Organon F 33 (2) 2026: 162–181 ISSN 2585-7150 (online)
https://doi.org/10.31577/orgf.2026.33203

RESEARCH ARTICLE

The Paradoxes of Time Travel: 50 Years Later

Ryan Wasserman*

Received: 6 October 2025 / Accepted: 17 December 2025

Abstract: David Lewis set the agenda for philosophical theorizing about time travel fifty years ago with his landmark paper "The Paradoxes of Time Travel". In honor of that anniversary, I review some of the central themes of his paper and discuss a few of the ongoing debates inspired by that work. In particular, I address some worries that have been raised for Lewis's characterization of time travel, his treatment of backward causation, and his proposed solution to the grandfather paradox.

Keywords: Time travel; David Lewis, backward causation; Grandfather Paradox.

"The Paradoxes of Time Travel" is a philosophical gem with many beautiful facets that reflect the brilliance of its creator. Originating as an informal talk for students and later developed into a book project, the work was eventually published in *American Philosophical Quarterly* in 1976.[1] It would go on to inspire a generation of philosophers and remains the starting

[1] For a history of the project and unpublished book chapters, see Lewis (2023).

* Western Washington University
https://orcid.org/0000-0003-2202-2131
Department of Philosophy, Western Washington University, Bellingham, WA 98225, USA
ryan.wasserman@wwu.edu

point for all serious work on the topic.[2] I cannot hope to do justice to the whole paper here, so I will limit my attention to three central themes: Lewis's characterization of time travel (section 1), his comments on backward causation (section 2), and his treatment of the grandfather paradox (section 3).[3] In each case, I will explain Lewis's position, review what has been said in response, and provide my own assessment of the debate. I will then conclude (in section 4) by reflecting on the continuing impact of Lewis's paper.

1. The Definition of Time Travel

Time travel can seem incoherent by its very definition. After all, ordinary trips through space involve being in different *places* at different times. So, time travel would presumably require being in different *times* at different times. But it is unclear what this would even mean. Worse yet, as soon as we begin to describe a case of time travel, we appear to led into contradictory statements. Suppose, for example, that Marty gets into his time machine in 1985, hits 88 mph, and then–at the very next moment–finds himself in 1955. In that case, Marty's arrival comes just *a moment after* his departure, but it also happens *thirty years before*. This seems impossible. "How can it be that the same two events, his departure and his arrival, are separated by two unequal amounts of time?" (Lewis 1976, 145)

Many thinkers–including philosophers, scientists, and science fiction authors–have tried to avoid these problems by reimagining time.[4] Normally, we think of events as unfolding along a single dimension where any two points are separated by a unique distance.[5] (See the figure below, where the 'D' marks Marty's departure and the 'A' his arrival.)

---

[2] For an overview of the literature, see Wasserman (2018).

[3] One thing I will not be discussing are Lewis's comments on persistence and personal identity. For a more developed statement of his views, see Lewis (1983). For a discussion of how these things relate to time travel, see Wasserman (2018, Chapter 6).

[4] See, e.g., Meiland (1974), Deutsch (1997), and Daniels (1935).

[5] Lewis understands time to be one dimension of a four-dimensional manifold (ibid., 145–46); for our purposes, a single line will do.

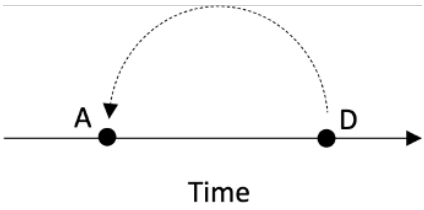

But we could instead think of time as a plane with room for multiple universes to run in parallel. (See Figure 2.):

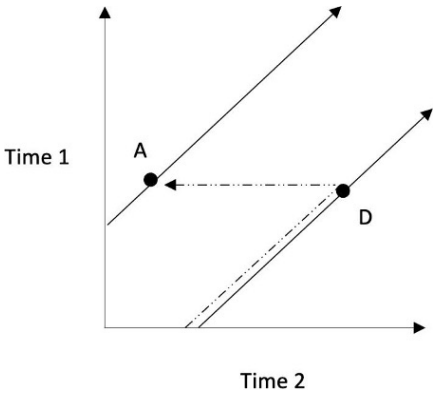

In this image, the solid angled lines depict different world-histories, and the dashed line represents Marty's journey through time. Importantly, this figure includes two different temporal dimensions, which allows us to say that Marty's departure and arrival are separated by unequal amounts of time– the departure is 30 years after the arrival in Time 2, but just a moment earlier according to Time 1. It also allows us to say that Marty's journey involves being at different times (of Time 1) at different times (of Time 2). In this way, we can avoid both of the initial worries outlined above.[6]

[6] Lewis credits the two-dimensional model to Meiland (1974) and Asimov (1955), but related ideas go back at least to Daniels (1935). For more recent developments of this approach, see Goddu (2003), van Inwagen (2010), and Law (2018).

In his paper, Lewis acknowledges the possibility of this model but questions its usefulness. When a time traveler visits his past, he is sup- posed to visit *his* past–that is, *the past where he previously existed*. Jumping to a parallel universe might provide a similar experience, but it is not the real thing. For this reason, Lewis claims that it does not give us "time travel as we know it from the stories" (ibid., 145).

More importantly, Lewis argues that the two-dimensional model is unnecessary, since we can make sense of temporal discrepancies within a single timeline. The key is to introduce two different ways of ordering events on that line and measuring the distances between them:

> I…distinguish…time itself, *external time* as I shall also call it, from the *personal time* of a particular time traveler: roughly, that which is measured by his wristwatch. His journey takes an hour of his personal time, let us say; his wristwatch reads an hour later at arrival than at departure. But the arrival is more than an hour after the departure in external time, if he travels toward the future; or the arrival is before the departure in external time (or less than an hour after), if he travels toward the past. (Ibid., 146)

Looking back at the first figure above, Lewis would say that, according to external time, Marty's arrival in 1955 comes first. Then, many years later, Marty gets into the time machine and hits 88 mph. That is the ordering of those events in real time. From Marty's perspective, however, the adventure begins with his getting into the machine in 1985. After that, he hits 88 mph and arrives in the past. That is the order of the events–*the very same events*–according to his personal time.

Lewis's picture, then, is one in which all the elements of a time-travel story occur along a single timeline. Moreover, he argues that this timeline is immutable. If Marty *is* (according to his personal time) going to arrive safely in 1955, then it is *already* the case (according to external time) that he arrived. If Marty *will* (personally) interact with his teenage father in 1955, then that is *already* (objectively) a fact before he gets into the time machine.

Importantly, Marty's actions in 1985 will still have an *effect* on the past–his hitting 88 mph will, for example, be a *cause* of his arrival in 1955–but

nothing Marty does will *alter* the past. Backward time travelers, in this picture, simply help to bring about the history that was.

Lewis concludes that the distinction between external and personal time allows us to give a consistent description of time travel without invoking parallel universes, rewritten histories, or other theoretical posits. Moreover, he suggests that the distinction provides for a perfectly general definition of time travel: for someone to be a time traveler is just for there to be a discrepancy between their personal time and the objective temporal relations that hold between events (ibid., 146–47). Marty's 1985-experiences, for example, come *before* his 1955-experiences in his personal time, but *later* according to external time. Or consider H.G. Well's famous example of a time traveler who journeys continuously into the future. That character experiences events in the same order as everyone else, but, from his perspective, things are sped up, so that a journey lasting thousands of external years takes less than a day of personal time.[7] According to Lewis, it is discrepancies of this kind that constitute time travel.[8]

Lewis's picture has proven extremely popular.[9] Nonetheless, worries remain. Here, I will focus on three problems related to Lewis's description of personal time, his definition of time travel, and his appeal to external time.

Lewis initially equates personal time with "that which is measured by [a] wristwatch" (ibid., 146). But, as he himself notes, the mere fact that someone's watch has gotten out of step does not make that person a time traveler. He therefore suggests the following account:

---

[7] Yet another example would be the case of the time traveler who takes an hour of personal time to travel an hour into the past. This would involve a discrepancy of or- der without any difference in distance.

[8] Lewis also famously characterizes time travelers in terms of their four-dimensional shapes (ibid., 146). Marty, for example, would be a "broken streak" with one segment ending in 1985 and another beginning in 1955. Wells's time traveler, meanwhile, would be a "continuous streak" which is "stretched out" relative to ordinary individuals. (The stretching, in this case, means that the mental and biological processes of the time traveler take longer to complete than those of an ordinary individual–see below for further details).

[9] Lewis's definition of time travel has been endorsed by Horwich (1987, 114), Dowe (2000, 441–42) and many others. The main alternative is to characterize time travel using the tools of relativity–see below.

> If you take the stages of a common person, they manifest certain regularities with respect to external time…Memories accumulate. Food digests. Hair grows. Wristwatch hands move. If you take the stages of a time traveler instead, they do not manifest the common regularities with respect to external time. But there is one way to assign coordinates to the time traveler's stages… so that the regu- larities that hold with respect to this assignment match those that commonly hold with respect to external time (1976, 146).

In Marty's case, his 1955 arrival comes after his 1985 departure (according to his personal time), since that sequence shows the hair growth, memory accumulation, etc. that we would normally see in the life of an ordinary person. That is what makes the arrival the "next" event in Marty's personal history. In contrast, imagine that Marty returns to 1985 after spending a week in the past and that he arrives at the precise time and place of his earlier departure. In that case, Marty's 1985 arrival-stages will follow immediately after his 1985 departure-stage in external time. However, to an outside observer, it will look as if Marty has instantaneously aged by a week (with longer hair, fingernails, etc.). Obviously, we do not typically see discontinuities like this in the lives of ordinary people. That is why the 1985 arrival-stages do not come "next" in Marty's personal time.

Lewis's suggestion may seem promising, but it is also problematic. The most obvious issue is that his characterization of personal time only applies to persons and other biological entities that undergo such changes. Moreover, there is no way to generalize Lewis's approach, since some potential time travelers are intrinsically unchanging. To make the worry as vivid as possible, imagine a single persisting electron in an otherwise empty world. Is that particle moving normally toward the future or continuously into the past? Whichever way you look at it, the pattern will remain the same–there is just one unit of negative charge after another. As a result, there is not enough structure to tell us whether there is a discrepancy between external time and the "personal" time of the electron.

A better option–suggested by Dowe (2000, 441), Wasserman (2018, Chapter 1, section 1), and others–is to characterize personal time in terms

of causation.[10] When Marty arrives in 1955, his hair, memories, and other features will closely match those of his 1985 departure. But those features are not *just* similar–they also depend, causally speaking, on what Marty was like when he hit 88 mph. In fact, Marty's very existence in 1955 is the direct causal result of what happened 30 years in the future. The same thing would be true for a lonely electron that is traveling continuously toward the past–at each moment, its existence and features will depend, causally speaking, on what it is like at later moments of external time. This would distinguish it from an ordinary, forward-moving electron, in which each stage depends on those in the external past. Something similar will presumably be true for any other kind of entity, since it is commonly thought that persistence through time (in either direction) requires causal dependence. We can therefore think of personal time, in general, as the order and metric provided by whatever causal relations are constitutive of identity through time.

Even if we grant Lewis his notion of personal time, we might still worry about his definition of time travel. One objection concerns the case of suspended animation. In the Marvel Cinematic Universe, Captain America crashes into the Arctic in the 1940s and is encased in ice for nearly 70 years. When he awakes, Cap is largely unchanged since his biological processes were "frozen" in time. Since these processes are constitutive of personal time for Lewis, he will have to say that there is a discrepancy between external time and Cap's personal time. But that seems problematic, since we would not normally take this to be a case of time travel.

It is less clear what to say about this case on my characterization, since I have not given a precise definition of "causal distance." When a time traveler discontinuously jumps into the future, there are no intermediate stages, so the causal distance from departure to arrival is zero. In Cap's case, his body is present throughout the 70-year period, and causal connections hold at each stage. So, perhaps the causal distance is 70 years as well. That would eliminate any discrepancy with external time, in which case Cap would not be a time traveler.

[10] Lewis himself thinks causation is required for personal identity, but he does not mention it in his characterization of personal time.

But now go back to the Wellsian time traveler. Like Captain America, he is continuously present throughout his journey. But unlike Cap, he is conscious during his trip and sees everything else sped up. However, if an external observer were to pass by the time machine, she would see something different: everything would, from her perspective, be moving at the normal pace, whereas the *time traveler* would appear to be in slow motion. Indeed, depending on the details, the time traveler might look just as "frozen" as Captain America. But, in that case, we seem to have a problem with *any* definition of time travel since we want to say that Wells's character is a time traveler and that Captain America is not. At this point, one might conclude that suspended animation *is* time travel or that Wellsian travel is *not*. But it would be even better if we could identify a relevant difference between the cases. One suggestion is made by Lewis himself in a letter to Jonathan Bennett:

> I'm not sure about cryogenic procedures. They slow down different processes to different degrees, and the rates of the fundamental atomic processes… are unchanged. So it's not clear to me that if we take a functional analysis of personal time… we get a discrepancy between personal and external time. (Quoted in Wasserman 2018, 13.)

Applied to the case of Captain America, the thought would be that his *biological* processes are frozen, but his basic subatomic particles are acting normally. If personal time is characterized in terms of these more fundamental processes, there would no longer be any discrepancy with external time. Moreover, if Wellsian time travel slows down *every* physical process, we would then be able to draw a distinction between the two cases.[11]

A third and final worry for Lewis concerns his appeal to external time. Lewis says little about this notion in his paper, beyond calling it "time itself." But for his definition of time travel to work, the external temporal relations between events would have to be independent of a particular individual's perspective, since that would be the only way to get a mismatch between personal and external time. The problem is that, on the standard

[11] For further discussion of this worry, see Wasserman (2018, Chapter 1, section 2).

interpretation of relativity, there is no objective temporal ordering for all events. In this sense, there is no (global) notion of real time. But if there is no real time, then there cannot be a mismatch between *it* and personal time.

One response to this worry is to say that non-relativistic worlds are possible and that Lewisian time travel could still occur in such worlds. As Lewis himself notes, a time-travel world might have to be "different in fundamental ways from the world we think is ours" (ibid., 145). However, many have argued that relativity allows for journeys through spacetime that appear to be time travel. One of the most famous examples involves the so-called *time dilation effect* of special relativity. To illustrate, suppose that two twins–Alice and Bob–start off at rest with respect to each other on the surface of the Earth. Bob then blasts off in a rocket ship, traveling near the speed of light. In a situation like this, relativity tells us that Bob's clock will appear to be running slow from Alice's perspective (and that Bob will see Alice's clock as being sped up). Importantly, this will be true for both wristwatches *and* biological clocks. So, if Bob travels far enough away and then returns to Earth (maintaining his high speed throughout), he will be visibly younger than Alice–perhaps even years younger–when he arrives. In this respect, he is like a Wellsian time traveler. But Bob's youthful appearance would not be due to a discrepancy between personal time and "real" time. Rather, the two individuals will simply have taken different paths through spacetime, and Bob's interval will be shorter, due to a difference in acceleration. This gives us a mismatch between the "proper time" of the two individuals, but it does not create a discrepancy with external time. Hence, this would not be a case of Lewisian time travel.

It is unclear what Lewis should say about this example. He could try to revise his definition, or he could maintain his account and insist that Bob's journey is not genuine time travel. More plausibly, he could simply say that there is semantic indeterminacy in this case, owing to our limited and sometimes conflicting talk of time travel. Since nothing important hinges on the issue, we can live with that indeterminacy.[12]

[12] For a classic statement of Lewis's view on linguistic use and semantic indeterminacy, see Lewis (1986a, 212). For more on relativity and time travel (including a discussion of general relativity), see Wasserman (2018, Chapter 2, sections 5 and 6).

2. The Problems of Backward Causation

If time travel involves discrepancies between personal time and external time, and the former is understood in terms of causation, then backward time travel would reverse causal sequences. This idea raises several concerns. Here, I will focus on three issues from Lewis's paper: his proposed analysis of causation, his attitude toward causal loops, and his discussion of self-undermining acts.

In his earlier work, Lewis (1973) argues that one event is a cause of another just in case those events are linked by a chain of counterfactual dependence (that is, a series of events in which each member would not have occurred without the previous event). For example, the engine wouldn't have started without a signal from the ignition, and the ignition wouldn't have signaled if I hadn't turned the key. So, my turning the key is a cause of the engine's starting.

To apply this to the case of time travel, we must first say a bit about Lewis's theory of counterfactuals. In general, a counterfactual of the form "A > C" is true just in case all the closest worlds where A is the case are also worlds where C is the case. Closeness, in turn, is taken to be a similarity relation governed by the following set of rules:

(1) It is of first importance to avoid big, widespread, diverse violations of law.
(2) It is of second importance to maximize the spatiotemporal region throughout which perfect match of particular fact prevails.
(3) It is of third importance to avoid even small, localized, simple violations of law.
(4) It is of little or no importance to secure approximate similarity of particular fact, even in matters that concern us greatly. (1986b, 47–48)

To illustrate, let w0 be the world in which Marty gets into his time machine, hits 88 mph, and appears in 1955. Let w1 be a world very much like w0, but where–contrary to the laws of w0–a few extra neurons fire in Marty's brain after he gets into the time machine and he chooses *not* to hit 88 mph. In that world, no one shows up in 1955. Finally, let w2 be a world like w1, but where a perfect duplicate of Marty miraculously shows up out of nowhere in 1955 (at the exact time and place of Marty's actual appearance in

w0). Because of that miraculous appearance, w2 will match w0 up until the time at which the extra neurons fire in 1985. So, w2 does a better job than w1 with respect to (2). But we can suppose that the sudden appearance of Marty's duplicate in that world would count as a "large miracle," relative to w0, since it would require billions of atoms to appear in a highly organized way–thus breaking Lewis's very first rule. If this holds for all worlds where a duplicate appears, it follows that no one would have appeared in 1955 had Marty not hit 88 mph. Lewis's theory of causation, therefore, says that Marty's hitting 88 mph was a cause of the appearance in 1955. For this reason, Lewis concludes that his theory of causation does not "rule out causal reversal *a priori*" (1976, 148).

This much is correct, but there is also a worry. Suppose that Marty's friend, Doc Brown, has a teleporter that he uses to send a single electron (all by itself) back to 1955. We can then ask: what would have happened had Doc not turned on his device? Let w3 and w4 both be worlds very much like w0, but where a few extra neurons miraculously fire in Doc's brain so that he does not turn on the machine. In w3, no extra electron appears in 1955; in w4, one does (at the exact time and place that one appears in w0). As in the previous case, the second world will score better with respect to (2), since it includes 30 extra years of exact match. However, unlike the case of Marty, this would *not* require a violation of (1), since a single electron's appearance would not constitute a big, widespread, diverse violation of law. That world *will* score slightly lower with respect to (3), since it will include one additional small miracle (corresponding to the electron's appearance in 1955). But this will not outweigh the decisive win on (2). Hence, Lewis will have to say that an electron would have appeared in 1955, even if Doc had *not* turned on the machine. Given his theory of causation, it would follow that the electron's appearance is not caused by future events. Assuming that causation is required for time travel, it would follow that the relevant electron is *not* a visitor from the future. Hence, Lewis's account seems to implausibly discriminate between would-be time travelers on the basis of their size.[13]

Even if one believes that causal reversals are not ruled out by the nature of causation, one might still worry about the phenomenon for other reasons.

[13] See Wasserman (2015) for further discussion.

One of the most common concerns is about the possibility of *causal loops*. Lewis gives the example of a traveler who goes back in time, meets his younger self, and passes on information about how to build a time machine (where this information is not available any other way). The younger self then uses that information to build a time machine, travel back in time, and meet his younger self. "But where did the information come from in the first place?" asked Lewis. "Why did the whole affair happen? There is simply no answer. The parts of the loop are explicable, the whole of it is not" (ibid., 149).[14]

Lewis's treatment of this case has always struck me as correct:

> Strange! But not impossible, and not too different from inexplicabilities we are already inured to. Almost everyone agrees that God, or the Big Bang, or the entire infinite past of the universe, or the decay of a tritium atom, is uncaused and inexplicable. Then if these are possible, why not also the inexplicable causal loops that arise in time travel? (Ibid., 149)

To this, I would add that a question like "Where did the information come from in the first place?" is ambiguous, since "first place" could refer to an initial point in external or personal time. On the first reading, the question has a perfectly good answer, since the information first shows up (in external time) with the older time traveler from the future. On the second reading, the question presupposes a false assumption, namely that there is a first moment in the information's personal history. From the information's perspective, there is no such time–it came from the older traveler, who got it from the younger man, who got it from the older traveler, who… In this respect, the information is like the "entire infinite past of the universe" that Lewis mentions in the quote above: it does not have a causal beginning.

[14] "Information loops" of this kind can be contrasted with "object loops." Suppose, for example, that Lewis's time traveler gives his younger self a set of printed out plans, which the younger self uses to build a time machine so that he can take back the very same sheets of paper to his younger self. Loops like this will require an object to revert back to its original state when traveling in time (e.g., coffee stains will have to disap- pear, torn corners will need to be replaced, etc.). This kind of process would be highly unusual, but not impossible. See Hanley (2004, section 4).

But, in that case, it seems strange to ask for a cause. If the series never started, why require a starter?[15]

Even if we accept the possibility of a causal loop (in which each event is a cause of the next and where the last is a cause of the first), we might still worry about the possibility of a *self-undermining series* (where each event is a cause of the next and the last is a *preventer* of the first). The standard example from science fiction literature is the traveler who goes back in time and kills her grandfather before her father is conceived. This kind of sequence would seem impossible, since removing the cause removes the effect, which, in this case, would mean removing the remover herself. *Perhaps* we could make sense of this by using parallel universes or rewritable histories (see above).[16] But Lewis rejects these things and therefore owes us a different explanation of what is going on in this case.

3. The Grandfather Paradox

Lewis's discussion of the grandfather paradox focuses on the following story:[17]

> Consider Tim. He detests his grandfather, whose success in the munitions trade built the family fortune that paid for Tim's time machine. Tim would like nothing so much as to kill Grandfather, but alas he is too late. Grandfather died in his bed in 1957, while Tim was a young boy. But when Tim has built his time machine and traveled to 1920, suddenly he realizes that he is not too late after all. He buys a rifle; he spends long hours in target practice; he shadows Grandfather to learn the route of his daily walk to the munitions works; he rents a room along the route; and there

[15] For more on this point (and causal loops in general), see Wasserman (2018, Chapter 5).

[16] For a fuller account of the possibilities, see Wasserman (2018, Chapter 3).

[17] We will focus on the classic version of the paradox, but Lewis also discussed the "epistemic" version of the puzzle in his early seminars. On the latter kind of argument, see Horwich (1987, Chapter 7), Smith (1997), and Wasserman (2018, Chapter 4, section 3).

> he lurks, one winter day in 1921, rifle loaded, hate in his heart, as Grandfather walks closer, closer… (Ibid., 149)

We may not know how the story turns out, but we do know that Tim will not succeed since Lewis stipulates that Grandfather died in 1957, rather than 1921. Still, it seems as if Tim has *the ability* to kill Grandfather:

> Conditions are perfect in every way: the best rifle money could buy, Grandfather an easy target only twenty yards away, not a breeze, door securely locked against intruders. Tim a good shot to begin with and now at the peak of training, and so on. What's to stop him? The forces of logic will not stay his hand! No powerful chaperone stands by to defend the past from interference… In short, Tim is as much able to kill Grandfather as anyone ever is to kill anyone. (Ibid.)

Yet there are equally good reasons for thinking that Tim *cannot* kill Grandfather:

> No Grandfather, no Father; no Father, no Tim; no Tim, no killing. And for good measure: no Grandfather, no family fortune; no fortune, no time machine; no time machine, no killing. So the supposition that Tim killed Grandfather seems impossible. (Ibid., 152)

In this way, the possibility of time travel seems to yield a contradiction: Tim both can and cannot commit *retrogrampatricide*. For Lewis, this is the central problem posed by the grandfather paradox.

His response to this puzzle goes as follows:

> To say that something can happen means that its happening is compossible with certain facts… Tim's killing Grandfather that day in 1921 is compossible with a fairly rich set of facts: the facts about his rifle, his skill and training… and so on… But his killing Grandfather is not compossible with another, more inclusive set of facts. There is the simple fact that Grandfather was not killed… You can reasonably choose the narrower delineation, and say that he can; or the wider delineation, and say that he can't. But choose. What you mustn't do is waver, say in the same breath that he both can and can't, and then claim that this contradiction proves that time travel is impossible. (Ibid., 150–51)

For Lewis, the proponent of the grandfather paradox is like someone who believes that her friend holds contradictory beliefs because she has heard him describe average Americans as both "rich" and "not rich." However, if the first of these descriptions occurred during a discussion of global poverty and the second happened in a conversation with global billionaires, then this inference would obviously be invalid–the two statements would not express contradictory beliefs since "rich" would mean different things on those occasions. In the same way, Lewis suggests that "can" means one thing when we are considering facts about Tim's skills and training, for example, and something else when we are focused on facts about his personal history (e.g., the fact that Grandfather was not killed in 1921). As a result, there are some contexts in which the sentence "Tim can kill Grandfather" would express a truth and others in which it would express a falsehood. But there is *no* context in which it would express *both* a truth and a falsehood, so there is no contradiction.[18]

Lewis is clearly correct about the context-sensitivity of "can." Indeed, "The Paradoxes of Time Travel"–along with his "General Semantics" (1970) and "Scorekeeping in a Language Game" (1979)–helped launch contemporary research into contextualism.[19] However, the above observations do *not* solve the grandfather paradox since we can state the problem using a single sense of "can."

To begin, consider the notion of ability that is typically at issue when undergoing deliberation. In such contexts, we usually assume that some things are genuine options for us and others are not. Walking across the room is an option for me right now; walking across the ocean is not. Arguably, this distinction is closely linked to our evaluation of agents. For example, if I promised you that I would walk across the room and did not, then you could justifiably hold me responsible for that failure. In contrast,

[18] Similar proposals were made, around the same time, by Fitzgerald (1974, 539–40), Horwich (1975, 435–37), and Thom (1975). However, Lewis's account has proven the most popular, having been endorsed by Dowe (2000, 448–50), Sider (2002), and many others.

[19] This includes contextualism in epistemology, where Lewis's "Elusive Knowledge" (1996) has played a major role. Among other things, that paper develops a response to skeptical arguments which mirrors his solution to the grandfather paradox.

it would make no sense for you to blame me for not walking across the ocean–even if I promised to–since that was never an option for me.

Kadri Vihvelin (1996, 2020) argues that this "agential" sense of "can" is what we often have in mind when discussing abilities. After all, we normally assume that something is an option for us only if we would (or at least might) succeed in performing the action, were we to try. But Tim would fail to kill Grandfather, no matter what he tried, since Grandfather's survival is causally necessary (in those circumstances) for Tim to do anything at all. Hence, Tim is not able to kill Grandfather in the ordinary sense of "able." For this reason, Vihvelin rejects Lewis's claim that "By any ordinary standards of ability, Tim can kill Grandfather" (ibid., 150). At the same time, she *agrees* with Lewis that backward time travel is possible. Her view is simply that time travelers are limited in ways that non-time-travelers are not.

I agree with Vihvelin that Tim cannot kill Grandfather in the ordinary, agential sense. But I still find the case puzzling. To see why, consider the following point from Lewis:

> Suppose that down the street another sniper, Tom, lurks waiting for another victim, Grandfather's partner. Tom is not a time traveler, but otherwise he is just like Tim: same make of rifle, same murderous intent, same everything. We can even suppose that Tom, like Tim, believes himself to be a time traveler. Someone has gone to a lot of trouble to deceive Tom into thinking so. There's no doubt that Tom can kill his victim; and Tim has everything going for him that Tom does. By any ordinary standards of ability, Tim can kill Grandfather. (Ibid., 149)

Lewis's analogy is compelling because Tim and Tom are subject to the exact same constraints–each experiences the same gravitational effect, the same psychological pressures, and so on. Moreover, neither of these agents is bound by shackles, subject to coercive threats, or watched over by chaperones from the future. Of course, there is one important difference between the two since Tim–unlike Tom–is ancestrally related to his target. But it is not as if the identity of the target interferes with Tim's decision-making or binds his arms to the wall. We are therefore left with the question: *what is holding him back?*

In my view, this is one of the most interesting puzzles surrounding the grandfather paradox. Here is one potential solution: the preceding argument assumes that causal constraints like threats and bindings are the only things that could limit an agent's abilities. But that seems like a mistake. I am not able to gain knowledge of a proposition without believing it, for example, but that is not due to causal forces. Rather, part of what it is to know something is to believe it, and believing a proposition is logically incompatible with *not* believing it. Hence, my lack of belief in a given proposition prevents me from knowing it by being constitutively incompatible with such knowledge. Perhaps something similar is true for Tim: part of what it is to have an ability (in the ordinary, agential sense) is for the performance of that action to be compatible with the agent's own causal history.[20] Since Tim's failure to kill Grandfather is part of his causal past, it is part of something constitutively incompatible with his success. That is how it renders him unable to kill Grandfather.[21]

4. Conclusion

In this paper, I have outlined some of the central moves in "The Paradoxes of Time Travel" and reviewed some of what can be said both for and against them. But one of the wonderful things about Lewis's work is that it continues to inspire new ideas about a variety of topics–including ones which lie outside the scope of the original paper. In just the last few years, for example, "Paradoxes" has been cited in works on the Christian doctrine of original guilt, the Aristotelian approach to understanding powers, the Molinist view of divine providence, the traditional debate over freedom and determinism, and various issues relating to ethics and moral responsibility.[22]

[20] See Rea (2015). For clarifications and further refinements, see Law and Wasserman (2020).

[21] For other discussions of Lewis's response to the grandfather paradox, see Wasserman (2018: Chapter 4), Garrett (2019), Effingham (2020: Chapter 10), Fernandes (2020), and Loewenstein (2022).

[22] See, respectively, Effingham (2022), Giannini and Donati (2023), Law (2024), Wasserman (2024), Bernstein (Forthcoming), and Cyr and Tognazzini (2024).

The diversity of these topics is a testament to the wide-ranging influence of Lewis's work, and I fully expect that influence to continue for the next 50 years and beyond.[23]

References

Asimov, Isaac. 1955. *The End of Eternity*. Garden City, New York: Doubleday.

Bernstein, Sarah. (Forthcoming.) "Ethical Puzzles of Time Travel". In *The Routledge Companion to Philosophy of Time*, edited by Nina Emery. New York: Routledge.

Cyr, Taylor, and Neal Tognazzini. 2024. "What Time Travel Teaches Us about Moral Responsibility." *Journal of Ethics and Social Philosophy*, 26: 523–46.

Daniels, David R. 1935. "The Branches of Time." *Wonder Stories*, 7: 294–303.

Deutsch, David. 1997. *The Fabric of Reality: The Science of Parallel Universes–and Its Implications*. New York: Penguin.

Dowe, Phil. 2000. "The Case for Time Travel." *Philosophy*, 75: 441–51.

Effingham, Nikk. 2020. *Time Travel: Probability and Impossibility*. Oxford: Oxford University Press.

Effingham, Nikk. 2022. "Fission Theories of Original Guilt." *International Journal for Philosophy of Religion*, 92: 15–30.

Fernandes, Alison. 2020. "Freedom, Self-Prediction, and the Possibility of Time Travel." *Philosophical Studies*, 177: 89–108.

Fitzgerald, Paul. 1974. "On Retrocausality." *Philosophia*, 4: 513–51.

Garrett, Brian. 2019. "Bulletproof Grandfathers, David Lewis, and 'Can't'-Judgments.'" *Acta Analytica*, 34: 177–80.

Giannini, Giacomo, and Donatella Donati. 2023. "Dispositionalism's (Grand)daddy Issues: Time Travelling and Perfect Masks." *Analysis*, 83: 40–49.

Goddu, G.C. 2003. "Time Travel and Changing the Past: (Or How to Kill Yourself and Live to Tell the Tale)." *Ratio*, 16: 16–32.

Hanley, Richard. (2004). "No End in Sight: Causal Loops in Philosophy, Physics and Fiction." *Synthese*, 141: 123–52.

Horwich, Paul. 1975. "On Some Alleged Paradoxes of Time Travel." *Journal of Philosophy*, 72: 432–44.

Horwich, Paul. 1987. *Asymmetries in Time: Problems in the Philosophy of Science*. Cambridge, Massachusetts: MIT Press.

Law, Andrew. 2018. "The Puzzle of Hyper-change." *Ratio*, 32: 1–11.

[23] I thank Brian Garrett for comments on earlier versions of this paper.

Law, Andrew. 2024. "If Molinism is True, What can you Do?" *International Journal for Philosophy of Religion*, 95: 307–22.

Law, Andrew, and Ryan Wasserman. 2022. "Lessons from Grandfather." *Philosophies*, 7: 1–10.

Lewis, David K. 1970. "General Semantics." *Synthese*, 22: 18–67.

Lewis, David K. 1973. "Causation." *Journal of Philosophy*, 70: 556–67.

Lewis, David K. 1976. "The Paradoxes of Time Travel." *American Philosophical Quarterly*, 13: 145–52.

Lewis, David K. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic*, 8: 339–59.

Lewis, David K. 1983. Survival and Identity. Reprinted with postscripts in *Philosophical Papers: Volume I*. Oxford: Oxford University Press.

Lewis, David K. 1986a. *On the Plurality of Worlds*. Oxford: Oxford University Press

Lewis, David K. 1986b. Counterfactual Dependence and Time's Arrow. Reprinted with postscripts in *Philosophical Papers: Volume II*. Oxford: Oxford University Press.

Lewis, David K. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy*, 74: 549–67.

Lewis, David K. 2023. *Philosophical Manuscripts*, edited by Frederique Janssen-Lauret and Fraser MacBride. Oxford: Oxford University Press.

Loewenstein, Yael. 2022. "Against the Standard Solution to the Grandfather Paradox." *Synthese*, 200, 172.

Meiland, J. W. 1974. "A Two-Dimensional Passage Model of Time for Time Travel." *Philosophical Studies*, 26: 153–73.

Rea, Michael. 2015. "Time Travelers Are Not Free." *Journal of Philosophy*, 112: 266–79.

Sider, Theodore. 2002. "Time Travel, Coincidences and Counterfactuals." *Philosophical Studies*, 110: 115–38.

Smith, Nicholas J. J. 1997. "Bananas Enough for Time Travel?" *British Journal for the Philosophy of Science*, 48: 363–89.

Thom, Paul. 1975. "Time Travel and Non-Fatal Suicide." *Philosophical Studies*, 27: 211–16.

van Inwagen, Peter. 2010. "Changing the Past". In *Oxford Studies in Metaphysics, Volume 5*, edited by Dean Zimmerman. Oxford: Oxford University Press.

Vihvelin, Kadri. 1996. "What Time Travelers Cannot Do." *Philosophical Studies*, 81: 315–30.

Vihvelin, Kadri. 2020. "Killing Time Again." *The Monist*, 103: 312–27.

Wasserman, Ryan. 2015. "Lewis on Backward Causation." *Thought*, 4: 141–50.

Wasserman, Ryan. 2018. *Paradoxes of Time Travel*. Oxford: Oxford University Press.

Wasserman, Ryan. 2024. "Time Travel, Freedom, and Incompatibilism." *Erkenntnis*, 89: 2953–66.