

CAN I KILL MY YOUNGER SELF? TIME TRAVEL AND THE RETROSUICIDE PARADOX*

Peter B. M. Vranas
vranas@wisc.edu
University of Wisconsin-Madison
25 October 2009

Abstract. If (backward) time travel is possible, presumably so is my shooting my younger self (YS); then apparently I can kill him—I can commit *retrosuicide*. But if I were to kill him I would not exist to shoot him, so how can I kill him? The standard solution to this paradox understands ability as *compossibility* with the relevant facts and points to an equivocation about which facts are relevant: my killing YS is compossible with his proximity but not with his survival, so I can kill him if facts like his survival are irrelevant but I cannot if they are relevant. I identify a lacuna in this solution, namely its reliance without argument on the hidden assumption that my killing YS is *possible*: if it is impossible, it is not compossible with anything. I argue that this lacuna is important, and I sketch a different solution to the paradox.

1. Introduction

If time travel is possible, presumably the following scenario is also possible: I travel back in time, I meet my younger self, and I point a loaded gun at him. It seems then that I can kill him; to put it more dramatically, I can commit “retrosuicide”. But it also seems that I cannot kill him: if I were to kill him, he would not grow up to become me, and I would not exist to travel back in time. This is a sketch of what I call the “retrosuicide paradox”; in §2 I provide a more detailed and rigorous formulation.

What I call the “standard solution” to the paradox is inspired by David Lewis (1976).¹ The solution understands ability as *compossibility* with the relevant facts and explains the paradox away as due to an equivocation about which facts are relevant. As an analogy, my speaking Finnish is compossible with the anatomy of my larynx but not with my lack of training, so in one sense I can speak Finnish but in another I cannot. Similarly, my killing my younger self is compossible with his proximity but not with his survival, so in one sense I can kill him but in another I cannot. It is reasonable to use ‘can’ in either sense, but one must choose; what one must not do is waver, say in the same breath both that I can and that I cannot (Lewis 1976: 150-1).

* I am very grateful to Allan Gibbard, John Hawthorne, Christopher Hitchcock, Aviv Hoffmann, James Joyce, Gerhard Nuffer, Peter Railton, Ted Sider, Nicholas J. J. Smith, David Velleman, and some anonymous reviewers for extensive discussion or comments. I am also grateful to Martin Barrett, Karen Bennett, Bob Brier, Phil Dowe, Molly Gardner, Thomas Hofweber, Eric Lormand, George Mavrodes, D. H. Mellor, Samuel Ruhmkorff, Lawrence Sklar, and J. J. C. Smart for helpful comments. Thanks finally to David Brink, Nancy Cartwright, John Earman, Antony Flew, Alan Hájek, Richard Hanley, Carl Hoefer, Louis Loeb, Alfred Mele, Graham Nerlich, Steven Quartz, Steven Savitt, Stephen Schulz, Paul Teller, Michael Tooley, and James Woodward for interesting questions or bibliographical help, and to my mother for typing part of the paper. Material from this paper was presented at the 96th Eastern APA meeting (December 1999), the California Institute of Technology (May 2000), the University of Michigan (September 2000), the London School of Economics and Political Science (December 2000), and the 96th Pacific APA meeting (March 2002).

¹ For (at least tentative or partial) endorsements of Lewis’s solution see: Brown 1992: 435-6; Dowe 2000: 448-51; Hanley 1997: 209-10; Kiourti 2008; MacBeath 1982: 411, 1983: 34; Richmond 2001: 310, 314, 317; Riggs 1997: 50-1; Sider 1997: 143, 2002: 116-7. For disagreements with Lewis’s solution see: Grey 1999: 66-7; Smith 1985: 60. Note that strictly speaking Lewis addresses the “grandfather paradox” rather than the retrosuicide paradox (see note 17); this is why I said that the standard solution is “inspired by” (rather than “due to”) Lewis.

In this paper I identify a lacuna in the standard solution, namely its reliance without argument on the *hidden assumption* that retrosuicide is *possible*: if it is impossible, it is not compossible with anything, and the standard solution fails because it supplies no sense in which I *can* commit retrosuicide (§2). I examine three possible responses on behalf of the standard solution, and I argue that they fail (§3). I conclude by sketching a different solution to the paradox (§4).

2. The retrosuicide paradox, the standard solution, and the hidden assumption

To formulate the retrosuicide paradox rigorously, call (*rigidly*) ‘my younger self’ an earlier temporal stage of mine, and consider an atom-for-atom duplicate of my younger self: the duplicate and my younger self have indistinguishable bodies but are temporal stages of different persons, whereas I and my younger self are different temporal stages of the same person.² Take any situation in which I *can* kill the duplicate: I have the ability (and the opportunity) to kill him.³ For example: he is asleep, I am pointing a loaded gun at him, there is no one around to help him, and so on.⁴ Suppose that I *don’t* kill the duplicate; for example, I shoot but I miss him. Now modify the situation by replacing the duplicate with my younger self: in the new situation *my younger self* is asleep, I am pointing a loaded gun at him, and so on.⁵ Say that the new situation is *retrosuicide-propitious*; more generally, say that a situation in which I am is retrosuicide-propitious exactly if it is identical to a situation in which I can (but I don’t) kill a duplicate of my younger self, except that (in the situation in which I am) my younger self is in the place of the duplicate. Given this terminology, to a first approximation the retrosuicide paradox can be formulated as the following argument:

- (P1) If time travel is possible, then it is possible for me to be in a retrosuicide-propitious situation.
- (P2) Necessarily, if I am in a retrosuicide-propitious situation, then in that situation I can kill my younger self.
- (P3) Necessarily, if I am in a retrosuicide-propitious situation, then in that situation I cannot kill my younger self.⁶

Thus: (C) Time travel is impossible.

(Throughout this paper I am talking about *backward* time travel and about *physical* necessity, possibility, impossibility, and compossibility.⁷)

² I understand person-stages as Sider (2001: 101) does (see also Lewis 1983a: 76-7), except that I take them to be extended rather than instantaneous (so I see no problem with my talk of “killing” a stage). The temporal boundaries between successive stages may be vague and arbitrary, but I see this as creating no problem for my reasoning. Three-dimensionalists may be unhappy with my talk of stages, but Effingham and Robson (2007) and Sider (2001: 101-9) have argued that three-dimensionalism is incompatible with time travel (on Sider’s argument see: Markosian 2004: 670-3; Sider 2004: 684-6; Simon 2005; Steen 2007: 190-202).

³ If you are a novice and I am an expert in chess but you refuse to play with me, then in one sense of ‘can’ I can defeat you but in another I cannot (cf. Austin 1956/1979: 229-30; Cross 1986: 54-7; Hasker 1989: 134-5); I am interested in the latter sense, so I require opportunity in addition to ability (cf. Kenny 1976: 219; Mann 1991: 620; Vihvelin 1996: 318; contrast Goldman 1970: 199). For the sake of simplicity, in general I omit mention of opportunity from now on.

⁴ I am *assuming* that in the above situation I can kill the duplicate; I am not saying that the features I list in the example guarantee I can kill him. The assumption that *in the above situation* I can kill the duplicate is compatible with a contextual understanding of ‘can’ (see note 9).

⁵ I don’t need to assume that it is possible for me to be in such a new situation; as it will be seen in a moment, the retrosuicide paradox assumes instead that this is possible *if* time travel is.

⁶ I understand “in that situation I cannot kill my younger self” as “it is not the case that in that situation I can kill my younger self”.

The argument is valid: from P2 and P3 it follows that it is impossible for me to be in a retrosuicide-propitious situation, and from this and P1 the conclusion follows that time travel is impossible. This conclusion is implausible because time travel is compatible with General Relativity (see, e.g.: Earman 1995a, 1995b: chap. 6; Gott 2001; Nahin 1999; Thorne 1994: chap. 14). On the other hand, the three premises of the argument are plausible. P1 is plausible because, if time travel is possible, presumably it is also possible that I travel back in time, I meet my younger self, I point a loaded gun at him, and so on. P2 is plausible because a retrosuicide-propitious situation and a situation in which I can kill a duplicate of my younger self are by definition indistinguishable in terms of local physical features, and it is plausible that ability *supervenes* on such features (see §3.3). P3 is plausible because it seems necessary that in *any* situation (retrocuicide-propitious or not) I cannot kill my younger self: it seems impossible that I kill him, since he grew up to become me (see §3).⁸ Of course the three premises are subject to objections, but the above preliminary considerations suggest that the above argument is indeed a *paradox*: an apparently sound argument with an apparently unacceptable conclusion (cf. Sainsbury 1995: 1).

The standard solution to the paradox understands ability in a given situation as compossibility with the relevant features of the situation and claims that “‘can’ is equivocal”:

To say that something can happen means that its happening is compossible with certain facts. Which facts? That is determined, but sometimes not determined well enough, by context. ... What I can do, relative to one set of facts, I cannot do, relative to another, more inclusive, set. Whenever the context leaves it open which facts are to count as relevant, it is possible to equivocate about whether I can [do something] (Lewis 1976: 150).⁹

Given this equivocation, according to the standard solution P2 and P3 amount to the following claims:

- (P2') Necessarily, if I am in a retrosuicide-propitious situation, then in that situation I can kill my younger self, in the sense that the relevant features of the situation—namely (for example) that he is asleep etc.—are compossible with my killing him.¹⁰
- (P3') Necessarily, if I am in a retrosuicide-propitious situation, then in that situation I cannot kill my younger self, in the sense that the relevant features of the situation—namely that *I don't kill him*, that *he is an earlier stage of mine*, and (for example) that he is asleep etc.—are not compossible with my killing him.¹¹

⁷ I am talking about physical rather than metaphysical impossibility because I will use the claim that resurrection is impossible (and I take resurrection to be physically but not metaphysically impossible).

⁸ Alternatively, one might argue that it is impossible for me to kill my younger self because it is impossible for me to change the past. I examine the corresponding variant of the retrosuicide paradox in another paper (Vranas 2009).

⁹ Agreeing with Kratzer (1977: 337-43), in a later piece Lewis regards ‘can’ as univocal but contextual, “unambiguous but relative” (1979/1983b: 246). For the sake of simplicity I will continue to speak of an alleged equivocation, but for my purposes only the alleged relativity matters.

¹⁰ To circumvent the objection that there is no such *fact* as my killing my younger self, take compossibility to be a relation between *propositions*; here, between the proposition that I kill my younger self and some other proposition.

¹¹ Are proponents of the standard solution *committed* to the claim that P2 amounts to P2' (and P3 amounts to P3')? They are if they accept the above quotation from Lewis: accepting that “[t]o say that something can happen means that its happening is compossible with certain facts” (Lewis 1976: 150) commits one to accepting that “I can kill my younger self” amounts to “my killing my younger self is compossible with certain facts”. One might object, however, that precisely because (as I go on to argue in the text) accepting that P2 amounts to P2' (and P3 amounts to P3') leads into trouble, a Lewis-style solution to the retrosuicide paradox should be charitably understood as denying that P2 amounts to P2'. I elaborate this objection and address it in detail in §3.

From P2' and P3' it does *not* follow that it is impossible for me to be in a retrosuicide-propitious situation, and the paradox is resolved—or so the standard solution goes. But what if it is necessarily *impossible* for me to kill my younger self? Then my killing him is necessarily *not compossible* with anything,¹² and via P2' it does follow that it is impossible for me to be in a retrosuicide-propitious situation; the paradox is restored, and the standard solution fails. So the standard solution works only if it is assumed that it is *not* necessarily impossible for me to kill my younger self—call this the *hidden assumption*. Equivalently, the hidden assumption is that my killing my younger self is *possible* (strictly speaking, *possibly* possible, but possible possibility is equivalent to possibility if the accessibility relation between possible worlds is reflexive and transitive¹³). In sum, the standard solution relies without argument on the hidden assumption that it is possible for me to kill my younger self (i.e., to commit retrosuicide).

3. The importance of the hidden assumption

How damaging to the standard solution is its reliance without argument on the hidden assumption? Proponents of the solution might see no damage: they might contend that the assumption (1) is obviously true and thus needs no argument, (2) can be dispensed with by charitably modifying the solution, or (3) must be true if a paradox is to arise at all. In the next three subsections I argue that all three contentions fail; it turns out that by relying without argument on the hidden assumption the standard solution misses the thrust of the paradox.

3.1. Is the hidden assumption obvious?

Consider this helpless boy, lying asleep within the range of my loaded gun. *Of course* it is possible for me to kill this boy; what is impossible is the *conjunction* of the propositions that I kill this boy and that this boy is an earlier stage of mine—or so a proponent of the standard solution might claim. In reply I grant that one can refer to my younger self by pointing to him and uttering “this boy” (and thus without mentioning that he is an earlier stage of mine), but I deny that it is obviously possible for me to kill this boy. It is possible for me to kill this boy only if there is a possible world in which I *do* kill this boy. But in any such world this boy *is* my younger self (given that ‘this boy’ and ‘my younger self’ are rigid designators), so what happens after I kill him—does he rise from the dead and grow up to become me? If resurrection is impossible, then apparently it is also impossible for me to kill my younger self.¹⁴ And resurrection, being a miracle, does seem to be impossible (recall that I am talking about *physical* impossibility).¹⁵ In §4 we

¹² One might respond that, even if my killing him is not *physically* compossible with anything, it may be *logically* or *metaphysically* compossible with the relevant features of the situation. I reply that the latter kinds of compossibility are not enough for ability: my killing someone is logically and metaphysically compossible with his being surrounded by bulletproof glass, but normally in such a situation I cannot kill him. (One could similarly argue that even physical compossibility is not enough for ability, but I don’t want to emphasize this problem because in §3.2 I examine a modification of the standard solution which does not rely on compossibility.)

¹³ Cf. Loux 1979: 22-4. Some people might deny that the accessibility relation is transitive (cf. Salmon 1989: 8-9), and thus might insist that the standard solution relies only on the *strict* hidden assumption that it is *possibly* possible for me to kill my younger self. For the sake of simplicity, in what follows I do not talk about the strict hidden assumption, but *mutatis mutandis* my reasoning would go through if I did.

¹⁴ Those who deny transworld identity but talk instead about counterparts can still accept this point (and my later points) by taking ‘this boy’ and ‘my younger self’ to be what Lewis calls ‘quasi-rigid’ designators (1986c: 6, 256).

¹⁵ It seems natural to understand the (arguably analytic) claim that death is permanent and irreversible as the claim that it is physically impossible to live again after one’s death, to be resurrected (cf. Feldman 1992: 64; contrast 66-71). But isn’t it physically—even if at present not technologically—possible to revive the brain dead? (Similarly, isn’t a quantum event possible in which someone who is brain dead spontaneously becomes fully functional again?)

will see that the story is more complicated (maybe in some possible world my younger self—and thus this boy—is not an earlier stage of mine); my present point is only that one cannot just take it for granted that it is possible for me to kill my younger self, so the current contention on behalf of the standard solution (namely that the hidden assumption is obvious) fails.

My reply to the contention that the hidden assumption is obvious hinged on the *rigidity* of the designator ‘this boy’; what if I refer to my younger self by using instead the *nonrigid* designator ‘the boy in front of me’? (Cf. Smith 1997: 379.) Then I grant that it is obviously possible for me to kill the boy in front of me: in some possible world the boy in front of me is (say) a duplicate of my younger self rather than my younger self, I kill the duplicate, and he does not rise from the dead. But from the fact that it is possible for me to kill *the boy in front of me* it does not follow that it is possible for me to kill *my younger self*, so it does not follow that the hidden assumption is true. A proponent of the standard solution might argue that (charitably understood) the solution relies not on the hidden assumption, but rather on the trivial assumption that it is possible for me to kill the boy in front of me, because P2’ (and P3’) should be modified by substituting ‘the boy in front of me’ for ‘my younger self’.¹⁶ In reply I grant that in response to a *variant* of the retrosuicide paradox (which corresponds to the *modified* P2’ and P3’) the standard solution does not rely on the hidden assumption (but relies instead on the trivial assumption that it is possible for me to kill the boy in front of me). But the fact remains that, *in response to the retrosuicide paradox as I formulated it*, the standard solution does rely on the hidden assumption. So it would be a Pyrrhic victory for a proponent of the standard solution to conclude that the solution works against the modified paradox but fails against the paradox I formulated. I am not saying that a single approach should resolve every paradox: it turns out, maybe surprisingly, that various variants of the retrosuicide paradox are subtly different. I am rather saying that, if it is conceded that the standard solution fails against the retrosuicide paradox as I formulated it, then it is conceded that this paradox requires some other solution.¹⁷

3.2. Is the hidden assumption dispensable?

One might try to get rid of the hidden assumption by charitably modifying the standard solution. The claim that ability in a given situation amounts to compossibility with the relevant features of the situation *grounds* in the standard solution the claim that ability can vary with the relevant

If it is, then brain death is not permanent and irreversible, so the *criterion* of brain death does not capture the *concept* of death (cf. Fischer 1993: 4).

¹⁶ The point would be that it is uncharitable to formulate the standard solution in terms of P2’ (and P3’): the use of the rigid designator ‘my younger self’ in P2’ smuggles in the facts that the boy in front of me is an earlier stage of mine and that I don’t kill him, facts which are irrelevant to the sense in which I can kill him. In §4, however, I suggest that my younger self is not *necessarily* an earlier stage of mine; if so, then the use of ‘my younger self’ in P2’ need not smuggle in any irrelevant facts.

¹⁷ Similar remarks apply to the *grandfather paradox*, which one gets by modifying the retrosuicide paradox so as to talk (not about my younger self, but rather) about my *grandfather*, namely a temporal stage of my paternal grandfather at a time “earlier” (see MacBeath 1982: 411) than my father’s conception. Even if resurrection is impossible and my father is born in 1933, it might seem obvious that it is possible for me to kill my grandfather in 1930 (suppose for the sake of argument that there is a possible world in which my father is born in 1929); so *in response to the grandfather paradox* the standard solution relies on an arguably trivial assumption rather than on the hidden assumption. But although Lewis (1976) proposed (a variant of) the standard solution in response to the grandfather paradox rather than the retrosuicide paradox and is thus in a sense immune to the criticism that his solution relies without argument on the hidden assumption, my point remains that, in response to the retrosuicide paradox as I formulated it, the standard solution does rely on the hidden assumption. This point does not hinge on Lewis’ exegesis: I am not arguing that Lewis relies on the hidden assumption, and I can grant that what Lewis says about the *grandfather paradox* is fine.

features, but one can reject the former claim and still accept the latter. More specifically, one can drop all talk of compossibility but claim that P2 and P3 amount (not to P2' and P3', but rather) to the following claims:

- (P2*) Necessarily, if I am in a retrosuicide-propitious situation, then relative to the relevant features of the situation—namely (for example) that my younger self is asleep etc.—I can kill my younger self.
- (P3*) Necessarily, if I am in a retrosuicide-propitious situation, then relative to the relevant features of the situation—namely that *I don't kill my younger self*, that *he is an earlier stage of mine*, and (for example) that he is asleep etc.—I cannot kill my younger self.

Even if the hidden assumption is false, so that my killing my younger self is necessarily impossible and thus is necessarily not compossible with anything, P2* (unlike P2') involves no compossibility claim, so from P2* it does not follow that it is impossible for me to be in a retrosuicide-propitious situation, and the paradox is resolved—or so the *modified* standard solution goes. It seems thus that the modified standard solution dispenses with the hidden assumption. But what if, in addition to its being necessarily impossible for me to kill my younger self, it is necessary that if it is impossible for me to kill him then relative to the relevant features of any situation I cannot kill him? Then it is necessary that relative to the relevant features of any situation I cannot kill my younger self, and via P2* it does follow that it is impossible for me to be in a retrosuicide-propitious situation; the paradox is restored, and the modified standard solution fails. More rigorously, even if the modified standard solution does not rely on the hidden assumption that P4 (see below) is false, it still relies on the *modified* hidden assumption that P4&P5 is false:

- (P4) Necessarily, it is impossible for me to kill my younger self.
- (P5) Necessarily, if it is impossible for me to kill my younger self, then relative to the relevant features of any situation I cannot kill him.

P5 is impervious to equivocation: according to P5, *no matter* what the relevant features of a situation are, relative to these features I cannot kill my younger self if it is impossible for me to kill him.¹⁸ So P5 suggests a general reason why solutions to the retrosuicide paradox that appeal to an equivocation about relevance fail if the hidden assumption is false: even if *normally* whether I can make a proposition true depends on what counts as relevant, in the *special* case in which the proposition is necessarily impossible I just cannot make it true—regardless of what counts as relevant.

Proponents of the (modified) standard solution might wish to contest P5. But how could they do so convincingly?¹⁹ The original standard solution *presupposes* P5: it can be seen that an under-

¹⁸ Lewis might reply that, not only is it possible to equivocate about whether I can do something, “[i]t is likewise possible to equivocate about whether it is possible for me to [do something]” (1976: 150). But such a reply would not sit well with Lewis’s own use of compossibility: to say that two propositions are compossible is to say that their conjunction is possible, so if ‘possible’ is equivocal then so is ‘compossible’. More importantly, however, it seems that there *is* a fact of the matter about whether in some (physically) possible world I kill my younger self, so P5 can escape equivocation.

¹⁹ One might contest the more general claim (which underlies P5) that *impossibility entails inability* as follows: if I ask my lawyer whether I can construct a chamber in which the laws of physics are violated, my lawyer might reasonably answer “yes, you can”, given that the law of the land does not prohibit the construction of such chambers. I reply that the lawyer would be saying that I can in the sense that I am *legally permitted*, not in the sense that I have the ability (and the opportunity).

standing of ability as compossibility entails P5.²⁰ Moreover, and more importantly, it seems that *any* reasonable understanding of ability must satisfy P5: we common mortals cannot bring about (physically) impossible states of affairs—even if God, saints, or magicians can, because they can perform miracles.²¹ One might object that my inability to perform miracles does not entail P5: even if it is *impossible* for me to kill my younger self, it does not follow that my killing him is a *miracle*. The idea is that my killing my younger self may be impossible by virtue of *entailing* (rather than *being*) a miracle, namely the miracle that my younger self is resurrected (maybe by someone else, who *can* perform miracles). But then, I reply, P5 follows from my inability to perform actions which *are or entail* miracles—to use a label, my inability to *guarantee* miracles.²² One might respond that, if the world is deterministic but I have free will, then I *can* guarantee miracles; for example, I can whistle although the laws of physics together with the current state *S* of the world entail that I will not. I reply that this need not count as “guaranteeing” because my whistling need not *entail* a miracle: even if in every world *with state S* in which I whistle a miracle occurs, maybe in some world with state *different* from *S* I whistle but no miracle occurs.²³ (From the claim that my whistling is *not compossible* with the current state of the world it need not follow that my whistling is *impossible*; cf. van Inwagen 1993: 188.) Let me emphasize, however, that I am not committed to the claim that I cannot guarantee miracles or to the *truth* of P5. My point is rather that P5 is *plausible*, so we have good reason to believe that even the above charitable modification of the standard solution relies without argument on the hidden assumption. In any case, the modified standard solution relies without argument on the modified hidden assumption that P4&P5 is false.²⁴

²⁰ The claim that, necessarily, if in (i.e., relative to the relevant features of) some situation I can kill my younger self then my killing him is compossible with the relevant features of the situation entails the claim that, necessarily, if in some situation I can kill my younger self then my killing him is possible, and the latter claim by contraposition entails P5.

²¹ Assuming that “[a] miracle is a violation of the [actual] laws of nature” (Hume 1748/1975: 114; contrast Clarke 1999; cf. Lewis 1979/1986a: 44-5). If God, saints, and magicians can perform miracles, then my objection to the modified standard solution does not work for a *variant* of the retrosuicide paradox formulated with respect to these kinds of beings; nevertheless, my objection works against the solution in response to the retrosuicide paradox as I formulated it.

²² More precisely, P5 follows from the conjunction of the claims that (1) necessarily, if it is impossible for me to kill my younger self, then my killing him guarantees a miracle, and (2) necessarily, if my doing something guarantees a miracle, then relative to the relevant features of any situation I cannot do it.

²³ A similar reply can be given to the claim that, if the world is deterministic but I have free will, then I can do something X such that, if I were to do it, then some miracle or other *would have occurred* (cf. Lewis 1981/1986b). The reply is that this need not count as “guaranteeing” a miracle: from the claim that in every *closest* world in which I do X some miracle occurs it does not follow that in *every* world in which I do X some miracle occurs.

²⁴ Another attempt to dispense with the hidden assumption starts with the claim that my killing my younger self is not a miracle; to use a label, it is *locally (physically) possible*—even if it is (“globally”) impossible. (For distinctions similar to the one between local and global possibility see: Brown 1992: 440; Earman 1995a: 286, 1995b: 174; Kutach 2003: 1099-100.) If P2' and P3' are modified by replacing compossibility with local compossibility, the resulting modification of the standard solution (call it the *local* standard solution) relies not on the hidden assumption but rather on the (trivial) *local* hidden assumption that it is *locally* possible for me to kill my younger self. I reply that, if (1) it is necessarily impossible for me to kill my younger self, (2) impossibility entails inability, and (3) ability amounts to local compossibility, then my killing my younger self is necessarily *not locally compossible* with the relevant features of any situation, and via the modified P2' it does follow that it is impossible for me to be in a retrosuicide-propitious situation; the paradox is restored, and the local standard solution fails. So even the local standard solution relies without argument on the modified hidden assumption that P4&P5 is false.

3.3. Is the hidden assumption innocuous?

If retrosuicide is impossible, then *of course* we common mortals cannot commit retrosuicide. But in such a case not much of a paradox is left: the proper reaction to the retrosuicide paradox is simply that, contrary to initial appearances, I cannot kill my younger self (Malament 1985: 98; cf. Vihvelin 1996). If so, then the standard solution cannot be faulted for relying on the hidden assumption: the assumption must be true if a paradox is to arise at all, and is thus innocuous—or so one might reason on behalf of the standard solution. I have three points in reply.

(1) Suppose I am presented with two individually plausible but mutually incompatible claims. Even if one of these claims is in fact true (so that the other one is false), I might not be justified in believing that it is true, so it does not follow that my proper reaction is simply to reject the other claim. Similarly, even if the hidden assumption is in fact false (so that it is impossible for me to kill my younger self), one might not be justified in believing that it is false, so it does not follow that “the proper reaction to the retrosuicide paradox is simply that, contrary to initial appearances, I cannot kill my younger self”. It is thus not the case that “the assumption must be true if a paradox is to arise at all”: a paradox can arise even if the assumption is false but one is *not justified* in believing that it is false.

(2) Even if, because somehow one *is* justified in believing that the hidden assumption is false, “the proper reaction to the retrosuicide paradox is simply that, contrary to initial appearances, I cannot kill my younger self”, it does not follow that “not much of a paradox is left”: the question may be left of what *explains* the failure of the reasoning to the effect that in a retrosuicide-propitious situation I *can* kill my younger self. That reasoning relies on the plausible claim that ability supervenes on local physical features.²⁵ Even if one is justified in believing that this supervenience claim is false, a paradox persists until it is satisfactorily explained *why* it is false, and pointing out that the supervenience claim conflicts with some claim that one is justified in believing may be too indirect an explanation to count as satisfactory. (What *would* count as a satisfactory explanation? Maybe one appealing to a general and well-supported theory of ability.)

(3) Finally, and most importantly, even if it is granted that “the assumption must be true if a paradox is to arise at all”, it does not follow that the assumption is innocuous. This is because different conclusions on whether (in a retrosuicide-propitious situation) I can kill my younger self may be warranted depending on whether the hidden assumption is true or false. If the assumption is *false* (so that it is *impossible* for me to kill my younger self), then—let me grant—one should conclude that I *cannot* kill my younger self. But if the assumption is *true* (so that it is *possible* for me to kill my younger self), then it may well be reasonable to conclude that I *can* kill my younger self after all. If so, then it does matter whether the hidden assumption is true or false, and the assumption is not innocuous.

²⁵ Somewhat less roughly, the supervenience claim is: for any situations *S* and *S'* that have the same local physical features and any person-stage *P*, *P* can make true in *S'* every proposition that *P* can make true in *S*. For similar claims see Deutsch and Lockwood's “autonomy principle” (1994: 53) and especially its reformulations by Sider (1997: 142-3) and Chambers (1999: 298; cf. 300). If the supervenience claim is true, it is not a *relevant* (to ability) difference between two situations that my younger self is an earlier stage of mine whereas the duplicate of my younger self is not. One might object that duplicates need not have the same abilities: “Suzy's duplicate can see Suzy's face without looking at a mirror, but Suzy cannot” (Vihvelin 1996: 330 n. 4; cf. 322, 327). I reply that we don't have *two* situations here: to a situation in which Suzy's duplicate is facing Suzy, no situation corresponds in which Suzy is facing Suzy (as opposed to facing Baby Suzy).

4. Conclusion: Beyond the standard solution

I said in §3.1 that, if resurrection is impossible, then *apparently* it is also impossible for me to kill my younger self. I also said, however, that the story is more complicated; here is why. The impossibility of resurrection entails that there is no possible world in which I kill a person-stage who *in that world* is an earlier stage of mine; but there may still be a possible world in which I kill my younger self, namely a person-stage who *in the actual world* is an earlier stage of mine. As an analogy, the impossibility of killing someone before she dies entails that there is no possible world in which I kill in 1990 someone who *in that world* dies in 2010; but there is still a possible world in which I kill in 1990 someone who *in the actual world* dies in 2010. The point is simply that someone who actually dies in 2010 could have died in 1990; is it similarly the case that my younger self, who is actually an earlier stage of mine, could have failed to be an earlier stage of mine? If it is, then it may be possible for me to kill my younger self even if resurrection is impossible, and the retrosuicide paradox can be resolved without appealing to an equivocation about relevance.

In another paper (Vranas 2009) I elaborate and defend the solution I just sketched: I argue that, if “origin essentialism” is false, then my younger self could have failed to be an earlier stage of mine. Here let me briefly examine the implications for the standard solution of accepting my proposed solution. If it is possible for me to kill my younger self, so that the hidden assumption is true after all, is the standard solution thereby vindicated? No, because it remains a significant flaw of the standard solution that it relies *without argument* on the hidden assumption. Moreover, my proposed solution does not appeal to an equivocation about relevance, so if my solution is accepted then the appeal of the standard solution to such an equivocation is superfluous. But then why is the standard solution intuitively appealing? Maybe because it seems to explain our vacillation between saying that I can and that I cannot kill my younger self: the fact that I don’t kill him (or the fact that he is an earlier stage of mine) seems to pop in and out of our mind. This explanation is superficial, however. I may believe that, *regardless* of what counts as relevant, in a retrosuicide-propitious situation I cannot kill my younger self because it is impossible for me to kill him; and I may *still* vacillate because, given that in a similar situation I can kill a duplicate of my younger self, it seems that in a retrosuicide-propitious situation I *can* kill my younger self after all. There is thus a good explanation of our vacillation (between saying that I can and that I cannot kill my younger self) which does not appeal to an equivocation about relevance. Once this is realized, the standard solution should lose even its superficial attraction.

REFERENCES

- Austin, J. L. (1979). Ifs and cans. In J. L. Austin, *Philosophical papers* (3rd ed., pp. 205-232). New York: Oxford University Press. (Originally published 1956.)
- Brown, B. (1992). Defending backwards causation. *Canadian Journal of Philosophy*, 22, 429-443.
- Chambers, T. (1999). Time travel: How not to defuse the principal paradox. *Ratio*, 12, 296-301.
- Clarke, S. (1999). Hume’s definition of miracles revised. *American Philosophical Quarterly*, 36, 49-57.
- Cross, C. B. (1986). ‘Can’ and the logic of ability. *Philosophical Studies*, 50, 53-64.
- Deutsch, D., & Lockwood, M. (1994). The quantum physics of time travel. *Scientific American*, 270 (3), 50-56.
- Dowe, P. (2000). The case for time travel. *Philosophy*, 75, 441-451.
- Earman, J. (1995a). Recent work on time travel. In S. F. Savitt (Ed.), *Time’s arrows today: Recent physical and philosophical work on the direction of time* (pp. 268-310). New York: Cambridge University Press.
- Earman, J. (1995b). *Bangs, crunches, whimpers, and shrieks: Singularities and acausalities in relativistic spacetimes*. New York: Oxford University Press.
- Effingham, N., & Robson, J. (2007). A mereological challenge to endurantism. *Australasian Journal of Philosophy*, 85, 633-640.

- Feldman, F. (1992). *Confrontations with the reaper: A philosophical study of the nature and value of death*. New York: Oxford University Press.
- Fischer, J. M. (Ed.). (1993). *The metaphysics of death*. Stanford, CA: Stanford University Press.
- Goldman, A. I. (1970). *A theory of human action*. Englewood Cliffs, NJ: Prentice-Hall.
- Gott, J. R. (2001). *Time travel in Einstein's universe: The physical possibilities of travel through time*. Boston: Houghton Mifflin.
- Grey, W. (1999). Troubles with time travel. *Philosophy*, 74, 55-70.
- Hanley, R. (1997). *The metaphysics of Star Trek*. New York: Basic Books.
- Hasker, W. (1989). *God, time, and knowledge*. Ithaca, NY: Cornell University Press.
- Hume, D. (1975). *Enquiry concerning human understanding*. In P. H. Nidditch (Ed.), *Hume's enquiries* (3rd. ed., pp. 5-165). Oxford: Clarendon Press. (Originally published 1748.)
- Kenny, A. (1976). Human abilities and dynamic modalities. In J. Manninen & R. Tuomela (Eds.), *Essays on explanation and understanding* (pp. 209-232). Dordrecht: Reidel.
- Kiourti, I. (2008). Killing Baby Suzy. *Philosophical Studies*, 139, 343-352.
- Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1, 337-355.
- Kutach, D. N. (2003). Time travel and consistency constraints. *Philosophy of Science*, 70, 1098-1113.
- Lewis, D. K. (1976). The paradoxes of time travel. *American Philosophical Quarterly*, 13, 145-152.
- Lewis, D. K. (1983a). Postscripts to "Survival and identity". In D. K. Lewis, *Philosophical papers* (Vol. 1, pp. 73-77). New York: Oxford University Press.
- Lewis, D. K. (1983b). Scorekeeping in a language game. In D. K. Lewis, *Philosophical papers* (Vol. 1, pp. 233-249). New York: Oxford University Press. (Originally published 1979.)
- Lewis, D. K. (1986a). Counterfactual dependence and time's arrow. In D. K. Lewis, *Philosophical papers* (Vol. 2, pp. 32-66). New York: Oxford University Press. (Originally published 1979.)
- Lewis, D. K. (1986b). Are we free to break the laws? In D. K. Lewis, *Philosophical papers* (Vol. 2, pp. 291-298). New York: Oxford University Press. (Originally published 1981.)
- Lewis, D. K. (1986c). *On the plurality of worlds*. Oxford: Blackwell.
- Loux, M. J. (Ed.). (1979). *The possible and the actual: Readings in the metaphysics of modality*. Ithaca, NY: Cornell University Press.
- MacBeath, M. (1982). Who was Dr. Who's father? *Synthese*, 51, 397-430.
- MacBeath, M. (1983). Communication and time reversal. *Synthese*, 56, 27-46.
- Malamet, D. B. (1985). "Time travel" in the Gödel universe. In P. D. Asquith & P. Kitcher (Eds.), *Proceedings of the 1984 biennial meeting of the Philosophy of Science Association* (Vol. 2, pp. 91-100). East Lansing, MI: Philosophy of Science Association.
- Mann, W. E. (1991). Jephtha's plight: Moral dilemmas and theism. *Philosophical Perspectives*, 5, 617-647.
- Markosian, N. (2004). Two arguments from Sider's *Four-dimensionalism*. *Philosophy and Phenomenological Research*, 68, 665-673.
- Nahin, P. J. (1999). *Time machines: Time travel in physics, metaphysics, and science fiction* (2nd. ed.). New York: Springer.
- Richmond, A. (2001). Time-travel fictions and philosophy. *American Philosophical Quarterly*, 38, 305-318.
- Riggs, P. J. (1997). The principal paradox of time travel. *Ratio*, 10, 48-64.
- Sainsbury, R. M. (1995). *Paradoxes* (2nd ed.). New York: Cambridge University Press.
- Salmon, N. (1989). On the logic of what might have been. *Philosophical Review*, 98, 3-34.
- Sider, T. (1997). A new grandfather paradox? *Philosophy and Phenomenological Research*, 57, 139-144.
- Sider, T. (2001). *Four-dimensionalism: An ontology of persistence and time*. Oxford: Clarendon Press.
- Sider, T. (2002). Time travel, coincidences and counterfactuals. *Philosophical Studies*, 110, 115-138.
- Sider, T. (2004). Replies to Gallois, Hirsch and Markosian. *Philosophy and Phenomenological Research*, 68, 674-687.
- Simon, J. (2005). Is time travel a problem for the three-dimensionalist? *The Monist*, 88, 353-361.
- Smith, J. W. (1985). Time travel and backward causation. *Cogito*, 3, 57-67.
- Smith, N. J. J. (1997). Bananas enough for time travel? *British Journal for the Philosophy of Science*, 48, 363-389.
- Steen, I. K. (2007). *Composition, vagueness, and persistence*. Doctoral dissertation, Syracuse University.
- Thorne, K. S. (1994). *Black holes and time warps: Einstein's outrageous legacy*. New York: Norton.
- van Inwagen, P. (1993). *Metaphysics*. Boulder, CO: Westview.
- Vihvelin, K. (1996). What time travelers cannot do. *Philosophical Studies*, 81, 315-330.
- Vranas, P. B. M. (2009). Time travel: Two kinds of consistency paradoxes. In preparation.