



Purely Theoretical Explanations

Giacomo Andreoletti¹  · Jonathan Tallant² · Giuliano Torrenco^{3,4} 

Received: 18 February 2020 / Revised: 14 July 2020 / Accepted: 28 July 2020 /

Published online: 2 September 2020

© Springer Nature B.V. 2020

Abstract

This paper introduces a new kind of explanation that we describe as ‘purely theoretical’. We first present an example, E, of what we take to be a case of purely theoretical explanation. We then show that the explanation we have in mind does not fit neatly into any of the existing categories of explanation. We take this to give us *prima facie* motivation for thinking that purely theoretical explanation is a distinctive kind of explanation. We then argue that it can earn its keep via application to two existing literatures: the literature on how we explain the truth of true negative existential propositions and the literature on how we explain the truth of true propositions about the past. We reply to some possible concerns regarding the introduction of purely theoretical explanations. We conclude that there is nothing obviously wrong with them and explore the ramifications for particular debates in metaphysics.

Keywords Metaphysics · Explanation · Truth-making · Presentism

1 Introduction

Our overarching aim in what follows is to articulate a particular type of explanatory mechanism and to defend its applicability in a variety of contexts. Before continuing, we need to make some clarificatory remarks.

First, as we understand them, explanations can be right or wrong, true or false. The obvious vehicles for bearing these features (truth, falsity, etc.) are propositions. Thus, explanation is propositional; explanations are either propositions or pluralities thereof.

✉ Giacomo Andreoletti
giacomo.andreoletti@protonmail.com

¹ University of Tyumen, Tyumen, Russian Federation

² Department of Philosophy, University of Nottingham, Nottingham, UK

³ Department of Philosophy, University of Milan, Milan, Italy

⁴ Department of Philosophy, Autonomous University of Barcelona, Barcelona, Spain

Second, explanations require an internal structure; a division into what is to be explained (*explanandum*) on the one hand, and what explains (*explanans*) on the other. To keep matters simple, we focus on the structure of propositions of the form $\langle p \text{ because } q \rangle$. Within this simplification, the structure of an explanation is such that an explanatory proposition is one that *contains* other propositions as constituents. More precisely, an explanation of the form $\langle p \text{ because } q \rangle$, itself a proposition, also contains the propositions p and q , which are identified as the *explanandum* and the *explanans* (respectively).

As we understand them, propositions are representational entities that are constituted by concepts.¹ The conceptual and representational nature of propositions plays an important role in the way explanation works. In the first place, which concepts can play an explanatory role depends on contextual factors.² But whilst context is important to individuate *what* is to be explained, to make explicit the *explanandum*, this is not its only role. Context is also relevant in determining what kind of explanation is appropriate. In other words, given a context and an *explanandum*, certain requirements on the explanatory mechanisms are thereby set. For instance, in the context of explaining behaviour in terms of naïve psychology, asking for an explanation of “why John got angry” requires an *explanans* containing reference to concepts such as behaviours, desire, beliefs, and so on. Somewhat crudely, the context is such that it determines that the *kind* of explanation required is an explanation given in terms of naïve psychology. Contrariwise, in the context of explaining behavior in terms of neuropsychology, asking for an explanation of John’s anger would require an explanation that makes reference to neuropsychological concepts.

In order to introduce purely theoretical explanations we must demonstrate the existence of pairs of *explanandum* and contexts, such that it is clear that we have a distinctive kind of explanation on our hands and that the context really is demanding an explanation of a novel kind. In order to help us achieve this, we want to present an example of an explanation, E. We’ll then introduce a (brief, incomplete) taxonomy of different kinds of explanation and their distinctive commitments and show that E does not fit into any of these kinds of explanation. We suggest that this motivates the claim that there is a distinctive kind of explanation—what we will call *purely theoretical explanation*. The example and discussion of the taxonomy will also help us to situate purely theoretical explanation within a wider theoretical framework.

So to the example. To introduce it, we need some background. The prevailing view in late nineteenth century physics was that there existed an ether; a substance-like medium permeating the whole universe that allowed for the propagation of light-waves. Bodies moving across this supposed ether would generate an ether-wind, which in turn would affect the speed of light because the medium velocity does affect the wave velocity that propagates through the medium. The Earth rotating around the Sun would be one of those bodies that generate an ether-wind. However, the Michelson-Morley experiment, aimed at detecting the effects of the ether-wind on the speed of light, famously provided a null result. No such effect was found. One of the explored explanations of this unexpected null result was that the Earth, where the Michelson-Morley experiment was conducted, somehow drags the ether thereby blocking the formation of an ether-wind on its surface. However, this was in contradiction with other

¹ We remain neutral as to how to cash this out.

² Indeed, the context sensitivity of explanations is often noted and discussed in the literature; see for instance van Fraassen (1980: 126).

known phenomena, like stellar aberration, which indicated that the Earth cannot drag this supposed ether. Eventually, the whole idea of there being an ether was discarded. This story provides us with the following example of what we take to be a novel kind of explanation.

E: The Michelson-Morley experiment gave a null result because there is no ether.

E is certainly a (true) explanation of why the Michelson-Morley experiment provided a null result. Our question in the next section: what *kind* of explanation?

2 The Varieties of Explanation and their Commitments: How to Classify E?

The first kind of explanation that we'll mention is ontological explanation. Ontological explanation is familiar from the literature on truthmaking (among other places). Truthmaker theorists hold (roughly) that the explanation of the truth of a true proposition is that there *exists* some entity or entities. Again very roughly, we may say that an ontological explanation is one where the *explanandum* mentions a particular occurrence or state, and the *explanans* makes ineliminable reference to *existing entities*, though does not require there to be any kind of causal connection between that which is described in the *explanandum* and what is referred to in *explanans*. To illustrate: <that *there exists a red rose* explains why *it is true that* <the rose is red>>. Ontological explanations are existentially committing and, as an explanatory kind, *call for the existence* of some entity or entities. The kind of entities to which one is committed by endorsing an ontological explanation may depend on what further metaphysical assumptions are accepted. For instance, someone who maintains that the truthmakers for the predicative truth <the rose is red> are tropes, would be required to give an ontological explanation of the truth of <the rose is red> that committed them to the existence of redness tropes.³

We do not see that E is ontological. There is no obvious appeal here to existing entities appearing in the explanans. Quite the opposite, in fact. It is crucial to E that the explanans is the *negation* of an existential claim, and this alone suffice to dismiss any question about whether E is an ontological explanation.

The second kind of explanation that we will mention is causal explanation.⁴ Examples of causal explanation are legion: the glass broke *because* it was dropped; Tom threw the stone *because* Alice told him to, etc. Very roughly, a causal explanation is one where the *explanandum* mentions a particular occurrence or state, and the

³ See, inter alia, Cameron (2008) who treats the Truthmaking as a measure of ontological commitment. For a discussion of those who see Truthmaking in terms of this kind of explanation, see Tallant (2017: 141-3).

⁴ One thorny issue concerns scientific explanation. So far as we know, the question of how to understand scientific explanation is undecided, with some favouring a causal account (see, e.g. Skow (2014)) and other favouring the view that scientific explanation is subsumption under the laws (e.g. Jansson (2015)). We are not too concerned here with the question of how to analyse the notion of scientific explanation. We believe that there are cases of causal explanation in the world. We believe that there are cases of explanation that require subsumption under the laws. Whether either of these kinds of explanation perfectly analyses scientific explanation is not, then, a question that we feel that we need to broach. In any case, we note again that we do not intend this survey of kinds of explanation to be exhaustive.

explanans describes features of the environment that are causally relevant to bringing about the occurrence or state described in the *explanandum*. Causal explanations commit us to there being what we can describe as *causal links* in the world. What exactly such a commitment amounts to depends on your background understanding of causality. For instance if you construe causal links as causal relations between events, then by accepting the explanation that glass broke *because* it was dropped, you will be committed to there being two events (a dropping and a breaking) standing in a causal relation; whereas if you construe causal relations as obtaining between facts, you will be committed to there being two facts concerning the glass.

Is E a causal explanation? Not so far as we can tell. First, we side with the popular view of event causation that construes causation as a relation between events (cf. Schaffer (2016)). As it is common to think that there are no negative events, viz. no events that essentially feature the absence of a property or particular, event causation rules out causation by absence (cf., inter alia, Beebe 2004: 291–92 and Lombard and Hudson 2020: 638–39). In the case of E, given that there is no such thing as the *event* of there being no ether, there simply is no causal *relation* between the ether’s absence and the experimenter’s result. What caused the Michelson-Morley experiment to give the result it did is not the ether’s absence, but rather all sorts of events regarding the set up of the interferometer used to make the experiment, the events regarding actual external influences, and so on. This is a *prima facie* reason to think that E is not a causal explanation.

Second, we think that even if someone is open to cases of causation by absence, for instance because she accounts for causation in terms of causal facts (cf. e.g., Mellor 1995), she would not consider E as reporting a causal fact or any other type of causal link. Suppose that Sarah promised to water Eric’s orchid during his vacation, but she forgot to do so. The orchid dried up and died. Someone who is sympathetic to causation by absence considers this as an example of an absence being a cause. Had Sarah watered the plant, the plant would have not died; Sarah’s omission to water the orchid *caused* its death.⁵ However, introducing the possibility of causation by absence comes with the well-known challenge of proliferation of causation. Consider Jennifer, who does not know Eric and lives miles away from him. Like Sarah, Jennifer did not water Erik’s plant *and* had she done so, the plant would not have died. Or, consider Marco, who lives on another continent, has never heard of Eric, did not water Eric’s plant either, and so on. Did *Jennifer’s* omission cause the plant’s death too? And Marco’s, as well as all those of many other people who did not water the orchid? To avoid this unwelcomed proliferation and still retain causation by absence, one needs a principled reason why some absences are causes, whereas others are not. That is, one needs to explain why, to use the example, it was Sarah’s omission that caused the orchid’s death, whereas Jennifer did not have any causal role in the plant’s death.

To meet the challenge, one can say that some absences are causes when they violate some sort of norm, for instance moral or legal norms. Sarah broke a promise, thereby violating a moral norm, and thus her omission counts as a cause of the plant’s death, whereas the same does not hold in the case of Jennifer. Or, one can follow Sarah McGrath in saying that an absence of an event-type *c* is a cause of some effect *e* only when the occurrence of an event-type *c* would be a *normal* would-be preventer of *e*. To characterize the notion of normality, McGrath (2005: 138) writes that “it is normal for x

⁵ We here understand omissions as being special cases of absences, i.e. absences of actions.

to φ iff x is supposed to φ ". Regularities that derive from behaviors, norms, and nature set standards for what is normal. In this vein, given that people are *supposed to* keep their promises, Sarah's omission caused the plant's death. In the case of Jennifer, however, she was not supposed to water the plant, hence her omission is not a cause of the death. Or, to use another example that is not about agents and hence more similar to our ether case: it can be correct to say that the complete *lack* of April's rain *caused* the linden's death. It is *normal* for the rain to come in April, but this year it did not. Hence the lack of rain, an absence, caused the linden's death. Or, another option to further discriminate between absences consists in appealing to reasonably close possible worlds.⁶ An absence of an event-type c is a cause of some effect e only when an event-type c occurs in a world that is reasonably close to the actual one. The event of *Jennifer* watering the plant does not occur in a reasonably close world, as she never does it nor she even know Eric, whereas it does rain in April at reasonably close worlds, since it normally does rain in April in the actual world. In accordance with common sense, the analysis in terms of reasonably close worlds implies that Jennifer did not cause the death of Eric's plant and that the lack of April rain did cause the linden's death. Could it be that according to one of these accounts of causation by absence, the absence of ether *caused* the Michelson-Morley experiment's result?

Without entering into the intricate details of these analyses, it should be clear that the ether case in E does not involve a case of causation by absence. The absence of ether does not violate any norm, be it legal, moral, or physical in terms of laws of nature. Moreover, appealing to normality and standards does not help either. It does not seem correct to say that it was normal for the ether to be there when Michelson and Morley performed the experiment. Nor is there some standard according to which the ether *should* have been there.⁷ Moreover, it does not seem correct to say that some reasonably close world features the ether. Surely there are possible worlds featuring ether, but in those worlds all sorts of physical phenomena would be profoundly different from the actual ones, as the ether was considered to be a substance permeating the *whole* universe. Hence worlds featuring ether would be rather distant from the actual one. Thus even if someone, unlike us, is open to causation by absence, the case of E does not seem to be one of those cases where an absence causes something to happen.⁸

If what we said so far is correct, then the claim that the ether's absence *caused* the null result is false, independently of whether or not one is open to causation by absence. Hence E does not report an episode of causation. However, this might not be enough to claim that E is not a causal explanation. In fact, Helen Beebe (2004), following Davidson (1967) and Lewis (1986), argues that not all causal explanations are reports

⁶ Beebe (2004:298-99) considers, but does not endorse, a definition of causation by absence in terms of worlds that are reasonably close to the actual world.

⁷ Unless we take into consideration epistemic standards. The scientific community at the time of the experiment *expected* the ether to exist. However, we do not want to let subjective elements such as human expectations to discriminate between what is and what is not a cause in the world.

⁸ One might here object that causation can be understood as counterfactual dependence and thus consider E a case of causal explanation. After all, had there been ether in our universe, the Michelson-Morley experiment would have detected its effects, given that it was designed precisely to do so. Yet, other explanations feature this counterfactual dependence. Had there not existed a red rose, it would not have been true that \langle the rose is red \rangle . This however does not make the corresponding ontological explanation of why it is true that \langle the rose is red \rangle a causal one. Moreover, a counterfactual condition is not sufficient to discriminate between absences that appear to be causes and those which do not, e.g. Jennifer's failure to water the plant.

of causation. That is, there can be causal explanations where what is mentioned in the explanans is not a cause of what is mentioned in the explanandum. If E were one of these cases, then we would be wrong to claim that E is not a causal explanation.

Is E one of those cases of causal explanation in which the explanans does not do any causing? We do not think so. Beebee observes that a causal explanation is one that provides *causal information* about the occurrence of an event. Crucially for our purposes here, she thinks that an absence can figure as an explanans in a causal explanation, as absences can provide causal information about the occurrence of an event.⁹ For instance, Sarah's *failure* to water the plant provides some causal information about why the plant died. That information tells us something about the causal chain of events that led to the plant's death—the causal chain did not contain Sarah's watering. It might be then that if someone wants to follow Beebee here, she should consider E to be akin to <the orchid died because Sarah did not water it> (a causal explanation). However, we do not think that this move would be correct. Beebee stands with the camp that takes causation to be a relation among events, hence she does not believe in causation by absence. When she argues that absences can figure as explanans in causal explanations without being causes, she is arguing for an account that is compatible with event causation, as event causation excludes absences from causes, and still gives some causal role to absences, as they can figure in causal explanations (cf. 2004: 304). Those absences that in accounts of causation featuring causation by absence would count as *causes* are treated by Beebee's account as explanans in causal explanations.¹⁰ Notice, though, that we argued above that the ether's absence should not count as a cause of the null result, *even if* someone believes in causation by absence. Hence, parallel reasoning should here hold for the claim that the ether's absence is *not* part of the causal history of the null result. Certainly, E gives us information about the null result, as any explanation gives some kind of information, but the information here is not causal in nature. Thus, if what we said so far is correct, E is not a causal explanation.

The third kind of explanation worthy of our attention is structural explanation. This takes some spelling out. The notion of ontological structure (of so-called 'joint carving') is sometimes appealed to in metaphysics. There are various ways to flesh out the details of the idea, and connect it to the notion of explanation. One is to appeal to Lewis' perfectly natural properties, namely the properties that account for objective similarities between things in the world.¹¹ Particular objects are said to resemble one another in particular ways. Given that resemblance is just sharing of properties, and (abundant) properties are easy to get by, any two objects will resemble one another in many ways. Yet certain resemblances seem to be objective, independent from the way we characterize things and from irrelevant contingencies. For instance, the reader and Barack Obama are similar in that they share the property of being a person; but the reader is also similar to a certain pen that now lies on a desk in Tokyo; they share the property of inhabiting the planet Earth. Clearly, the first similarity seems to us more important, and — crucially — appealing to it can be exploited as a mechanism to

⁹ See also Tang (2015) for an extension of Beebee's account where an absence can figure in a causal explanation not only as the explanans but also as an explanandum.

¹⁰ This has the merits, in her view, that we do not need to commit ourselves to the problematic claim that absences can cause.

¹¹ See Dorr and Hawthorne (2013), who individuate ten more roles that the notion of naturalness should play according to Lewis (and cast some doubts on whether anything can play all such roles at once).

generate explanations of various kinds; for instance, a causal explanation of why both the reader and Obama need to inhale oxygen to survive. That qualifies the property of being human as natural (or at least as more natural than the property of inhabiting the same planet¹²).

Structural explanations of similarities of this kind commit those who endorse them to reality having a certain structure — to there being certain joints in reality rather than others. In order to accept a structural explanation we do not *need* to reify the similarities or the properties they are based on. That is, we do not need to treat them as the subject of an *explanans* in an ontological explanation. Furthermore, the commitment of structural explanations is not confined to the structure that can be captured by attribution of natural properties to reality. At least as is understood by Sider, the notion of structure is a *generalization* of Lewis' notion of perfectly natural property to other features or aspects of reality, such as quantification, modality and tense.

As an illustration consider quantification, which may be thought of as “structural” and “joint carving”, and follow Sider (2009: 405), who invites us to consider two sentences:

–Ted is human.

–Ted is located in North America.

These two sentences do not obviously point to any shared natural properties, and yet there seems to be some form of genuine similarity between the *facts* expressed by each. Indeed, naturalness, in his case, is a product of the similarity between the facts expressed by the two sentences—the particular ‘Ted’ appearing in both facts.

Analogously, Sider then invites us to consider the two sentences:

–Something is human.

–Something is located in North America.

As Sider (2009: 405) then puts it:

If the existential quantifier carves at the joints, we again have fact-similarity. There is some genuine commonality between cases in which something is human and cases in which something is located in North America. Each is a case of something being a certain way, and that is a genuine similarity.

Let us grant Sider that there is a sense in which there is a genuine similarity. Let us also grant that this similarity is not to be reified; there is not an *entity* that *is* this similarity. Nonetheless, the similarity between these facts can be put to explanatory work: that we are able to properly execute the tasks involved in metaphysical enquiry is to be explained (in part) by the fact that there is ‘structure’ of a certain

¹² Although perfectly natural is a notion that accept of degree, here – to keep things simple – we don't discuss such an aspect.

kind (viz. quantificational) in the world. Thus, again only roughly, we can say that an explanation is structural just in case it makes ineliminable reference to structure in the *explanans*. And, as in the case of perfectly natural properties or quantificational structure, someone endorsing a structural explanation is committed to the structural features her explanation relies on. The commitment is not to a particular ontology, but it is “serious” commitment to certain metaphysical structure. Again, it’s then straightforward to see that against this context there and a particular explanandum, only an explanans that makes reference to structure will suffice. For instance, <the facts *something is human* and *something is in North America* are similar *because* both expressions carve reality at its ontological structure>. A causal explanation, for instance, simply wouldn’t cut it here.

That was a long introduction to the notion of structural explanation. Sadly, E is not a structural explanation. Since the ether does not exist, there is no metaphysical structure of the kind that Sider describes that can be doing work in this case. There is no ‘fact similarity’ in a case like this where there is no existent at all in the explanans. Note that although both structural explanations and E do not carry a specific *ontological* commitment, there is a crucial difference between the two cases. By accepting that there is genuine similarity between the fact that something is human and something is in North America, we are committed to quantification being part of (or an aspect of) the metaphysical structure of reality. We are not committed to the existence of quantification as an entity along North America (or what North America is reducible to), yet we are committed to the structure of reality being characterized by quantification. Contrariwise, accepting that there being *no* ether explains the null result does not commit us to the structure of the world being characterised in some specific way, for instance containing some “negative” structural element.

The penultimate kind of explanatory mechanism that we wish to consider is mathematical/logical explanation. (We concede that these are likely distinct, though we run them together here and focus on the logical case.) There are cases where deductive patterns of inference can be explanatory. Thus: that $\langle p \rangle$ is true explains why $\langle p \vee q \rangle$ is true; that $\langle p \rangle$ is true and that $\langle p \rightarrow q \rangle$ is true, explains why $\langle q \rangle$ is true, and so on. Logical explanations appeal to *sound* inferential patterns and logical laws. It is a further (albeit interesting) question to ask what “grounds” such soundness, whether it is a matter of psychological “laws” or not, for instance.¹³ But we don’t (we assume) need to *reify* the logical laws. Nonetheless the commitment is serious and makes a difference for overall complex of one’s beliefs. And again, it’s easy to see that there will be contexts in which this kind of logical explanation will be *demand*ed by the context and explanandum. To illustrate, suppose that we are in a 1st year logic class. Our student is trying to solve a problem. The problem presents them with two premises: p and $p \rightarrow q$. They are told that q is true and asked to explain *why* q is true. Only one explanation looks viable given that explanandum and context: q is true because p and $p \rightarrow q$.

Once again, though, E is not an explanation of this kind. No deductive patterns of inference are used by E. Nor does the explanandum follow from the explanans as a matter of logical laws. Logic by itself does not capture the explanatory mechanism carried by E. Perhaps E can be logically regimented so that it looks like a logical explanation. The experiment gave a null result because *if* there is no ether *then* the experiment gives a null

¹³ There may also be ‘extra-logical’ explanations. See Baron and Colyvan (2016) for discussion.

result *and* there is no ether. Yet, many (if not all) good explanations display some kind of logic (in a broad sense) and can be regimented in a logical form, but this does not suffice to make them explanations of a logical kind. Moreover, E as it stands does not have such regimented form. Thus, whatever E *is*, it is not a logical explanation.

The last kind of explanation we will introduce from the literature is conceptual. Examples here are easy. Suppose that Mark is unmarried and that this is known to Alice. Suppose that Sally asserts that ‘Mark is a bachelor’. Alice then asks Sally to explain why this is so. Sally first notes that Mark is unmarried, but Alice retorts that she knows this already. Sally’s next line of response, then, is to say to Alice that what it is to be a bachelor *just is* to be an unmarried man, and this is why Mark is a bachelor. Thus, in this case the explanation that Sally provides Alice with is one where the explanation is one that connects two concepts together. It seems to be the explanation that is demanded by the context. We can, again roughly, then say that an explanation is conceptual just in case the *explanans* analyses at least one concept in terms of another (or others). While — as we stressed before — all explanations in general make use of concepts (being their vehicle a proposition), only conceptual explanations need to provide us with an analysis of one concept in terms of another. Thus, only conceptual explanations are committal to there being certain conceptual links among such and such concepts.

E is not a conceptual explanation. In offering E as an explanation of why the Michelson-Morley experiment gave a null result, we are simply not analysing a concept in terms of another concept. To be sure, there are concepts *involved*. Someone lacking the concept ‘ether’ would not be able to make sense of what is being said. But that is not quite the point of what we are after, here. The fact that concepts are involved in an explanation is not to say that the explanation *is* conceptual (were that the case, then there would be but one kind of explanation, for every explanation will trade upon concepts). Again, whatever E is, it is not a conceptual explanation.

We don’t think that these different kinds of explanatory mechanism are always distinguished as carefully as they could be. We concede that not everyone in the various literatures over which we are skimming would accept the legitimacy of the different explanatory kinds mentioned here. We also note that we have made no attempt to give any analyses of these different kinds of explanatory projects. We concede this as a project that could be attempted, but since spelling out the various different kinds of explanation that are available is not the principle aim of our project, we hope that the reader will permit us this lack of detail.

We don’t take the list to be complete. There may well be other kinds of explanation that we’ve not considered. But what we can’t see is that any other kind of explanation is going to help us here. E seems to be a distinctive kind of explanation. We suggest that this ‘seeming’ gives us a *prima facie* motivation for thinking that there is a distinctive kind of explanation here. We will call this kind of explanation—explanations like E—*purely theoretical explanation*.

3 Introducing Purely Theoretical Explanations

Suppose we accept that there is a *prima facie* case for thinking that E is a distinctive kind of explanation that we’re calling purely theoretical. There remains a question of how to understand the natures of such explanations. When introducing E, we did

nothing to clarify its nature. What can we say about purely theoretical explanation? We're going to make four remarks that, we hope, do enough to clarify the notion.

First, the explanatory claim made by a purely theoretical explanation (PTE) is of the right syntactic form to count as an explanation: by way of reminder, the Michelson-Morley experiment gave a null result *because* there is no ether. Merely syntactically, then, PTEs have an appropriate syntactic structure.

Second, we cannot simply substitute *anything* into the explanans. It seems clear that although the following would be purely theoretical (in the sense that it might satisfy the other criteria adumbrated), it would not be the right explanation: the Michelson-Morley experiment gave a null result *because there are no unicorns*. Certainly, we do not see that this is a good explanation at all. That being so, we think that PTEs are sensitive to the right kinds of contextual features to function as explanations: the contexts demand that some claims that we might make in the explanans get to count as good explanations, others as bad, and others as not explanatory at all (as in our unicorn case).

Third, purely theoretical explanations are not committing in any of the senses that ontological, structural, conceptual, etc. explanations are. To be maximally clear: if *x* is a PTE, then *x* does not make reference to any existing entity (or structure, etc.). Certainly, that seems to be a natural reading of *E*. *There is no ether* does not seem to be an entity of any kind. More to the point perhaps, accepting that there is no ether does not come with an endorsement concerning a particular ontology, a particular metaphysical view on how reality is structured or how concepts are related to one another. Nonetheless, *E* still seems perfectly explanatory.

Fourth, though related, a PTE seems to come with its own kind of commitment, although it may be more difficult than in the other cases to pin down exactly what the commitment of a PTE is. We will call this a commitment to a certain *doctrine* and we do our best, in what follows, to capture the notion.

To get some traction on the notion of doctrine, consider *E* once again. When we offer *E* as an explanation, we are stating something about a particular subject matter, that is ether. In so doing, we are making a substantial commitment towards a certain view about ether: there is no such thing—that is, the world lacks ether. A commitment to such a theory of ether is a doctrinal commitment. Now, the claim that explanations come with commitments with respect to background views or theories about specific subject matter is something unsurprising and likely true of any explanation. An explanation that commits to the existence of entities of kind *K* will usually also commit to the theory that at least in part is about the *K*s, and so will an explanation that commits to the world having certain structural features. In this sense, doctrinal commitment is a minimal commitment of *any* theoretical explanation. However, it is the *typical* commitment of PTE because PTE do not come with any further commitment to a specific class of entities, natural properties, causal or conceptual links, etc. as other kinds of theoretical explanations do. Thus, although different from all kinds of theoretical explanations seen before, purely theoretical explanations have their own specific commitments.

However, whether an explanation comes with a commitment to a doctrine without thereby also coming with a commitment to a specific ontology or class of natural properties is not always something so clear cut. Examples that involve negative existentials as *explanans*, such as *E*, can be easily non committal to anything further except the denial of certain entities (when they are not exemplifications of causal explanations). Also, the examples given below in terms of a minimal theory of truth are relatively

uncontroversial. However, other cases are more nuanced, and we do not exclude the fact that there may be a *spectrum* with respect to how “pure” a theoretical explanation is; or that there may be cases in which it is ultimately quite arbitrary to classify the explanation as purely theoretical rather than ontological, structural, or causal; or cases in which whether the explanation comes only with a doctrinal commitment or with a more substantive commitment is one of the points at issue in the philosophical debate (see below the discussion of presentist explanation of past-tensed truth, especially in 4.2 and 4.5).

This does not mean that doctrinal commitments do not come with constraints on how we take the world to be. A doctrine represents the world as being some particular way. We do not wish to commit, here, as to the vehicles that provide that representation, though we suspect that they will be propositional. By representing the world as being some particular way, these representations can be faithful, accurate representations (of the way that the world was, of the way that the world is not, and so on). Even if the lack of further (ontological, structural, etc.) commitments, beside the acceptance of the theory, means that a merely doctrinal commitment does not posit any *direct* constraint on our conception of reality, a doctrinal commitment will typically put limitations and restrictions on such further commitments.

Consider again E. Accepting E commits us to the “no ether” doctrine about ether, namely the very simple doctrine that there is no ether. Being committed to the “no ether” doctrine does not directly lead us to accepting any specific ontology (we are not committed to the existence of entity of a particular kind), or metaphysics (we are not committed to the world being structured in some way rather than another), etc. And yet accepting the “no ether” doctrine is *not* compatible with just any ontology or metaphysics whatsoever. For one, it is not compatible with an ontology that contains ether, and plausibly also with a metaphysics that posits structural features of reality that are characteristic of a world inhabited by ether. In other words, the constraints that doctrinal commitments *specifically* put on how we conceive of reality are indirect constraints on the ontology, structure, conceptual connections, etc.

More pressingly, perhaps, we want to present a challenge. We have given examples of doctrines. In particular, we said with E that *there being no ether* explains the null result of the Michelson-Morley experiment. This PTE comes with a commitment toward a doctrine that represents the world as being ether-free. Suppose that our opponent wants to insist that they really do not understand what such doctrines are supposed to be. In that case, we worry that their inability to comprehend runs deeper than merely not seeing what a doctrine is supposed to be. If our opponent does not understand how E works, then we do not think that they understand how explanation works. And if they do not understand that, then, since we aren’t offering an analysis of elucidation of the general notion of explanation, we do not see that anything we have said can be of use to them.

Another way to point down the idea of a doctrinal commitment is to connect it with “explanatory deductions”. As Tallant notes (Tallant 2017: 147–55), elaborating on a point in Dummett (1978: 296), sometimes a deduction is explanatory, and the fact that a deduction can be provided it may be the only reason for us to believe a certain conclusion. In cases like that, it may turn out that one of the premises we endorse to construct the explanatory deduction is what we call a (mere) doctrine. Think for instance of a “minimal” positions about a subject matter such as pain:

- 1 X is in pain iff X's C-fibres are firing
- 2 There is nothing more to pain than the satisfaction of 1
- 3 X's C-fibres are firing

Therefore,

- 4 X Is in Pain

If we accept an explanation such as the above of the reason why X is in pain, we are surely taking on board ontological commitments (to the existence of C-fibers) and possibly structural (firings of C-fibers) commitments. Yet, in accepting 2 we commit ourselves to a certain view about the world, namely that when someone is in pain *nothing more than* the firing of their C-fibers goes on, which does not come with any further ontological or structural (or conceptual, etc.) commitment. 2 forms a part of the explanation, here, as to why x is in pain. For instance, consider the following context.

Elizabeth: 'I know that the brain and mind are supposed to be connected, but I'm not sure quite how and I really don't know what's going on with pain.'

Katherine: 'Ah, well, in my philosophy course we learned that someone is in pain if and only if their C fibres are firing'.

Elizabeth: 'Oh, ok. Well, in our Physics class we were looking at MRI scans. We saw some c-fibres firing. I guess that means that the subject was in pain. But why were they in pain? You told me that their c-fibres would be firing if and only if they are in pain, so I know that they're in pain, but *why*? What is this stuff 'pain'?'

Katherine: 'MRI? Cool! But, well, I've said everything that there is to say: the reason that x is in pain when their c-fibres are firing is that there's nothing more to pain than that someone is in pain if and only if their c-fibres are firing. That's it! Pain isn't a *thing*. Pain is just every instance of that bi-conditional.'

In the final section, where Katherine notes that 'the reason that x is in pain when their c-fibres are firing is that there's nothing more to pain than that someone is in pain if and only if their c-fibres are firing', we think that she's offered Elizabeth an explanation. (Perhaps it's not a good one; perhaps there is more to pain than just the satisfaction of every instance of the bi-conditional; in that case it's a bad explanation, but an explanation nonetheless.) That is a purely theoretical explanation.¹⁴

¹⁴ To see why this is not an instance of conceptual explanation, it makes sense for us to foreshadow some things that we say later on. The explanatory mechanism at the core of a purely theoretical explanation is not that of a conceptual explanation, relying on certain connections between concepts. Thus, it is not by pondering the link between the *concept* of pain, and the *concept* of c-fibres firing, that we can come out with an explanation of why it is true a person who is in pain is a person who's c-fibres are firing. The explanation in question is purely theoretical because its background theory is what we call a (mere) doctrine: it does not come with a commitment concerning specific conceptual link; one can perfectly well grasp the concept of pain without grasping the concept of a c-fibre (let alone a firing c-fibre), and one can perfectly well grasp the concept of a c-fibre firing without grasping the concept of being in pain.

Another example is the case of truth. Minimalists or deflationists about truth seem more likely candidate for philosophers who resorts to purely theoretical explanations. An explanation such as E' below (adapted from Tallant (2017: 152)) does not commit the minimalist to any specific entity or structural feature of reality connected to the nature of truth, but nonetheless seems to offer us an explanation of why a given proposition is true.

(E') $\langle p \rangle$ is true because (i) $\langle p \rangle$ is true iff p; (ii) Truth's nature is exhausted by every instance of (i); (iii) p.

This seems to be an explanation of why $\langle p \rangle$ is true. Endorsing E' does come with a commitment, namely the (deflationary) doctrinal commitment to the view that there is nothing more to truth than the satisfaction of every instance of the T schema. And the claim that 'there is nothing more to truth than the satisfaction of every instance of the T schema' is itself a doctrinal commitment. It tells us that there is no N such that N is both a part of the nature of truth and anything other than the satisfaction of every instance of the T schema. That being so, we have a purely theoretical explanation.

So, in sum: a purely theoretical explanation is an expression that: is of the right syntactic form to be regarded as an explanation; appears explanatory (and can be contrasted with putative explanations that are clearly *not* explanatory; commits us to (what we call) a doctrine, but makes no further ontological, structural, conceptual, etc. commitment in the explanans.

All of which is well and good (perhaps), but there may at this point be a little uncertainty as to why we are supposed to care. To be sure, there is some intrinsic interest in clarifying and classifying the different kinds of explanatory projects in which we might engage, and so it is good to have some grip on PTEs; nonetheless, we might also retain some further expectation that a philosophically interesting distinction in explanatory kinds will bring with it something of consequence. So, yes: there are PTEs. What can we do with them? Tentatively, though it is not our primary focus in this paper, we think that there is an obvious *potential* outlet here via the truthmaker literature.

Like ontological explanation, we think that (though not recognised as a distinctive explanatory kind) purely theoretical explanation is sometimes appealed to in the truthmaking literature. To illustrate, consider the true negative-existential proposition: \langle there are no unicorns \rangle . How are we to explain the truth of this true proposition? How does \langle there are no unicorns \rangle get to be true? According to the truthmaker maximalist, the answer here is ontological: for every true proposition there is some existent that explains the truth of the true proposition in question—even negative existentials.

According to some, however, the correct explanation for the truth of \langle there are no unicorns \rangle is that *there are no unicorns*. According to the views currently under consideration, that *there are no unicorns* is not intended to denote an existing thing (for example, a reified absence). Nor, on the views under consideration, is *there are no unicorns* intended to denote a 'negative fact', a 'negative state of affairs', a 'negative property', or any other kind of entity or structural feature of reality. According to the particular view (or family of views) that we have in mind, there are some truths that simply require no truthmakers. All that there is to the explanation of how the true negative existential propositions get to be true is that some thing or thing *not* exist.

Thus, once again: the truth of <there are no unicorns> is to be explained very simply; it is true because there are no unicorns.¹⁵

The explanans in the case of <there are no unicorns>, and other negative existentials, is such that the explanation cannot be classified as any of ontological (the explanans makes no reference to an existent), causal (no causal explanation is being given), structural (no similarity is being appealed to), logical (it is not just the inferential import of the quantifier and the negation that does the explanatory work), or conceptual (merely that we have a concept ‘unicorn’ and the concept of negation is no guide to the existence of unicorns, or otherwise). So if there is explanation here, it seems to be a PTE. And, it seems reasonable to us that we may explain the truth of the true proposition <there are no unicorns> by recourse to an explanans that states that *there are no unicorns*.¹⁶

Another case has also received quite extensive discussion. Presentists claim that only present objects exist. They also claim that there are truths about the past. There is no *ontological explanation* that the presentist is (easily) able to give for the truth of propositions about the past, since the relevant past-entities do not exist. For instance, though it’s true that <there were dinosaurs>, there exist no past-dinosaurs to explain the truth of the propositions. Nonetheless, claim some presentists, there is still an explanation of the truth of these true past-tensed propositions that can be given: <there were dinosaurs> is true because *there were dinosaurs*.¹⁷ And, once more, this seems to be an explanation that is demanded by the context. We are told that Presentism is true and that there are no dinosaurs, though there were. We are then asked to explain why it’s true that <there were dinosaurs>. This context/explanandum pair makes it clear enough that the explanation for the truth of <there were dinosaurs> is that *there were dinosaurs*. There is no entity, no structure, no concept and no cause, etc., that is being appealed to as the source of the explanation. If PTEs can do a good explanatory job in the other contexts we have seen, a presentist can resort to this PTE to explain truths about the past.

In sum: in the opening section we identified an explanation, E, and we suggested that, at least prima facie, this kind of explanation does not fall naturally into any of the pre-existing kinds of explanation described in the wider literature. We’ve made some attempts to spell out what’s involved in offering a purely theoretical explanation like E, we’ve shown how this kind of explanation is sometimes appealed to in the philosophical literature on truthmaking, and we showed how PTEs can be used by presentists. What we want to do now is turn our attention to some objections to purely theoretical explanation.

4 Objections to Purely Theoretical Explanation

We’ve now said a little bit about what purely theoretical explanations are, how they would work, and the kinds of context in which we might appeal to them. The objections

¹⁵ We have in mind here views defended by, inter alia, Simons (2008) and Tallant (2010a).

¹⁶ Tallant (2017: chapter 7) is representative, here.

¹⁷ See, inter alia, views defended by Sanson and Caplan (2010) and Tallant (2009, 2010b).

that we pursue here allow us to say a bit more about purely theoretical explanation and the notion of a doctrine.

4.1 It Is Unclear why E Should Count as an Explanation

One might have some worries about the initial example we used, i.e. E. We claimed that ‘the Michelson-Morley experiment gave a null result because there is no ether’ is a PTE, thereby assuming it is an explanation. Is it an explanation in the first place, though? Maybe E by itself does not provide an explanation of why the experiment gave the result it did. Perhaps we need to know more about the set up of the experiment, its initial conditions, what work the ether was supposed to do, and so on.

Agreed. Probably almost any explanation needs a background. Take ‘the argument whose premise is (P and not-P); the conclusion Q is valid because (P and not-P) is a contradiction’. One cannot even understand the explanation unless she knows that an argument is valid if and only if it’s impossible for its premises to be true and the conclusion false, that ‘and’ is the truth-functional conjunction, how P and Q are to be understood, and so on and so forth. Yet, the requirement of a proper background does not make it the case that the explanation about the validity of the argument is not an explanation.

We think that matters are similar for E. To be sure, any explanation requires a context and background. We agreed to such a thing at the outset. Nonetheless, at least so far as we can tell, there are numerous conceptual backgrounds against which E can properly be regarded as an explanation.

Another worry might be the following. One might explain the null result of the Michelson-Morley experiment in terms of laws of nature and claim that E is nothing else than a shorthand for an explanation in terms of laws of nature. For instance, in this case the Special theory of Relativity (STR) and the constancy of the speed of light. Agreed. One can explain the null result by saying that STR implies that the ether as it was conceived *cannot* exist. For if it existed, the ether would provide an absolute frame of reference, which conflicts with STR. Certainly this would be a further correct explanation of why the experiment gave the result it did. In fact, any physical process, as an experiment is, can be explained via its initial conditions and the laws of nature, provided one thinks that laws of nature have explanatory power.

Yet, this does not imply that E is not an explanation *too*. Nor does it tell us that E is not an explanation of a different kind than the one in terms of laws of nature. For instance, E does not require any commitment to the existence of laws of nature as a certain kind of necessary relation among universals, or a commitment to the laws of nature being a sort of structure over mere regularities. The only commitment E requires is a commitment to a doctrine that represents the world as ether-free. Furthermore, it is interesting to observe that the scientific community was quite puzzled with the Michelson-Morley experiment result. The experiment was cleverly designed to detect the effects of the ether but it did not do so. E was a possible explanation of the null result, but physicists were not eager to go in that direction. For light was taken to be a wave and any wave needs a medium. The ether was the medium for light. In fact, possible explanations other than E were devised and considered. It was thought for instance that the ether was there and that there was the ether wind too, which as an effect it would have slightly slowed down the light in one direction. Yet, it was thought that there was

also a slight contraction of the experimental apparatus in just one direction, due perhaps to the fact that the apparatus itself was going through the ether in one direction. If so, this would have explained the null result because the two effects would cancel out.

Only later on, with the full development of STR, which as said rules out the possibility for the ether to exist, it became clear that E was not only a possible explanation of the null result, but *one correct* explanation of the null result.

4.2 It Is Unclear why E Is a PTE

Having suggested that E really is explanatory in the Mickleston-Morley case, we should also say something about whether or not that kind of explanation is really of the same kind as the sorts of explanation present in cases of negative existential propositions, or in cases of past-tensed propositions where we couple that to a presentist metaphysic. We turn to the case of negative existentials first. We think that the case of negative existentials is straightforwardly of the same kind as E. By way of a reminder, here is E again.

E: The Michelson-Morley experiment gave a null result because there is no ether.

In the case of negative existentials, we have a structurally similar explanatory pattern. Thus:

<there are no unicorns> is true because there are no unicorns.

To the left, in the explanandum, we have some feature of the world (a truth; a experimental result) and to the right we have the explanans, that commits us to the non-existence of a particular entity: ether; unicorns. As above, what we have in the explains is a representational device that places constraints on the world and then uses that to provide the explanation of the explanandum. To be sure, there are differences in kind between the two, for one is an explanation of an experimental result and the other is the explanation for the truth of a true proposition. Thus, whilst there may be differences between the two, it seems nonetheless reasonable to treat both as instances of purely theoretical explanations and to make space for further individuation within the category of purely theoretical explanations.

The presentist case is slightly more complex. For the presentist does not prima facie commit to saying that, for instance:

<There were dinosaurs> is true because *there are no dinosaurs*.

That would be a very odd attempt at an explanation indeed! Rather, what the presentist will say is that:

<There were dinosaurs> is true because *there were dinosaurs*.

We concede, then, that this is different from E and the case of negative existentials. Nonetheless, just as we suggested that E and explanation of negative existential truths may both be regarded as purely theoretical explanations, but as different kinds of purely theoretical explanation, we think that similar things can be said here. To the left, in the

explanandum is something that requires explanation. The explanans then consists solely of a representation of the world that places constraints on the world; *the world must be such as to have contained dinosaurs*. This is doctrine for the presentist. Now, to be sure, this is a different kind of doctrine to that posited in the case of negative existentials. Rather than the doctrine being such that *x not exist*, the doctrine is such that *x existed*, and some may be unsatisfied with the claim that the presentist doctrine does not come with further commitments (to Lucretian properties, uninstantiated haecceities or other amenities). But this is an *independent* theoretical worry, which concerns the very general issue of what kind of commitment comes from endorsing a given theory (for a given explanatory purpose). Indeed, the worry could be reformulated in terms of whether a presentist can resort to a purely theoretical explanation to account for the truth of past-tensed claims or she needs to resort to an ontological (by committing herself to uninstantiated haecceities, for instance), or metaphysical (by committing herself to Lucretian properties, for instance) kind of explanation. Our point here is more limited, it is that (i) doctrinal commitment is a cost per se (to be gauged against other costs, such as ontological and ideological costs), and (ii) an explanation is purely theoretical (as opposed to various other kinds of explanation) if it is based on an explanatory mechanism that comes with a mere doctrinal commitment (and no further kind of commitment). Thus, even if the presentist commitment that there were dinosaurs is different from the commitment that comes with explanation of the truth of negative existential, insofar as it does not come with further commitment it is of the right kind for being a commitment distinctive of PTE.

To be sure, then, there may be different kinds of purely theoretical explanation and future work on their natures would do well to map their different kinds and explore the precise relationship between them. Nonetheless, for the reasons given we think it intuitive enough to regard them as all instances of purely theoretical explanation.

4.3 Purely Theoretical Explanation Requires us to Posit Primitive Concepts

Here is a further objection. As we noted, one kind of purely theoretical explanation takes the form: <there are no unicorns> is true, because *there are no unicorns*. But, as we also saw above, this is not the only instance of purely theoretical explanation. Presentists (who believe that only present objects exist) look to explain the truth of true propositions about the past by saying things like ‘<there were dinosaurs> is true because *there were dinosaurs*’. In this case, though, the *explanans* features a tensed claim: *there were dinosaurs*. This is ineliminably tensed. It cannot be analysed. The past tense, here, is expressed by a primitive operator that stands for a primitive concept: for the presentist, some things *were* the case; they exist no longer. Similarly, we might raise a concern for the explanation of the truth of <there are no unicorns>. The *explanans*, here, is that *there are no unicorns*. But this *explanans* ineliminably features negation; there is *no* entity that is a unicorn. And, much like tense is for the presentist, negation may be treated as an operator. It’s also (arguably) an operator that stands for a primitive concept (negation) that does not admit of a revealing analysis. We therefore face a problem. These PTEs are explanations that feature primitive concepts. Our explanations ought to not feature primitive concepts. Thus, PTEs are to be rejected.

Whilst it is true that PTEs may feature primitive concepts, it’s also highly likely that other forms of explanation feature primitive concepts: causal and ontological (certainly

ontological) are likely forms of explanation that may feature primitive concepts. Consider: an ontological explanation will make an existence claim in the explanans: it will say that say something of the form $\langle q \rangle$ because there exists some $x \dots$. The concept of ‘existence’ cannot be analysed—or, at least, it’s plausible to think so. (Similar remarks may apply to causation as it features in causal explanation unless we can provide a robust and satisfactory analysis of causation.) Nonetheless, we still have an explanation here.

Note also that while it is true that purely theoretical explanations commit those who endorse them to the primitive concepts that it makes use of, it is not thereby to be considered as a type of conceptual explanation. As it should be clear by what we said in the previous session, we distinguish the commitment to the concept we *make use of* when we endorse a certain explanation, and the commitment to what we resort to in the *explanans* — that is, the specific commitment that individuate the kind of explanation at issue. The explanatory mechanism at the core of a purely theoretical explanation is not that of a conceptual explanation, relying on certain connection between concepts; it is not by pondering the link between the concept of unicorn, the concept of negation, and the concept of existence that we can come out with an explanation of why it is true that there are no unicorns. When we appeal to a purely theoretical explanations and claim that the proposition \langle there are no unicorn \rangle is true in virtue of reality *not* being in certain ways, we commit ourselves to a certain *doctrine* about reality. Such a doctrine encompasses concepts (such as the primitive concept of negation), but it is the commitment to the doctrine and not to the (relations between) concepts that show up in the doctrine what characterizes it as a purely theoretical explanation. The explanation in question is purely theoretical because its background theory is what we call a (mere) doctrine: it does not come with a commitment concerning specific conceptual link, or a further piece of ontology or ideology.

4.4 There’s no Way to Distinguish Good Instances of PTE from Bad

In this section, as well as the next two, we want to move to discuss objections to (what we would describe as) PTEs being deployed in the philosophical literature. The thought behind this concern is that if our opponents in these literatures have identified good reasons to dislike PTEs *within those debates*, those concerns may well generalise to infect the broad notion of a PTE. In each case, we’ll argue that our opponents challenge can be met. First, we’ll consider two putative PTEs from the literature on presentism and truthmaking. Only one is a genuine explanation.

(1) \langle there were dinosaurs \rangle is true because *there were dinosaurs*.

(2) \langle there were dinosaurs \rangle is true *because Sherlock Holmes lived at 221B Baker street*.

The concern (from those who have raised this kind of worry) is that it’s hard to see what makes these cases different. Why is (1) a good PTE, but (2) not? If we cannot distinguish good PTEs from bad, then we cannot make use of them. It’s easy to see how this will infect the notion of a PTE generally (cf. Baron and Asay 2014: 323–4)).

Well, if the *explanandum* asks for an explanation of the truth of a past-tensed claim, and we are presentists who think that tense is primitive and unanalysable, then the *explanans* had better feature a past-tensed claim about how the world was. The proposition <Sherlock Holmes lived at 221B Baker street> is not a past-tensed claim about how the world was. It is a claim about a fiction. Thus, that different kinds of explanation are called for in different cases is unsurprising. It is a feature of explanation *in general* that the context determines the kind of explanation required. There is no special problem here. We say that (1) is true and that (2) is false. We are doctrinally committed to dinosaurs *having existed*. If the likes of Baron and Assay think that more needs to be said, we should be told what and why.

Indeed, very generally, we struggle to see that Asay and Baron *really* fail to understand what is going on. Consider the explanation of the truth of negative existentials propositions., and a putative alternative explanation.

E*: <there are no unicorns> is true because there are no unicorns.

E**: <there are no unicorns> is true because there are no centaurs.

It seems perfectly clear that whilst *there are no unicorns* is well suited to its position in the explanandum, *there are no centaurs* would not be. In order to explain why the proposition <there are no unicorns> is true, we need to say something relevant to its subject matter. E* makes reference to a doctrine about unicorn (a very simple one, there aren't any), whereas E** is about centaurs. Missing a connection between centaurs and unicorns (unicorns are neither "made out" of little centaurs, nor are just centaurs in disguise), E** makes reference to a doctrine that is explanatory irrelevant and so it does not count as an explanation, let alone a PTE. To sum up, so far as we can see, the way in which we handle E* is structurally analogous to the presentist's position. We take E* to be a paradigmatic example of a sensible explanation and, given its structural similarity to the presentist position, we see little to object to in the claim that the presentist offers us an explanation in cases like (1), but not like (2).

4.5 The Ostrich

Torrenco (2014) has criticised "ostrich presentism". An ostrich is a presentist who claims that by endorsing a primitively tensed language she can avoid paying a price in terms of ontology or ideology when she provides an explanation of why true past tensed existential claims are true. Although, as we have stressed, a purely theoretical explanation comes with its own kind of commitment, the ostrich's strategy is clearly very close to that of endorsing a purely theoretical explanations.

The core of Torrenco's anti-ostrich criticism is the following. Call a claim *p* that is true in a world *w* an ungrounded truth (in *w*) if (i) it is a brute truth (i.e. there is no non-circular explanation of why it is true), and (ii) the entities and properties that *p* is about do not exist and are not genuine¹⁸ (respectively) in *w*. Can ungrounded truths be

¹⁸ A "genuine" property is one that emerges in the substantial metaphysical doctrine that we accept (those expressed by its primitive predicates). Lewis (1983) talks of "perfectly natural properties", Fine (2001) of "real properties", and Sider (2011) of "joint carving properties" are ways to spell out the idea.

explanans? If ungrounded truths are accepted with *no restrictions*, one could gain all parsimony in ontology and primitive ideology with no loss of explanatory power. For instance, if we do not want *tables* in our domain of quantification, we could accept ungrounded truths that explain (or explain away) anything about tables that our adversaries (who accept the existence of tables) explain (see Torrenco 2014). But we do so with no cost to our ontology. Purely theoretical explanations may seem amenable to this criticism. PTEs about the past may seem to be explanations that invoke ungrounded truths, after all; we have seen on a number of occasions, ungrounded truths appeal in the *explanans* via being doctrinal.

Nonetheless, we wish to resist the impression that we are engaged here in ‘ostrich-like’ behaviour. Regardless of whether Torrenco’s criticism to the kind of presentism he discusses is correct, we are not claiming that PTE’s come with no cost at all. As we have stressed above, PTEs come with a cost; they are committal. Although they do not commit those who uphold them to new ontology or structure, they require a particular theoretical stance toward the world: that the world *is not* some way; that the world *was* some way, and so on. They commit us to a doctrine and that is a commitment. We don’t claim to know how to balance doctrinal commitment against ontological commitment, but that’s perhaps unsurprising. We don’t know how to weigh qualitative parsimony against quantitative parsimony. We don’t know how to weigh explanatory power against the pull towards explanatory unity. Simply: weighing up competing theoretical virtues is tough. We take it that every time we posit a doctrine, we’re incurring some kind of cost and so just as we weigh a theory with respect to the number of entities and the kinds of entity that it posits, so we should add to that weighing the number and complexity of doctrines posited.

We can imagine an opponent putting the point another way. The demand for ontological or structural explanations is something like a demand for *the world* (or something in the world) to *be* the explanation of what is true. Thus, when we seek an explanation for the truth of <there is an electron> we are looking for something, on the side of the world, that explains why this proposition is true. And, so goes the thought, matters are similar in the case of past-tensed propositions and negative existential propositions; we are looking for something, on the side of the world, that explains why these propositions are true.

At this point, though, we think that our opponent has erred. We want to refer back to some remarks that we made earlier. Explanations are propositional in form, or so we said. An explanation *is* a proposition—albeit quite plausibly a complex proposition. In its canonical form an explanation will say something of the form of <*p* because *q*>. Of course, on both sides of the ‘because’ claim we find further propositions. As a consequence, our opponent misspoke. There is nothing *in* the explanation that is ‘something on the part of the world’. Indeed, ‘the world’ is not ‘in’ the explanation at all. Rather, an explanation is a proposition of a particular kind. As the form of the proposition reveals, <*q*> is taken to explain <*p*>. One proposition is taken to explain another. Thus, the very claim that there *could* be anything ‘on the side of the world’ that explains why a proposition is true (or false) is simply misguided from the outset. What one should say is that explanations come with commitments of various sorts, and that these can place *constraints* on how the world must be for the explanation to be apt. But so long as there are no unicorns, it’s hard to see that much will have gone wrong with the explanation of the truth of the proposition <there are no unicorns> if we say that it’s

true because there are no unicorns. As we pointed out above (see 4.2), sometimes the philosophical point at issue is precisely whether a PTE is enough to account for an *explanandum*, or some more committal strategy is required. But our aim here is merely to point out that PTE are a distinct kind of explanation, which are used in philosophy and beyond, and whose existence is worth recognizing; we are not defending any specific example as the best explanatory strategy.

4.6 Semantic Indeterminacy

We consider the possibility that what we call “purely theoretical explanations” are not explanations at all. Rather, they merely provide (contextually) hints towards explanations of some theoretical kinds (e.g., ontological, structural, conceptual). In other terms, we can think of them as shorthand for disjunctions of possible explanations. For instance, if someone claims “<there are no unicorns> is true because there are no unicorns”, what she is stating is better paraphrased as “<there are no unicorns> is true because either there are unicorns-absence, or the world is structurally as such that there are no unicorns, or it exemplifies a totality fact that entails that there are no unicorns, or...” where the list of potential explanations is left open.

Our reply: while the phenomenon of exploiting some sort of semantic indeterminacy to achieve explanatory import may happen in the context of explanation in philosophy and in the sciences, as stated the objection does not do anything to *show* that purely theoretical explanations are not a genuine phenomenon. As stated at least the objector does not provide an argument to think that *every time* that someone appeals to an alleged PTE, something like E, she must be understood as expressing something else (an indeterminate disjunction of contextually salient eligible explanations). Thus we have reason to think that there are cases of purely theoretical explanations. We conclude that the case for purely theoretical explanations stands.

Acknowledgements An early version of this paper was presented at the Explananza workshop at the University of Western Australia. We thank all participants for feedback. We would also like to thank the project PID2019-108762GB-I00 of the Spanish Ministry of Science and Innovation. Finally, we thank two anonymous referees of this journal for helpful comments.

References

- Baron, S., & Asay, J. (2014). The hard road to Presentism. *Pacific Philosophical Quarterly*, 95, 314–335.
- Baron, S., & Colyvan, M. (2016). Time enough for explanation. *Journal of Philosophy*, 113(2), 61–88.
- Beebe, H. (2004). Causation and nothingness. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals*. Cambridge: The MIT Press.
- Cameron, R. (2008). Truthmakers and ontological commitment: Or how to deal with complex objects and mathematical ontology without getting into trouble. *Philosophical Studies*, 140, 1–18.
- Davidson, D. (1967). Causal relations. *Journal of Philosophy*, 64, 691–703.
- Dorr C. & Hawthorne J. (2013), “Naturalness”. In Karen Bennett & Dean Zimmerman (eds.), *Oxford Studies in Metaphysics*, vol. 8., Oxford University Press.
- Dummett M. A. E. (1978). *Truth and Other Enigmas* Harvard University Press.
- Fine, K. (2001). The question of realism. *Philosophers' Imprint*, 1, 1–30.
- Jansson, L. (2015). Explanatory asymmetries: Laws of nature rehabilitated. *Journal of Philosophy*, 112(11), 577–599.

- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61, 343–377.
- Lewis, D. (1986). “Causal explanation”, in philosophical Papers, vol II, New York: Oxford University Press.
- Lombard, L. B., & Hudson, T. (2020). Causation by absence: Omission impossible. *Philosophia*, 48(2), 625–641.
- McGrath, S. (2005). Causation by omission: A dilemma. *Philosophical Studies*, 123, 125–148.
- Mellor, D. H. (1995). *The facts of causation*. London: Routledge.
- Sanson, D., & Caplan, B. (2010). The way things were. *Philosophy and Phenomenological Research*, 81, 24–39.
- Schaffer, J. (2016). The metaphysics of causation. *The Stanford Encyclopedia of Philosophy* (fall 2016 edition), E.N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2016/entries/causation-metaphysics/>>.
- Sider, T. (2009). Ontological realism. In D. Chalmers, D. Manley, & R. Wasserman (Eds.), *Metametaphysics* (pp. 384–423). Oxford: OUP.
- Sider, T. (2011). *Writing the book of the world*. Oxford: OUP.
- Simons, P. (2008). Why the negations of false atomic sentences are true. In T. de Mey & M. Keinanen (Eds.), *Essays on Armstrong* (Vol. 84, pp. 15–36). Acta Philosophica Fennica.
- Skow, B. (2014). Are there non-causal explanations (of particular events)? *British Journal for the Philosophy of Science*, 65, 445–467.
- Tallant, J. (2009). Ontological cheats might just prosper. *Analysis*, 69, 422–430.
- Tallant, J. (2010a). There’s no existent like ‘no existence’ like no existent I know. *Philosophical Studies*, 148, 387–400.
- Tallant, J. (2010b). Still cheating, still prospering. *Analysis*, 70, 502–506.
- Tallant, J. (2017). *Truth and the world*. London: Routledge.
- Tang, Z. (2015). Absence causation and a Liberal theory of causal explanation. *Australasian Journal of Philosophy*, 93(4), 688–705.
- Torrenzo, G. (2014). Ostrich Presentism. *Philosophical Studies*, 170, 255–276.
- van Fraassen, B. (1980). *The scientific image*. Oxford: OUP.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.