

## KEARNS ON RULE A

*Philosophia* 43 (2015): 205 – 215

This is the penultimate version of the paper. Please do not cite this draft; rather, see <http://link.springer.com/article/10.1007%2Fs11406-014-9555-5> for the published version of the paper. You can cite *it* using the following DOI: **10.1007/s11406-014-9555-5**

Recall, from chapter V of Peter van Inwagen’s influential *An Essay on Free Will*, the following crucial inference rule which helps to undergird his so-called Direct Argument against the compatibility of moral responsibility and determinism.

Rule A:  $\Box p \vdash Np$

Given van Inwagen’s definition of the operator, ‘N’, the upshot of Rule A is that no human person is now, has been, or ever will be even partly directly morally responsible for a necessary truth. Moreover, most philosophers agree with van Inwagen that the validity of Rule A is “beyond dispute” (van Inwagen, 1983, p. 184); Stephen Kearns (2011), however, thinks that Rule A is *invalid*.

According to Kearns, Rule A is invalid because there are counterexamples that show that people are (or could be) partly directly morally responsible for a necessary truth. So, in what follows, I wish to examine three of Kearns’s attempts to provide a counterexample to Rule A. I’ll argue that these examples fail; and since these three alleged counterexamples exhaust the *types* of counterexamples Kearns wishes to give, if I’m right that these counterexamples fail, then it will follow that Kearns has not successfully shown that Rule A is invalid.

### II. Kearns’s First Alleged Counterexample to Rule A: Murder!

To begin to see how Kearns's objection to Rule A will go, consider the following scenario:

Murder!: Steven murders someone, and it is uncontroversial that he's responsible for his murdering since he does so knowingly and intentionally and could have done otherwise, etc. He is, therefore, responsible for the fact that he *actually* murders someone.

However, the fact that he actually murders someone is necessarily true. It is true in every possible world that, in the actual world, Stephen murders someone. (Kearns, 2011, p. 309)<sup>1</sup>

Kearns thinks that, given the set-up of Murder!, it's necessarily true that Stephen actually murders someone. Why think a thing like that? Well, Kearns thinks that *the actual world* is a rigid designator for a particular possible world (Ibid., p. 311). So, call the actual world  $\alpha$ . If *the actual world* rigidly designates  $\alpha$ , and Stephen murders in  $\alpha$ , then, necessarily, Stephen murders in the actual world. What Kearns thinks follows from this is that, necessarily, Stephen actually murders; and, of course, Stephen is morally responsible for actually murdering. And even though it's a necessary truth that Stephen actually murders, he is at least partly morally responsible for this fact because, as Kearns puts it, "in some sense, Stephen *makes it the case* that he actually murders someone" (Ibid., 309, my emphasis). The actual world is the way it is, in part, because of what Stephen does; and since this is a necessary truth, Stephen is morally responsible for a necessary truth. So, Rule A is invalid.

But I don't think that Murder! succeeds in showing Rule A to be invalid. To see why not, I think it prudent to remove some of the (what I think is) obfuscatory language that Kearns uses. For example, Kearns asks us to assume that Stephen murders in the actual world, where the

---

<sup>1</sup> For sake of expediency, I've made minor changes to the wording of Kearns's Murder! case. Nothing crucial rests on these changes.

locution ‘the actual world’ is a rigid designator. So, as above, I will continue to call ‘the actual world’  $\alpha$ , instead. Here is why I think calling *the actual world* ‘ $\alpha$ ’, will help get clear some, at least initially, confusing stuff. Suppose, along with Kearns that, in  $\alpha$ , Stephen murders. I think Kearns is right to conclude that if Stephen murders in  $\alpha$ , then it’s necessarily true *that Stephen murders in  $\alpha$* . But I wish to point out that Kearns thinks that it follows from its being necessarily true *that Stephen murders in  $\alpha$*  that he *actually* murders. This is confusing. For, does Kearns mean to argue that it’s necessarily true that *in fact* Stephen murders? Or does he mean simply to argue that it’s necessarily true that in *the actual world* (i.e.  $\alpha$ ) Stephen murders. I agree with the latter claim, but think that the former claim is obviously false. If it’s possible for Stephen to refrain from murdering (and Kearns grants that it is (Ibid., p. 310)), then it’s possible that Stephen refrains from murdering *in fact*. So, Kearns must mean to argue for the claim that it’s necessarily true that Stephen murders in  $\alpha$ .

Kearns, then, thinks that ‘actually’ denotes the *location* of Stephen’s murder. But, if that’s right, then I think that Kearns is mistaken. For, I think that ‘actually’ acts as an *intensifier* to ‘murders’ (i.e. ‘actually’ acts as a word that adds emotional context to ‘murder’, but nothing of propositional significance). What’s more, I think this reading of ‘actually’ is utterly plausible; in fact, I think it’s the correct reading. Moreover, it seems to me that that using ‘actually’ to denote the location of Stephen’s murder is a bit of a rhetorical trick to pull the reader’s intuition toward the conclusion that, if Stephen is morally responsible for actually murdering, and necessarily *that Stephen actually murders* is true, then Stephen is morally responsible for a necessary truth. To see better why it is that I say using ‘actually’ to denote a locator of Stephen’s murder is a rhetorical trick, allow me to restate the proposition *that Stephen actually murders* (a

proposition that is *ex hypothesi* a necessary truth) using my preferred name for *the actual world*, viz.,  $\alpha$ .

$S_\alpha$ : That Stephen murders in  $\alpha$

$S_\alpha$  is necessarily true.  $S_\alpha$  is necessarily true because, if Stephen murders in  $\alpha$ , then it's true in all possible worlds *that Stephen murders in  $\alpha$* . That is, it is true in all possible worlds *that Stephen murders in  $\alpha$* . So,  $S_\alpha$  is necessarily true.

Is Stephen morally responsible for  $S_\alpha$ ? I doubt it. To see why, suppose another possible world,  $\beta$ , had obtained instead of  $\alpha$ . And suppose that in  $\beta$ , Stephen refrains from murdering. Now, it's true in  $\beta$  *that Stephen murders in  $\alpha$*  (since that's necessarily true); but, is Stephen, in  $\beta$ , morally responsible for the fact *that Stephen murders in  $\alpha$* ?  $\alpha$  didn't obtain; and the world that obtained,  $\beta$ , is a world in which Stephen *doesn't* in fact murder. Given such a case, I think it's intuitive to conclude that Stephen is *not* morally responsible for  $S_\alpha$  since Stephen's murdering doesn't obtain.

Now, Kearns might respond that there it is still true *in some sense* that Stephen *makes it the case* that  $S_\alpha$ . And, he might continue, if this is true, then there's at least some sense in which Stephen is morally responsible for the fact *that Stephen murders in  $\alpha$* . Thus, Kearns might conclude that Stephen is at least partly morally responsible for the fact *that Stephen murders in  $\alpha$* .

But, in reply, I wish to cast doubt on the claim that Stephen makes it the case *that Stephen murders in  $\alpha$* . For, I think that in order for it to be true that Stephen makes it the case *that Stephen murders in  $\alpha$* , the fact *that Stephen murders in  $\alpha$*  has to depend, in the relevant way, on Stephen and what he does. So, let's consider the proposition *that Stephen murders in  $\alpha$* . By

hypothesis, this proposition is necessarily true. Now, it might help us figure out whether or not Stephen makes it the case *that Stephen murders in  $\alpha$*  by thinking about what this proposition is about. I say it's about Stephen and whether or not he murders in  $\alpha$ . Even so, I claim that Stephen is not responsible for the fact that the proposition *that Stephen murders in  $\alpha$*  is about him and whether or not he murders in  $\alpha$ . I say Stephen is not responsible for this fact because I think that the truth of *that Stephen murders in  $\alpha$*  does not depend on Stephen.

Following Merricks (2007), I think that some, but not all, truths require truthmakers. I claim that *that Stephen murders in  $\alpha$*  is one of those truths that does not require a truthmaker. Or, even if it does, its truthmaker is *not* Stephen. Why not? Because *that Stephen murders in  $\alpha$* 's truth is independent of whether or not Stephen is, in fact, actual (i.e. whether or not Stephen, in fact, exists), or whether or not Stephen, in fact, murders. And I think that in order for it to be the case that Stephen makes it the case *that Stephen murders in  $\alpha$* , it has to be the case that the truth of the proposition depends on Stephen (in the sense of 'depends on' that truth depends on the world<sup>2</sup>); that is, it has to be the case that Stephen *makes it true that Stephen murders in  $\alpha$* . But, even though the proposition is about Stephen and what he does, he does not make it true *that Stephen murders in  $\alpha$* ; for, this proposition is true regardless of what Stephen does, or whether or not he exists (and how can a non-existent thing make something the case?). So, I conclude that Stephen doesn't make it the case *that Stephen murders in  $\alpha$* . And since his making

---

<sup>2</sup> Very briefly, here's what I mean when I say that "the truth of the proposition depends on Stephen." I mean 'depends on' in a very trivial sense, the very trivial sense that truth (obviously) depends on the world and not the other way around. So, for example (and, again, borrowing from Merricks (2007)), it's true, *that dogs bark* because dogs bark. It's true *that there are no hobbits* because there are no hobbits. And so on. This is what I mean when I say that 'truth depends on the world'; I mean it in just this trivial way, setting aside questions as to whether or not there are more *substantive* things to say about how truth depends on (or is otherwise related to) the world. My claim is that, in order for Stephen to make it the case *that Stephen murders in  $\alpha$* , the truth of *that Stephen murders in  $\alpha$*  has to depend on Stephen in this trivial way.

it the case *that Stephen murders in  $\alpha$*  was supposed to show that he is at least partly directly morally responsible for a necessary truth, I conclude that nothing of the sort has been shown. Thus, I conclude Rule A is safe from Kearns's argument from Murder!.

But suppose someone were to object as follows:

Your argument seems to be assuming something like the following conditional:

If a proposition, P, will be true no matter what an agent, S, does, then S does not make it the case that P is true.

But this conditional is not obviously true. For, don't Frankfurt-style cases cast some doubt on this particular assumption? Given Black's presence, plans, and powers, Smith is going to be fatally shot no matter what Jones does. But in fact Jones, himself, freely chooses to shoot Smith, subsequently shoots Smith 'on his own' without Black's intervention, etc. By the above-assumed conditional, Jones doesn't make it the case that *that Smith is fatally shot* is true (since, again, that was going to be true no matter what). This is not obviously correct, though, given that the contribution that Jones's free choice (etc.) actually made to Smith's being fatally shot.<sup>3</sup>

In reply, I argue that Frankfurt cases, such as the one under consideration, do *not* show that the above conditional is false, and here is why. The truth of *that Smith is fatally shot* does not depend on, in the sense of 'depends on' that truth depends on the world, what Jones does. Rather, the truth of *that Smith is fatally shot* depends on Smith and whether or not he is fatally shot. But for Jones to make it the case *that Smith is fatally shot*, it has to be the case that this fact depends on (in the sense of 'depends on' that truth depends on the world) what he does. Since it

---

<sup>3</sup> Thanks to E. J. Coffman for the objection.

doesn't, Jones does not make it the case *that Smith is fatally shot*.<sup>4</sup> Moreover, we're now missing the alleged sufficient condition for Jones' being morally responsible for Smith's being fatally shot; thus, we haven't any reason to think that Jones is morally responsible for the fact *that Smith is fatally shot*. And so it goes with Murder!.

Before I finish with Murder!, however, I want to offer another reason to doubt that Stephen is morally responsible for S.. I think that Kearns, in his attempt to argue that Stephen makes it the case *that Stephen murders in  $\alpha$*  assumes something like the following principle:

ER: If S is directly morally responsible for p, and  $p \supset q$ , then S is directly morally responsible for q.

Here is why I think that Kearns's argument that Stephen makes it the case *that Stephen murders in  $\alpha$*  relies on something like ER. Consider another way Kearns might respond to my above argument that Stephen *does not* make it the case *that Stephen murders in  $\alpha$* :

By hypothesis, Stephen murders someone; that is, Stephen *in fact* murders someone. So, Stephen makes it the case that he murders someone. But this fact implies *that Stephen murders someone in  $\alpha$* . So, it follows that Stephen makes it the case *that Stephen murders someone in  $\alpha$* . And since all should agree that Stephen is directly morally responsible for the fact that he, in fact, murders someone, and this fact implies *that Stephen murders someone in  $\alpha$* , we should all agree that he is directly morally responsible for the fact *that Stephen murders someone in  $\alpha$* .

But such a response relies on two invalid inferences. One is ER, from above, and the other is something like it that I'll call:

---

<sup>4</sup> If Jones is morally responsible for anything it's the truth of *that Jones fatally shoots Smith* since this fact *does* depend on Jones and what he does.

MC: If S is morally responsible for making it the case that p, and  $p \supset q$ , then S is morally responsible for making it the case that q.

But, while ER is superficially plausible, it is false by counterexample. And so is MC, for the exact same reasons. For, suppose that Soren Kierkegaard is morally responsible for his being (or his making it the case that he is) a melancholy Dane. Kierkegaard's being a melancholy Dane implies that there is such a place as Denmark. But surely Kierkegaard is not morally responsible for (and neither does he make it the case that) there being such a place as Denmark, even if he is (we may assume) morally responsible for being (or making it the case that he is) a melancholy Dane.

So, if Kearns is relying on ER (or MC), then he is erroneously led to believe that Stephen is responsible for (or makes it the case that he is) actually murdering someone because he thinks that Stephen's actually murdering someone is a consequence of his murdering someone (something I am happy to agree that Stephen makes the case). Thus, he might conclude via something like ER, that Stephen is at least partly directly morally responsible for actually murdering. But, since ER (and MC) is false, the envisaged reasoning is unsound. I conclude, then, that Kearns's argument from Murder! fails to show that Rule A is invalid.

### III. Kearns's Second Alleged Counterexample to Rule A: Hey Jude

In this section, I wish to turn our attention to a second counterexample that Kearns gives, a counterexample that is of a different *type* than the first. Consider:

Hey Jude [HJ]: Paul McCartney composes the melody of Hey Jude and does so knowingly and intentionally, etc. He is morally responsible for various facts concerning Hey Jude like the fact that Hey Jude starts the way it does and for the fact that Hey Jude has a beautiful melody since he *created* Hey Jude. However, these facts that McCartney

is responsible for are necessary truths. It is necessarily true that Hey Jude starts in the distinct way it does and that it has a beautiful melody.

I say that HJ is a different *type* of counterexample to Rule A because it is designed to undermine Rule A in a way that doesn't rely on a person's being responsible for a world-indexed truth.

Murder!, you'll recall, relied on Stephen's being morally responsible for a world-indexed truth, and since world-indexed truths are necessary truths, it was supposed to follow that Stephen is morally responsible for a necessary truth. But, as Kearns rightly points out, an incompatibilist *could* (though, given my argument from the last section, I don't think that an incompatibilist *should*) concede that Murder! is a counterexample to Rule A, but that the Direct Argument needn't rely on Rule A. All the Direct Argument need rely on, the imagined incompatibilist might claim, is the following:

Rule A\*: Necessarily, for any proposition, p, if it is necessary that p (and this necessary truth is not world-indexed), then no one is even in part directly morally responsible for the fact that p. (Ibid., p. 311)

The argument *for* Rule A\*'s truth, then, can go as follows:

P1. An agent is in part directly morally responsible for a fact only if the agent makes this fact obtain.

P2. No agent can make a non-world-indexed necessary fact obtain.

C1. Therefore, no one is in part directly morally responsible for any non-world-indexed necessary fact (i.e. Rule A\* is true). (Ibid., p. 312)

It is this argument to which HJ is supposed to provide a counterexample. So, since Murder! would not have provided a counterexample to P1 – C1, HJ is a different *type* of counterexample to Rule A because it rules out Rule A\*, as well.<sup>5</sup>

Now, why think that HJ provides a counterexample to Rule A\*? Kearns thinks HJ shows that A\* is false because McCartney is morally responsible for some necessary truths. McCartney is responsible for, among other things, the melody and beauty of “Hey Jude”. And, since “Hey Jude” has these properties necessarily—that is, “Hey Jude” couldn’t have been any other way—McCartney is morally responsible for some necessary truths since he’s the one who created “Hey Jude”. So, Rule A\* (and Rule A) is invalid.

But I think that HJ fails to show that Rule A\* (and Rule A) is invalid. To see why, note, first, that Kearns thinks that “it is a necessary truth that Hey Jude starts a certain way, and is beautiful”, and that “if Hey Jude had had a different melody, it wouldn’t have been Hey Jude” (Ibid., p. 313). This is just what HJ says: “Hey Jude” *just is* a certain way, and McCartney is morally responsible for “Hey Jude” because he created it.

But this seems wrong to me. Why should we suppose that, if it’s true that “Hey Jude” is necessarily a certain way, that McCartney created it? It seems to me that supposing “Hey Jude” *just is* a certain way is to suppose that “Hey Jude” has some essential properties. But, if that’s right, then “Hey Jude” just seems to be an instantiation of something like *Hey Judeness* or *Hey Judaity*. I think it’s dubious whether or not there is such a thing as Hey Judaity, but I think that Kearns’s contention that “Hey Jude” *just is* a certain way (that is, “Hey Jude” is necessarily the

---

<sup>5</sup> To be clear, HJ (and the alleged counterexample that I consider in the next section) serves to perform two functions. The first is to show that P2 of the argument for Rule A\* is unsound. The second is to show that a person can be at least partly directly morally responsible for a necessary truth—and, insofar as this is the case, Rule A\* (and Rule A) is false.

way it is, whether or not Paul McCartney had ever existed, or if he *had* existed but never actualized it) commits him to such an ontology.

So, suppose that this is what Kearns thinks, and that his thinking this way about “Hey Jude” is right. Again, I think this means that “Hey Jude” is an *instantiation* of *Hey Judavity*. And, if that’s right, then McCartney *did not* create “Hey Jude”; rather, he *actualized* it, i.e., he created an instance of the property *Hey Judeness*.<sup>6</sup> So, what McCartney is morally responsible for is not his *creating* “Hey Jude”—*Hey Judavity* seems to exist necessarily as an abstract object—he’s morally responsible for *actualizing* it. And since he needn’t have actualized it, then he’s not morally responsible for a necessary truth. He’s not morally responsible for the fact that, necessarily, “Hey Jude” starts a certain way or has a certain melody. He’s morally responsible for the fact that “Hey Jude” is actualized; and that’s not a necessary truth. So, I conclude that HJ fails to show that Rule A\* (and Rule A) is false.

To see another reason to think that HJ fails to show that Rule A\* (and Rule A) is false, consider the following argument, one that is generated straight from HJ.

- (1) Paul McCartney is at least partly directly morally responsible for the fact that “Hey Jude” starts the way it does.
- (2) Necessarily, “Hey Jude” starts the way it does.
- (3) Therefore, Paul McCartney is at least partly directly morally responsible for a necessary truth (i.e. Rule A\* (and Rule A) is false)

---

<sup>6</sup> Why is it better to say that McCartney *actualizes* “Hey Jude” rather than *create* it? Assuming, as we are, that Kearns is committed to the existence of some abstract object, *Hey Judeness* (or *Hey Judavity*), Kearns accepts an ontology of songs such that songs exist necessarily as abstract objects. Thus, if McCartney want’s one of these songs (“Hey Jude”) to be actual, he can’t *create* it, strictly speaking; it already exists. What he *can* do is make it actual; that is, he can *actualize* it. So, Paul McCartney can’t, strictly speaking, *create* “Hey Jude”; what he can do is *actualize* it. For more on this, see Alvin Plantinga’s (1974).

Given (1) – (3), I think it's pretty clear that HJ fails as a counterexample to Rule A\* (and to Rule A); for, I think that premise (2) is clearly false. To see that this is so, imagine that you and I have just come from seeing Paul McCartney play a live music concert. And suppose, further, that I am no fan of McCartney's, I don't know anything about his music; but you are a fan, and you do know something about his music. Now, suppose that I ask you something like the following:

Hey, what was that third song McCartney played? The one that said something about taking a sad song and making it better. I liked that one; what was it called?

I think, obviously, your answer to my question should be that the song in question is called "Hey Jude". And I think this is true *even if* Paul McCartney (unbeknownst to me) adlibbed the beginning, or played the verse a step up (that is, in a higher key) from the way he normally plays it.

For, suppose that Paul McCartney did these things, that he adlibbed a bit. Would it be plausible for you to respond to my question in something like the following way?

Well, it's difficult to say what song McCartney played third. I know to which song you refer, and it *sounds* an awful lot like this other song of his, "Hey Jude", but it wasn't *that* song; for, he played the beginning a different way (or he played the verse a step up from normal). So, I don't know what to tell you, but it wasn't "Hey Jude".

If you answered me in this way, and I went and listened to the recorded version of "Hey Jude", I think I'd have every right to think you'd not been straightforward with me, or that you'd confused my question for some other. For, I think that "Hey Jude" does *not* necessarily start a certain way. There are many things about "Hey Jude" that can change, it seems to me, without it being the case that "Hey Jude" isn't being played. Now, this of course leads to sorities problems

with respect to the identity of things like songs, etc.; but, that's another issue entirely. The point is that, unless Kearns is willing to accept a dubious metaphysical claim like *that there is such a thing as Hey Judaity*, then I think he should concede that "Hey Jude" isn't *necessarily* a certain way, or, anyway, that "Hey Jude" doesn't necessarily start in a particular way. And if not, then, given that the intro to "Hey Jude" can change without the song's numerical identity being changed, premise (2) is false. So, I conclude that HJ fails as a counterexample to Rule A\* (and Rule A).

But, perhaps Kearns *isn't* committed to the existence of *Hey Judaity*, and that Paul McCartney really did *create* "Hey Jude". If so, then I think Kearns is guilty of basing his HJ argument on a premise like ER (or MC). Kearns erroneously believes that Paul McCartney is morally responsible for the fact that "Hey Jude" has the melody that it does because he reasons that "Hey Jude" has the melody that it does is a consequence of Paul McCartney's having composed it (something that I am happy to agree that McCartney makes the case). Thus, by ER (or MC), Paul McCartney is morally responsible for the fact that "Hey Jude" has the melody that it does. But we saw that ER (and MC) is false. So, I conclude that Kearns's argument from HJ fails to show that Rule A\* (and Rule A) is false.

#### IV. Kearns's Third Alleged Counterexample to Rule A: Torturing Babies

In this section, I'd like to consider another type of alleged counterexample to Rule A\* (and Rule A), namely:

Torturing Babies [TB]: Karen decides to make eating cake immoral by making it the case that when someone eats cake, thousands of people are killed. Moreover, Karen decides to make torturing babies for fun immoral. She does so by making it the case that thousands of people are killed when someone tortures a baby for fun. (Ibid., p. 314 – 315)

TB is a different type of counterexample to Rule A\* (and Rule A) because it purports to show that a person can be partly directly morally responsible for a necessary truth, even when that necessary truth's being true is *overdetermined*. To be clear, Kearns thinks TB shows that Rule A\* (and Rule A) is false because, though torturing babies is necessarily immoral—i.e. it's necessarily true that torturing babies is wrong—Karen is at least *partly* directly morally responsible for the fact that torturing babies is wrong because she, at least in part, *makes it the case that* torturing babies is wrong. Perhaps she overdetermines its being wrong to torture babies, but torturing babies is still wrong at least partly because of what Karen does (viz., insuring that thousands of people will die if someone tortures a baby). So, Rule A\* (and Rule A) is invalid because Karen is partly morally responsible for a necessary truth.

But as with the foregoing alleged counterexamples, I think that TB fails to show that Rule A\* (and Rule A) is invalid, as well. And I think the fact that Karen's actions overdetermine the wrongness of torturing babies helps make it clear why TB fails as a counterexample. Now, before I say why I think that the fact that Karen's actions overdetermine the wrongness of torturing babies helps make it clear why TB fails as a counterexample, it needs to be noted that Kearns thinks such overdetermination is *not* a problem for TB. He thinks overdetermination isn't a problem because he thinks Fischer and Ravizza's (1998) Erosion\* example shows that overdetermination doesn't rule out moral responsibility. Recall, Fischer and Ravizza's Erosion and Erosion\* cases:

Erosion: Imagine that Betty [a soldier charged with destroying an enemy fortress] plants her explosives in the crevices of the glacier and detonates the charge at T1, causing an avalanche that crushes the enemy fortress at T3. Unbeknownst to Betty and her commanding officers, however, the glacier is gradually melting, shifting, and eroding.

Had Betty not placed the dynamite in the crevices, some ice and rocks would have broken free at T2, starting a natural avalanche that would have crushed the enemy camp at T3.

(Fischer and Ravizza, 1998, p. 157)

Erosion\*: [The details are the same as Erosion except that, in this case,] the conditions of the glacier *do* actually cause the ice and rocks to break free, triggering an avalanche that arrives at the fortress precisely at the same time as the independent avalanche triggered freely by Betty. Each avalanche is sufficient for the destruction of the enemy fortress.

(Ibid., p. 160 – 161)

The conclusion of Erosion\*, with which Kearns agrees, is that Betty is morally responsible for the fact that the enemy base is destroyed, and this is true even though that fact is overdetermined. Thus, Kearns thinks overdetermination doesn't rule out moral responsibility.

But, this conclusion isn't at all clear to me. The destruction of the enemy fortress is overdetermined (that much is not up for debate); so, its being true *that the enemy fortress is destroyed* doesn't depend on Betty; for, the fact *that the enemy base is destroyed* doesn't depend on Betty or what she does (for, given the case, it would be true *that the enemy base is destroyed* no matter what Betty did). And I think in order for Betty to be morally responsible for some fact or other, it needs to be the case that that fact depends on her in the relevant way. Thus, it seems to me that Betty *isn't* morally responsible for the fact *that the enemy base is destroyed*. And the same will be true for *anyone* with respect to an overdetermined fact.

Moreover, I think that what Betty plausibly *is* responsible for is the fact *that Betty causes an avalanche that destroys the enemy base*. And the same can be said, *mutatis mutandis*, for Karen and her overdetermining the immorality of torturing babies. She's not *at all* responsible for its being immoral to torture babies; for its being immoral to torture babies (which is *ex*

*hypothesi* necessarily true) doesn't depend in the relevant way on Karen and what she does. That is, the fact *that it's immoral to torture babies* doesn't depend on Karen because it would be true no matter what she did. Thus, it seems to me that she is not at all directly morally responsible for the fact *that torturing babies is immoral*. What she plausibly *is* responsible for is her *attempt*, or, perhaps, her *intention*, to make torturing babies immoral. But this latter bit is not a necessary truth, while the former is. And since Karen is not at all morally responsible for the former—but *is* the latter—she is not at all morally responsible for a necessary truth. I conclude that TB fails to show that Rule A\* (and Rule A) is false.

But I don't want to move too hastily here, so here I'll provide the same argument I ran against Kearns's argument from Murder!, and his argument from HJ. In his argument from TB, Kearns seems to me erroneously to believe that Karen makes it the case that it's wrong to torture babies because Kearns reasons that its being wrong to torture babies is a consequence of Karen's insuring that thousands of people will die if a baby is tortured. Thus, by MC, Karen makes it the case that it's wrong to torture babies; and if she makes the case that it's wrong to torture babies, then it's plausible that she's at least partly directly morally responsible for having done so. But as we've seen, MC is false. So, I conclude that Kearns's argument from TB fails to show that Rule A\* (and Rule A) is invalid.

## V. Conclusions

I have considered, and rejected, three alleged counterexamples to Rule A. More precisely, I have considered one alleged counterexample to Rule A *in particular*, namely, Murder!, and two alleged counterexamples to Rule A\*, a revised version of Rule A. I argued that all of these counterexamples fail, in the first place, because each case fails to show that the truth of the fact for which the agent is allegedly morally responsible *depends on* the agent in the

relevant way. After all, necessary truths are true regardless of whether or not the agent in question exists. And if that's true, then it cannot be the case that the truth of the fact in question depends on the agent. Thus, I conclude that none of the agents in question are even in part directly morally responsible for the truth of the facts under consideration.

Moreover, I argued that it seems as if Kearns's argument, in each case, relies on one or both of the fallacious inference rules, ER and MC. I argued that in each alleged counterexample, Kearns attempts to show that someone is (or could be) directly morally responsible for a necessary truth because the relevant agent in each case *makes it the case* that some necessary truth or other is true by virtue of the fact that the agent makes it the case that some fact that *implies* the necessary fact is true. But, this reasoning doesn't follow. For it's perfectly possible, as my Soren Kierkegaard example shows, to be morally responsible for making it the case that p, while p implies q, without, thereby, being morally responsible for making it the case that q. And if I'm right that Kearns's alleged counterexamples rely on this reasoning, this is all the more reason to conclude that his counterexamples to Rule A fail. Thus, I conclude that Kearns has failed to provide a successful counterexample Rule A.<sup>7</sup>

#### References

Fischer, John Martin and Mark Ravizza 1998, *Responsibility and Control: A Theory of Moral Responsibility*, Cambridge: Cambridge University Press.

Kearns, Stephen 2011, "Responsibility for Necessities," *Philosophical Studies* 155: 307 – 324.

Merricks, Trenton 2007, *Truth and Ontology*, New York: Oxford University Press.

Plantinga, Alvin 1974, *The Nature of Necessity*, Oxford: Oxford University Press.

van Inwagen, Peter 1983, *An Essay on Free Will*, Oxford: Oxford University Press.

---

<sup>7</sup> I wish to thank E. J. Coffman, John Martin Fischer, Markus Kohl, and David Palmer for helpful and instructive comments on earlier drafts of this paper.