

KADRI VIHVELIN

WHAT TIME TRAVELERS CANNOT DO

Those of us who believe that time travel is possible tend to think that objections to it trade on various sorts of modal confusions. Consider, for instance, the objection which is based on the claim that it is *metaphysically impossible* for the time traveler to kill a helpless sleeping baby. It goes something like this:

Suppose, for the sake of *reductio*, that time travel is possible. Then the time traveler can journey back to her own past and meet up with her younger self. Suppose that the time traveler, unhappy with her life, decides that it would have been better never to have lived it. She packs a gun and travels back through time, determined to kill her infant self. She picks a time when she knows that the baby will be alone. She checks carefully to make sure her gun is loaded. She fires. She misses. She fires again. This time the gun jams.

We don't need to go on with the story. We know that the time traveler's efforts will all fail. No matter how many times she tries, something will go wrong— her gun will jam, the bullet will be deflected, it will turn out to have been the wrong baby. We know that the time traveler's baby self will live— *must* live— for her survival is what makes possible the journey of the time traveler. So we know that the time traveler *cannot* kill her baby self; it is metaphysically impossible for her to do so. On the other hand, we can stipulate that our time traveler has the ability to shoot close-range targets and plenty of opportunity to shoot and kill the baby. By our ordinary standards of what people can do, the time traveler *can* kill the baby who is her younger self. So if time travel is possible, it's true both that the time traveler can kill her baby self and false that she can do so. Therefore, time travel is not possible.

The standard reply to this objection (let's call it the 'naive objection') goes something like this: Of course the time traveler (let's call her 'Suzy') *will not* kill the baby who is her younger self (let's name her 'Baby Suzy'). But that doesn't mean she *can't*. We may be tempted to think she can't because we know that all her attempts to do so will fail. Since we know that Suzy in fact lived until 1995 (when she stepped into the time machine), we also know that any earlier attempts— by time-traveling Suzy or by anyone else— to kill

Baby Suzy failed. But this is just a special case of our knowledge of the general truth that no one ever succeeds in killing anyone before the date of their actual death. It would be a mistake to conclude that no one *can* kill us before the day that we die. And it would be a mistake to conclude either that no one can kill Baby Suzy or that timetravelling Suzy cannot do so.¹

The standard reply in effect accuses the naive objector of being guilty of the kind of mistake the fatalist makes. The fatalist is someone who thinks that if it's true now that something will happen in the future, that its happening is metaphysically necessary and unavoidable. The fatalist thinks that if it's true now that I will eat cornflakes for breakfast tomorrow, then I *cannot* have pancakes instead. And she thinks that if it's true now that no one will kill me tomorrow, then no one *can* kill me tomorrow. This is a mistake. When we say that I can eat pancakes for breakfast tomorrow, we mean that my eating pancakes is consistent with some subset of the actual facts — facts about my competence as a pancake cook and facts about what's in my kitchen, for instance. We don't mean that my having pancakes for breakfast is consistent with the fact that I actually end up eating cornflakes instead. The same thing is true of the time traveler. When we say that Suzy can— though she won't— kill Baby Suzy, we mean that her killing the baby is consistent with some subset of the actual facts — facts about Suzy's shooting skills, her proximity to the unprotected baby, and so on. We don't mean that Suzy's shooting Baby Suzy is consistent with the fact that Baby Suzy lives to be the adult who travels back to visit her own past.

That's the standard reply. For a long time, I was persuaded by it, but I no longer think it's adequate. I agree, of course, that there is a difference between 'will not' and 'cannot', and I agree that the argument for the *metaphysical* impossibility of autoinfanticide is based on fatalistic confusions. And I agree that there are many ways in which time travelers are like the rest of us; there are things which they do not do which they nevertheless *can* do. But I now think that no time traveler can kill the baby who in fact is her younger self, given what we *ordinarily* mean by 'can'. This is a surprising conclusion, one which I resisted for some time, and one which some of my readers have taken as a new argument against the possibility of time travel. But I think it's a mistake to reject the logical (or even physical) possibility of time travel on the grounds that it cramps our style. I will argue that the time traveller's inability to committ autoinfanticide is just

what we should expect, given the time traveler's somewhat peculiar situation, and given some very general considerations about how we evaluate counterfactuals.

Before I give my argument, we need to consider the relevance of counterfactuals to our ordinary sense of 'can'.

'CAN' AND COUNTERFACTUALS

Let's begin by noting that there is a sense of 'can' in which everyone, except perhaps the fatalist, will agree that the time traveler *can* kill her baby self. It is logically possible for her to do so. There is a possible world at which the following things happen: the time traveler shoots and kills her baby self; the baby is buried but three days later is resurrected from the dead; the baby grows up to become the time traveler who journeys back to the past where she shoots and kills her baby self. This world has laws which are different from ours and is strange in other ways— we may wonder why the time traveler does not know these extraordinary facts about her past history. But it is a possible world; we do not contradict ourselves when we describe what happens there.

But this is not the sense of 'can' we ordinarily use when we talk about what people can and cannot do. It's logically possible for me to run faster than the speed of light, but I cannot do so. It's logically possible for me to walk on water, but I can't do that either. On the other hand, I can swim and I can ride a bicycle.

What this ordinary sense of 'can' commits us to is a long-standing philosophical issue, and what we say depends in part on what position we take on the free will/determinism question. Some people think that the disagreement between compatibilists and incompatibilists is just a disagreement over which sense of 'can' is the appropriate one. The incompatibilist, they say, is someone with a more demanding sense of 'can' than the compatibilist. The incompatibilist holds that 'can do otherwise' means that the person's doing otherwise is consistent with *all* the facts about the laws and past history, whereas the compatibilist says that 'can do otherwise' means that the person's doing otherwise is consistent with a proper subset of these facts— facts about the person's abilities and opportunities. On this analysis, the debate about free will and determinism dissolves when we recognize that these two senses of 'can' are in play. If determinism is true, then no one can ever do otherwise, given the incompatibilist's sense of 'can', but we can often do otherwise, given the compatibilist's less demanding sense of 'can'.

I think that this analysis of the debate is wrong. I don't think that the compatibilist and the incompatibilist *mean* something different by 'can'. I think that they both mean by 'can' what we ordinarily mean. And what we ordinarily mean when we say that someone can do something is that she has both the ability and the opportunity to do it. More precisely, we mean that she has the ability to do an act of the relevant kind and that nothing prevents her from exercising this ability. The debate is not about the meaning of 'can', but about a substantive modal question: Would deterministic causes deprive us of the opportunity (prevent us from exercising our ability) to do anything other than what we in fact do?

We won't answer that question here, but putting it this way can help us to see that there is something about which all sides can, or at least should, agree. We should agree that someone can do something, in the *relevant* sense, only if it's true that if she tried to do it, she would or at least might succeed. And everyone should agree that if someone would fail to do something, no matter how hard or how many times she tried, then she cannot do it. I can swim and ride a bicycle at least partly because it's true that I would succeed at both of these activities if I tried to do them. I can hit a bullseye on a dart board at least partly because it's true that if I tried to do it, I might succeed. But I *cannot* walk on water or run faster than the speed of light because if I tried to do these things, I would fail, no matter how often or how hard I tried.

The standard reply to the naive objection to time travel says that the time traveler can kill the baby who is her baby self because doing so is consistent with facts about the time traveler's abilities (e.g. her shooting skills) and opportunities (the gun in her hand, the proximity of the baby, and so on). I say that it is not at all obvious what the time traveler's abilities and opportunities are. If someone wants to defend the claim that Suzy *can* kill Baby Suzy, despite the fact that all her attempts to do so have failed and will continue to fail, then she has to defend the claim that there is at least one occasion on which it's true that if Suzy had tried (one more time) to kill Baby Suzy, she would, or at least might, have succeeded.

But I don't think this claim can be defended. The fact that someone tries and fails to do something is consistent with its being true that she could have done that thing, for it is consistent with its being true that if she had tried, just one more time, she would, or at

least might, have succeeded. But in the case of the time traveler, I think that this counterfactual is always false.

Consider Suzy after she has pulled the trigger for the fifth time and once again failed to kill her infant self. Would she have succeeded if she had tried one more time? I think that we have good reason to believe that her sixth attempt would also have failed. More generally, I say that it's true not only that Suzy's attempts to kill Baby Suzy *will* all fail; I say that it's also true that *if* Suzy had made any further attempts to kill Baby Suzy, these attempts *would* also all have failed. And because she wouldn't succeed no matter how often or how hard she tried, I don't think the time traveler can kill her baby self

WHY SUZY CANNOT KILL HER BABY SELF

Let's remind ourselves of how we evaluate counterfactuals. We evaluate counterfactuals by considering what happens at the *closest* worlds at which the antecedent is true. 'If P, it would be the case that Q' is true just in case Q is true at *all* the closest P-worlds. 'If P, it would or might be the case that Q' is true just in case Q is true in *at least some* of the closest P-worlds.

How do we decide which antecedent-worlds are closest? Well, that depends on your theory of counterfactuals, but most theories agree about this much: the closest antecedent-worlds are those antecedent-worlds which are most similar, in the relevant respects, to our own. So, for instance, when we ask what would have happened had I tried to walk on water, we do not consider worlds where I'm as light as a feather, worlds at which the earth has a much smaller mass, worlds at which I have wings, and so on. Rather, we consider worlds which have the same (or nearly the same)² laws of nature as ours, where the earth has the same mass, and where I am a human being weighing what I actually weigh.

If it is true that Suzy can kill her baby self then it must be the case that there are at least *some* occasions on which it's true that:

- (S) If Suzy had tried to kill Baby Suzy, she would or at least might have succeeded.

But that is so only if on some occasions at least some of the closest worlds at which Suzy tries to kill Baby Suzy are worlds at which she succeeds. Whether this is so will depend, of course, on what the time traveler's world is like. I assume, as discussion of time travel standardly assume, that the time traveler's world is much like the actual world. When we ask "What would things be like if there were time travel?", we are asking a

subjunctive question. We evaluate it like any counterfactual. We imagine a world as similar to our world as can be, consistent with its laws allowing time travel. Maybe, as some scientists tell us, the actual world is already a world like that.

I claim that at any world remotely like ours, (S) is always false. I say that *all* the closest worlds at which Suzy tries to kill Baby Suzy are worlds at which she fails. That is, I say:

(F) If Suzy had tried to kill Baby Suzy, she would have failed
is invariably true.

And if (F) is always true, and if I'm right about what we mean by 'can', then it's always false that:

(C) Suzy can kill Baby Suzy.

When we send Suzy back in the time machine we know that she will not kill her baby self. We know that she will not because we know that she did not. Why didn't she? There are two possible explanations. Either she tried and failed, or she didn't try. If she tried and failed there will be some explanation for her failure: the gun jammed, the baby moved, the bullet was a dud Something happened that prevented her from succeeding. On the other hand, if she didn't even try to kill her baby self then we don't have to suppose that any of these factors prevented her from doing so. The gun may have been well-oiled, the ammunition live, and the baby a sitting duck. But if she didn't try to kill the baby, it's not surprising that she didn't succeed. Our question is: Is it true, on any of these occasions, that if she had tried she might have succeeded?

I don't think so. Take any case you like: Suppose that the gun was greased, the bullets live, and the baby a sitting duck. I say that even in this case (S) is false and (F) is true. Despite appearances, Suzy cannot kill her baby self.

Of course, it's *logically possible* for Suzy to kill Baby Suzy; that is, there are worlds at which Suzy tries to kill Baby Suzy and her attempt succeeds. But consider what those worlds must be like. They are worlds at which Baby Suzy dies and is buried, but is later resurrected from the dead and grows up to be the adult Suzy who travels back through time and kills her baby self.

Is there any other logically possible way to work the trick? Well, *maybe*. Maybe there are worlds at which Suzy kills Baby Suzy, and Baby Suzy stays dead, but Baby Suzy

is somehow not Suzy's younger self. (Whether this is really possible is something we will take up shortly.)

Either way we are talking about possible worlds very different from our own. So different, it seems to me, that they are not relevant to counterfactuals about what would have been the case had Suzy tried to kill Baby Suzy. If Suzy had tried to kill Baby Suzy—if she had squeezed the trigger of the gun she was in fact holding in her hand—the laws of nature would still have been the same (or nearly the same) and there would still have been no resurrections from the dead. And if Suzy had tried to kill Baby Suzy, Baby Suzy would still have been Suzy's younger self.

Compare Suzy's case to more prosaic cases. Take any real life actual baby you like. Don't kill it! Please, don't even *try* to kill it! But do consider the counterfactuals. Do you think it is true that if you tried to kill it, it would be resurrected from the dead? Of course not. Do you think that if you tried to kill it, then someone else would, in later life, somehow take the place of the adult this baby will become? Of course not. There are possible worlds where there is resurrection from the dead, and maybe there are worlds where babies have stand-ins waiting in the wings ready to take their place. But neither of these worlds is remotely like the world that would come about if you tried to kill that baby. If you tried to kill that baby, then either you would succeed and the baby would be gone forever or you would fail. If you think that worlds where you fail are more like the actual world than any world where you succeed, then you think that you cannot kill the baby. Precisely the same is true about Suzy and Baby Suzy. But in Suzy's case there is this difference. The worlds where Suzy tries and *fails* are worlds where the gun jams, the bullets are duds, the baby moves The worlds where Suzy tries and succeeds are worlds which either have resurrection from the dead or some sort of system of ontological understudies. (There's no other way for her to succeed). But worlds of the former sort—where guns jam or bullets misfire — are always more similar to the actual world than worlds of the latter sort. Which is why it is always true that if Suzy tried to kill Baby Suzy, she would fail.

The reason the counterfactuals stack up this way has everything to do with the fact that Suzy is a time traveler and Baby Suzy is her baby self. Other people can kill Baby Suzy. Suzy's mother could have killed Baby Suzy. Suzy can kill other babies. But she can't kill *that* baby. If Suzy's atom-for-atom duplicate were in Suzy's shoes, she would be

able to kill Baby Suzy. If Baby Suzy were replaced by a baby who is the atom-for-atom duplicate of Baby Suzy, then Suzy would be able to kill the duplicate. But if that baby is Baby *Suzy*, Suzy cannot kill it.

Objection: You've argued that if Suzy killed Baby Suzy, Baby Suzy would be— would have to be— resurrected from the dead. But then the standard defense of time travel is right after all. Suzy *can* kill Baby Suzy — no contradictions or paradoxes arise if we suppose that she does.

Reply: I agree that if Suzy *succeeded* in killing Baby Suzy, Baby Suzy would be resurrected from the dead. But that's not the counterfactual relevant to our ordinary sense of 'can'. What's relevant is what would happen if Suzy *tried* to kill Baby Suzy. And it's false that if Suzy tried to kill Baby Suzy, then Baby Suzy would be resurrected from the dead. Compare: If I succeeded in running faster than the speed of light, then the laws of nature would be— would have to be— very different. But it's false that if I *tried* to run faster than the speed of light, the laws of nature would be very different.

Objection: Doesn't your argument show that time travel is possible only at worlds whose laws provide for resurrection from the dead? And since our world is not such a world, you've shown that time travel is *nomologically* impossible.

Reply: Not so. I think that time travel is possible at worlds very much like ours, maybe exactly like ours. But in any case there is no reason to suppose that there is any connection between time travel-permitting laws and resurrection-permitting laws. My claim is that *if* time travel occurred, then the laws would still be much the same; there would still be no resurrection from the dead.

Objection: Suzy *can* kill *that* baby— the one you're calling 'Baby Suzy'. If she tried to kill that baby, she would succeed, but in that case it would turn out that the baby is not Baby Suzy, but some other baby who looks just like Baby Suzy.

Reply: You are changing the subject. That baby *is* Baby Suzy and our question is whether Suzy can kill Baby Suzy, not whether she can kill Baby Suzy's twin or any other baby. The only worlds relevant to the evaluation of the counterfactual are worlds at which Suzy tries to kill Baby Suzy.

Objection: But what's at issue is not whether Suzy can kill a baby who satisfies the description 'Suzy's baby self'. Maybe by definition it is impossible for Suzy to kill any baby

which satisfies that description. But the interesting question is whether Suzy can kill a particular baby — *that* baby, however we describe it.

Reply: The issue is not terminological. ‘Baby Suzy kills her baby self’ isn’t false “by definition”, there are worlds where it is true — worlds where there is resurrection from the dead. The impossibility I am arguing for is neither analytic nor logical; it’s nomological. Given our laws, no one can kill her younger self. Laws sustain counterfactuals, and one of the counterfactuals sustained by our laws is this one: If anyone tried to kill the baby who is her younger self, she would fail.

Objection: But surely you agree that it’s a contingent fact that *this* baby grew up to be *that* person— our grown-up Suzy, the time traveler. *This* baby could have died in infancy and Suzy could have had a different past; she could have been born to a different mother. Surely there are worlds where *this* baby and *that* adult are two different people. So surely there are worlds with our laws (no resurrection from the dead) where *that* adult kills *this* baby. Why not say that a world thus described is what the world would be like if Suzy tried to kill Baby Suzy?

Reply: Let’s take this one step at time. Certainly there are worlds where *this* baby fails to become an adult— worlds where this person dies in infancy. But it is less clear that Suzy— *that* adult— could have failed to come from this child. I take it that Kripke would argue that this is not possible on the grounds of the “necessity of origins”. Kripke’s intuition seems to be that there is no possible world where Suzy grows from anything other than Suzy’s zygote. If that’s so, then it would seem to follow that any world with an adult Suzy and a Baby Suzy must be a world where the latter is the origin of the former.

But even if we reject the “necessity of origins” thesis, I don’t think that this objection can be made to work.

The relevant question is whether that adult can kill this child. But we have to be careful because the question is potentially ambiguous. When we talk about “this baby” or “Baby Suzy”, we might be talking about a person-stage or we might be talking about the temporally extended person of whom the stage is only an early part. We have the same ambiguity with “that adult” or “Suzy”. When we ask what would have happened had Suzy (that adult) tried to kill Baby Suzy (this baby), we might be asking what would have happened had that person-stage tried to kill this person-stage. Or we might be asking what would have happened had that person tried to kill this person.

I take the latter question, the question about persons, to be the relevant one. Our question is about what *time travelers* can do. To say that someone is a time traveler presupposes that she is a person who extends through time. A time traveler is a temporally extended person with one stage in the future and a later stage in the past. But it's easy to get confused here because we are dealing with identities in several dimensions. There are questions about the identity of Suzy over time, questions about the trans-world identity of Suzy; and questions about the trans-world identity of Suzy-stages.

To avoid confusion, think of persons as four-dimensional objects. Suzy, the person, is an ensemble of Suzy-stages, strung out across time and connected by chains of causal dependency. Are there worlds where the stages which are temporal parts of the actual person, Suzy, occur as parts of different persons? Let us suppose so. Let us suppose, moreover, that at one such world the ensemble that embeds the adult-stage kills off the one containing the baby-stage. We might describe this by saying that at this world, Baby Suzy doesn't become the adult Suzy or by saying that the adult Suzy comes from a different baby.

Well and good. But what does this show? Note that worlds like this are irrelevant to deciding what Suzy, the *person*, can or cannot do. Suppose that at the world just described the adult succeeds in killing the baby. Is this a world where Suzy, the person, tries and succeeds in killing Baby Suzy, the person? I think not. Remember Suzy isn't a stage; she is an ensemble. If Suzy exists at this world at all she must be identical with some ensemble in it, either the ensemble that is killed or the different ensemble that does the killing. Is Suzy at this world the killer or the victim? In either case this is not a world where Suzy kills Baby Suzy, it is a world where she either kills or is killed by *someone else*.³

Can we stipulate that Suzy is *both* the killer and the victim — that our Suzy, the whole temporally extended being, is identical to *both* of these two different persons in this other world? No. Here I think we have to draw the line. Suzy cannot be identical to both of these non-identical persons. Maybe this is metaphysical short-sightedness on my part. Maybe, out there on the far fringes of logical space, Suzy is a multitude engaged in trans-temporal auto-genocide. Even so, I'm sure that these worlds are far too dissimilar to our world to be relevant to the counterfactual that concerns us.

Objection: Maybe you are right about persons. But surely that shows that the relevant question *isn't* about what persons can do; it's about what person- stages can do. And haven't we just seen that there are worlds where this adult stage kills that baby stage — worlds where the stages are stages of different people?

Reply: I am not denying that the world as described is logically possible (though I have misgivings). And I'll allow (for the sake of this argument anyway) that this world is more like ours than a world whose laws allow resurrection from the dead. If so, then it is true that if that adult stage *tried and succeeded* in killing this baby stage, the world *would* be as described: the stages would be stages of different persons, and this baby would not be that adult's younger self. But this is beside the point.

Our question is not about what the world would be like if that adult-stage *tried and succeeded* in killing this baby stage. Our question is about what that adult stage *can* do, and that is a question about what whether that adult stage would have succeeded if it had *tried* to kill this baby stage. And I think the answer is clearly "No". It is false that if that adult stage had pulled the trigger, then the adult stage and the baby stage would have been stages of two different persons. Any world like this is far less like ours than a world where that adult stage tries, but fails to kill this baby stage

To see this, consider what the world would have to be like for it to be true that if that adult stage had pulled the trigger, then that adult stage and this baby stage would have been stages of two different persons. Let's pick a specific occasion -- say Sept. 3, 1960, when Baby Suzy is one year old and adult Suzy is thirty-six and is standing a few feet from Baby Suzy, gun in hand and murder on her mind. As a matter of fact, that adult stage and this baby stage have matching DNA and are linked by chains of causal dependencies stretching from 1960 to 1995 and then back again to 1960. We are being asked to suppose that if that adult stage had pulled the trigger, then this baby stage would have died and these causal chains would not have obtained. But then how would that adult stage have come to be there in 1960, gun pointed at this baby-stage with matching DNA? There seem to be just three possibilities: Either that adult stage would have miraculously come into existence out of thin air, or the DNA of some actually existing person would miraculously have changed so that it matches the DNA of that adult stage, or the world would have included another baby stage with DNA that matches the DNA of that adult stage and thirty-six years worth of the appropriate causal connections would have linked this extra

baby stage to that adult stage. But worlds like these are far less like ours than worlds where the gun jams, or the bullet misses, or... and that adult stage fails to kill this baby stage. So if that adult stage had tried to kill this baby stage, the stages would still have been stages of the same person, and the attempt would have failed.

Objection: Look, suppose that over on Putnam's Twin-Earth, Suzy's doppelganger gets into a time machine to kill her baby self, but wires get crossed and she materializes right beside our Baby Suzy. You agree that Twin Suzy can kill Baby Suzy. Yet you are committed to saying that Suzy cannot do what her atom-for-atom duplicate can do.

Reply: I am committed to that, and I agree it seems odd.⁴ It seems odd because, ordinarily, when we are considering what someone is able to do right now, the past is irrelevant. Or rather, the past is irrelevant except insofar as it affects current present-tensed properties. Of course, Suzy can't swim if she hasn't taken lessons, but the lessons enable her to swim only because of the difference to the way her brain is set up now. It's her current brain set-up that gives her that ability, however she acquired it; and if some doppelganger had the same set-up, she would have the same ability to swim, however she had acquired that set-up. So, normally we think we can, at least in principle, read a person's abilities off her present-tensed properties — the way she is now. And we think we can read facts about what a person *can* do from present-tensed facts about the person — her abilities — together with present-tensed facts about her immediate environment — her opportunities. But then normally we are not dealing with time travelers. Normally causation runs in one direction only -- from past to future. And the past constrains what we can do in the future only by affecting the present.

But with time travelers it is different. The way Suzy and her surroundings are now is caused not only by the past but also by the future-tensed fact that Baby Suzy will survive to be the adult Suzy. So we shouldn't be surprised that what Suzy *can* do cannot be read off present-tensed facts about Suzy and her surroundings. We should not be surprised that the future-tensed fact of Suzy's *not* killing Baby Suzy turns out to be one of the facts which must be held constant in our evaluation of counterfactuals about what would have happened had Suzy *tried* to kill Baby Suzy. Suzy's failure to kill Baby Suzy is what makes Suzy's attempt possible. The relevant sense of possible is nomological; given our laws, Suzy tries to kill Baby Suzy only if her attempt fails. (The laws of our world dictate that no one tries to do anything in later life if she was killed as a child.)

Objection: But if your argument works, doesn't it show that time-travelers are able to do — can do— only what they in fact do?

Reply: No. There are other things time travelers cannot do; for instance, they cannot kill their ancestors before they pass on their genes. But they are able to do lots of other things; they have plenty of abilities which they do not exercise, but could have exercised. Suzy could have pinched her baby self. She could have gone next door and killed some other baby. There are worlds with our laws at which Suzy tries and succeeds in pinching her baby self or in killing that other baby. So we have no reason to doubt the truth of the relevant counterfactuals: If Suzy had tried to pinch her baby self (kill that other baby), she would have succeeded.

But this leaves open lots of interesting questions. For instance: Can a time traveler kill the infant Hitler? Here is an argument that no time traveler can do so.

If a time traveler is in the past, gun aimed at Baby Hitler, this can only be because she's read history books reporting the baby's future's crimes. But the history books will report these crimes only if the baby survives the time traveler's visit and grows up to commit the crimes. So if the time traveler is in the past with Baby Hitler this can only be because Baby Hitler survived her visit. We should hold this fact constant when we evaluate counterfactuals about what would have happened if the time traveler had tried to kill Hitler. So if the time traveler had tried to kill Baby Hitler, she would have failed.

But the first premise of this argument is false. It's false that the *only* reason a time traveler might try to kill Baby Hitler is because she knows the baby's future deeds. A time traveler might travel back to the past, pick a victim at random, and kill for sport. If the victim is Baby Hitler, the time traveler will not succeed. But there is no reason to doubt that if the time traveler had tried (one more time) to kill that baby, she would or might have succeeded.

Still, we might think that this argument works for time travelers who travel back to the past with the deliberate intention of killing Baby Hitler. Let's suppose that Suzy is such a time traveler and that she has, by our *ordinary* (present-tensed) standards, both the ability and the opportunity to kill Baby Hitler. Does my argument show that Suzy cannot kill Baby Hitler? More generally, does it show that no time traveler can *intentionally* affect the past? I'm not sure. I think that in these sorts of cases the counterfactuals are

not always clear. But it is at least arguable that Suzy *can* kill Baby Hitler (though of course she won't).

The laws of nature, our laws, dictate that adults do not arise either *ex nihilo*; a person's adulthood is caused by her childhood. The particular counterfactuals these laws sustain make it false, I have argued, that time travelers can kill their baby selves. On the other hand, the laws of nature do allow memories and beliefs to arise *ex nihilo*, or at least without being true. There are worlds with our laws where Suzy tries to kill Baby Hitler because of what she believes this child will become, and she succeeds, and Hitler is not resurrected from the dead, nor does anyone take his place and commit his crimes. At these worlds, Suzy's beliefs are false; her memories are hallucinations. So it's false that if Suzy had tried to kill Baby Hitler, this would *have to have been* because her attempt failed and Baby Hitler survived; there are other ways in which Suzy might have ended up trying to kill Baby Hitler. Can Suzy kill Baby Hitler? That depends on what *would or might* be the case if she tried, and that depends on what happens at the *closest* worlds where Suzy tries. Arguably, worlds where Suzy tries, succeeds, and turns out to have radically mistaken beliefs about Baby Hitler's future are closer to the actual world than worlds where Suzy has true beliefs but something or other goes wrong, and her attempt fails. If so, then there is reason to believe that Suzy *can* kill Baby Hitler, even though we know that she never will.⁵

NOTES

- ¹ For versions of what I'm calling 'the standard reply', see David Lewis, "The Paradoxes of Time Travel", *American Philosophical Quarterly* 12 (1976), 145-152 (also in Lewis's *Philosophical Papers*, vol.2, Oxford University Press, 1986) and Paul Horwich, "Time Travel", Chapter 7 of his *Asymmetries in Time* (Bradford Books, MIT Press, 1989).
- ² David Lewis is the best-known defender of the view that the closest antecedent-worlds may be worlds with somewhat different laws. But on his view, the laws are only as different as they have to be to accommodate the event which, by the standards of our laws, is a "divergence miracle". See Lewis, "Counterfactual Dependence and Time's Arrow", *Nous* 13 (1979), 455-76 (also in his *Philosophical Papers*, Volume 2, *ibid.*).

- ³ We get the same result if we substitute talk of counterparts for trans-world identity. There are worlds where counterparts of the baby stage and the adult stage are parts of different persons. We might choose to treat either of these persons as the counterpart of our Suzy, but we cannot simultaneously treat both as Suzy counterparts. We can imagine Suzy as one of a pair of identical twins. She might have been the older; she might have been the younger. But she could not have been *both* twins.
- ⁴ But perhaps not that odd. We should all agree that Suzy's duplicate can do some things that Suzy cannot do. For instance, Suzy's duplicate can see Suzy's face without looking at a mirror, but Suzy cannot see Suzy's face without looking at a mirror. Of course, neither Suzy nor her duplicate can see *her own* face without looking at a mirror. My claim is that the situation with respect to autoinfanticide is exactly the same. Neither Suzy nor her duplicate can kill her younger self. But just as Suzy and her duplicate differ with respect to what particular faces they can see, Suzy and her duplicate differ with respect to what particular babies they can kill. Suzy's duplicate can kill Baby Suzy, but Suzy cannot kill Baby Suzy (though she can kill Baby Suzy's duplicate).
- ⁵ I'm grateful to Jonathan Bennett, Bob Bright, Bryson Brown, Curtis Brown, John Fischer, Jeff Foss, Mark Heller, Brian Loar, Ray Martin, George Mavrodes, Adam Morton, Alex Rosenberg, Ted Sider, Paul Teller, and especially Terrance Tomkow for helpful discussion and comments on this paper.

School of Philosophy
University of Southern California
Los Angeles, California 90089-0451
USA