

Duality and Ontology*

Baptiste Le Bihan^{a1} and James Read^{b,c2}

^a Philosophy Department, University of Geneva
2 Rue de Candolle, 1205 Geneva, Switzerland

^b Faculty of Philosophy, University of Oxford
Radcliffe Humanities, Oxford OX2 6GG, United Kingdom

^c Department of Philosophy, University of Illinois at Chicago
1423 University Hall (M/C 267), 601 South Morgan Street
Chicago, IL 60607, United States

Abstract

A ‘duality’ is a formal mapping between the spaces of solutions of two empirically equivalent theories. In recent times, dualities have been found to be pervasive in string theory and quantum field theory. Naïvely interpreted, duality-related theories appear to make very different ontological claims about the world—differing in e.g. space-time structure, fundamental ontology, and mereological structure. In light of this, duality-related theories raise questions familiar from discussions of underdetermination in the philosophy of science: in the presence of dual theories, what is one to say about the ontology of the world? In this paper, we undertake a comprehensive and non-technical survey of the landscape of possible ontological interpretations of duality-related theories. We provide a significantly enriched and clarified taxonomy of options—several of which are novel to the literature.

*Forthcoming in *Philosophy Compass*.

¹baptiste.lebihan@unige.ch

²james.read@philosophy.ox.ac.uk

Contents

1	Introduction	2
2	Review	5
3	Taxonomy	7
4	Antirealism and structuralism	10
5	Underdetermination	11
6	Discrimination	13
7	Nihilism	16
8	Alternatives to the common core	17
9	Conclusion	22

1 Introduction

Contemporary physics is built upon two central—and *prima facie* incompatible—frameworks: the theory of general relativity on the one hand, and the standard model of particle physics—a certain quantum field theory—on the other. Although these two frameworks are strikingly effective at describing the actual world in their relevant domains (*viz.*, macroscopic, astrophysical and cosmological scales for general relativity, and atomic, subatomic and molecular scales for the standard model), they rest upon very different assumptions. For example, one central feature of general relativity is that spacetime is rendered dynamical (that is, it is not a fixed background, but rather ‘curves’ in response to its matter content); by contrast, spacetime remains a fixed background in the standard model of particle physics.

This notwithstanding, various fields of physics—such as the study of the early universe, or of black holes—lie at the intersection of the domains of these two theories, and thereby call for a *quantum theory of gravity*.³ Constructing such a theory capable of overcoming

³Strictly, one might distinguish different senses of a ‘quantum theory of gravity’—this could be e.g. (i) a quantised version of a theory which describes gravity, such as a quantised version of general relativity; or (ii) a theory which unifies general relativity and the quantum-mechanical standard model of particle physics; or perhaps (iii) something else. (Sometimes, theories of type (ii) are called ‘theories of everything’, since they encompass all four fundamental interactions.) For further discussion on these matters, see

the tensions between general relativity and the standard model is an ongoing matter of profound difficulty; at present, there exist several candidate options which remain the subject of active research. (Cf. [Huggett and Wüthrich, 2013] for a philosophical overview of such options.)

According to the naïve ontological picture presented by *string theory*—arguably the most popular extant research programme in quantum gravity—reality is constituted by one-dimensional strings, as well as by other higher-dimensional entities called ‘branes’. Moreover, reality is not made up of four spacetime dimensions (three spatial and one temporal), but rather of ten, or eleven. But string theory embodies another intriguing notion that should be of interest to philosophers and metaphysicians—the notion of *duality*.

Associate with every theory a class of ‘models’, equivalently ‘solutions’. As van Fraassen puts it, a model is “Any structure which satisfies the axioms of a theory” [van Fraassen, 1980, p. 43]. (Roughly speaking, physical theories are specified by certain dynamical equations; each model represents one possible solution to those equations—hence the latter choice of nomenclature.) In turn, take it that two solutions are ‘empirically equivalent’ just in case they agree on all ‘physically observable data’, i.e. on empirical substructures, in the sense of van Fraassen [van Fraassen, 1980, p. 64]. Then, a duality is a mapping—i.e., a systematic correspondence—between the spaces of solutions of two theories, such that models related by that map are empirically equivalent.⁴

There exists not just one string theory, but rather five, related by an intricate web of dualities. That is to say, each model of a given string theory possesses (via the duality map under consideration) a dual model, which *prima facie* makes very different ontological claims

[Huggett and Callender, 2001, Wüthrich, 2005].

⁴Such is the definition of dualities presented in e.g. [Matsubara, 2013, Read, 2016], which will suffice for our purposes. Note that one might augment the criterion of ‘empirical equivalence’ by requiring that all quantities regarded as being physically meaningful (whether observable or not) be preserved under the duality map. We agree that the most striking examples of dualities—including examples of dualities from string theory—satisfy this stronger condition. However, in this paper we choose to work with the weaker notion of a duality proceeding in terms of empirical equivalence alone, for this will suffice to make all necessary points regarding the interpretation and ontology of duality-related theories. For more detailed and comprehensive approaches to the definition of dualities, see e.g. [De Haro, 2017b, De Haro and Butterfield, 2018]. Note also that it need not be the case that the duality map is one-one—it might instead be that a class of solutions of the first theory is mapped to a single solution of the second theory under the duality. This will be of relevance below.

about the world, while nevertheless being empirically equivalent.^{5,6} Thus, at first blush, dualities instantiate the *underdetermination of theory by empirical evidence* familiar from the philosophy of science. In the case of dualities, however, this underdetermination is peculiar, as the empirical equivalence of the solutions under consideration was often not expected *ab initio*, but rather came as a profound surprise, in light of their apparently diverging ontological pictures.

In this paper, we undertake a comprehensive yet non-technical survey of the terrain of possible interpretative options for ascertaining the ontology of duality-related models of physical theories in such a way as to resolve any threat of underdetermination, introducing several novel observations and options along the way; our hope is that this work will both open this area of research to an expanded class of philosophers of all stripes, bring enhanced clarity to the literature on these matters, and show the way to novel ontological interpretations

⁵Here is a more precise description of the situation, for the most technical-minded reader. The four best-known examples of dualities arising in string theory are ‘T-duality’, ‘mirror symmetry’, ‘S-duality’, and the ‘AdS/CFT correspondence’. In the case of T-duality, type IIA superstring theory (one of the five superstring theories) on a product manifold $M \times S^1$ with radius of the periodic dimension R is found to be dual to type IIB superstring theory (another of the five superstring theories) on the product manifold $M \times S^1$ with radius of the periodic dimension proportional to $1/R$ [Becker et al., 2007, ch. 6]. Mirror symmetry is a generalisation of T-duality to the case of topologically inequivalent manifolds. (For a philosophical introduction to mirror symmetry, see [Rickles, 2013b].) S-duality relates solutions of one superstring theory with string coupling constant g_s to solutions of another superstring theory with string coupling constant $1/g_s$; it is thus a so-called ‘strong/weak’ duality. For example, strongly/weakly coupled type I superstring theory is dual under S-duality to weakly/strongly coupled $SO(32)$ heterotic string theory [Becker et al., 2007, §8.2]. Finally, in the AdS/CFT correspondence, a string theory in so-called ‘AdS spacetime’ (the ‘bulk theory’) is dual to a conformal field theory (CFT) in a lower number of spacetime dimensions (the ‘boundary theory’) [Becker et al., 2007, ch. 12]. The AdS/CFT correspondence was originally introduced in [Maldacena, 1998]. In this paper, the relevant aspects of these dualities will be introduced when needed in the course of the dialectic.

⁶Some care regarding the notion of a ‘model’, or ‘solution’, of string theory is needed. In each of the five versions of perturbative string theory, one posits a background spacetime manifold; one may then write down certain fields (ultimately understood to be ‘coherent states’ of strings—cf. [Callan et al., 1985, Huggett and Vistarini, 2015]) on this manifold, obeying (to a certain order of approximation) certain dynamical equations, consistent with the string theory under consideration (for details here, see e.g. [Green et al., 1987, §3.4]). These are the ‘models’ of the string theories under consideration, which are related by dualities (for example, the background fields in T-duality are related by the so-called ‘Buscher rules’—cf. [Blumenhagen et al., 2013, §14.2]). Note that these models are often understood to be the ‘ground states’ of a more fundamental theory, called ‘M-theory’. (That is, they are states of the more fundamental theory, holding only at some minimal energy.) We will have more to say about M-theory below.

of duality-related theories.⁷

2 Review

In the recent literature, one finds the above-mentioned claim that string-theoretic dualities present a case of underdetermination of theory by evidence—that is, a situation in which there exist multiple theories, each of which (*prima facie*) makes different ontological claims about the world, yet which are all adequate to exactly the same stock of (possible) empirical data. Such underdetermination is typically understood to be problematic for the scientific realist, for how can one plausibly maintain that one’s preferred theory is *true*, if a range of other theories are also consistent with the data? In e.g. [Matsubara, 2013, Read, 2016], authors began to compile a taxonomy of interpretative options available to the realist, in order to ‘break’ the putative underdetermination arising in the case of dualities. The two (allegedly) most plausible options in this regard were claimed to be the following:

- (*Discrimination.*) Privilege the ontological claims of just one of the two dual theories. That is, consider two dual theories, \mathcal{T}_1 and \mathcal{T}_2 , with (respectively) solutions \mathcal{M}_1 and \mathcal{M}_2 related by the duality map. Naïvely interpreted,⁸ \mathcal{M}_1 and \mathcal{M}_2 represent two *distinct* worlds, respectively W_1 and W_2 (hence a case of underdetermination). However, according to this *discriminatory* strategy, only one of \mathcal{M}_1 and \mathcal{M}_2 is a legitimate description of the actual world.⁹ Though coherent, this approach faces an obvious problem: principled reasons for privileging the ontological

⁷For other philosophically-oriented introductions to dualities—including all the theories and their respective dualities mentioned in this work—see e.g. [Polchinski, 2017, Rickles, 2011]. For an introduction to the philosophy of quantum gravity more generally, see [Matsubara, 2017].

⁸What do we mean, when we speak of the ‘ontological claims’ of (solutions of) a theory, or of the ‘naïve interpretation’ of (solutions of) that theory? In this paper, we take the ‘ontological claims’ of a theory to be given by its ‘naïve interpretation’; in turn, we understand this to be an interpretation of the theory in question (*a fortiori* its solutions) such that the worlds represented by the solutions of that theory are ‘isomorphic’ to those solutions. (Here, we set aside legitimate concerns that speaking of isomorphism between mathematical structures and worlds constitutes a category error—cf. [van Fraassen, 2002, ch. 1].) More technically, our focus is upon *internal* interpretations, in the sense of [Dewar, 2017a, De Haro, 2017b].

⁹Throughout this paper, by the ‘actual world’, we mean a hypothetical world in which the empirical data consistent with two dual solutions are observed. Though this use is non-standard, it will simplify our discussion. It is also worth clarifying what we mean by ‘legitimate description’. Suppose that one observes certain empirical data, and one has to hand a certain range of mathematical models consistent with that data. In spite

claims of just one of the two dual solutions appear (in general) to be lacking (cf. [Read, 2016, Teh, 2013]).

- (*Common core.*) ‘Break’ the underdetermination by interpreting only the ‘common core’ of the solutions related by the duality map as representing physical states of affairs. In more detail, consider again two dual theories, \mathcal{T}_1 and \mathcal{T}_2 , with (respectively) solutions \mathcal{M}_1 and \mathcal{M}_2 related by the duality map. On this position, the ‘naïve’ interpretation of \mathcal{M}_1 and \mathcal{M}_2 , according to which these solutions represent distinct worlds W_1 and W_2 , is not correct. Rather, we should identify the mathematical structure *common* to those solutions, and interpret \mathcal{M}_1 and \mathcal{M}_2 in terms of only that common structure—call it \mathcal{M}_c . In so doing, the underdetermination is (apparently) broken, for in so interpreting \mathcal{M}_1 and \mathcal{M}_2 , these solutions may be regarded as representing the *same* world—call it W_c —the ontology of which is taken to be represented by \mathcal{M}_c .

In this paper, our concerns are twofold: (1) We contend that both the discriminatory and common core approaches are more subtle than has hitherto been appreciated—and in fact, both approaches are consistent with a number of *distinct*, more fine-grained views, only some of which overcome the putative underdetermination in the case of dualities. (2) We maintain that there exist (at least) two *further* approaches for addressing the underdetermination which arises in the case of dualities—these we call ‘nihilism’ and ‘pluralism’. Roughly speaking, nihilism is the view that *no* solutions of dual theories constitute legitimate descriptions of the actual world;¹⁰ pluralism is the view that *all* dual solutions may be taken to represent *the same* actual world—but not because such a world is represented by the common core of those solutions, but rather because the structure of all dual solutions may be instantiated *simultaneously*.¹¹

In the remainder of this paper, we undertake the following tasks. In §3, we propose an expanded taxonomy of options for the interpretation of dualities. In §4, we present and set aside various ‘antirealist’ and ‘structuralist’ approaches to dualities—for our concern in this paper

of all such models being consistent with the data, it may be that one discounts certain models, for certain super-empirical, philosophical reasons (more on this below). Such models are, then, taken to *not* constitute legitimate candidates for representing the actual world. The complement of this set of available models is the set of legitimate candidates for representing the actual world.

¹⁰One might wonder how nihilism can be compatible with scientific realism; we shall see in §7 multiple senses in which this could be the case.

¹¹We agree that such a view is strange. However, as we argue in §8.2, it is *coherent*, and therefore deserves to be mentioned explicitly.




Figure 1: Five dual solutions $\mathcal{M}_1, \dots, \mathcal{M}_5$, and their common mathematical core, \mathcal{M}_c . ‘Naïvely interpreted’, such solutions may be understood as representing certain worlds (to which they are ‘isomorphic’—cf. footnote 8). If one such world is green, this indicates that this world is regarded as being a legitimate candidate for being the actual world. If the world is red, on the other hand, this indicates that this world is not regarded as being a legitimate candidate for being the actual world. Options (1)-(6) categorise different verdicts on which of these worlds are regarded as being candidates for being the actual world in this sense.

is to interpret dual theories realistically; that is, to get a handle on what dual theories tell us about *what the world is really like*. In §§5-8, we discuss each of the above-mentioned realist interpretative options in turn, and assess whether they succeed in resolving the putative underdetermination arising from dualities.

3 Taxonomy

What we call above ‘discrimination’ and the ‘common core’ approach form just two of a substantially broader range of interpretative options *vis-à-vis* dualities. To make this explicit, consider figure 1. Here, $\mathcal{M}_1, \dots, \mathcal{M}_5$ represent (respectively) five solutions of five theories $\mathcal{T}_1, \dots, \mathcal{T}_5$, which are dual to one another.¹² The solution \mathcal{M}_c to the right of the $\mathcal{M}_1, \dots, \mathcal{M}_5$ consists in the common mathematical structure of each of the five dual solutions.¹³ (The rightward arrow

¹²We here include five solutions with an eye to the five superstring theories. However, any number of dual solutions greater than or equal to two would suffice for our purposes.

¹³What exactly does this ‘common mathematical structure’ consist in? This is a question worthy of considered attention; thankfully, authors in the philosophy

indicates that \mathcal{M}_c is typically constructed once the $\mathcal{M}_1, \dots, \mathcal{M}_5$ are given.) Beneath each of the $\mathcal{M}_1, \dots, \mathcal{M}_5$ are sets of possible worlds (1)-(6), to which the solutions are interpreted as corresponding.¹⁴ If such a world is to be regarded as being a legitimate candidate for being the the actual world, we colour it green; otherwise, we colour it red.¹⁵

Clearly, a range of interpretative options are possible. If each of the worlds associated with $\mathcal{M}_1, \dots, \mathcal{M}_5$ are regarded as being legitimate candidates for being the actual world (cases (1) and (2)), then we have a case of underdetermination: for given a set of empirical data compatible with one of these solutions, we do not know which of these five worlds is the actual world. Within this situation, two further options are possible: either the world represented by \mathcal{M}_c is also a legitimate candidate for being the actual world (case (1)), or it is not (case (2)). We discuss the underdetermination approach in §5.

Suppose instead that just one of the original dual solutions is regarded as being a legitimate candidate for being the actual world (cases (3) and (4)); this is the discriminatory approach. (Clearly, our decision to regard \mathcal{M}_1 as representing this world in figure 1 is made without loss of generality.) Within this scenario, two sub-scenarios again arise: either W_c is regarded as also being a legitimate candidate for being the actual world (case (3)), or it is not (case (4)). In §6, we reappraise the discriminatory approach.¹⁶

of dualities have already done much to clarify these matters. See in particular [De Haro and Butterfield, 2018, §2], in which this ‘common core’ is taken to consist of a ‘bare theory’—in itself understood to be a triple $\langle \mathcal{S}, \mathcal{Q}, \mathcal{D} \rangle$ of (respectively) states, quantities, and dynamics, isomorphic representations of which are contained in the structure of each of the dual theories under consideration (along with, potentially, further theory-specific structure). For our purposes, it suffices to know that the common mathematical core of two dual theories *can* be constructed in a well-defined manner.

¹⁴Via ‘naïve interpretation’, in the sense of footnote 8.

¹⁵Whether or not one regards a certain empirically adequate model as being a legitimate candidate for representing the actual world will, to repeat, depend upon one’s prior philosophical predilections—cf. footnote 9. We will see below a range of explicit examples of such predilections.

¹⁶Let us demonstrate that (3) and (4) are distinct options by way of an analogy with a well-known case in the literature on symmetry transformations. Consider models of Newtonian gravitation theory, set in Newtonian spacetime (cf. [Earman, 1989, ch. 2]); let $\mathcal{M}_1, \dots, \mathcal{M}_5$ correspond to solutions of this theory which differ only with regard to the absolute velocity of (the centre of mass of) the entire material content of the universe; let \mathcal{M}_1 , in particular, be the model which states that the absolute velocity of (the centre of mass of) the entire material content of the universe is zero; and let \mathcal{M}_c be the model of Newtonian mechanics set in Galilean spacetime corresponding to $\mathcal{M}_1, \dots, \mathcal{M}_5$, in which the notion of absolute velocity has been excised. Note that \mathcal{M}_1 and \mathcal{M}_c make different ontological claims about the world: the former states that the absolute velocity

Finally, suppose that *none* of the worlds represented by $\mathcal{M}_1, \dots, \mathcal{M}_5$ are regarded as being legitimate candidates for being the actual world (cases (5) and (6))—this is the nihilist gambit. Again, within such a scenario, either W_c may be regarded as itself being a legitimate candidate for being the actual world (case (5)), or it may not (case (6)). We discuss nihilism in §7.

Setting things up in the above manner is significantly more nuanced than extant taxonomies of interpretative options. For example, this framework illustrates that the discriminatory approach—represented in cases (3) and (4)—is not to be regarded as disjoint from the common core approach—represented by cases (1), (3), and (5).

After discussing options (1)-(6) in §§5-7 of this paper, we consider in §8 whether there exist any *alternatives* to the common core approach (note that these alternatives are not represented in figure 1). We identify two: (i) a tactic—common in the physics literature—of attempting to embed the two dual theories into some ‘deeper’ theory (§8.1); and (ii) the ‘pluralist’ strategy indicated above, in which the structures of all dual solutions are regarded as being *jointly* instanti-

of the material content of the universe is zero; the latter states that this is undefined. (Cf. [van Fraassen, 1980, pp. 45-46].)

With this in mind, let us now consider (3) and (4), beginning with the latter. According to this position, only \mathcal{M}_1 of $\mathcal{M}_1, \dots, \mathcal{M}_5$ is a legitimate candidate for representing the actual world. Why would one think this? One reason would be on the grounds of the principle of sufficient reason (PSR)—i.e. (roughly speaking) Leibniz’s principle that God must have a sufficient reason to realise one of a class of symmetry-related models of a given theory (more on Leibniz’s principles below). Maudlin articulates the point clearly, in the context of the same example:

But on reflection, the PSR argument cannot get off the ground in the case of absolute velocity. Suppose that God wishes to create the material world in a heretofore empty absolute space. God could give the material world, as a whole, any absolute velocity in that space without affecting the relative positions and motions of bodies. Among all of these possible velocities, one stands out as special: absolute rest. For if God should give the material world any nonzero velocity, He would have to choose a *direction* in absolute space for that velocity to point. [Maudlin, 2012, p. 48]

If one accepts such reasoning, one may embrace option (4) in our taxonomy of options. In addition, however, there still remains the question whether \mathcal{M}_c —which, recall, represents (naïvely interpreted) a *distinct* world to \mathcal{M}_1 , in which there is no meaningful notion of absolute velocity—is a legitimate candidate for representing the actual world. If one *does* think this, one will be pushed to (3)—one will think that either the universe is at rest, or that the actual world does not contain facts about absolute velocity at all. If one does *not* think that \mathcal{M}_c is a legitimate candidate for representing the actual world (perhaps if one is strongly wedded to a notion of absolute velocity in order to maintain the coherence of one’s metaphysics), one will remain with option (4).

ated (§8.2).

4 Antirealism and structuralism

Before we discuss the above-outlined realist strategies for the interpretation of dualities, it is worth taking some time to present and set aside certain positions in the philosophy of science, which bear upon the interpretation of dualities, but which we shall not consider further in this paper.

The first class of such positions consists in ‘antirealist’ views, which (in some sense) deny that questions of ‘ontology’ are meaningful, or at least central to scientific practice. Stronger versions of antirealism are to be found in logical positivism and logical empiricism. Loosely, advocates of these views prioritise the observable, and claim that ‘metaphysical’ questions regarding the (unobservable) ‘ontology’ of the world are, strictly speaking, nonsensical. (For some contemporary discussion of the logical positivist movement, and its transmutation into logical empiricism, see e.g. [Creath, 2017, Friedman, 1999, Godfrey-Smith, 2013].) A weaker view is the constructive empiricism of van Fraassen [van Fraassen, 1980], according to which we should be agnostic about claims read off from our theories of science regarding the unobservable, and have credence only in the statements read off from our theories regarding observable phenomena. (Thus, there is a sense in which the antirealism of van Fraassen is ‘epistemological’, whereas the antirealism of the logical positivists and empiricists is ‘metaphysical’.)

Since both strands of antirealism deny (in some sense) that it is of value to be concerned with answers to questions of ontology, we set them aside in the remainder of this paper—for it is precisely such questions which we investigate, and which are regarded (at least *prima facie*) as being legitimate and worthwhile in the context of investigations into the ontological commitments of theories related by dualities.

As a halfway house between realism and antirealism, one may embrace a ‘structural realist’ thesis. There are many brands of structural realism—see e.g. [French, 2014, Ladyman, 2014, Ladyman, 1998, Worrall, 1989]. One version of structuralism of interest is the following: (i) the solutions of our best theories of physics are constituted formally by certain mathematical objects; (ii) we should take the world represented by that solution to be ‘isomorphic’ to the solution itself;¹⁷

¹⁷Cf. footnote 8. Note that there are some subtleties here. First, such a claim can only strictly hold true for fundamental theories, in which we believe that the totality of the world is modelled by that solution of the theory in question. Second, such a claim effaces

(iii) such a world may or may not be amenable to description in terms of extant ontological categories—but regardless of whether that be so, we should *still* believe that the solution of the theory in question describes some possible world (potentially the actual world, if the empirical substructures of that solution match the empirical data from the actual world).

This brand of structuralism is certainly an interesting thesis, which merits much investigation. However, we elide further such discussion in this paper, for two reasons. First (and to repeat), we are interested in saying something, in light of our best theories of physics (and especially in light of dualities), about *what the world is really like*; this does not appear possible in general on the above structuralist view. Second, we incline towards a more cautious attitude, according to which we may *only* regard a particular solution of a particular theory as representing a possible world (rather than being mere mathematics) when we have in hand a clear picture of *what that solution is supposed to represent*. In this latter regard, cf. [Møller-Nielsen, 2017, Read and Møller-Nielsen, 2018].¹⁸

5 Underdetermination

We return now to the taxonomy of realist options for the interpretation of dualities presented in §3; we begin with the underdetermination view. One might not regard the putative underdetermination presented by dualities as being problematic. Indeed, according to the underdetermination interpretation, solutions of dual theories are in fact all legitimate candidates for representing the actual world—and we cannot ascertain *which* of the worlds associated with those dual solutions (naïvely interpreted) is the actual world.

Should one rest satisfied with the underdetermination interpretation? Arguably *no*, for in cases of underdetermination, it is impossible

important issues regarding gauge redundancies—i.e., aspects of the formalism of a given theory which we do not believe to have representational capacities.

¹⁸A second version of structuralism called ‘ontic structural realism’ is also worthy of mention. (We do not comment on whether the version of structuralism presented above should also qualify as ‘ontic’, in some sense.) This latter view finds presentation in e.g. [Esfeld and Lam, 2008, French, 2014, French, 2010], and appeals to what one might call an ‘ontology of physical structures’. There is a sense in which, on this view, it *is* legitimate to make claims about ‘what the world is really like’; however, such an ontology is to be explicated in terms of ‘structures’, rather than ‘objects’. Since our concern in the ensuing is not with structuralism, we again elide further discussion. For more detailed reflections on structural realism in the context of dualities, see [Dawid, 2013, pp. 178-189] and [Matsubara, 2013, §3].

to ascertain which of a class of empirically equivalent worlds—all empirically adequate to the actual world—is, in fact, the actual world. Thus, underdetermination gives rise to a sceptical challenge: absent a means of determining which of a class of worlds is the actual world, *we have at hand no determinate picture of what the world is really like.*

One might seek to overcome such underdetermination in the following way: identify the ‘common mathematical core’ of the duality-related solutions under consideration, and take the actual world to be represented by this common mathematical core.¹⁹ This, in effect, takes us from scenario (2) of figure 1, to scenario (1). Note, however, that in itself such a move is *insufficient* to resolve the underdetermination under consideration. Indeed, there is a sense in which, absent further philosophical details, such a move has made the situation *worse*: we have, in effect, identified a *further* world which is empirically adequate to the actual world.

If such a ‘common core approach’ is to constitute a viable route to resolving the underdetermination, then it must be augmented by some *philosophical* reasoning, such that the original dual solutions do not in fact constitute legitimate descriptions of the actual world. One option here would be to appeal to putative super-empirical virtues of just one of the solutions under consideration (typically, though not necessarily, the common core solution \mathcal{M}_c)—e.g., simplicity, explanatory power, etc.—in order to break the underdetermination. However, there remains a gap to be bridged between such putative virtues, and truth-conduciveness (cf. [van Fraassen, 1980, §4.1]). Only if such a bridge is constructed will appeal to such virtues move us from scenario (1) to scenario (3)/(5) of figure 1. We return to this issue of breaking the underdetermination in §6, on the discriminatory approach.

In this connection, one might think that another philosophical principle delivering scenario (3)/(5) from scenario (1), and thereby breaking the underdetermination, is Occam’s razor—i.e., the principle that, all else being equal, otiose structure should not be introduced into one’s ontology. Note, though, that as with the virtues discussed

¹⁹Such a position is common in the literature on dualities—for an explicit, paper-length defence of this view, see [Rickles, 2017]. In this vicinity, cf. also [De Haro, 2017b, De Haro and Butterfield, 2018, Huggett, 2017, Matsubara and Johansson, 2018]. While it is true that the common core approach is sometimes identified with structural realism, note that the approach is *distinct* from the two versions of structuralism articulated in §4—for (i) this approach still attempts to explicate the ontology of a given model in terms of extant ontological categories (unlike the version of structuralism considered in the body of §4), and (ii) this approach is not necessarily committed to an ‘ontology of physical structures’ (unlike the version of structuralism considered in footnote 18). Of course, one might still choose to *call* the common core approach a version of structuralism—but we take this, ultimately, to be a semantic matter.

above, such a principle is first and foremost a *practical* principle, reminding us that it is (in general) preferable to work with (solutions of) more parsimonious physical theories.²⁰ Again, however, such a principle in itself does nothing to rule out as candidates for representing the actual world the other dual solutions under consideration.

What is needed is a more robust metaphysical principle to exclude all but one of the dual solutions under consideration as being candidates for representing the actual world. One obvious option here is Leibniz’s ‘principle of the identity of indiscernibles’ (PII), which (applied to worlds) can be understood (for our purposes) to state that there can be no ‘distinctions without a difference’—that is, no empirically equivalent but physically distinct worlds, which vary with respect to unobservable structure. (The PII was famously presented by Leibniz in the *Leibniz-Clarke Correspondence* [Alexander, 1986]. For contemporary discussion of Leibniz’s principles, and in particular the PII, see [Rodriguez-Pereyra, 2014, Saunders, 2003].) Embracing such a metaphysical thesis seems to deliver us from scenario (1) to scenario (5)—since only W_c does not possess the variant undetectable structure under consideration, and is thereby compatible with this principle. Of course, however, the question naturally arises at this juncture: *why* should we commit ourselves to such a metaphysical principle?

6 Discrimination

In this section, we examine the discriminatory approach, according to which only one of the dual solutions under consideration is a legitimate candidate for representing the actual world. This approach is represented in scenarios (3) and (4) of figure 1. Claims appearing to be consonant with this approach are found in particular in the literature on the ‘AdS/CFT correspondence’—one particularly famous duality, in which a string theory in so-called ‘AdS spacetime’ (the ‘bulk theory’) is dual to a ‘conformal field theory’ (CFT) in a lower number of spacetime dimensions (the ‘boundary theory’) [Becker et al., 2007,

²⁰In fact, there is a sense in which this is not true in the case of dualities. For example, electromagnetism formulated in terms of a vector potential A^a is ‘theoretically equivalent’ (for our purposes: dual) to electromagnetism formulated in terms of the Faraday tensor F_{ab} (cf. [Weatherall, 2016a, Weatherall, 2016b, Weatherall, 2018]). Though the former theory has more structure (i.e., degrees of freedom) than the latter, and so is in this sense ‘less parsimonious’, there nevertheless exist many practical virtues of using the former over the latter—e.g. its amenability to variational principles; locality principles; etc. (cf. [Dewar, 2017b, Nguyen et al., 2018]). Of course, strictly speaking this is compatible with Occam’s razor, for in the foregoing language, all else is *not* equal.

ch. 12]. As Oriti points out, many string theorists speak as if the four-dimensional spacetime of the boundary theory is real, with the bulk spacetime appearing only as an auxiliary construction [Oriti, 2009]. For example, Horowitz and Polchinski note that the AdS/CFT correspondence is a little different from other dualities in that the conformal field theory side is exactly understood, whereas the string theory side is only approximately understood. Building on this, they write,

In the AdS/CFT case, the situation may not be so symmetric, in that for now the gauge side has an exact description and the string/gravity side only an approximate one: we might take the point of view that strings and spacetime are “emergent” and that the ultimate precise description of the theory will be in variables closer to the CFT form. [Horowitz and Polchinski, 2009, p. 230]

Here, Horowitz and Polchinski appear to claim that since one of a pair of dual theories (here, the conformal field theory) is better understood, we should privilege the ontological claims associated with solutions of that theory over those of its dual (here, the bulk string theory). Faced with such a passage, the question arises naturally: why should the *epistemological* fact about what human beings happen to currently know about two dual theories relative to one another warrant the *metaphysical* conclusion that the theory about which we currently know more must give the correct description of the world? Such worries have been expressed by Teh [Teh, 2013, §4], Rickles [Rickles, 2013a, p. 317], and Dieks, van Dongen, and De Haro [Dieks et al., 2015, De Haro, 2017a]. What is needed is some argument to the effect that one description of the world has metaphysical priority over its dual; such authors, however, treat this with suspicion—for example, Teh writes, “We have no good reason to think of the gravitational side of the duality as metaphysically emergent from the gauge theory side, or vice versa” [Teh, 2013, p. 310].

Is this response to Horowitz and Polchinski reasonable? In fact, there is perhaps room to defend the discriminatory approach to dualities such as the AdS/CFT correspondence in the face of such criticism. Here is an alternative way to read Horowitz and Polchinski: it is not that we simply better *understand* one of the two dual theories (namely, the CFT); rather, it is that we only *have available* the full mathematical structure of the CFT, whereas the other dual theory—the AdS string theory—is only *partly constructed* (our understanding of the AdS side of the AdS/CFT duality is inherently perturbative—cf. [Ammon and Erdmenger, 2015]). The question here is not (*pace* the argument above) one of our only investing with ontological im-

port those physical theories which we best understand, but rather one of our only investing with ontological import those physical theories which *which we actually have to hand*. Absent a completion of the AdS string theory, why take seriously its ontological claims (naïvely interpreted)—or even think that such a coherent completion and subsequent interpretation is to be had?²¹

Though it seems to us that such a defence of the discriminatory approach is reasonable in context of dualities such as the AdS/CFT correspondence, in which the full structure of only one of the duality-related theories is available, let us now consider other ways in which the metaphysical primacy of one of the dual theories might be established. Such principles might be e.g. the super-empirical, or metaphysical, principles introduced in §5. However, as we have already seen, such super-empirical principles do not go far enough, for they do not *preclude* certain worlds from being legitimate candidates for being the actual world; moreover, while metaphysical principles such as the PII arguably do not face such difficulties, we appear to lack independent reason to embrace such principles.

It is worth dwelling a little longer upon how a defence of the discriminatory approach based upon the PII might proceed. Consider the situation in which two theories \mathcal{T}_1 and \mathcal{T}_2 are dual, with each solution of \mathcal{T}_1 corresponding to a class of solutions of \mathcal{T}_2 .²² If one also embraces the PII, then one should not regard the elements of this class of solutions of \mathcal{T}_2 as representing distinct worlds. But which (unique) world *should* one take these solutions to represent? One natural answer is that one should take this world to be that represented by the unique solution of \mathcal{T}_1 (naïvely interpreted) to which this class of solutions of \mathcal{T}_2 corresponds.²³

Note that here the PII is being put to a *different* use to that in §5. While in that case, the principle was used to establish that *neither* dual solution represents the actual world (rather, only their mathematical common core does so), here the principle is used to establish, in light of ‘gauge redundancy’ in one of the two dual theories under consideration, that the ontological claims of the *other* dual theory are to be preferred. These are, then, inter- versus intra-theoretic applica-

²¹We are grateful to an anonymous referee for helpful discussion on this point. Note that there is some parallel between the caution advanced in the above paragraph, and that presented in the context of dualities in [Read and Møller-Nielsen, 2018].

²²Cf. footnote 4.

²³This, indeed, is exactly the standard answer given in the case of e.g. the theoretical equivalence of Newtonian gravitation theory set in Galilean spacetime, and Newton-Cartan theory—cf. [Friedman, 1983, Malament, 2012, Read and Møller-Nielsen, 2018, Weatherall, 2016a].

tions of the PII; in our view, both are in principle legitimate.

7 Nihilism

If all dual theories are to be considered on a par *vis-à-vis* the legitimacy of their ontological claims, then one must advocate either interpretations (1) and (2) in figure 1 (*viz.*, underdetermination interpretations), or interpretations (5) and (6). It is these latter two approaches which we now discuss; we call these ‘nihilist’ strategies. In particular, we focus on (6), for we have seen above circumstances in which one may be led to interpretation (5)—e.g., through embracing both the common core interpretation and the PII, the latter in order to exclude the legitimacy of the original dual solutions *qua* descriptions of the actual world.

As we see it, one might endorse nihilism for two reasons. The first applies in cases in which the dual theories under consideration are not expected to be final theories. In this case, one might claim that solutions of non-final theories (naïvely interpreted) simply cannot be understood to represent *any* possible world. There is something to be said for this view, for consider e.g. solutions of Newtonian mechanics—an uncontroversial case of a non-final theory. It is well-known that the stability of matter cannot be accounted for in this theory; for this, one must proceed to some quantum-mechanical successor. But in that case, how *could* solutions of this non-final theory constitute legitimate candidates for representing the actual world? Thus, in this regard, one might be a nihilist about the ontological claims of any non-final dual theory, while still remaining a scientific realist—for in this case, one might still maintain that solutions of a *final* theory *could* constitute legitimate candidates for representing the actual world.

The second sense in which one might be a nihilist is the following. Suppose that, for antecedent reasons, one is unsympathetic to a particular research programme, e.g. string theory. In that case, one might reject (for said to-be-articulated antecedent reasons) solutions of all e.g. dual string theories as being legitimate candidates to describe the actual world, while remaining a realist, for one might think that solutions of *other* theories (e.g. loop quantum gravity, or extensions thereof) may legitimately describe the actual world.

In brief, then: nihilism is compatible with scientific realism just so long as there is *some* solution of *some* theory which one takes to constitute a legitimate description of the actual world. Perhaps such a solution is provided by looking to the mathematical common core of the dual solutions under consideration—in which case, one finds oneself in situation (5) of figure 1. Otherwise, one finds oneself in

situation (6) of figure 1. But, what solutions *could* describe the world, if not the common core, and not the original dual solutions under consideration? It is to this question that we now turn.

8 Alternatives to the common core

So far, we have considered underdetermination, discriminatory and nihilist approaches to the interpretation of dualities—these correspond respectively to options (1) and (2), (3) and (4), and (5) and (6) of figure 1. In this section, we consider whether there exist any alternatives to the ‘common core approach’ (options (1), (3), and (5) in figure 1). In our view, there are (at least) two such options: (i) rather than construct a theory which represents the ‘common core’ of the original dual theories, construct a new theory in which each of the dual theories are *embedded* (§8.1); (ii) argue that each dual theory describes, correctly but partially, the same *one* world (§8.2).

8.1 Overarching theories

The common core approach purports to identify a possible world ‘isomorphic’ to the mathematical structure common to the dual solutions under consideration. Though this approach is popular in the philosophy of physics literature (see e.g. [De Haro and Butterfield, 2018, Huggett, 2017, Matsubara, 2013, Matsubara and Johansson, 2018, Rickles, 2011, Rickles, 2017]), it is not the only live interpretative option purporting to break the underdetermination. Indeed, a distinct position—widely embraced in the physics community in the context of string-theoretic dualities—is to embed the spaces of solutions of the two dual theories under consideration into that of some deeper, ‘overarching’ theory.

What is meant by an ‘overarching’ theory? The answer to this question is best given by way of example. It is sometimes claimed that the five superstring theories are certain ‘limits’ of some deeper, ‘M-theoretic’ structure, in exactly the sense that their solutions spaces can be embedded into that of this deeper theory—cf. figure 2. Such an ‘M-theory’ is conceptually *distinct* from the common core of the dual string theories under consideration.

A parallel example of this manoeuvre can be drawn from the history of physics: consider the relation between Heisenberg matrix mechanics and the (putatively) empirically equivalent Schrödinger wave mechanics, in the 1920s and 30s.²⁴ Ultimately, the unification of

²⁴For a detailed study of the relations between these two early approaches to quantum mechanics, see [Muller, 1997a, Muller, 1997b].




Figure 2: The spaces of two dual theories \mathcal{T}_1 and \mathcal{T}_2 , embedded into the solution space of an ‘overarching’ theory, $\tilde{\mathcal{T}}$.

these two theories did not proceed by finding their mathematical common core; rather, the two theories were embedded into a *deeper* theory—what we now refer to as orthodox quantum mechanics (see [Muller, 1997b, §5]). This latter theory has a *richer* space of solutions than that of the two original theories—again, the structure is as per figure 2.

What to make of this approach with regard to the problem of underdetermination arising in the context of dualities? Merely embedding the spaces of solutions of the two dual theories into that of some ‘deeper’ theory does *not* in itself resolve the underdetermination—for again, given the empirical evidence compatible with one dual solution, it is not clear whether one should embrace the ontological claims of that solution, or of its dual, or of the overarching theory.

In fact, the situation here is even more subtle, and merits more detailed consideration; two scenarios are possible. First, if this ‘embedding’ of (the solution spaces of) two dual theories \mathcal{T}_1 and \mathcal{T}_2 into that of some overarching theory $\tilde{\mathcal{T}}$ is such that solutions of \mathcal{T}_1 and \mathcal{T}_2 *just are* solutions of $\tilde{\mathcal{T}}$, then it is clear that, for those solutions, $\tilde{\mathcal{T}}$ (naïvely interpreted) does not offer a distinct ontological picture over and above that of \mathcal{T}_1 and \mathcal{T}_2 . In that case, we cannot interpret the ontology of these solutions in terms of the ontology of $\tilde{\mathcal{T}}$; that is, the introduction of $\tilde{\mathcal{T}}$ does nothing to resolve the problem of underdetermination which arises in the context of string-theoretic dualities. In such a case, one seeking this end would need recourse to (for example) one of the other interpretative options elaborated in this paper.²⁵

Second, suppose that (somehow, in some sense not fully captured

²⁵In this sense, the overarching theory approach can be combined with the common core approach. For further detailed discussion of the differences between these two approaches, and the extent to which they can be combined, see [De Haro, 2018].

by e.g. figure 2) solutions of $\tilde{\mathcal{T}}$ (naïvely interpreted) *are* taken to offer a distinct ontological picture from solutions of \mathcal{T}_1 and \mathcal{T}_2 . In that case, one *may* say that the problem of underdetermination is resolved by interpreting solutions of \mathcal{T}_1 and \mathcal{T}_2 in terms of the ontology associated to the solutions of $\tilde{\mathcal{T}}$ —however, as before, the problem is only truly resolved if some further, metaphysical principle is embraced—to the effect that worlds represented by solutions of \mathcal{T}_1 and \mathcal{T}_2 are *not* legitimate candidates for being the actual world, whereas those represented by solutions of $\tilde{\mathcal{T}}$ *are* legitimate such candidates. In this case, however, it is unclear that e.g. the PII could deliver this verdict—for it need not be the case that the solutions of $\tilde{\mathcal{T}}$ have *less* structure than those of \mathcal{T}_1 and \mathcal{T}_2 (as was the case in the common core approach)—but this is a crucial assumption of PII-style arguments, which are used to conclude that the ontology of the ‘new’ theory under consideration is to be preferred. Perhaps some *explanatory* thesis may be advocated here instead: ‘since the overarching theory can account for more physical scenarios in its enriched space of solutions, that approach is to be preferred.’ However, again, it is not clear whether touting of such super-empirical virtues truly resolves the *metaphysical* problem of underdetermination. Moreover, it is not clear that such explanatory reasoning is sound, for recall that in the context of classical gravity, the fact that the solution space of general relativity is enriched over that of the alternative programme of ‘shape dynamics’ is sometimes taken to constitute an advantage of the latter theory over the former.²⁶

8.2 Pluralism

The second alternative to the common core approach we dub ‘pluralism’. On this approach, we consider the (distinct) structures of each of the dual solutions under consideration as describing *co-instantiated* structures in the actual world. That is, on this view, dual solutions may be taken to represent parts of *one* world. Each dual describes a seemingly different reality, but each of the structures under consideration represents a numerically distinct part of one world. Thus, according to pluralism, the physical world is in a certain sense fragmented.²⁷

²⁶The claim here is that, since shape dynamics has a restricted space of solutions as compared with general relativity, it is ‘more predictive’. See e.g. [Barbour, 2010, Mercati, 2018] for further discussion.

²⁷Pluralism is incompatible with our definition of the common core approach, since according to the latter we should read the ontology in (roughly speaking) the intersection of the mathematical structures of the two duals under consideration, rather than the union of those structures as on the pluralist approach. Nevertheless, a ‘fragmented’ ontology could, in principle, be consistent with the common core approach. Our thanks to Keizo

The pluralist strategy must pass a number of hurdles, if it is to be regarded as being successful. First, if pluralism is to resolve the putative underdetermination arising in cases of dualities, then, as before, some principled argument according to which the original dual solutions are *not* legitimate candidates for representing the actual world must be issued. Note, though, that this is not a problem *particular* to pluralism.

As we see it, there exist two central particular issues for pluralism: an overdetermination problem, and an ontological problem. On the former, since each of the individual dual solutions has certain empirical substructures, taken to correspond to the body of empirical data in the actual world, it seems that the pluralist strategy has, in a certain sense, swapped a problem of underdetermination for a problem of *overdetermination*, for now all of the dual structures may be taken to account for that body of empirical data. On the latter, the ontological problem is to understand what such a ‘fragmented’ world might look like, and whether the notion is consistent.

Let us focus upon the overdetermination problem. There exist two central options available by way of response to this issue: either the pluralist may claim (a) that only one of the co-instantiated dual structures gives rise to the *observed* empirical data in the actual world,²⁸ or she may claim (b) that that all such structures account *non-redundantly* for that data. In our view, there exist legitimate concerns regarding both proposals; let us discuss them in turn.

On the former view—(a)—the observed empirical data in the actual world may be accounted for by appeal to just *one* of the dual structures—though we do not know which one. Clearly, such an option is problematic, for it merely pushes the putative underdetermination arising in the case of dualities from the question of which of a number of worlds (i.e., those corresponding to the dual solutions, naïvely interpreted) could be the actual world, to the question of which of a number of distinct structures *within* a world could be that which accounts for the observed empirical data.

In order to maintain the latter view—(b)—the pluralist will argue that, in fact, all of the co-instantiated dual structures are *necessary* for accounting for the observed body of empirical data in the actual

Matsubara for discussion on this point.

²⁸The emphasis here is on ‘observed’. Both of the dual structures are, by construction, compatible with the same empirical data. However, in a pluralist world, it could (*prima facie*) be the case that the ontology which gives rise to the *observed* empirical data is either that of the first theory, or that of the second. This is the scenario countenanced in case (a).

world.²⁹ It is not, however, clear that such a view is compelling, for it certainly appears that each of the dual structures could account for the observed empirical data *in and of themselves*. This point is best illustrated by way of example.

Consider electromagnetism, formulated in terms of the vector potential A^a .³⁰ It is well-known that the space of solutions of this theory partitions into equivalence classes of solutions, elements of each of which are related by a $U(1)$ gauge transformation, and are regarded as being empirically equivalent (since each of these gauge-equivalent solutions gives rise—up to isomorphism—to the same Faraday tensor F_{ab} , which is taken to encode the observable data in the theory). In this case, we face an apparent problem of underdetermination analogous to that arising in the case of dualities,³¹ for one might ask: ‘which of the gauge-related solutions of electromagnetism instantiates the structure of the actual world?’

Of course, in this case it is standard to maintain that the ‘true’ ontology of solutions of the vector potential formulation of electromagnetism is that represented by the associated solution of the Faraday tensor formulation of electromagnetism.³² Nevertheless, let us consider how the pluralist strategy would pan out in this case. On the analogue of the pluralist view here, *each* of the gauge-related solutions in a particular equivalence class of the theory, the empirical substructures of which correspond to the empirical data in the actual world, is instantiated in the actual world. But—and here is our response to this position—since each of these structures could *individually* give rise to that observed empirical data, it simply does not seem correct to state that all structures *together* must be co-instantiated in order to account for this data. To claim otherwise appears metaphysically otiose.³³ In our view, the central challenge for the pluralist is to articulate a sense in which the above analogy does *not* hold, the co-instantiated structures *do* non-redundantly account for the observed empirical data, and

²⁹One recent notion in metaphysics to which the pluralist might appeal here is that of ‘multiple grounding’—cf. [Dasgupta, 2014].

³⁰Cf. footnote 20.

³¹Indeed, one might treat the $U(1)$ gauge symmetry of the vector potential formulation of electromagnetism as giving rise to a ‘self-duality’, in which case this example *just is* a case of underdetermination of the kind considered in this paper.

³²This is, in a sense, the analogue of the common core approach discussed above; for the (heterodox) analogue to the discrimination approach, see Maudlin’s suggestion that there is “one true gauge” [Maudlin, 1998, p. 367].

³³Our mention of the co-instantiation of all gauge-related structures may remind the reader of Pitts’ approach to gravitational energy in general relativity [Pitts, 2010]. Exploring the overlap—if ultimately any—between this view and pluralism would constitute an interesting task for future pursuit.

therefore the problem of overdetermination is evaded.

Thus, pluralism faces a number of challenges, if it is ultimately to be regarded as being a compelling resolution to the problem of underdetermination in the case of dualities. Nevertheless, for the purpose of fully mapping the terrain on this topic, and since the view is *prima facie* consistent, it certainly deserves to be studied further.

9 Conclusion

In this paper, we have cut the issue of the interpretation of dualities along two distinct axes. First, which of the dual solutions under consideration should be taken to constitute a legitimate description of the actual world. The options here divide into three categories: underdetermination, discrimination, and nihilism. Second, candidate replacements for the ontology represented by the dual solutions under consideration, naïvely interpreted. Though the best-known approach in the philosophical literature in this regard is the common core approach, we have identified in this paper two others: (a) appeal to an overarching theory; and (b) what we have dubbed ‘pluralism’. This latter position offers a novel avenue for the interpretation of dualities—albeit one that currently faces difficulties.

Acknowledgements

Very many thanks to Nick Huggett, Niels Linnemann, Keizo Matsubara, Tushar Menon, Oliver Pooley, Dean Rickles, Christian Wüthrich, and an anonymous referee for this journal, for valuable comments on earlier drafts of this article. We are also grateful to Augustin Baas, Rasmus Jaksland, and Vincent Lam for feedback, and to Sumana Sharma for constructing the figures. The work of B.L.B. was performed under a collaborative agreement between the University of Illinois at Chicago and the University of Geneva, and was made possible by grant number 56314 from the John Templeton Foundation. J.R. is supported by an AHRC graduate scholarship at the University of Oxford, a senior scholarship at Hertford College, Oxford, and a visiting fellowship at the University of Illinois at Chicago. The latter was again funded by grant number 56314 from the John Templeton Foundation. The contents of this paper are solely the responsibility of the authors, and do not necessarily represent the official views of the John Templeton Foundation.

References

- [Alexander, 1986] Alexander, H. G., editor (1986). *The Leibniz-Clarke Correspondence: Together with Extracts from Newton's Principia and Optiks*. Manchester University Press.
- [Ammon and Erdmenger, 2015] Ammon, M. and Erdmenger, J. (2015). *Gauge/Gravity Duality: Foundations and Applications*. Cambridge University Press.
- [Barbour, 2010] Barbour, J. (2010). Shape dynamics: an introduction. In *Proceedings of the Conference: Quantum Field Theory and Gravity*, Regensburg.
- [Becker et al., 2007] Becker, K., Becker, M., and Schwarz, J. (2007). *String Theory and M-Theory: A Modern Introduction*. Cambridge University Press.
- [Blumenhagen et al., 2013] Blumenhagen, R., Lüst, D., and Theisen, S. (2013). *Basic Concepts of String Theory*. Springer International Publishing.
- [Callan et al., 1985] Callan, C., Friedman, D., and Perry, M. (1985). Strings in background fields. *Nuclear Physics B*, 262(4):593–609.
- [Creath, 2017] Creath, R. (2017). Logical empiricism. In *The Stanford Encyclopedia of Philosophy*.
- [Dasgupta, 2014] Dasgupta, S. (2014). On the plurality of grounds. *Philosophers' Imprint*, 14(20):1–28.
- [Dawid, 2013] Dawid, R. (2013). *String Theory and the Scientific Method*. Cambridge University Press.
- [De Haro, 2017a] De Haro, S. (2017a). Dualities and emergent gravity: Gauge/gravity duality. *Studies in History and Philosophy of Modern Physics*, 59:109–125.
- [De Haro, 2017b] De Haro, S. (2017b). Spacetime and physical equivalence.
- [De Haro, 2018] De Haro, S. (2018). The heuristic function of duality. *Synthese*.
- [De Haro and Butterfield, 2018] De Haro, S. and Butterfield, J. (2018). A schema for duality, illustrated by bosonization. In Kounieher, J., editor, *Foundations of Mathematics and Physics One Century After Hilbert*, pages 305–376. Springer International Publishing.
- [Dewar, 2017a] Dewar, N. (2017a). Interpretation and equivalence; or, equivalence and interpretation.

- [Dewar, 2017b] Dewar, N. (2017b). Sophistication about symmetries. *British Journal for the Philosophy of Science*.
- [Dieks et al., 2015] Dieks, D., van Dongen, J., and de Haro, S. (2015). Emergence in holographic scenarios for gravity. *Studies in History and Philosophy of Modern Physics*, 52:203–216.
- [Earman, 1989] Earman, J. (1989). *World Enough and Spacetime*. MIT Press.
- [Esfeld and Lam, 2008] Esfeld, M. and Lam, V. (2008). Moderate structural realism about space-time. *Synthese*, 160(1):27–46.
- [French, 2010] French, S. (2010). The interdependence of structure, objects and dependence. *Synthese*, 175(1):89–109.
- [French, 2014] French, S. (2014). *The Structure of the World: Metaphysics and Representation*. Oxford University Press.
- [Friedman, 1983] Friedman, M. (1983). *Foundations of Space-Time Theories*. Princeton University Press.
- [Friedman, 1999] Friedman, M. (1999). *Rethinking Logical Positivism*. Cambridge University Press.
- [Godfrey-Smith, 2013] Godfrey-Smith, P. (2013). *Theory and Reality: An Introduction to the Philosophy of Science*. University of Chicago Press.
- [Green et al., 1987] Green, M., Schwarz, J., and Witten, E. (1987). *Superstring Theory: Volume 1*. Cambridge University Press.
- [Horowitz and Polchinski, 2009] Horowitz, G. and Polchinski, J. (2009). Gauge/gravity duality. In Oriti, D., editor, *Approaches to Quantum Gravity: Toward a New Understanding of Space, Time and Matter*. Cambridge University Press.
- [Huggett, 2017] Huggett, N. (2017). Target space \neq space. *Studies in History and Philosophy of Modern Physics*, 59:81–88.
- [Huggett and Callender, 2001] Huggett, N. and Callender, C. (2001). Why quantize gravity (or any other field for that matter)? *Philosophy of Science*, 68:382–394.
- [Huggett and Vistarini, 2015] Huggett, N. and Vistarini, T. (2015). Deriving general relativity from string theory. *Philosophy of Science*, 82(5):1163–1174.
- [Huggett and Wüthrich, 2013] Huggett, N. and Wüthrich, C. (2013). Emergent spacetime and empirical (in)coherence. *Studies in History and Philosophy of Modern Physics*, 44(3):276–285.

- [Ladyman, 1998] Ladyman, J. (1998). What is structural realism? *Studies in History and Philosophy of Modern Physics*, 29(3):409–424.
- [Ladyman, 2014] Ladyman, J. (2014). Structural realism. In *The Stanford Encyclopedia of Philosophy*.
- [Malament, 2012] Malament, D. (2012). *Topics in the Foundations of General Relativity and Newtonian Gravitation Theory*. University of Chicago Press.
- [Maldacena, 1998] Maldacena, J. (1998). The large N limit of superconformal field theories and supergravity. *Advances in Theoretical and Mathematical Physics*, 2:231–252.
- [Matsubara, 2013] Matsubara, K. (2013). Realism, underdetermination and string theory dualities. *Synthese*, 190(3):471–489.
- [Matsubara, 2017] Matsubara, K. (2017). Quantum gravity and the nature of space and time. *Philosophy Compass*, 12(3).
- [Matsubara and Johansson, 2018] Matsubara, K. and Johansson, L.-G. (2018). Spacetime in string theory: A conceptual clarification. *Journal for General Philosophy of Science*.
- [Maudlin, 1998] Maudlin, T. (1998). Healey on the Aharonov-Bohm effect. *Philosophy of Science*, 65(2):361–368.
- [Maudlin, 2012] Maudlin, T. (2012). *Philosophy of Physics: Space and Time*. Princeton University Press.
- [Mercati, 2018] Mercati, F. (2018). *Space Dynamics: Relativity and Relationalism*. Oxford University Press.
- [Møller-Nielsen, 2017] Møller-Nielsen, T. (2017). Invariance, interpretation, and motivation. *Philosophy of Science*, 84(5):1253–1264.
- [Muller, 1997a] Muller, F. (1997a). The equivalence myth of quantum mechanics—part I. *Studies in History and Philosophy of Modern Physics*, 28(1):35–61.
- [Muller, 1997b] Muller, F. (1997b). The equivalence myth of quantum mechanics—part II. *Studies in History and Philosophy of Modern Physics*, 28(2):219–247.
- [Nguyen et al., 2018] Nguyen, J., Teh, N. J., and Wells, L. (2018). Why surplus structure is not superfluous. *British Journal for the Philosophy of Science*.
- [Oriti, 2009] Orti, D., editor (2009). *Approaches to Quantum Gravity: Toward a New Understanding of Space, Time and Matter*. Cambridge University Press.

- [Pitts, 2010] Pitts, J. B. (2010). Gauge-invariant localization of infinitely many gravitational energies from all possible auxiliary structures. *General Relativity and Gravitation*, 42(3):601–622.
- [Polchinski, 2017] Polchinski, J. (2017). Dualities of fields and strings. *Studies in History and Philosophy of Modern Physics*, 59:6–20.
- [Read, 2016] Read, J. (2016). The interpretation of string-theoretic dualities. *Foundations of Physics*, 46(2):209–235.
- [Read and Møller-Nielsen, 2018] Read, J. and Møller-Nielsen, T. (2018). Motivating dualities. *Synthese*.
- [Rickles, 2011] Rickles, D. (2011). A philosopher looks at string dualities. *Studies in History and Philosophy of Modern Physics*, 42:54–67.
- [Rickles, 2013a] Rickles, D. (2013a). AdS/CFT duality and the emergence of spacetime. *Studies in History and Philosophy of Modern Physics*, 44(3):312–320.
- [Rickles, 2013b] Rickles, D. (2013b). Mirror symmetry and other miracles in superstring theory. *Foundations of Physics*, 43(1):54–80.
- [Rickles, 2017] Rickles, D. (2017). Dual theories: ‘same but different’ or ‘different but same’? *Studies in History and Philosophy of Modern Physics*, 59:62–67.
- [Rodriguez-Pereyra, 2014] Rodriguez-Pereyra, G. (2014). *Leibniz’s Principle of Identity of Indiscernibles*. Oxford University Press.
- [Saunders, 2003] Saunders, S. (2003). Physics and Leibniz’s principles. In Brading, K. and Castellani, E., editors, *Symmetries in Physics: Philosophical Reflections*. Cambridge University Press.
- [Teh, 2013] Teh, N. J. (2013). Holography and emergence. *Studies in History and Philosophy of Modern Physics*, 44(3):300–311.
- [van Fraassen, 1980] van Fraassen, B. (1980). *The Scientific Image*. Oxford University Press.
- [van Fraassen, 2002] van Fraassen, B. (2002). *The Empirical Stance*. Yale University Press.
- [Weatherall, 2016a] Weatherall, J. O. (2016a). Are Newtonian gravitation and geometrized Newtonian gravitation theoretically equivalent? *Erkenntnis*, 81(5):1073–1091.
- [Weatherall, 2016b] Weatherall, J. O. (2016b). Understanding gauge. *Philosophy of Science*, 85(5):1039–1049.
- [Weatherall, 2018] Weatherall, J. O. (2018). Categories and the foundations of classical field theories. In Landry, E., editor, *Categories for the Working Philosopher*. Oxford University Press.

- [Worrall, 1989] Worrall, J. (1989). Structural realism: The best of both worlds? *Dialectica*, 43:99–124.
- [Wüthrich, 2005] Wüthrich, C. (2005). To quantize or not to quantize: Fact and folklore in quantum gravity. *Philosophy of Science*, 72(5):777–788.