

Research Bank

Journal article

The end of mystery

Baron, Sam and Colyvan, Mark

This is the accepted manuscript version. For the publisher's version please see:

Baron, S. and Colyvan, M. (2019). The end of mystery. *American Philosophical Quarterly*, 56(3), pp. 247-264. <https://doi.org/10.2307/48570634>

The End of Mystery

Abstract

Tim travels back in time and tries to kill his grandfather before his father was born. Tim fails. But why? Lewis's response was to cite 'coincidences': Tim is the unlucky subject of gun jammings, banana peels, sudden changes of heart and so on. A number of challenges have been raised against Lewis's response. The latest of these focuses on explanation. This paper diagnoses the source of this new disgruntlement and offers an alternative explanation for Tim's failure, one that Lewis would not have liked. The explanation is an obvious one, but controversial and so it is defended against all the objections that can be mustered.

Keywords: Time · Time Travel · Explanation · Logic · Paradox

1. Introduction

Tim the time traveler travels back in time, hell-bent on murdering his grandfather before his father was conceived. As we learnt at mother's knee: Tim fails. But why? Lewis's (1976) (in)famous response was to cite 'coincidences': Tim is the unlucky subject of gun jammings, accidents with banana peels, sudden changes of heart and so on. A number of challenges have been raised against Lewis's response. The latest of these focuses on explanation. In this paper we diagnose the source of this new disgruntlement and offer an alternative explanation for Tim's failure— one that Lewis would not have liked. Very roughly, Tim fails because contradictions are impossible. We think that the explanation is an obvious one, but controversial so we defend it against all of the objections we can muster. We begin, in § 1, with Tim's tale along with Lewis's account of it, and the perceived explanatory woes of that account before, in § 2, diagnosing the source of these woes and offering an alternative explanation. In § 3 we consider objections.

2. The Grandfather Paradox

Why does Tim fail to kill his grandfather? According to Lewis (1976),

Tim fails to kill his grandfather for some ‘commonplace’ reason. Tim’s gun might jam. Or he might slip on a banana peel. Or he might have a sudden change of heart. Or he might get distracted at the last minute or ... And so it goes. Call the explanation that Tim fails because of some commonplace reason: the *Lewis* explanation. The Lewis explanation for Tim’s failure has long worried philosophers. Resistance to the Lewis explanation comes in two forms. First, some deem the Lewis explanation to be unsatisfactory on probabilistic grounds (see Horwich (1987; 1975)). This concern, however, has been adequately answered (cf. Smith 1997). So set it aside. Second, some deem the Lewis explanation to be unsatisfactory on explanatory grounds. The Lewis explanation does not, some have argued, constitute a complete explanation; the case of Tim the time traveller remains mysterious. Arntzenius and Maudlin (2002, p. 180) express this sentiment when they write of Tim’s failed schemes to kill his grandfather:

One worry is the question as to **why** such schemes always fail. Doesn’t the necessity of such failures put *prima facie* unusual and unexpected constraints on the actions of people, or objects, that have travelled in time? Don’t we have good reason to believe that there are no such constraints (in our world) and thus that there is no time travel (in our world)? [our emphasis]

Smith (2017, p. 157) does a nice job of outlining the problem:

The worry is as follows. The would-be autoinfanticidal time traveller is attempting to do something impossible ... We accept that she will not succeed. We also accept that what will stop her succeeding is a succession of commonplace occurrences ... We are bracketing any worry that such a succession is *improbable* ... Yet still there is a problem (so the worry goes). The problem is that the exclusion of the time traveller from successfully committing autoinfanticide seems mysteriously inexplicable. Each particular event that foils the time traveller is explicable in a perfectly ordinary way; but the *inevitable combination* of these events amounts to a ring-fencing of the forbidden zone of autoinfanticide—and this ring-fencing is mystifying. It’s like a grand conspiracy to stop the time traveller doing what she wants to do—and

yet there are no conspirators: no time lords, no magical forces of logic. This is profoundly perplexing.

Call the worry that Lewis's explanation is incomplete: *the explanation problem for backwards time travel*. We will have more to say about Smith's solution to the explanation problem in due course. For now, we wish to focus on the problem itself. As Smith shows, the explanation problem can be traced to the work of a number of Lewis's critics. In addition to Arntzenius and Maudlin, Dowe (2007, p. 724), Gorovitz (1964, pp. 366–367), Horwich (1987, pp. 119–121), Riggs (1997, p. 52), Ismael (2003, p. 308) and Carroll (2010, p. 86) all consider some version of the explanation problem. The problem is compelling. Tim fails to kill his grandfather because his gun jams? Surely there's more to the story than that!

Such an explanation is unsatisfactory in three respects. First, the explanation is not particularly unified. Suppose Tim tries to kill his grandfather repeatedly. He fails but each time for a different reason: his gun jams, he is distracted at the last moment, he slips on a banana peel, he has a sudden change of heart. All of these events are utterly disparate. And yet the repeated failure is not as disjunctive as its explanation. Second, the explanans and the explanandum have very different modal profiles. There is something inevitable about Tim's failure, but there is nothing inevitable about gun jammings, distractions and changes of heart. Such events could easily have been otherwise; but not so for Tim's failure. Third, the Lewis explanation does not yield much by way of understanding. Sure, Tim fails because he slips on a banana peel, but merely citing this fact gives us no 'Aha!' moment.¹

So the Lewis explanation leaves us with a sense of mystery. Shortly we will offer an alternative to the Lewis explanation or, rather, a supplement to it. First, however, we believe that the explanatory structure of the grandfather paradox needs to be unpacked. It is imperative to distinguish three explananda that are apt to be conflated. Here they are:

[EXP1] Tim fails to kill grandfather.

[EXP2] All of Tim's actual attempts result in failure.

[EXP3] Tim is doomed to fail in his attempts to kill grandfather.²

A bit about the difference between these two explananda: [EXP1] concerns a particular event in which Tim fails to kill his grandfather. So, for instance, suppose that on the 12th November 1955 at 2:00pm, Tim takes aim at his grandfather, pulls the trigger and fails to make his mark. Why? What exactly is it that happens on the 12th November 1955 at 2:00pm? That's the kind of question that [EXP1] raises. [EXP2] by contrast, are not about any event in particular. It is, rather, about the manner in which, no matter how hard he tries, Tim fails to kill his grandfather. All of his attempts are stymied. [EXP2], then, is about the sum-total of Tim's failures; it is about Tim's life story. Why, we might wonder, is Tim's life so beset by defeat? What accounts for the string of fails? [EXP3] is about the possibility of what Tim is trying to do. [EXP3] is not just about Tim's life story, but about any way that Tim's life story might have gone. No matter how Tim might possibly try to kill grandfather, Tim must fail. What accounts for this modal fact?

To highlight the difference between [EXP1], [EXP2] and [EXP3] it is useful to consider a different case altogether. In Königsberg there were once seven bridges connecting a series of islands. Suppose that Bridget attempts to cross all seven bridges, crossing over each bridge exactly once. She fails to do so. Suppose she tries again and again and she fails each time. Why? As with Tim the time traveler, there are two explananda that demand attention:

[EXP4] Bridget fails to cross the Seven Bridges of Königsberg.

[EXP5] Bridget fails to cross the Seven Bridges of Königsberg repeatedly.

[EXP6] Bridget is doomed to fail in her attempts to cross the Seven Bridges of Königsberg.

As with [EXP1], [EXP4] is about a particular event. So, for instance, suppose that on the 12th November 1755 at 2:00pm, Bridget attempts her

crossing and is foiled. Why? What exactly is it that happens on the 12th November 1755 at 2:00pm? [EXP5] like [EXP2] is not about any event in particular. It is, rather, about the manner in which, no matter how hard she tries, Bridget fails to cross the Seven Bridges of Königsberg. [EXP4] is about Bridget's bridge crossings taken together. Why is Bridget, like Tim, such an actual failure? That's what we want to know. [EXP6] is about any possible attempt that Bridget might make. No matter how she actually tries to cross the seven bridges, any possible alternative will also meet with failure. Why? What accounts for the fact that she must fail, no matter what she tries to do?

[EXP4], [EXP5] and [EXP6] call for different kinds of explanation. [EXP4] calls for a causal explanation: on the 12th November 1755 at 2:00pm Bridget fails to cross the Seven Bridges because she becomes distracted and accidentally doubles-back over a bridge she forgot she had already crossed. [EXP5] calls for an explanation of the string of causal histories corresponding to Bridget's attempts; it calls for an explanation that ties together these actual attempts under a common banner and explain why she always fails. [EXP6] calls for an explanation that abstracts away from the actual causal details entirely.

The abstracted explanation for [EXP6] is mathematical in nature, and lies in graph theory. If we treat each of the seven bridges as an edge, and each of the land masses as vertices, the seven bridges can be treated as a connected graph. The resulting graph is non-Eulerian, which means that it provably lacks both an Eulerian path and an Eulerian circuit. An Eulerian path is just a continuous path through a graph that passes over each edge exactly once; an Eulerian circuit is a continuous path through a graph that passes over each edge exactly once, starting and ending at the same vertex. The lack of Eulerian paths explains [EXP6]: there just is no successful crossing available to Bridget so she *must* fail. The lack of Eulerian paths also explains [EXP5]: Bridget repeatedly fails to cross the seven bridges because there is no possible way of succeeding. And, of course, this impossibility also explains [EXP4].

As with [EXP4], [EXP5] and [EXP6], [EXP1], [EXP2] and [EXP3]

demand different kinds of explanation. [EXP1], like [EXP4] calls for a causal explanation: on the 12th November 1955 at 2:00pm, Tim fails to kill his grandfather because his gun jams. [EXP2], like [EXP5], calls for an explanation that sets aside any particular gun jammings, changes of heart and so on, and tells a general story about Tim's actual failures. [EXP3], like [EXP6], requires an explanation that abstracts away from Tim's actual attempts, either taken individually or together, and tells a general story about why, no matter how Tim tries, no matter what possible attempt he makes on grandfather's life, he must fail.

Before we outline the correct explanation for [EXP3], it is worth pausing to consider how conflating [EXP1], [EXP2] and [EXP3] might bring discomfort. Suppose one is perplexed about Tim the time traveler and asks: "Why does Tim fail to kill grandfather?" and we give back the answer: because Tim's gun jams. In giving this answer, we have understood the question being asked to be a causal why-question. We took the question to be asking after [EXP1], and so we answered accordingly. If that's what was really being asked, if all one wanted to know is what gets between Tim and homicide *on a particular occasion, in the past, when Tim fails* then one ought to be satisfied. The question was asked and answered. One has the explanatory goods, so to speak.

If, however, one was *really* asking after [EXP2] or [EXP3], then the answer we have provided will seem utterly unsatisfying. Suppose one is asking after [EXP2]. What is sought is an explanation for why Tim repeatedly fails. We have explained why he fails on a particular occasion. That does nothing to explain the string of failures to which he succumbs. Similarly, suppose one is asking after [EXP3]. What is sought is an explanation for why Tim was doomed to fail. We offered the causal story about his failure. But that causal story does nothing to explain the inevitability of Tim's failure. Indeed, when the explanation for [EXP1] is offered up as an explanation for either [EXP2] or [EXP3], that explanation should seem unsatisfactory in exactly the respects discussed above. An account of the causal history of a particular event does nothing to unify Tim's various failures, and it is a unified account of why he always fails

that one ultimately wants when asking after [EXP2]. Similarly, an account of the causal history of a particular event is out-of-step with the modal force of the ‘why’ question one is asking when asking after [EXP3]. It was the inevitability of Tim’s failure that was asked after. We responded with a contingent could-easily-have-been-otherwise causal story. No wonder the answer isn't mollifying. Finally, the causal story does little to help one understand [EXP2] or [EXP3]. To be sure, it helps one understand why Tim fails on a particular occasion, but no understanding of why he always fails or why he is doomed to fail has been given (apart from the conjunction of all the particular causal stories).

The explanation problem arises, we contend, because the Lewis explanation answers at best one ‘why’ question. It answers a ‘why’ question posed at [EXP1]. But that why question is of little interest. It is [EXP2] and [EXP3] that really matter. The Lewis explanation can therefore seem unsatisfying in one of two important ways. First, one might mistakenly take the Lewis explanation to be an explanation for either [EXP2] or [EXP3], a task to which it is unsuited. Second, one might correctly take the Lewis explanation to be an explanation for [EXP1] but maintain—and we think rightly—that the important fact in need of explanation is [EXP3]. Again, we can look to the Seven Bridges of Königsberg for guidance. Exactly the same confusion would arise in that case if the causal explanation were mistakenly treated as an explanation for [EXP5] or [EXP6]. Imagine that someone claims to have solved the Bridges of Königsberg case once and for all. Why does Bridget repeatedly fail to cross the seven bridges? And why is she doomed to fail? For some commonplace reason, is the answer offered. She keeps doubling back over bridges, or losing her concentration. Once she even slipped on a banana peel and fell headlong into the canals. Such a common-place solution is clearly missing something. Euler and the graph theory explanation has it right. But if we didn’t have Euler to guide us, and we were confused about the particular explanatory question being asked, we might well think that there’s something to this common-place account.

So what then *is* the explanation for [EXP3]? The answer is one that Lewis

would not have liked: logic. Tim is doomed to fail in his attempts to kill grandfather because any successful attempt would result in a logical contradiction, and contradictions are impossible given the law of non-contradiction. In short, just as mathematical facts are needed to explain why Bridget cannot complete the bridge walk, logical facts are needed to explain why Tim's homicidal tendencies must go unfulfilled. Moreover, just as the mathematical explanation for [EXP6] also explains [EXP5], so too does the logical explanation for [EXP3] also explain [EXP2]. It is because Tim must fail that he fails in all of his attempts actually. The comparison with the mathematical case is no accident. In both cases, the explanation operates in a very similar fashion, via abstraction: the particular causal details of the scenario are set aside in order to focus on the broad structure of the two cases. In one situation, the structural facts most salient are mathematical ones, in the other case the facts are logical. In both cases, though, abstraction is necessary for the explanatory task at hand.^{3,4}

We say that Lewis would not like this explanation because he seems to rule it out. Here's Lewis (1976, p. 149):

Tim can kill Grandfather. He has what it takes. 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.

To be clear: we are not suggesting that there is any logical chaperone.⁵ But we do believe that, in a certain (non-causal) sense, the forces of logic stay Tim's hand. It is because contradictions are impossible that Tim is doomed to failure. Of course we don't mean that the law of non-contradiction *causes* Tim to fail, and we don't think that logic is a *thing* and that the thing gets in the way. Our claim is just that the best, systematic explanation for why Tim is doomed to fail in his attempts to kill grandfather is one that appeals to logical facts in the explanans.

Our response is unpopular, it seems. Smith (2017) has, to some degree, argued against it, and we can muster a few objections against the response ourselves. Before addressing objections, however, it is useful to pause and consider the power of the logical explanation we are proposing. First, the logical explanation provides the sought-after unity that we described above. Tim's various failures are unified in virtue of what they are attempts to do: they are all attempts to do something impossible. Second, the logical explanation is modally robust: why is Tim doomed to fail? Because success is impossible. The modal strength of the explanans and the explanandum are, finally, in accord. Third, the logical explanation yields understanding: once we learn that success would be impossible, we just *get* why Tim must fail whenever he tries: "Aha!", we say, "but of course!"

3. Objections

As we said, there are objections.

3.1. Nothing New!

Here's the first objection: what we have said is nothing new! Lewis's account of Tim the time traveler appeals to consistency in a number of places. At best, then, we have simply restated Lewis's position.

Even if this is correct, we believe that the restatement is useful and important in two respects. First, we have offered a diagnosis of where the explanation problem goes wrong, and a solution to that problem. Maybe the solution and the diagnosis were already there in Lewis. If so, then the present paper can be considered as a timely reminder that there is a perfectly adequate explanation available for why it is that Tim is doomed to fail. Second, Lewis was clearly dubious about the idea that the laws of logic have a role to play in accounting for Tim's failed attempts to kill his grandfather. So if the position we are defending is Lewis's then a further puzzle arises: what's Lewis playing at? Suppose he accepts that the correct explanation for Tim's failure is a logical one, then why cast shade on the idea that logic has anything to do with it?

The apparent tension in Lewis's thought can be explained by looking to Lewis's wider views about explanation. Lewis (1986) held the view that all explanation is causal. Accordingly, the only way that Lewis can see the laws of logic being explanatorily relevant is through the lens of causal reasoning. Through such a lens, the only way for the law of non-contradiction to do explanatory work is for it somehow cause Tim the time traveler's repeated failure. This way of thinking about the explanatory power of the laws is bizarre indeed. The central move we are making then, and the move that we deem to be novel is to give up the shackles of a causal theory of explanation and see the time travel case in the simplest of terms. Tim is doomed to fail because the law of non-contradiction is inviolable.⁶ That's an explanation, and it is a non-causal one. So on the one hand we agree with Lewis: the laws of logic do not stay Tim's hand, if 'to stay' is 'to prevent causally'. But we disagree in so far as we think that Tim's hand is stayed (in the relevant non-causal sense) by logic nonetheless; it is stayed because it must be, on pain of contradiction. And, of course, the specific means of staying his hand is via various common-place occurrences. The laws of logic work in concert with or, rather, *via* commonplace causal explanations.

One might seek to press the 'nothing new' objection a bit further. In a recent paper, Smith (2017, pp. 161–162) takes a line that is quite similar to the one that we are taking.

It should now be clear what to say about [the explanation problem]: it isn't a sound one! This is a case of the type in which no (further) explanation of failure is required. There are *no scenarios at all*—no points in logical space—satisfying the description 'a time traveler commits autoinfanticide'. There *is no* forbidden zone and hence, no need or even possibility of an explanation of why the time traveler does not enter 'it'. *Whatever happens* it won't be autoinfanticide because *no scenario at all* satisfies that description. The reason for this is that the description is self-contradictory ... So, the crucial point here is that there *is no* forbidden zone. This is completely different from saying that there is one, but 'laws of logic' prevent us from entering it... The key point is that the purported descriptions of scenarios in the forbidden zone

(involving autoinfanticide—or indeed *any* successful changing of the past) are all self-contradictory and describe *no sequence of events at all*. Hence, there just is no forbidden zone and thus no mysterious ring-fencing of ‘it’: there is *nothing to ring-fence*.

In order to understand Smith’s response to the explanation problem, let’s consider Smith’s idea of ring fencing. Suppose that, all of a sudden, an invisible barrier appears around Smith’s office. No one can get in or out. The office is, as it were, ring-fenced by a mysterious force. This ring-fencing stands in need of explanation. For it is clearly possible, in some broad sense, to get into the office. People have been in there before. All of a sudden, however, the possibility of entering the office is off-limits. This case differs sharply from the time-travel case. With respect to Smith and his office, there are points in logical space—namely ones where he is in his office—that are closed off. This corresponds to a ‘forbidden zone’ that Smith talks about in the quotation above. With respect to poor Tim and his travels through time, there are no points in logical space that are closed off. It is not possible for Tim to kill his grandfather. There just is no corresponding ‘forbidden zone’ of logical space.

This difference between the two cases makes all the difference for Smith. When there is a forbidden zone, simply pointing to the ring-fence surrounding the forbidden zone does nothing to alleviate the mystery. If Smith asks “why can’t I get into my office?” and we say “because you can’t!” No steps have been taken toward alleviating the mystery at hand. When there is no forbidden zone, however, pointing to the absence of such a zone is enough to remove all mystery. This is how Smith sees the time travel case: when we point to the fact that what Tim is trying to do is impossible, there is no longer anything mysterious about the fact that Tim is doomed to fail in his attempts to kill grandfather. As Smith notes, pointing to the fact that there is no forbidden zone is quite different to pointing to a forbidden zone—some set of points in logical space that are closed off—and remarking that laws of logic keep us from reaching those points in logical space.

But isn’t that what we’re saying? Not exactly. Our claim is not that there

is some forbidden zone of points in logical space that the laws of logic keep us from entering. At least, not if what we mean by ‘points in logical space’ is ‘logical possibilities’. Our claim is that there are no logical possibilities in which Tim succeeds in killing grandfather, because the law of non-contradiction is true. If, however, what one means by ‘points in logical space’ is something much broader, something that allows the relevant points to be impossibilities then we are guilty as charged: there absolutely is a forbidden zone that the laws of logic prevent us from entering, and that forbidden zone is the zone of impossibility. But we see nothing untoward about the laws of logic cleaving the possible from the impossible, and we see no reason why one cannot appeal to the shape of that ring-fence—the precise divide between the possible and the impossible—in the course of offering an explanation.

Ultimately, we are in agreement with Smith on the following crucial point: when we realize that what Tim is trying to do is impossible, the mystery surrounding Tim’s repeated failure to kill his grandfather dissolves. But we don’t find that at all surprising. Indeed, this is precisely what we should expect if what we have just done is offer a complete explanation of the fact that Tim is doomed to fail. We believe that Smith is offering the same account that we are offering, and that’s why we both see the dissolution of mystery. The difference, however, is that we are willing to call the appeal to the impossibility of what Tim is trying to do “an explanation”, whereas Smith (2017, p. 160) wants to leave this open. Whether or not this is an explanation, remarks Smith, depends on one’s preferred theory of explanation.

We will return to this difference between our view and Smith’s in a moment. First, it is useful to consider a second point of difference between Smith and us. While Smith is willing to leave it open that appealing to the impossibility of what Tim is doing amounts to an explanation, ultimately he does not think that an explanatory question aimed at [EXP3] is a question that should be asked and answered.

So ‘nothing’ and ‘commonplace occurrences’ are both, in a way, correct answers to the question ‘What stops him?’—but a much less misleading

response to the question is to point out that the question itself is out of place, that it really should not be asked at all. There simply *is nothing* for him to succeed at—there is no such thing as a scenario that satisfies the description ‘autoinfanticide’—and so there is no question as to why he fails to do ‘that’ ... I am suggesting that, rather than try to answer it, we should reject the question ‘Why did X not occur?’ (‘Why does the time traveller fail?’, ‘What stops him?’ etc.) as out of place when the X in question is autoinfanticide or some other example of changing the past... we should reject such questions when, respectively, it is impossible for Y not to occur or impossible for X to occur. (Smith, 2017, pp. 165–166)

Smith’s view is that ‘why’ questions about impossibility are not legitimate explanatory questions to ask. So while one *can* provide an answer to such a ‘why’ question, thereby (potentially) providing an explanation, one *shouldn’t* do so; the proper thing to do is to reject the initial question as being somehow out of place. Moreover, it is clear that Smith takes this to be the correct response in general when the phenomenon of interest is impossible. We disagree, both with respect to the general claim and with regard to its particular application to the time travel case.

We think there is a slippage in Smith’s reasoning between two sorts of cases. By differentiating them we will see more clearly that a why question aimed at [EXP3] is a reasonable explanatory question to ask. In the first case, one knows that it is impossible for Tim to kill his grandfather, because one knows that success would lead to a contradiction and contradictions are impossible. In the second case, one does not know that it is impossible for Tim to kill his grandfather, because one does not know that success would lead to a contradiction, though one does know that contradictions are impossible. In the second case, it is perfectly reasonable to ask why Tim is doomed to fail. Indeed, asking such a question seems like exactly the right thing to do in this situation. Denying that this is a reasonable question to ask is risky business. There are many cases where a particular outcome is impossible, but we can and should ask explanatory ‘why’ questions in order to reveal the impossibility at issue. A case in point is the Bridges of Königsberg case described above. But there are others. Indeed, in any case in which the laws of nature rule out a

particular outcome, it is reasonable to ask why the outcome at issue does not occur. And that's precisely because identifying what is and is not impossible is a very important part of explaining the world. In order to develop the best explanatory theories we can, we must limn the boundary between the possible and the impossible, and the only way to do that is to ask 'why' questions.

In the first case, by contrast, when one knows that Tim's success would lead to contradiction, and that contradictions are impossible, we agree that asking 'but *why* is Tim doomed to fail?' is a silly thing to do. But the silliness has nothing to do with the fact that what Tim is trying to do is impossible. The silliness is just the same silliness involved whenever one has a complete explanation before them of some phenomenon, and then stubbornly refuses to take that explanation seriously. For instance, suppose that someone throws a rock at a window and breaks it. The broken window is discovered, and the breakage is explained in terms of the thrown rock. If one were to ask, in the face of that information, 'but *why* did the window break?', then one has just failed to grasp the explanation at issue. In that situation, there is something untoward about re-asking the 'why' question, at least if one asks the question expecting there to be some further explanation available for the breaking of the window. The question is an unreasonable one, and it is unreasonable in exactly the same way as asking why Tim is doomed to fail, once one has the logical explanation for his failure in hand. In both cases one is either failing to understand the explanation, or failing to understand the game of giving and receiving explanations. But one can fail in these ways for any explanation.

This brings us back to the first point of difference between us and Smith: we do not think it should be left open whether an appeal to logic should count as an explanation. The account we have offered for why it is that Tim is doomed to fail appears to have the necessary features to be an explanation. First and foremost, it is an answer to a 'why' question. The why question at issue being, of course, 'why is Tim doomed to fail in his attempts to kill grandfather?' One might seek to press the objection by arguing that the account we have offered does not answer a *contrastive*

why question, and all explanatory ‘why’ questions are contrastive, of the form ‘why P and not $Q_1 \dots Q_n$?’ Where the Q_i designate some salient range of alternatives. On the contrary, the ‘why’ question that logic answers in the time travel case can be framed contrastively, as follows: ‘why is Tim doomed to fail in his attempts to kill grandfather rather than succeeding even once?’ The answer to this question is just that: Tim is doomed to fail and not succeed even once because if Tim were to succeed *even once*, then a contradiction would follow and contradictions are ruled out by the law of non-contradiction.

In addition to being an answer to a ‘why’ question, the account we have offered for why it is that Tim is doomed to fail removes mystery, as all good explanations should. Smith concedes as much: by pointing to the fact that Tim is trying to bring about a contradiction, and noting that contradictions are impossible, all mystery surrounding Tim’s repeated failures has been removed. There is nothing more to be said; a full explanatory account of [EXP3] has been given. The logical account of Tim’s failure also provides us with understanding of the situation, as any good explanation should. When we work out that what Tim is trying to do would yield a contradiction, and we reflect on the fact that contradictions are impossible, we fully understand why it is that Tim is doomed to fail. We can see that understanding is in fact gained by contrasting the logical account of Tim’s failure with the causal explanation of [EXP1]. That Tim’s gun jams does not help us to understand why he is doomed to fail. As soon as we build in the logical constraints and the inevitable contradiction associated with Tim’s success, we come to understand why [EXP2] is the case.

Finally, the account we have provided has two salient features that, for many philosophers, are deeply connected to explanation: unification and prediction. As already noted, appealing to the law of non-contradiction serves to unify all of Tim’s various failures to kill his grandfather under a common banner. They are all attempts to do something impossible; a commonality that cannot be seen by attending to the causal facts alone, since the sum total of failures are quite disparate with regard to their

causal profiles. The law of non-contradiction also has predictive power: it predicts that for any further attempts that Tim might make on grandfather's life, Tim will fail.

The account we have provided walks, talks and looks like an explanation; we think it is one. Still, one might remain unconvinced. In the next sub-section we therefore consider three objections against the idea that we have offered a genuine explanation of the fact that Tim is doomed to fail in his murderous efforts.⁷

3.2. Not an Explanation!

The first objection takes the form of a dilemma. Either the law of non-contradiction causes Tim to fail repeatedly or it does not. If it does, then what we have offered counts as an explanation, but the explanation is horrifying. The explanation is horrifying because, if true, then it really would seem as though there are powerful 'logical guardians' or 'magical laws of logic' that reach down into the world at the last moment, just before Tim is about to pull the trigger, and stay his hand. If the law of non-contradiction does not cause Tim's repeated failures, then it is not an explanation, because all explanations are causal.

We have already made it clear that we are averse to the idea of logical guardians. We don't think that the law of non-contradiction somehow causes anyone to do anything. The objection, however, trades on a false presupposition, namely that all explanations must be causal. There are some, such as Lewis, who have held this view (which, as already discussed, is why he wouldn't like the logical explanation we have proposed). But the idea that all explanation is causal has been the subject of sharp criticism in recent times. There has been a proliferation of examples of non-causal explanation within the philosophy of science and philosophy of mathematics literatures.⁸ Accordingly, without an argument for the view that all explanation is causal, the assumption simply begs the question against our position. Arguments in favor of the view that all explanation is causal are hard to come by. In so far as such arguments have been offered, they are arguments in favor of the view that all

explanation for the occurrence of particular events is causal explanation. Philosophers who put forward these arguments tend to concede that *regularity* explanations—explanations for why certain regularities appear—are likely to be non-causal explanations. But that is precisely the kind of explanation we take ourselves to be offering: we are confronted with a regularity, the repeated failure of Tim's attempts to kill his grandfather. The explanation for this regularity lies in logic.

The explanation we are offering, then, is similar to the explanation of why one law holds, in terms of another. For instance, one can explain a number of conservation laws by appealing to rotational symmetries within spacetime.⁹ These symmetries are part of the nomic structure of the geometric manifold that our universe is constituted by. It seems very strange, however, to say that the rotational symmetries cause the conservation laws to be thus and so. The explanation, then, is non-causal. Our explanation is like this: there is a regular happening in our universe, time travelers never kill their grandfathers. The explanation lies in logic: it is because of the law of non-contradiction that time-travel events are thus and so. But the law of non-contradiction no more causes Tim's failure than do the deep rotational symmetries of spacetime cause the conservation of mass/energy across the manifold.

The second objection focuses on the concepts of unification and prediction. According to Frost-Arnold (2010) explanations must be unifying and predictive in a particular way. They must be capable of unifying and predicting novel phenomena, where a 'novel phenomenon' is (roughly) any phenomenon that one did not set out to explain. So, for example, consider the following putative explanation for why it is that opium puts one to sleep: because it has a dormative virtue. This is no explanation at all, one might argue. Why? Because the fact that opium has a dormative virtue does not unify the properties of opium with anything other than exactly the thing we are trying to explain. It does not, for instance, explain why ketamine or marijuana also have doping effects. Similarly, the fact that opium has a dormative virtue does not predict anything other than the fact that opium puts one to sleep. It does not, for

instance, predict the fact that alcohol at suitably high doses will also put one to sleep.

Our explanation for why it is that Tim is doomed to fail, one might argue, is just like the dormative virtue explanation. Its unificatory and predictive powers are restricted to all and only those cases of Tim failing. The explanation does not unify Tim's failure with anything other than exactly that which we set out to explain; and it does not predict anything other than the regularity that we set out to explain.

But that's not right. The explanation we have offered has novel unificatory and predictive powers. The explanation unifies Tim's failure with the failure of every single attempt by a time traveler to initiate a self-defeating causal chain. In addition, as we shall see presently, the explanation provides a unified account of grand- father paradoxes that arise in higher temporal dimensions (indeed, the explanation works for paradoxes in n-dimensional temporal spaces).¹⁰ The predictive power of the explanation is also apparent. The law of non-contradiction predicts the presence of *grey state* solutions in time travel cases (cf. Dowe 2007). Grey-state solutions arise for physical systems involving time-traveling particles and the like. Some such systems can be placed into paradoxical set-ups: set-ups that seem bound to produce a contradiction. The discovery that such systems always have grey state solutions shows that there will always be some consistent state for the system to enter into, despite its apparent tendency to engender paradox. The fact that there is always a consistent state for the system to enter into is exactly what we should expect if the law of non-contradiction is true. Notice, however, that the existence of grey state solutions was no part of the initial explanatory target.

The third objection focuses on the link between explanation and difference- making. For some, explanation is to be analyzed in terms of difference-making (see, for instance, Strevens (2008)). For others, explanation is not to be analysed in terms of difference-making but is nonetheless deeply connected to that concept (see Woodward (2003) and Woodward and Hitchcock (2003)). Difference-making, in turn, is often understood counterfactually: x makes a difference to y when if x had not

been the case, y would not have been the case. Given all this, one might argue that the law of non-contradiction cannot do explanatory work as it is not appropriately connected to difference-making. For there to be an appropriate connection, the following counterfactual would need to be true:

[1] If the law of non-contradiction had been false, Tim would have succeeded in killing grandfather.

This is a counterpossible. As Lewis (1973b; 1973a), Williamson (2007) and others have argued, however, all counterpossibles are *trivially* true. Accordingly, on this approach to counterfactuals, the following counterfactuals are also trivially true:

[2] If the law of excluded middle had been false, Tim would have succeeded in killing grandfather.

[3] If $2 + 2 = 5$, Tim would have succeeded in killing grandfather.

[4] If bachelors had not been unmarried males, Tim would have succeeded in killing grandfather.

Now there's trouble. If [2]–[4] are true, and counterfactual dependence is indicative of difference-making and thus explanation, then it turns out that all manner of junk explains Tim's repeated failure to kill his grandfather. We don't believe that all counterpossibles are trivially true. Moreover, we are not alone here. A number of philosophers maintain, contra Lewis and Williamson, that counterpossibles have non-trivial truth-values. Indeed, it is only if one adopts a particular semantics for counterfactuals—namely the Lewis-Stalnaker semantics¹¹—that all counterpossibles are trivially true. Fortunately, there is a straight-forward extension of the Lewis-Stalnaker semantics available that yields non-trivial truth-values for counterpossibles.¹² So whether or not all counterpossibles are trivially true will depend on the arguments that one can muster for one semantic approach over the other. We have no intention of mustering such arguments here. Rather, we will just admit that our view is beholden to thinking that counterpossibles have non-trivial truth-values and leave the

matter at that.

Besides, there is a more troubling concern in the neighborhood. Let's suppose that counterpossibles take non-trivial truth-values. If the law of non-contradiction is false, then *absolutely everything* would follow, one might think. So while it might be true that if the law of non-contradiction had been false, Tim would have succeeded in killing his grandfather, the following counterfactuals are also true, and many more besides:

[4] If the law of non-contradiction had been false, Caesar would not have crossed the Rubicon.

[5] If the law of non-contradiction had been false, the Apollo missions would not have successfully landed on the moon.

[6] If the law of non-contradiction had been false, then the Allies would not have won World War II.

As before, if difference-making is indicative of explanation, then it would seem that the explanatory power of the law of non-contradiction is great indeed. The law of non-contradiction is explanatorily implicated in just about everything. It simply explains too much.

This worry trades on a misunderstanding about how counterfactuals with logically impossible antecedents ought to be evaluated. In order to get the result that *everything* would follow from denying the law of non-contradiction, one must accept that the laws of logic license *explosion*: $A \wedge \neg A \vdash B$, for any arbitrary B . For it is only if the laws of logic are explosive that the failure of the law of non-contradiction implies everything whatsoever. We admit that it is tempting to imaginatively engage in counterfactuals with logically impossible antecedents in this fashion. What one is effectively doing is attempting to hold classical logic fixed when imagining a scenario in which the law of non-contradiction is false. One way to do this is to imagine a scenario that is closed under classical consequence and then imagine *of that scenario* that the law of non-contradiction fails because of the presence of some contradiction or other. But obviously if the scenario one imagines is closed under classical

consequence (or closed under any explosive consequence relation), then all hell will break loose when a contradiction is added. So just don't imaginatively engage with counterpossibles in that way. Or, rather, when one attempts to so imaginatively engage, one typically fails to appreciate the scenario one is imagining.

Instead, imagine a scenario that is closed under some non-explosive consequence relation. One option—and the one that we are attracted to—is to imagine that the scenario is closed under the consequence relation of a paraconsistent logic such as Priest's LP.¹³ Assuming the scenario to be closed under LP is useful because LP permits non-trivial reasoning with contradictions. One can, as it were, unravel the consequences of a given contradiction using the inferential capacity of LP. And let's face it: if Tim manages to succeed in his endeavor, then the outcome will be logically disastrous. A slew of contradictions will quickly follow. But not every claim and its negation will be true: claims about the past before grandfather survived should remain untouched. LP has the capacity to vindicate that suggestion.¹⁴

In sum, then, if we use the right logic to think through an inconsistent time-travel scenario, there is no reason to think that counterfactuals [4]–[6] will be true. Still, one might wonder: what can be said in favour of the truth of counterfactual [1]? Well, if the law of non-contradiction were false, then there would be ways for Tim to kill his grandfather. That's because the outcome that Tim is so persistently trying to pursue would no-longer be impossible. So, assuming that Tim has what it takes, that he's the best person for the job, that he's a crack shot, then he should be able to succeed. The principal barrier to Tim's success has been removed, and the whole point of these time travel events is that Tim would succeed but for the paradox he would create. So let there be paradoxes. Let contradictions thrive through time. Then Tim will have his day.

3.3. The Wrong Explanation!

The third objection we wish to consider denies that the explanation we have offered is the right explanation for why Tim is doomed to fail. It

could be argued that if there are two time dimensions, then Tim is perfectly capable of killing his grandfather, by traversing hypertime. Of course, our universe is one in which there is only one time dimension. So Tim is doomed to fail in his attempts to kill his grandfather in the actual world and, indeed, in any world like it, temporally speaking. But there are possible worlds in which Tim succeeds. So it is at best an empirical or metaphysical limitation that prevents Tim from killing his grandfather, not a logical one. If that is correct, then the correct explanation for why Tim is doomed to fail is that time has only one dimension in our world. It is not, as we have suggested, because the law of non-contradiction is true.¹⁵

What is death? Here's one answer: death is that point in time at which one's last temporal part is located. Which is to say that death is the cessation of further propagation through time of one's temporal parts. Now, let us suppose that in addition to the normal temporal dimension, there is an additional dimension: hypertime. Time can be modelled as a line. Time and hypertime can be modelled as a plane. Suppose further that Tim's grandfather is both temporally and hypertemporally extended. Which is to say that Tim's grandfather has temporal parts as well as hypertemporal parts. Effectively, what this means is that Tim's grandfather exists along multiple timelines. Now, suppose that Tim wants to kill grandfather. To do so, he travels backwards in time but forwards in hypertime. He then murders grandfather and returns home. This 2D time travel story is perfectly consistent. Tim has killed his grandfather without plunging the universe headlong into chaos via contradiction. The case may be modelled as follows:

<insert figure 1 here>

Figure 1. A Two Dimensional Travel Story

Tim steps into the time machine at $\langle t_4, ht_0 \rangle$, travels to $\langle t_1, ht_4 \rangle$ and kills grandfather. He then gets back into the time-machine at $\langle t_4, ht_0 \rangle$ and travels forward through time to $\langle t_4, ht_4 \rangle$.

So it would seem that by adding a second time dimension Tim can kill his

grandfather. So Tim is only doomed to fail if there is a single time dimension. It follows, one might argue, that Tim's repeated failure to kill grandfather is best explained by appealing to the fact that our world has only a single time dimension.

The addition of a second time dimension makes it possible for Tim to kill grand- father in some sense. But there is still something that Tim cannot do: namely, travel to a particular time, hypertime coordinate and kill grandfather *at that coordinate*. In particular, even in two-dimensional time, Tim cannot travel backwards in his own timeline and kill his grandfather in that timeline. At best, he can travel to a different timeline, one that is indexed to a distinct hypertemporal coordinate to the coordinate that indexes his own timeline, and kill grandfather in that timeline. But now ask the following question: why is Tim doomed to fail in his attempts to kill grandfather in his own timeline? It cannot be because there is only one time dimension. And it won't help to add a third time-dimension, since even if we add a third time dimension Tim still can't kill his grandfather in his own timeline. The only explanation, that we can see, for why Tim is doomed to fail to kill grandfather is the logical explanation that we have offered. If that's right, then the logical explanation is still required to explain certain constraints on time travelers. Given that we need the logical explanation, however, then it seems natural to use the logical explanation to explain all of the constraints that time travelers are under, where those constraints are needed to prevent paradoxes from forming.^{16,17}

Indeed, this is the heart of the matter. It does not matter how many dimensions of space and time there are: fix an event in that space-time manifold (where the event is located at a space-time point, or region, in the manifold in question) and that event can't be other than the way it is without inducing a contradiction. As has been noted before (Baron and Colyvan, 2016, p. 77) there is nothing special about time in all of this. There are spatial analogues of the paradoxes of time travel, except that no one has ever managed to get themselves worked into a state of confusion over such spatial cases.¹⁸ In short, any claim to alter the past by appeals to

hypertime in the manner discussed here are equivocations. Such changes are no more changing the event in question, than an individual going grey somehow brings it about that they were grey haired as a child.

4. Conclusion

The explanation problem for backwards time travel arises because it seems as though there is something inexplicable about the repeated failure of a time traveler to perform certain actions. We have argued for a particular solution to the explanation problem. The solution begins, first, by differentiating between two distinct explananda: why does Tim fail? And: why is Tim doomed to fail? We suggested that the confusion between these different explananda is partly responsible for the seeming plausibility of the explanation problem. We noted, however, that an explanation for why Tim is doomed to fail is needed. We offered one such explanation: Tim is doomed to fail because the law of non-contradiction is true. We went on to defend this proposed explanation from a range of objections. The explanation, we submit, holds its own. And so we herald the end of mystery: there is nothing inexplicable about what time travelers cannot do. The universe is consistent and consistency is a serious and non-trivial constraint. There are limits to what one can do in consistent universes.

Sam Baron

*Department of Philosophy
University of Western Australia
samuel.baron@uwa.edu.au*

Mark Colyvan

*Department of Philosophy
University of Sydney
mark.colyvan@sydney.edu.au*

Acknowledgements

Thanks to Nicholas J. J. Smith for a very helpful discussion on the topic of this paper and for comments on an earlier draft. Mark Colyvan's work on this paper was funded by an Australian Research Council Future Fellowship (grant number FT110100909).

Notes

¹ Although we have identified these three shortcomings of the Lewis explanation separately, we have no doubt that they are all connected. The lack of understanding, in particular, seems likely to be a function of the other two shortcomings.

² Note that “doomed to failure” is to be understood as follows: Tim’s failure is (logically) necessary because success would lead to a (logical) contradiction. I.e. there is no possible world in which Tim succeeds to kill his grandfather. More on this below.

³ See Baron and Colyvan (2016) for more on the relationship between mathematical and logical explanations and for a defense of logical explanations.

⁴ It is important to note that the mathematics and logic *on their own* do not explain Bridget’s or Tim’s failure. Rather, it is the mathematics/logic in combination with various physical facts about the universe. For instance, we must hold fixed that there are no wormholes that Bridget can utilize to get around the seven bridges, or that the Bridges are not situated in an exotic geometry. Similarly, in the time travel case, we must hold fixed various facts about the linearity, dimensionality and existence of time. But all explanations are subject to the enforcement of background conditions in this manner; there is nothing special about the logical and mathematical cases under consideration.

⁵ The “End of Eternity” by Isaac Asimov (1955) is the tale of such chaperones.

⁶ Even dialetheists do not countenance contradictions such as these. So even if one is open to the idea that some contradictions might be true, there is no reason to entertain changing the past as one of them.

⁷ As Smith rightly notes (in private communication) there are, in fact three views in the vicinity and it is worth distinguishing them.

(a) There is a substantive explanation: the forces of logic *causally* stay Tim’s hand. All parties (Lewis, Smith, and we) reject this view.

(b) There is a substantive explanation but it is *not* a causal one. This is our view.

(c) There is no substantive explanation. This is Smith’s view. There is no substantive explanation of the time traveler’s failure (beyond the low-level facts about banana peels etc.) and, more- over, no such explanation is required. (Smith is happy to call this pointing out that there are no scenarios satisfying a certain description ‘an explanation’ of sorts but he rejects that it’s any kind of *substantive* explanation.)

⁸ See, for instance, Baker (2005); Baron (2014); Colyvan (2002, 2010); Lyon (2012); Lyon

and Colyvan (2008); Pincock (2015); Rice (2015).

⁹ See Noether's (1918) theorem.

¹⁰ The logical explanation also unifies other, related cases of attempting to change the past. As is well known, the Grandfather Paradox is merely a particularly dramatic way to draw attention to puzzles about changing the past. The logical explanation treats all such cases the same: the past cannot be changed because given that the past is thus and so, it cannot be otherwise. Once put like this, we see that time travel is something of a red herring. Whenever things are thus and so, they cannot be *at the same time* not thus and so. This is the case whether we're talking about the past, the present, or the future. Time travel is just a nice way of dramatizing the alleged problem (Baron and Colyvan 2016).

¹¹ See Lewis (1973b); Stalnaker (1968).

¹² For discussion of this extension see Baron et al. (2017); Beall and van Fraassen (2003); Bjerring (2014); Mares (1997); Nolan (1997); Priest (2002); Restall (1997).

¹³ LP is a non-explosive logic. In this sense it is "contradiction friendly" yet it has the following version of the law of non-contradiction as a theorem: $\neg(P \wedge \neg P)$. See Priest (2008) for an introduction to LP.

¹⁴ To reason fully about time dialetheically, something like the dialethic tense logic developed by Tanaka (1998) may be needed.

¹⁵ For discussion of hypertime models of time-travel see Goddu (2011, 2003), Loss (2015), Meiland (1974) and van Inwagen (2010). Note that we do not attribute the two-dimensional view discussed here to any of these authors.

¹⁶ One might think that there is *some* sense in which Tim can kill his grandfather if there are n time dimensions but cannot if there are $n - 1$ time dimensions. Even so, the dimensional explanation is a poor one. For a start, it is disunited. In one case it is the fact that there is only one time dimension that explains the relevant impossibility, in another it is that there only two time dimensions, and so on: it is a different explanation each time. More importantly, the dimensional explanation fails in the infinite limit. For if there are infinitely-many time dimensions, the only explanation for why a time traveler cannot perform a certain action—and there *are* actions that time travelers cannot perform in spaces with infinitely-many time dimensions—is the logical one. Why not embrace the logical explanation from the get go?

¹⁷ A proponent of the two-dimensional model may well accept that the law of non-contradiction is a component of the dimensional explanation. She may then point out,

however, that the law of non-contradiction is a part of *many* explanations. Indeed, it might be thought of as a background constraint on all explanations: all explanations presume consistency. One might worry, then, that the law of non-contradiction is too widespread to be explanatorily relevant. There are two things to say in response. First, that the law of non-contradiction is widespread does not imply it is explanatorily idle. Perhaps it really does do this explanatory work. The reason we don't usually focus on the law, however, is because it plays the role of a background constraint in most contexts. There are few contexts in which the law is raised to salience for us. Second, not all explanations presume consistency. There are explanations in and about non-classical logic that do not presume that consistency is in play. So it is not true that all explanations presume consistency. This makes the law of non-contradiction non-trivial as a component of best explanatory practice.

¹⁸ For example, given that an individual is at a particular spatial location at a given time, she must have failed to reach a different spatial location by the time in question. See (Baron and Colyvan, 2016, p. 77) for discussion of such cases and their relevance to the time travel case.

References

- Arntzenius, Frank and Tim Maudlin. 2002. "Time Travel and Modern Physics," In *Time, Reality and Experience*, ed. Craig Callender (Cambridge MA: Cambridge University Press), pp. 169–200.
- Asimov, Isaac. 1955. *The End of Eternity* (Doubleday).
- Baker, Alan. 2005. "Are There Genuine Mathematical Explanations of Physical Phenomena?" *Mind*, vol. 114, pp. 223–238.
- Baron, Sam. 2014. "Optimization and Mathematical Explanation: Doing the Lévy Walk," *Synthese*, vol. 191, pp. 459–479.
- Baron, Sam and Mark Colyvan. 2016. "Time Enough for Explanation," *Journal of Philosophy*, vol. 113, pp. 61–88.
- Baron, Sam, Mark Colyvan, and David Ripley. 2017. "How Mathematics Can Make a Difference," *Philosophers' Imprint*, vol. 17, pp. 1–19.
- Beall, Jc and Bas C. van Fraassen. 2003. *Possibilities and Paradox: An Introduction to Modal and Many-valued Logic* (Oxford: Oxford University Press).
- Bjerring, J. C. 2014. "On Counterpossibles." *Philosophical Studies*, vol. 168, pp. 327–353.
- Carroll, John W. 2010. "Context, Conditionals, Fatalism, Time travel and Freedom," in *Time and Identity*, eds. Joseph Keim Campbell, Michael O'Rourke, and Harry S.

Silverstein (Cambridge MA: MIT Press), pp. 79–93.

Colyvan, Mark. 2002. “Mathematics and Aesthetic Considerations in science.” *Mind*, vol. 111, pp. 69–74.

Colyvan, Mark. 2010. “There is No Easy Road to Nominalism.” *Mind*, vol. 119, pp. 285–306.

Dowe, Phil. 2007. “Constraints on Data in Worlds With Closed Timelike Curves.” *Philosophy of Science*, vol. 74, pp. 724–755.

Frost-Arnold, Greg. 2010. “The No-miracles Argument for Realism: Inference to an Unacceptable Explanation.” *Philosophy of Science*, vol. 77, pp. 33–58.

Goddu, G. C. 2003. “Time Travel and Changing the Past: (or How to Kill Yourself and Live to Tell the Tale).” *Ratio*, vol. 16, pp. 16–32.

Goddu, G. C. 2011. “Avoiding or Changing the Past.” *Pacific Philosophical Quarterly*, vol. 92, pp. 11–17.

Gorovitz, Samuel. 1964. “Leaving the Past Alone,” *Philosophical Review*, vol. 73, pp. 360–371.

Horwich, Paul. 1975. “On Some Alleged Paradoxes of Time Travel,” *Journal of Philosophy*, vol. 72, pp. 432–444.

Horwich, Paul. 1987. *Asymmetries in Time* (Cambridge MA: MIT Press).

Ismael, Jennan. 2003. “Closed Causal Loops and the Bilking Argument.” *Synthese*, vol. 136, pp. 305–320.

Lewis, David. 1973a. *Counterfactuals* (Oxford: Blackwell).

Lewis, David. 1973b. “Counterfactuals and Comparative Possibility.” *Journal of Philosophical Logic*, vol. 2, pp. 418–446.

Lewis, David. 1976. “The Paradoxes of Time Travel.” *American Philosophical Quarterly*, vol. 13, pp. 145–152.

Lewis, David. 1986. “Causal Explanation,” in *Philosophical Papers II* (Oxford: Oxford University Press), pp. 214–240.

Loss, Roberto. 2015. “How to Change the Past in One-dimensional Time.” *Pacific Philosophical Quarterly*, vol. 96, pp. 1–11.

Lyon, Aidan. 2012. “Mathematical Explanations of Empirical Facts, and Mathematical

Realism.” *Australasian Journal of Philosophy*, vol. 90, pp. 559–578.

Lyon, Aidan and Mark Colyvan. 2008. “The Explanatory Power of Phase Spaces.” *Philosophia Mathematica*, vol. 16, pp. 227–243.

Mares, Edwin D. 1997. “Who’s Afraid of Impossible Worlds?” *Notre Dame Journal of Formal Logic*, vol. 38, pp. 516–526.

Meiland, Jack W. 1974. “A Two-dimensional Model of Time for Time Travel.” *Philosophical Studies*, vol. 26, pp. 153–173.

Noether, Emily. 1918. “Invariante Variationsprobleme.” *Gott. Nachr.*, 1918, pp. 235–257.

Nolan, Daniel. 1997. “Impossible Worlds: A Modest Approach.” *Notre Dame Journal of Formal Logic*, vol. 38, pp. 535–572.

Pincock, Christopher. 2015. “Abstract Explanations in Science.” *British Journal for the Philosophy of Science*, vol. 66, pp. 857–882.

Priest, Graham. 2002. *Beyond the Limits of Thought* (Oxford: Oxford University Press).

Priest, Graham. 2008. *An Introduction to Non-Classical Logic: From If to Is*, 2nd edition (Cambridge: Cambridge University Press).

Restall, Greg. 1997. “Ways Things Can’t Be.” *Notre Dame Journal of Formal Logic*, vol. 38, pp. 583–596.

Rice, Collin. 2015. “Moving Beyond Causes: Optimality Models and Scientific Explanation.” *Noûs*, vol. 49, pp. 589–615.

Riggs, Peter J. 1997 “The Principal Paradox of Time Travel.” *Ratio*, vol. 10, pp. 48–64.

Smith, Nicholas J. J. 1997. “Bananas Enough for Time Travel?” *British Journal for the Philosophy of Science*, vol. 48, pp. 363–389.

Smith, Nicholas J. J. 2017. “I’d Do Anything to Change the Past (but I Can’t Do ‘That’).” *American Philosophical Quarterly*, vol. 54, pp. 153–168.

Stalnaker, Robert C. 1968. “A Theory of Conditionals.” in *Studies in Logical Theory*, ed. Nicholas Rescher (Oxford: Blackwell), pp. 98–112.

Strevens, Michael. 2008. *Depth: An Account of Scientific Explanation* (Cambridge MA Harvard University Press).

Tanaka, Koji. 1998. “To Be Something and Something Else: Dialethic Tense Logic.” *Logique Et Analyse*, vol. 16, pp. 189–202.

van Inwagen, Peter. 2010. "Changing the Past," in *Oxford Studies in Metaphysics Volume 5*, ed. Dean W. Zimmerman (Oxford: Oxford University Press), pp. 1–28.

Williamson, Timothy. 2007. *The Philosophy of Philosophy* (Oxford: Oxford University Press).

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. (Oxford: Oxford University Press).

Woodward, James and Christopher Hitchcock. 2003. "Explanatory Generalizations, Part I: A Counterfactual Account." *Noûs*, vol. 37, pp. 1–24.