Physical Theory and Physical Possibility

Sam Baron[*], Baptiste Le Bihan[†] & James Read[‡]

Abstract

It is plausible that the models of our scientific theories correspond to possibilities. But exactly which models of which scientific theories stand in this correspondence? The answers to this question hinted at so far in the literature are too restrictive: they don't support the idea that the models of many of our best scientific theories correspond to physical possibilities. The paper thus provides a novel proposal for guiding belief about physical possibilities based on physics. The proposal draws on the notion of an effective theory: a theory that applies very well to a particular, restricted domain. We argue that it is the models of effective theories that we should believe correspond, at least in part, to physical possibilities. It is thus effective theories that should guide modal reasoning in science.

0. Contents

---

[*] s.baron@unimelb.edu.au

[†] baptiste.lebihan@unige.ch

[‡] james.read@philosophy.ox.ac.uk

1. Introduction

The question of what is possible and what is not is at the heart of all philosophy and has been approached in various ways from different directions. The aim of this article is to contribute to the empiricist tradition in the philosophy of modality, which attempts to discover genuine physical possibilities in our scientific theories—and not in a faculty of *a priori* conceivability (see, e.g., Berenstain and Ladyman 2012; Ruyant 2019; Wilson 2020). Following this approach, and looking at a scientific theory, it is very natural to suppose that its models, here regarded as solutions to its equations, correspond to possibilities with, if the theory is to be empirically adequate, at least one model corresponding to what is actual. Van Fraassen articulates this idea clearly, when he writes that:

> To believe a theory is to believe that one of its models correctly represents the world. You can think of the models as representing the possible worlds allowed by the theory; one of these possible worlds is meant to be the real one. To believe the theory is to believe that exactly one of its models correctly represents the world (not just to some extent, but in all respects). (van Fraassen, 1980, p. 47)

This natural thought gives rise to an important question: if genuine possibilities are to be found in theories' models, exactly *which models* of *which theories* should we believe correspond to genuine possibilities? For one sense of 'possibility' the answer appears straightforward. If possibility means logical possibility, then presumably a model need only be consistent for it to correspond to a genuine possibility. We should thus believe that all consistent models of all theories correspond to logical possibilities, where 'correspond' here is understood as in the quote above: as the representation of worlds by models.

For a more restricted notion of physical possibility, however, the answer is less obvious. The physical possibilities, we might assume, are those situations that are consistent with the actual laws of nature (Earman 1986, p. 13; Maudlin 2007,

p. 18).[1] It is thus tempting to think that all theories that accurately capture the laws have models that correspond to physical possibilities (exactly what it means for a theory to accurately capture the laws is an issue to which we return below). As we will argue here, however, this way of specifying physical possibilities via the actual laws of nature faces serious problems. We thus suggest an alternative approach based on the notion of an *effective theory*, which aims to capture *effective laws*: general principles that apply only to a specific domain. We use those terms as generalisations of the way they are used in the context of effective field theories where 'specific domains' are identified to energy scales.[2] It is effective theories whose models we should believe correspond to physical possibilities.

The paper is structured as follows. We begin, in §2, by considering a law-based approach to linking theories to physical possibilities and explaining why this doesn't work. In §3 we consider an alternative approach, offered by Hazen. We show that Hazen's principle fares no better than the law-based approach. In §4, we consider an approach based on approximation. This is a step in the right direction, but ultimately faces difficulties of its own. In §5, we propose a way to overcome the problems facing other approaches, using effective theories. In §6, we address some potential difficulties with our approach, and in §7 we wrap up.

2. Laws

Physical possibility, we assume, is a mind-independent, objective modality. As noted, the physical possibilities are represented by those worlds that are compatible with the actual laws of nature. This suggests a straightforward method for determining which models of which theories correspond to physical possibilities: just look to the models of the theories that accurately capture the laws. We can capture this thought with the following principle:

---

[1] While we focus our attention on the laws, we don't presuppose any particular kind of realist interpretation of laws. What we do presuppose, however, is a kind of minimal realism about physical modality compatible with a variety of radically different approaches (such as, e.g., Cartwright 1983 and Berenstain and Ladyman 2012). See §6.6 for discussion.

[2] Although we borrow the term from the context of effective field theories, our purpose is tangential to the recent discussions of the relation between effective field theories and scientific realism (Williams 2019; Rivat 2021). We use the label 'effective' to refer to the fact that all empirically confirmed scientific theories are effective in a broader sense—i.e., that does not reduce the notion of validity domain of theories to the notion of energy domain.

> For any theory that accurately captures the laws, every model of that theory corresponds to a physically possible world.

The difficulty with this straightforward proposal, however, is that it is doubtful that *any* of our past or present scientific theories manage to accurately capture the actual laws of nature. First: what is it for a theory to capture the actual laws? A law of nature, presumably, is an exceptionless, non-accidental regularity that applies universally. A theory can be regarded as a class of models.[3] This class of models can generally be described using a set of equations or general principles. A theory accurately captures the actual laws of nature when the sets of equations or general principles correspond to the actual laws. In other words, the actual laws, when written down, just are the equations or general principles of a given theory that defines a class of models.

Now, while many of our theories imply universal claims that are candidates to be laws, none of them imply such claims that are true. The generalisations that we find in even our best theories do not hold universally or have exceptions. For example, consider general relativity. This is one of our most successful theories to date and does seem to provide universal principles for gravity. These principles, however, are known to break down at very high energies. Something similar is true for the Standard Model of particle physics—our most empirically successful quantum theory—which (being an effective field theory) is also only valid at certain energy scales. Developments in quantum gravity notwithstanding, at present we simply do not have a theory which can lay any claim to being exceptionless and/or universal.

We are not the first to point out that our scientific theories fall short of accurately capturing the actual laws of nature, whatever they may be. Both Cartwright (1999) and Lange (2012) offer similar examples. In both cases, they point to apparently universal claims, and show that they have exceptions. Cartwright, for instance, focuses on Newton's law of gravitation, and shows that it applies only under certain *ceteris paribus* conditions (e.g., at relatively low speeds, or in the absence of electromagnetic interference). Lange, by contrast, considers general principles of thermal expansion found in thermodynamics, and shows that these principles apply at least conditionally. For instance, general principles of thermal expansion apply to metal rods only in the absence of external factors influencing

---

[3] Here, we sidestep the debate between semantic versus syntactic formulations of physical theories. In light of the work of Lutz (2017), it is not obvious that this is a substantive debate in any case.

the system, such as mechanical impact on the rod (e.g., someone hammering the rod into place).

One might respond that while our current scientific theories don't manage to capture the actual laws, future theories will, and it is the models of those theories that correspond to physical possibilities. The difficulty with this suggestion, however, is that we need guidance on what is or is not physically possible *right now*, not at some putative 'end of science'. The reason for this is that physical possibility already plays a substantial role in scientific reasoning. To see this, let us briefly pause to consider some examples.

First, physical possibilities support counterfactual reasoning. Indeed, within the domain of science this kind of reasoning just is reasoning about, and in terms of, physical possibilities. Counterfactual reasoning is important for explanation. A prominent approach to explanation within the philosophy of science analyses explanation, at least partly, in terms of counterfactuals (see e.g. Woodward 2003). It is thus very natural to suppose that explanatory connections within science must be underpinned by counterfactuals.

Even if explanation is not analysed in counterfactual terms, counterfactual dependence is usually treated as a necessary condition on explanation. For x to explain y, it is plausible that x must make a difference to y, in the sense that if x had not occurred, y would not have occurred. At the very least, counterfactual dependence is a mark of the explanatory, and thus reasoning with counterfactuals is a powerful heuristic for reasoning with explanations.

Setting aside counterfactual dependence, and focusing purely on explanation, a further role for physical possibilities can be discerned. For there appear to be some explanations that appeal directly to underlying modal facts. Consider, for instance, equilibrium explanations, in the sense of Sober (1983). Very roughly, these explanations show that some actual state of a system is to be expected because all other possible states of the system are unstable. The gaps in the Rings of Saturn provide a case in point. The explanation for why there are gaps in the Rings relies on a harmonic analysis of the matter distribution in our solar system. This analysis reveals certain regions in the solar system within which nothing can stably orbit, without falling into some massive body or other. The explanation for the gaps is just that there is no possible way for an object to remain in orbit within the relevant region. This kind of explanation relies crucially on an underlying space of physical possibilities involving orbitals within the solar system.[4]

Symmetry explanations are another case in point. Some explanations appeal to the fact that certain features of a physical system are invariant under some salient transformation. The notion of invariance so crucial to symmetry explanations is essentially modal. Invariance of the kind needed for symmetries thus really only makes sense against the backdrop of some space of physical possibilities. For some, the fundamental quantities of physics just are those that remain invariant under changes, whereas non-fundamental quantities are variant (see e.g. Saunders 2003).[5] This notion of fundamentality forges a tight connection between physical and modal notions, making it quite important to specify the link between physics on the one hand, and physical possibility on the other.

Setting aside explanation, physical possibility is also important to the distinction between determinism and indeterminism. Earman, for instance, specifies a notion of determinism as follows:

> Letting $\mathcal{W}$ stand for the collection of all physically possible worlds, that is, possible worlds which satisfy the natural laws obtaining in the actual world, we can define Laplacian variety of determinism as follows. The world $W \in \mathcal{W}$ is *Laplacian deterministic* just in case for any $W' \in \mathcal{W}$ if W and W' agree at any time, then they agree for all times. (Earman, 1986, p. 13)

Earman's notion of determinism presupposes a collection of possible worlds. It thus relies, at least tacitly, on a principle that connects scientific theories to possibilities. Statements of determinism quite generally are likely to be similar in this respect. Determinism roughly speaking relies on the idea that there is only one possible way for the world to be, given certain pre-specified conditions. In order to determine whether we should endorse determinism, a space of physical possibilities is required, even if just to narrow the number of elements of that

---

[4] Tegmark (1997), Wolf and Thébault (2023), and Swanson (2022) deploy something resembling equilibrium explanations to account for the four-dimensionality of our universe, the explanatory depth of inflationary models in cosmology, and the ubiquity of second derivatives in physics, respectively.

[5] Think also of Noether's theorems: conserved quantities are associated with symmetries of an action. See Read and Teh (2022) for recent work on the philosophy and physics of Noether's theorems.

space down to one. Again, some link between theory and physical possibility is required.[6]

Physical possibilities also show up in statistical mechanics. For instance, according to the popular 'Boltzmannian' approach, one can characterise entropy as follows: the (Boltzmann) entropy of a macrostate is a measure of the number of possible microstates with which it is compatible. This Boltzmannian notion of entropy clearly makes use of a notion of possibility, one that it seems natural to call physical possibility. Indeed, the alternative 'Gibbsian' approach to statistical mechanics countenances ensembles of physical systems that evolve in a probabilistic way, and thus would also appear to trade in notions of physical possibility. Thus, contemporary statistical mechanics would appear to presume that one can make sense of these notions of physical possibility.

It is easy to multiply examples. As Wilson puts it:

> It is a truism that physics involves reasoning about alternative physical possibilities. To explain and predict actual observations, we construct models of the phenomena which represent alternative physically possible histories; perhaps we also assign a probability distribution over these histories. This essentially modal character of physics has been emphasized by authors as varied as Sellars, Suppes, Cartwright, Ladyman and Ross, and Maudlin. (Wilson, 2021, p. 1113)

We won't labour the point any further. We rely extensively on reasoning about alternative physical possibilities, and we therefore need to find a procedure for guiding belief about physical possibilities on the basis of our current theories. What's clear is that we cannot focus on those theories that get the actual laws right, since none of our theories provide true, universal exceptionless generalisations. Some other approach to working out which models of which theories correspond to physical possibilities is therefore required.

3. Hazen's Principle

In correspondence with Lewis, Hazen offers an alternative to a law-based approach to guiding beliefs about possibilities. Hazen writes:

---

[6] Readers familiar with the hole argument—see Pooley (2021) for a recent review—won't need to be convinced of the significance of physical possibility in discussions of determinism.

> Every mathematically standard model of any coherent comprehensive physical theory corresponds to a genuine possible world. (Hazen in Lewis 2020)

Call this *Hazen's Principle*. Hazen's Principle faces some obvious difficulties. To begin with, it is unclear what a 'mathematically standard' model of a theory is supposed to be. This could mean any number of things, depending on the standards being used. A mathematically standard model could be the one most commonly used in science. Or it could be the one that deploys a widely-used mathematical approach. Or it could simply be a model that is most commonly discussed. Without a way to determine what it is for a model to be mathematically standard, there is little hope of identifying the models of a theory that correspond to possibilities. More than this, however, it is difficult to see what justification there might be for focusing only on the mathematically standard models when it comes to possibility. Why should the fact that a model is mathematically standard be a necessary condition for its corresponding to a possibility?

A second, more serious problem with Hazen's Principle concerns the notion of a 'complete and comprehensive' theory. This seems troubling, for a couple of reasons. First, it's unclear what 'complete' means and, anyway, it's not clear that we have any such theories. 'Complete' seems to suggest some kind of final theory, or at least final theory of a particular topic. E.g., a final theory of physics. If that's what 'complete' means, though, then it is plausible that none of our current theories qualify, for the reasons which we have already articulated. This means that current science provides no guide to what is possible.

It could be that 'complete' does not reference a final theory, but only a theory that provides a total description of some phenomenon. This would effectively rule it out that any of the models of either general relativity or quantum field theory correspond to possibilities. For while these are our two most successful physical theories to date, they are known to be incomplete. For instance, as already indicated, general relativity fails to provide an account of gravity at small scales (i.e. high energies), and so in that sense is not a complete theory. We take it to be a problem for a principle linking models to possibilities if it cleaves incredibly successful theories from physical possibility.

What about 'comprehensive'? This seems to signal generality. A comprehensive theory will be a maximally general theory, a putative theory of everything. But, again, we have no such theory, so we'd be waiting for the end of science to get guidance on what we should believe the physical possibilities to be. And, again, neither general relativity nor the Standard Model would provide a guide to the physical possibilities, since neither theory is fully general.

4. Approximation

We are thus still in need of a method for working out which models of which theories correspond to physical possibilities. An obvious way forward is to use the notion of approximation. Consider, again, the law-based approach. There the idea was that all the models of those theories that accurately capture the laws correspond to physical possibilities.

A weaker principle that uses approximation, such as the following, would seem to do better:

> For any theory that approximately captures the laws, every model of that theory approximates a physically possible world.

This weaker principle involves approximation twice over. First, instead of requiring that our theories accurately capture the laws, the weaker principle requires only that our theories approximately capture the laws (by which we mean: the principles that define a class of models approximate the actual laws of nature). This is an easier condition to satisfy, and one that will generally avoid the kinds of problems discussed in §2. Given said approximate representation, however, the correspondence to physical possibilities can itself at best be approximate. A model merely approximates a world when there is some degree of mismatch between the model and the world. We take it that this could happen in many different ways, depending on the model and the world.

While a step in the right direction, the approximation principle faces problems of its own. The main problem is that we have no idea how close our current theories are to approximating the laws, in the above sense. This is a problem, because the looser the connection between the theory and the actual laws of nature, the less assured we can be that the models of the theory approximate physical possibilities. We thus need to know how close we are to the actual laws before we can come to a reasonable view about whether the models approximate physical possibilities.

This uncertainty about whether the models approximate physical possibilities gets in the way of modal reasoning in science. For presumably we don't just want to be able to draw modal inferences using our scientific theories about what is or is not possible, but in addition we want those inferences to be reliable. Until we know how closely our theories approximate the laws, however, we have no way of determining the reliability of our modal inferences drawn on the basis of those theories. This introduces a high level of uncertainty into the use of models as a guide to possibility.

One might respond that we can be sure that some of our theories closely approximate the laws because those theories are approximately true. So, for instance, one might say that general relativity is approximately true, and because of this we should believe that any universal principles issuing from the theory closely match the actual laws. But this merely shifts the bump in the carpet. For we simply do not know how close general relativity is to being true and, in particular, we don't have a precise measure of how close any universal claims it implies might be to the actual laws. Thus, the shift to focusing on theories that are approximately true does nothing to address the uncertainty around possibility that arises from the approximation principle.

A further problem with the approximation principle is that there seem to be at least some scientific theories that aren't even approximately true, but whose models we commonly take to correspond to physical possibilities. Take, for instance, Newtonian theories of gravity. *Prima facie*, Newtonian gravitation does not seem to be an approximately true theory, if what we are looking for is a theory that captures the actual universal laws. This is because Newtonian gravity is superseded by general relativity, which is considered to be a much more accurate account of gravity; moreover, of course, Newtonian mechanics is not a quantum mechanical theory. Nonetheless, we still use the models of Newtonian gravity for a variety of purposes, to give us information about physical possibilities. Some (indeed, most) models of Newtonian gravity correspond to at least parts of the actual world. Consider, for instance, the solar system (perhaps *sans* Mercury): this system is modelled well by Newtonian gravity, and indeed different models of Newtonian gravity can provide possible ways for regions of the world like this (i.e., other stellar systems, of which there are of course many!) to be.[7]

The problem, then, is that even if we can somehow work out which theories approximate the actual laws, and then claim that the models of those theories correspond to physical possibilities, this would leave us with no way to accommodate the apparent success of less accurate theories that somehow still seem to provide good guides to physical possibility.

---

[7] In fact, there are well-known problems with treating Newtonian gravity as a cosmological theory: see Malament (2012, ch. 4) and Wallace (2017).

5. Effective Theories

Thus, the move to approximation doesn't fully solve the problem. The trouble is that we don't have a good way of measuring the distance between our theories and the actual laws. Perhaps, however, the focus on laws as standardly understood is the problem. Thus far we have been focusing on exceptionless regularities that apply universally. While these laws may determine the scope of physical possibility, we don't have scientific theories that accurately track them or, at least, we are uncertain as to whether any of our theories do.

Our suggestion, then, is to focus on *effective laws* instead. An effective law is not a universal, exceptionless regularity. Rather, an effective law is a regularity that holds exceptionlessly *for a specific domain*. Domains are specified in terms of the conditions that hold within them. For instance, a region in which the gravitational field is very weak, and everything is moving at low speeds relative to the speed of light, corresponds to a particular domain. Similarly, a particular energy scale corresponds to a domain, such as the extremely high energies at which general relativity is thought to break down, or the comparatively low energies at which general relativity applies.

Effective laws are grounded in universal laws. Which is to say that the universal laws plus the conditions which specify a given domain ground the truth of the effective law for that domain. Effective laws thus correspond to only a *part* of the world, and thus not to the world as a whole.

An effective theory is a theory that captures the effective laws for a domain either exactly or to a high degree of approximation. Our proposal, then, is to focus on effective theories. The idea is that the models of effective theories approximate parts of genuine physical possibilities.

Thus, we offer the following principle:

> For any model M of an effective theory T, M approximates a part of a physically possible world.

This principle is an improvement on the principles considered thus far. For while it is difficult to work out whether our theories approximate the actual laws, it is easier to determine whether our effective theories approximate the effective laws. If an effective theory is extremely well-supported empirically for a given domain, and there are no known counterexamples to the general principles that issue from that theory as applied to just that domain, then we have good evidence that the theory is managing to fully capture the effective laws for the domain in question. For if the theory were somehow failing to capture these effective laws, then we should expect there to be empirical data within that domain that the laws cannot accommodate.

Given this, we have strong evidence that some of our current theories capture actual effective laws. Take, for instance, general relativity. The principles of gravitation from general relativity don't hold universally, for all domains. However, we have good evidence that they hold universally for a particular domain, namely certain energy scales. The same point applies for the Standard Model of particle physics, for Newtonian gravity, or indeed for essentially any other extant empirically successful theory of physics to which one can point.

We are claiming that strong accord with empirical data is good evidence that a theory is an effective theory: one that approximates, to a high degree, the effective laws for a particular domain. Given this, why not just say that strong accord with empirical data is good evidence that a theory accurately captures the actual universal laws, not just the actual effective laws? The answer is that in the case of effective laws, we can be confident that there is very little, if any, extra empirical data to be gathered that would undermine the accuracy of the theory for a given domain.[8] Thus, we have good reason to suppose that the principles articulated by the theory won't need to be modified much if at all to capture the actual effective laws. In the case of a theory and the universal laws, however, we have no similar assurance. There may well be a great deal of empirical data that would require us to change the general principles articulated by a theory in order to get it to match the actual laws.

Take general relativity again. This theory does not capture fully the actual laws, as already explained. We know that there are facts about high-energy domains that will need to be accommodated. We currently have no idea about what those facts might be or about the scope of the revisions to general relativity which might be required in order to bring it into line with the actual laws. As a result, we have significant uncertainty about how close general relativity is to capturing laws of gravitation. If, however, we focus just on the low energy domain, we know that most of the data relevant to general relativity in this domain has been gathered already. Thus while general relativity might not get the effective laws for this domain exactly right, we can be comparatively confident that it is quite close. In particular, we can be confident that no new information about this domain will

---

[8] Cf. Hoefer (2020): "*the experiments have all been done already, and current theory is confirmed by them*" (emphasis in original).

come in and radically undermine the theory.[9]

Our principle thus avoids one of the problems facing the approximation principle. It also overcomes the other problem, namely that of explaining how the models of (say) Newtonian gravity might correspond to physical possibilities despite Newtonian gravitation not being true to a very high degree. For even though Newtonian gravitation might not provide a very good picture of the actual universal laws, as already mentioned it does provide an accurate picture of effective laws for a specific domain. In particular, for low-energy domains in which the gravitational field is weak and velocities are low, Newtonian gravitation is highly accurate, and very well-supported empirically. Moreover, we can be confident that new data won't overturn this fact, and so we can be confident that Newtonian gravitation is accurately capturing effective laws. It is for this reason that its models correspond to physical possibilities, albeit approximately and only in the sense of corresponding to parts of possible worlds.

The shift from correspondence with possible worlds to correspondence with parts of worlds might seem troubling. Granted, we are able to reduce uncertainty about whether the parts of worlds at issue are parts of physical possibilities, but we do so apparently at the cost of our capacity to draw reliable inferences about complete worlds. We admit as much: our principle provides no way to work out which models of which theories correspond to complete worlds. But we don't see this as a problem. For we rarely if ever use science to draw modal inferences about worlds as a whole. We are typically interested in possibilities in a narrower sense. That's because, when we draw modal inferences from science, we generally draw those inferences with respect to specific domains. We are therefore confident that our principle is sufficient to guide beliefs about which models of which theories correspond to physical possibilities, at least insofar as we want those beliefs to be both reliable and useful in modal reasoning within science.

Before we consider some difficulties facing our principle, it is worth identifying another advantage. So far, we have focused primarily on theories within physics. It is plausible, however, that the models of theories outside of physics also correspond to physical possibilities, at least some of the time. Take biology, for instance. It seems plausible that models of evolution by natural selection correspond to physical possibilities. Principles that focus on universal, exceptionless laws will find it hard to accommodate this fact. That's because it is unlikely that

---

[9] Leading on from the previous footnote: this is a point which has been made with great elegance by Hoefer (2020).

biological theories capture laws of this kind even approximately. Indeed, there may not even be such laws for biology. It is plausible, however, that there are effective laws of biology: laws that apply to certain domains (such as for all biological organisms) that are ultimately grounded in the universal laws at some level.

As such, our principle applies equally well to biological theories, or indeed to any effective theories in any domain, as to theories from physics. For extremely well-confirmed theories of any domain where we can be confident that new data won't undermine the theory, we can be relatively certain that the models of the theory approximate parts of physical possibilities. For theories that are less well-confirmed, we can be less confident of the modal inferences drawn from those theories. But that is as it should be: it seems right for theories that enjoy less evidential support to be less of a guide to what is possible.

6. Problems

So far we have introduced a new principle that aims to determine which models of which theories correspond to physical possibilities. In this section we consider and respond to some potential problems.

6.1. Not Every Model

The principle, as stated, is totalising: *every* model of an effective theory approximates part of a physically possible world. This, one might argue, is a problem. Even for effective theories, one might argue, we should allow that some of their models don't correspond to any genuine physical possibilities. Here are a few examples from the philosophy of spacetime in particular:

i. In the case of the hole argument of general relativity, all three of Butterfield (1989), Maudlin (1988), and Rynasiewicz (1994) hit upon the idea that, having fixed a representational context, only one of a class of models of general relativity related by hole diffeomorphisms represents a physical possibility, while the other models in that class represent nothing at all.

ii. Those inclined towards relationalism about spacetime might aver that vacuum solutions of general relativity do not represent physical possibilities, for there can be no spacetime in the absence of matter. (Pooley 2001, for example, explores this idea.)

iii. One might maintain, *à la* Earman (1995), that solutions of general relativity which admit of closed timelike curves and which violate certain consistency conditions (designed to rule out grandfather-type paradoxes) do not represent physical possibilities.

iv. Relatedly, 'Machian relationalists' such as Barbour (2012) maintain that not all solutions of Newtonian gravity represent genuine physical possibilities—rather, only those for which the total angular momentum of the universe vanishes do so. (For philosophical discussion of such approaches, see Pooley and Brown (2002).)

These kinds of cases can be handled, to some extent, using the notion of approximation already built into our principle. As noted, the models of an effective theory correspond to parts of physical possibilities by approximating them. This means that there can be some measure of mismatch between the models and the physical possibilities. This mismatch leaves it open that some aspects of a model don't correspond to anything in a given physical possibility.

This is, perhaps, what is happening with case (iii), if we follow Earman. Solutions to general relativity that feature closed timelike curves would only approximate physical possibilities insofar as they get the topology wrong. One way in which they may get the topology wrong is through the inclusion of closed timelike curves. Nevertheless, such solutions may still approximate physical possibilities to some degree, for (say) the local physics described in such solutions may still agree with the local physics in (a region of) a world devoid of closed timelike curves.

The other cases are harder. Take, for instance, Pooley's view on vacuum solutions to general relativity. If he's right, then at least some of the models of general relativity don't approximate even parts of physical possibilities.[10] If this can happen, then uncertainty about the reliability of our modal inferences starts to creep back in. For if there are some models that don't approximate any physical possibilities at all, how confident should we be that any particular model approximates a physical possibility?

In order to address this issue, we propose the following modification of our principle: we should believe that all models of an effective theory approximate

---

[10] One could push back against this line, by e.g. arguing that vacuum solutions still represent *parts* of physical possibilities—e.g., Minkowski spacetime can still represent parts of solutions to e.g. Maxwell's equations. In which case, the worries and responses to follow are moot.

physical possibilities, in the absence of countervailing philosophical principles or evidence from a more fundamental theory.[11] Which philosophical principles one endorses is, of course, a matter of choice. The thought, then, is that the models of an effective theory are innocent until proven guilty: our default position is to assume that they all approximate physical possibilities.[12] This default position, however, is subject to specific defeaters, though what those defeaters might be is considered on a case-by-base basis.

What if there are always some defeaters? Then we can slightly weaken our principle by including a second dimension of approximation. We can thus add that the space of models is itself an approximation to a space of models that captures the physical possibilities exactly. A weakened principle along these lines can be stated as follows:

> For most models M_i of an effective theory T, each of the M_i approximates a part of a physically possible world.

This weakened principle reduces the reliability of modal inferences drawn on the basis of science, since it introduces a second source of error. The first source of error consists in the fact that models now merely approximate parts of physical possibilities. The new source of error is based on the fact that some models may not approximate parts of physical possibilities at all. While this second source of error reduces the reliability of modal inferences drawn using science, it doesn't undermine them completely. For it is still the case that most models of effective theories approximate parts of physical possibilities, and so our inferences drawn using those theories are still generally reliable. Moreover, one can still get a handle on the subset of models which (approximately) represent (parts of) physical possibilities, once one is explicit about the above-mentioned countervailing philosophical principles in play. That is still enough to capture the importance of modal inference to science.

---

[11] What is it for one theory to be more fundamental than another? Whatever the account of relative fundamentality might be, we can plug it in here. See Crowther and Linnemann (2019) for recent discussion.

[12] What we propose is in the spirit of (a modalised version of) the 'totalitarian principle', popularised in physics by Gell-Mann (1956): "everything not forbidden is compulsory".

6.2. Approximation

Our principle uses a notion of approximation. However, we haven't said much about what this is. One might object, then, that our principle is intolerably vague.

As we see it, there are at least three ways in which models may 'correspond' to worlds:

1. A model corresponds to a world if it is isomorphic to a complete possible world.[13]
2. A model corresponds to a world if it is isomorphic to a part of a complete possible world.
3. A model corresponds to a world if it represents either a part of a complete possible world or a complete possible world to some degree.

Our principle uses the weakest notion of correspondence. Thus, when we talk of *approximation* we mean *representation to some degree*. What is it for a model to represent to a degree? It depends on how models represent.[14] This is a difficult issue, and not one we aim to discuss here for there are many accounts available. The point is just that whatever one's account of the representation relation, it is a desideratum on such accounts that they be able to make sense of more or less accurate representations. That, however, is all we need for approximation relations between models and (parts of) worlds.

If approximation is understood in terms of degree of representation, then a further question arises: to what degree should we believe that the models of an effective theory represent and thus approximate physical possibilities? Our answer is that we should believe that the models of an effective theory approximate physical possibilities to a high degree, in the absence of any countervailing considerations (of the kind that might lead us to believe that the models of a theory don't approximate physical possibilities at all, as discussed above).

What justifies this claim? The claim is justified by the fact that the models at issue are models of an effective theory, which is highly accurate for a domain. When a theory manages to capture the effective laws for an actual system, it seems

---

[13] There are serious questions regarding what it could mean for models to be 'isomorphic' to worlds—see e.g. van Fraassen (2002). We register these as serious and set them aside.

[14] See Frigg and Nguyen (2017) for an overview of various approaches.

plausible to suppose that its models will also do a good job of capturing possibilities for that system (again, in the absence of reasons to think otherwise). In short, the actual is a good guide to what is possible. We admit that this connection between actuality and possibility could be rejected. If one rejects the connection, however, then it is unclear why we should take any theory to be a guide to possibility in any sense. Rejecting the connection would thus lead to a version of modal scepticism about physical possibility. Since we reject scepticism of this kind, we take it to be a reasonable assumption that the models of effective theories approximate physical possibilities to a high degree.

6.3. Effective and *Ceteris Paribus* Laws

Thus far we have employed the notion of effective laws. One might argue, however, that effective laws are nothing but *ceteris paribus* laws. The *ceteris paribus* conditions on laws, however, are notoriously difficult to spell out and, in the end, may collapse into triviality. There is thus a real risk that our account will inherit these problems.

Effective laws are different from *ceteris paribus* laws, however. A *ceteris paribus* law is generated by taking a universal exceptionless regularity and adding extra conditions that constrain the application of the law. An effective law, by contrast, is not a way of modifying an existing universal law with extra conditions. It is a non-universal law, one that holds as an exceptionless regularity for a specific domain. The effective laws are grounded in the universal laws, whereas *ceteris paribus* laws are supposed to be replacements for universal laws.

Unlike *ceteris paribus* laws, effective laws are not difficult to spell out. In order to spell out an effective law, we simply identify a particular domain of application, and then state a general principle for that domain. Since we can specify domains of application by identifying properties of physical systems (such as energy scales, gravitational strength, and so on) there is no problem with stating effective laws. There is thus no reason to suppose that the problems that plague *ceteris paribus* laws will apply to effective laws as well.

6.4. Beyond Effective Theories

The principle that we have outlined focuses on effective theories. This principle represents an advance over existing principles connecting theories to possibilities (such as Hazen's Principle) in part because it makes room for theories like Newtonian gravitation to correspond to physical possibilities. One might worry, however, that the principle doesn't go far enough. For, one might argue, there are theories that aren't even effective theories but which still manage to correspond to physical possibilities.

One way to press this point is by appealing to theories outside of physics. We have already shown how some biological theories can be accommodated. But, one might argue, there are many biological theories that cannot be accommodated, because they can't even be thought of as highly accurate effective theories. Thus, our principle is not sufficiently liberal.

We admit that we struggle to come up with plausible examples that might support this objection. For note that effective laws can be *highly specific*. An effective law is highly specific if it applies only to a very small part of the universe. One example might be a general principle that captures the breeding cycle of a particular species of finch in the Galápagos. An effective theory for this domain can accurately capture a highly specific effective law, that nonetheless generalises across the whole domain.

Because effective theories can be highly specific in this sense, we find it hard to see how many theories will be left out of the picture, regardless of which particular area of science they might come from. For keep in mind that effective theories must be very well-confirmed for their domain. Indeed, the level of confirmation is quite demanding: we must be fairly confident that no new data will require significant revisions to the theory. Any theory like that, no matter how specific, will count as an effective theory.

Thus, the only theories that our principle won't capture will be theories that are not well-confirmed empirically. But we don't see that as a problem. For we really should be careful about drawing inferences from such theories (modal or otherwise), and so it seems reasonable to be sceptical that the models of these theories correspond to physical possibilities.

6.5. Grounding of Effective Laws

We have claimed that effective laws are grounded in universal principles. We have also said that effective laws can be highly specific. One might object, however, that these two points are in tension, or at least that it's not completely clear how this can be done. Take, for instance, the example of the finches mentioned above. Suppose we come up with an effective theory of the finches, which accurately captures an effective law. It's implausible, one might argue, to suppose that this effective law is grounded in a set of universal laws. The universal laws, whatever they might be, will be far too general, and the effective finch law will be far too specific. Effective and universal laws would involve too different sorts of things, jeopardising the possibility of grounding the first into the latter.

We admit that there might be some implausibility to the idea that the finch law is *directly* grounded in the universal law *alone*. But, firstly, effective laws can, and likely will be, indirectly grounded in universal laws, via legal intermediaries. In particular, each effective law will be grounded in another effective law that is perhaps more general, applying to a domain that is broader than but includes the domain covered by the highly specific effective law. Indeed, as we see it there will be a chain of effective laws that issue from the most basic, universal laws, whatever they might be to the most specific effective laws. The models of theories that capture effective laws at each step of the chain will correspond to physical possibilities, as and when they are developed.

Secondly, that effective laws are grounded in universal laws doesn't mean that universal laws are their *sole* ground. Rather, effective laws are *partially* grounded in the universal laws. Other factors contribute to the grounding of effective laws as well: the concepts used to articulate the target domain of investigation, and the various properties of the effective domain itself. Consider our finches again: effective laws will only be partially grounded in the universal laws. In addition, they will be partially grounded in the taxon of finches, and in the specifics of their ecological niche.

6.6. Realism About Laws

On the face of it, our principle presupposes realism about the laws of nature. That is, it assumes that there are universal, exceptionless generalisations which (i) define the set of physically possible worlds, and (ii) provide the grounding basis for effective laws. Not everyone accepts realism about the laws, however. For those who don't (e.g., Cartwright 1983 or van Fraassen 1989), neither the metaphysical picture of the set of physically possible worlds nor the principle we have provided will seem very plausible.

Ultimately, nothing we have said here relies on any kind of realism about the laws of nature. Realism about the laws has just been a convenient foil for considering the question of whether and how scientific theories guide beliefs about physical possibilities. Our principle works even if one rejects realism about the laws of nature.

To see this, note that if one gives up on realism about the laws, then one must provide an alternative metaphysical account of what physical possibility is, since one can no longer rely on the idea that physical possibility is just consistency with the actual laws of nature. As we see it, there are then two broad options for how to specify physical possibility.

First, one could take the set of physically possible worlds to be the set of worlds which are consistent with facts about the actual world. On this picture, the relevant facts are not laws, because (of course) there aren't any. Second, one could take the set of physically possible worlds to be primitive. On this view, it is not the case that the physically possible worlds are those that are consistent with actual facts. There is just a primitive modal fact that these are the physically possible worlds.

The first option is entirely compatible with our principle. For even if one rejects realism about laws, one can still say that there are effective theories and that these effective theories manage to capture effective principles (perhaps we don't want to call these effective laws either). Then these effective principles can describe facts about the world that are grounded in whatever facts define the space of physical possibilities. On this view, the models of effective theories provide the same guide to physical possibilities as in the case of realism about the laws.

The second option—in which there are primitive modal facts about what is physically possible—is harder to square with our approach. These primitive modal facts may come apart from the actual facts entirely. It is thus unclear how effective theories that capture actual facts could provide much of a guide to primitive physical possibilities. But notice that if that's right, then it is doubtful that science could guide belief about physical possibility in a reliable way at all. That's because our scientific theories are developed in terms of actual facts, and so the only way that they could track physical possibilities is by capturing actual facts about the world. This, we submit, provides a reason to reject primitivism about physical possibility, since it would make a mystery of the reliability of modal inferences that are based on science.

This is hardly conclusive, and there is likely more to say about primitivist approaches to physical possibility. What matters is that realism about the laws of nature is not essential to our proposal, and so that proposal is open to anti-realists about the laws as well.

6.7. Change Over Time

New theories are developed every day. Some of these end up being novel effective theories. It follows that over time we must revise our views about what we take the physical possibilities to be. But, one might continue, that's implausible: what is physically possible should not change over time.

Our response is to emphasise that the principle we seek, and indeed take ourselves to have identified, is supposed to be a guide to belief. It is a method for working out what we should believe the physical possibilities to be, by scrutinising our scientific theories. It is not supposed to be some magical modal telescope that allows us to peer into modal space at some fixed class of possible worlds.

As a guide to what to believe based upon our best extant science, we should in fact expect the principle to yield different results at different times. Moreover, this seems quite natural for physical possibility: we can and do change our views about what is physically possible as we learn more about the world. It is thus a feature of our principle that it allows for belief revision in this manner.

Of course, we grant that there is a distinction to be drawn between what one believes to be physically possible and what really is physically possible. The best we can hope for is to develop a principle that gives us reliable beliefs about what is physically possible at each point in time. As far as we can tell, our principle achieves this.

6.8. Epistemic Possibility

Our principle is supposed to guide belief. In particular, it tells us which models of which theories we should believe correspond to physical possibilities. By providing a principle that is a guide to belief, one might point out, we are at best characterising a notion of epistemic possibility. Our goal, however, is to provide a principle for specifying non-epistemic, physical possibilities. So we have failed to achieve our goal.

But just because something is a guide to belief about possibility does not automatically turn the possibilities at issue into epistemic possibilities. To think otherwise is to suppose that there is no way to guide belief about non-epistemic possibility, which seems implausible.

Moreover, note that epistemic and non-epistemic possibilities play different roles in theorising. Thus, whether our principle delivers a guide to belief about non-epistemic possibilities depends on what role the possibilities being delivered are playing. So long as our principle guides belief about possibilities that can play the role of non-epistemic possibilities in theorising, that seems like reason enough to suppose it is guiding belief about possibilities of the right kind.

One role that non-epistemic possibilities play is in supporting counterfactuals of a certain type, namely counterfactuals about how matters would be, were certain changes to be made. The possibilities that are delivered by our principle support counterfactuals of this kind, so far as we can tell, which is reason enough to suppose that our principle is not a guide to mere epistemic possibility.

7. Conclusion

It's time to take stock. In this article, we've sought to identify a principle that can guide beliefs about physical possibility based on scientific theories. We concluded that current offerings (e.g. Hazen's Principle) are by and large too restrictive: they present a rather narrow picture of when the models of theories correspond to physical possibilities. We thus proposed an alternative. On our alternative, it is the models of effective theories that correspond to physical possibilities by approximating parts of worlds.

This alternative is superior for three reasons. First, it does not require theories that capture universal laws. It only requires theories that capture effective laws, which is a less demanding requirement. Second, unlike other options, our principle allows that the models of theories like Newtonian gravitation and evolution by natural selection correspond to physical possibilities. Third, our principle can, with minimal adjustments, accommodate cases in which the models of a theory fail to correspond to any physical possibility, even in part.

Note that this is a first stab at trying to settle on an appropriate principle to guide belief about physical possibility, and so work remains to be done. The principle needs to be tested against a larger range of theories in both physics and science more generally to ensure that it continues to deliver plausible results. We also need to consider the principle in light of the full range of ways in which physical possibilities are used in science, to ensure that it can appropriately scaffold scientific practice. It is plausible that the principle will require further refinement going forward. In this way, what we have said provides a foundation for future work on this topic.

8. Acknowledgements

We thank Al Wilson for exposing us to Hazen's Principle and for some discussions related to this topic.

8. References

Barbour, J. (2012). Shape dynamics. An introduction. In F. Finster, O. Müller, M. Nardmann, J. Tolksdorf, and E. Zeidler (Eds.), *Quantum Field Theory and*

Gravity, pp. 257–297. Basel: Springer.

Berenstain, N. and J. Ladyman (2012). Ontic structural realism and modality. In E. Landry and D. Rickles (Eds.), *Structural Realism: Structure, Object, and Causality*. Springer.

Butterfield, J. (1989). The hole truth. *British Journal for the Philosophy of Science 40*(1), 1–28.

Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford: Oxford University Press.

Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.

Crowther, K. and N. Linnemann (2019). Renormalizability, fundamentality, and a final theory: The role of uv-completion in the search for quantum gravity. *British Journal for the Philosophy of Science 70*(2), 377–406.

Earman, J. (1986). *A Primer on Determinism*. Dordrecht: Springer.

Earman, J. (1995). *Bangs, Crunches, Whimpers, and Shrieks: Singularities and Acausalities in Relativistic Spacetimes*. New York: Oxford University Press.

Frigg, R. and J. Nguyen (2017). Models and representation. In L. Magnani and T. Bertolotti (Eds.), *Springer Handbook of Model-Based Science*, pp. 25–48. Dordrecht: Springer.

Gell-Mann, M. (1956). The interpretation of the new particles as displaced charge multiplets. *Il Nuovo Cimento (1955-1965) 4*(2), 848–866.

Hoefer, C. (2020). Scientific Realism without the Quantum. In S. French and J. Saatsi (Eds.), *Scientific Realism and the Quantum*, pp. 19–34. Oxford University Press.

Lange, M. (2012). There sweep great general principles which all the laws seem to follow. *Oxford Studies in Metaphysics 7*, 154.

Lewis, D. (2020). Letter from Allen P. Hazen to David Lewis, 6 march 1983. In H. Beebee and A. R. J. Fisher (Eds.), *Philosophical Letters of David K. Lewis: Volume 1: Causation, Modality, Ontology*, pp. 320–321. Oxford: Oxford University Press.

Lutz, S. (2017). What was the syntax-semantics debate in the philosophy of science about? *Philosophy and Phenomenological Research 95*(2), 319–352.

Malament, D. B. (2012). *Topics in the Foundations of General Relativity and Newtonian Gravitation Theory*. Chicago: Chicago University Press.

Maudlin, T. (1988). The essence of space-time. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association 1988*, 82–91.

Maudlin, T. (2007). *The Metaphysics within Physics*. Oxford: Oxford University Press.

Pooley, O. (2001). Relationism rehabilitated? II: Relativity. Unpublished manuscript. `http://philsci-archive.pitt.edu/221/`.

Pooley, O. (2021). The hole argument. In E. Knox and A. Wilson (Eds.), *The Routledge Companion to the Philosophy of Physics*, pp. 145–159. New York: Routledge.

Pooley, O. and H. R. Brown (2002). Relationalism rehabilitated? I: Classical mechanics. *British Journal for the Philosophy of Science 53*(2), 183–204.

Read, J. and N. J. Teh (2022). *The physics and philosophy of Noether's theorems*. Cambridge: Cambridge University Press.

Rivat, S. (2021). Effective theories and infinite idealizations: A challenge for scientific realism. *Synthese 198*(12), 12107–12136.

Ruyant, Q. (2019). Structural realism or modal empiricism? *British Journal for the Philosophy of Science 4*(70), 1051–1072.

Rynasiewicz, R. (1994). The lessons of the hole argument. *British Journal for the Philosophy of Science 45*(2), 407–436.

Saunders, S. (2003). Physics and Leibniz's principles. In K. Brading and E. Castellani (Eds.), *Symmetries in Physics: Philosophical Reflections*, pp. 289–307. Cambridge: Cambridge University Press.

Sober, E. (1983). Equilibrium explanation. *Philosophical Studies 43*(2), 201–210.

Swanson, N. (2022). On the ostrogradski instability; or, why physics really uses second derivatives. *British Journal for the Philosophy of Science 73*(1), 23–46.

Tegmark, M. (1997). On the dimensionality of spacetime. *Classical and Quantum Gravity*, L69–L75.

van Fraassen, B. C. (1980). *The Scientific Image*. Oxford: Clarendon Press.

van Fraassen, B. C. (1989). *Laws and Symmetry*. Oxford, England: Oxford University Press.

van Fraassen, B. C. (2002). *The Empirical Stance*. New Haven: Yale University Press.

Wallace, D. (2017). More problems for Newtonian cosmology. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics 57*, 35–40.

Williams, P. (2019). Scientific realism made effective. *British Journal for the Philosophy of Science 70*(1), 209–237.

Wilson, A. (2020). *The Nature of Contingency: Quantum Physics as Modal Realism*. Oxford: Oxford University Press.

Wilson, A. (2021). Counterpossible reasoning in physics. *Philosophy of Science 88*(5), 1113–1124.

Wolf, W. J. and K. P. Y. Thébault (2023). Explanatory depth in primordial cosmology: A comparative study of inflationary and bouncing paradigms. *British Journal for the Philosophy of Science*.

Woodward, J. (2003). *Making Things Happen*. New York: Oxford University Press.