

A defense of idealism

*submitted to the Central European University in partial fulfillment of
the requirements of the degree of Doctor of Philosophy*

by Daniel Kodaj

Supervisor: Professor Howard Robinson

Abstract

This dissertation defends idealism. Chapter 1 defines idealism as the thesis that physical objects are ideal, with ideal objects defined as objects the existence of which necessitates the existence of subjects who can observe them. The definition is fine-tuned in various ways, and it is related to recent theories of metaphysical fundamentality. Chapter 2 reconstructs three contemporary arguments for idealism, Howard Robinson's argument against real matter and John Foster's arguments against real space. I claim that Robinson's argument is open to the objection that spacetime is real but Foster effectively challenges that idea. Chapter 3 presents a new argument against the reality of spacetime from a puzzle about relativity. Chapter 4 faces two metaphysical objections to idealism, the truthmaker objection, which says that the idealist cannot supply truthmakers for physical truths, and the nomic objection, which says that the idealist cannot explain laws. I argue that these objections can be deflected by conjoining idealism with theism.

Contents

Introduction	1
1 The concept of idealism	4
2 The ideality of matter and space: Three contemporary arguments	29
3 Real spacetime as excess structure	64
4 Truthmakers and laws in idealism	102
Conclusion	123
Appendix: Clock transport synchrony	125
References	130
Long contents	137

Acknowledgments

I would like to thank Howard Robinson for allowing me to be an elephant in his idealist china shop and for wading through an endless series of less than perfect drafts for this essay. Without his help and encouragement, I would have had no chance of completing this project. I also owe him a large general debt for illuminating insights about metaphysics and about the proper way of formulating philosophical arguments.

I would like to thank Mike Griffin for advice on Leibniz and Kant, comments on various preliminary drafts of the argument, lots of help in securing travel grants, and for buying me an embarrassing amount of beer and hot chicken wings to keep me going. I'd also like to thank Hanoah Ben-Yami for help on Chapter 3, Katalin Farkas for the basic idea behind the definition of ideality I propose, and Dean Zimmerman for comments on an early draft of Chapter 1 as well as for his help in making a research trip to Rutgers University possible.

The drawings of the quad in Chapter 3 are by Krisztina Óré.

Materialism is an erroneous way of life, deriving from an erroneous way of thought. Materialism derives from the habit of hanging the head and looking down to matter, instead of looking out with a level gaze on the given of sense, and then looking up to God who gave it.

A.A. Luce (1954: vii)

Introduction

This dissertation is a defense of idealism in the context of analytic metaphysics. It constructs a definition of ideal physical objects, it shows how the hypothesis that physical objects are ideal can be motivated by puzzles in the philosophy of physics, and it argues that there are no decisive reasons for preferring physical realism to idealism, provided that idealism is upheld in conjunction with theism.

Investigating idealism in the context of analytic metaphysics is worthwhile for two reasons. (i) Idealism is a venerable doctrine that was hugely influential in the history of modern philosophy. It is reasonable to ask whether it can be reformulated and defended in contemporary terms. (ii) Philosophers of mind usually embrace some form of physicalism nowadays, dismissing idealism out of hand and treating the dualist minority as an anomaly that is hard to eradicate. It is worth one's while to ask whether the establishment can feel safe about this attitude. To sum up, investigations of idealism are well in order because they can challenge the physicalist orthodoxy while reconnecting metaphysics with an important historical theme.

I emphasize that whenever “idealism” is used in this dissertation, I mean an ontological thesis about the relationship between the mental and the physical, the kind of thesis familiar from Berkeley and (on certain views¹) from Leibniz. Nothing that I say is meant to concern any other doctrine that is or was or could be called “idealism,” such as transcendental, absolute, objective, and Platonic idealism, or idealism conceived as a belief in noble principles etc.²

¹ E.g. Adams (1994: Part III). The theory I defend below also bears some resemblance to the doctrines of the Yogācāra school of Buddhism (see Tola and Dragonetti 2004).

² The meanings of “idealism” are so disparate that not even family resemblance unifies them. See Burnyeat (1982) for an argument that Greek idealism was a form of realism. For two unusually clear takes on Kant's less than pellucid doctrines, see van Cleve (1999) (who portrays Kant as a cautious Berkeley) and Langton (1999) (who makes him into a Locke). For opposing views on German Idealism, see Ameriks (1987) (who thinks that “idealism” has no meaning in Hegel's case), Pippin (1991) (who sees Hegel as a Kantian), Stern (2009) (who thinks that Hegel's idealism is a sort of Platonism), Beiser (2002: 351–5) (who portrays German Idealism as a form of vitalist Spinozism), and Brandom (2009) (who thinks that German Idealism prefigured his own brand of pragmatism). See Sprigge (1983, 1993, 2006: Ch. 5) for an interpretation of British Absolute Idealism as a form of panpsychism, and Quinton (1972) for the thesis that their core doctrine was holism. (Ewing (1934) offers a piecemeal analytic critique of Absolute Idealist tenets.) Putnam (1978: 18) calls verificationism a form of idealism. One sometimes also hears about “linguistic idealism,” a mythical postmodern doctrine on which language creates reality (see Fodor (1998: 36) for such a use of the word). “Idealism” is a very confused term.

I'll assume that, as far as our dialectic is concerned, idealism has two rivals: physicalist monism (the thesis that everything real is physical) and dualism (the thesis that there are real physical entities as well as real nonphysical entities or properties). I'll call the shared assumption of these two doctrines (the assumption that there are real physical objects) "physical realism."

I restrict my attention to the debate between physical realism and idealism because these are the basic options about the ontology of the physical once we accept that (i) something is real and (ii) whatever is real is either physical or mental (at actuality and nearby worlds, in any case). My defense won't address views that deny (i) or (ii). Neither will it address theories that deflate or dismiss ontology or question our epistemic access to ontological truth. I have neither the space nor the expertise to tackle various forms of anti-metaphysicalism. However, this restricted perspective is sufficient to meet my two main goals. The proper contemporary formulation of idealism is unaffected by issues about the viability of metaphysics, and challenging physical realism is enough to challenge the physicalist orthodoxy, because physicalism about the mental is closely allied with realism about the physical nowadays.

My defense of idealism will be a defense of the following argument:

- (I) If physical realism is true, then the particles and fields studied by fundamental physics are real.
- (II) If physical objects are real, they have real categorical properties.
- (III) The only categorical property of the particles and fields studied by fundamental physics is spatiotemporal position.
- (IV) Spatiotemporal position is not a real property.
- (V) Therefore, physical realism is false.

Chapter 1 clarifies the concept of real and ideal objects and properties and it defines idealism as the claim that physical objects are ideal. Chapter 2 defends premises (II) and (III) through a reconstruction of Howard Robinson's case for idealism, and it endorses two arguments for (IV) by John Foster. Chapter 3 is a new argument for (IV), based on a puzzle from the philosophy of relativity. Chapter 4 focuses on indirect reasons to reject the argument at hand, reasons to maintain that one of (I)–(IV) is guaranteed to be false because the conclusion is metaphysically unacceptable. I'll argue that there are no good reasons to treat idealism as metaphysically unacceptable, provided that idealism is upheld in conjunction with theism.

Premise (I) is taken up in the conclusion. I'll claim that denying (I) makes the physical realist committed to a theory that is unlikely to be compatible with current mainstream forms of physicalism. Denying (I) is therefore not a very good move in the present dialectic, the point of which is to see if contemporary philosophy of mind can find a solid metaphysical foundation in real matter.

This dissertation is a *defense* of idealism, not a proof of it. By a “defense of theory Θ ,” I mean an argument establishing that theory Θ is (i) consistent, (ii) well motivated by certain puzzles in its domain, and (iii) unlikely to be preferred to its rivals by ideal agnostics.

By “ideal agnostics,” I mean (hypothetical) addressees of philosophical arguments who have no convictions or intuitions whatsoever about the domain under scrutiny (the domain at hand being the ontology of the physical world). I envisage philosophical argumentation as akin to a trial: a case is presented for and against some theory (in our case, idealism), and a group of ideal agnostics pass a sentence solely on the basis of the strength of the arguments, without any regard to private convictions and intuitions. If, from this objective viewpoint, there is reasonable doubt that the theory on trial is false or that it is preferable to any of its rivals, the defense is successful.³

Ideal agnostics don’t exist. Worse, the participants of actual philosophical debates will have conflicting opinions about what an ideal agnostic would say in a given situation. Still, my dissertation will proceed under the pretense that ideal agnostics exist and that they are addressees of the following arguments. I have two reasons for working under that pretense.

The first reason is that two of my premises, (III) (the claim that all fundamental physical properties are dispositional) and (IV) (the claim that spacetime is ideal) involve difficult issues in the philosophy of physics that cannot be settled in a dissertation (or at all, perhaps). My aim will be to show that these premises are nonetheless defensible in the sense that they fit available data, they have explanatory value, and they are not open to any knock-down objections.

The second reason why I offer a defense and not a proof of idealism is that a defense is enough to address the two core issues that make idealism interesting nowadays. Constructing a defense is sufficient for showing that idealism can be reformulated in contemporary terms, because a defense must establish that the theory being defended is coherent, which, in turn, requires formulating it. And a defense of idealism is sufficient for challenging the physicalist orthodoxy, since a successful defense of idealism will show that idealism is not preferable to physical realism from an objective standpoint. If that much is established, then cocksure confidence in the truth of physicalism turns out to be rooted in prejudice, fashion, or dogma.

³ For more on ideal agnostics, see van Inwagen (2006: 44–49). Note, however, that on van Inwagen’s conception, a defense of theory Θ is successful iff ideal agnostics are converted into believers in Θ . I only require them to develop reasonable doubt that Θ is false.

1 The concept of idealism

The goal of this opening chapter is to clarify the concept of idealism. I'll define idealism as the thesis that physical objects are ideal, and I'll define ideal objects (roughly) as objects that necessarily coexist with subjects who can observe them. The chapter argues that the proposed definition captures the necessary and sufficient condition for the truth of idealism, that it is superior to rival formulations of idealism, and, when appropriately fine-tuned, it is immune to counterexamples.

Section 1.1 kicks off the dialectic by arguing that idealism has no standard analytic formulation today and the formulations that are on the table fail to meet the basic criterion for being a definition, the criterion that the definiens must be better understood than the definiendum.

Section 1.2 introduces the proposed definition of idealism. I'll define idealism as the thesis that physical objects are ideal and I'll define ideal objects (roughly) as objects that necessarily coexist with subjects who can observe them. I'll also construct a theistic variant of the definition, one on which ideal objects are objects that can only exist if God is disposed to make subjects observe them. I'll argue that the claim that physical objects are ideal in the sense indicated is necessary and sufficient for the truth of idealism, hence this claim captures the core metaphysical commitments of idealism.

Section 1.3 looks at a series of counterexamples to the proposed definition of ideal objects. The counterexamples involve objects that are classified as ideal by the proposed definition but are, or can be conceived as, real. By engaging with these puzzles, I'll construct a fine-tuned version of the definition of ideal objects, one that is immune to all relevant counterexamples.

Section 1.4 engages with the problem of grounding. Recently, concepts of fundamentality, grounding, and dependence have become prominent in analytic metaphysics, and some metaphysicians argue that theories like idealism and physicalism are to be framed as theories about grounding (fundamentality etc.) and not in modal terms. I'll argue that the modal criterion proposed here meets all known formal and substantive features of grounding, hence the charge that my proposed definition of idealism should be replaced by a grounding claim is not very well motivated.

1.1 Problems about defining idealism

Since idealism is virtually never discussed in contemporary metaphysics and philosophy of mind, it lacks a standard analytic formulation. Indeed, it even lacks a standard *non*-analytic formulation—there is no single slogan, awaiting analysis, that one could identify with idealism. Or so I’ll argue here. Consider the following proposals for a definition of idealism:

- (1) Fundamental reality is mental.
- (2) The physical world is mind-dependent.
- (3) There is no external world.

These slogans are certainly suggestive of idealism. But they have very little analytic cash value, or, at any rate, not enough to buy us a definition that can be used without further ado to develop arguments for idealism. A good definiens is better understood than the definiendum—there is no point explaining something through concepts that are harder (or just as hard) to grasp than the concept that is being explained. But (1)–(3) violate this rule.

Proposal (1) violates this rule because “fundamental reality” hardly wears its meaning on its sleeve. The phrase is deeply suggestive, but it is not backed by a familiar everyday concept, nor by a relatively well understood scientific notion that is applicable here. Even worse, its philosophical profile is unclear, because the formal and substantive characteristics of fundamentality are controversial. So (1) defines idealism through a concept that isn’t better understood than idealism itself. Similar remarks apply to (2), since the concept of ontological dependence is about as clear as the concept of fundamentality.

The worth of (3) depends on what one means by “external world.” On one reading of the phrase, even the idealist can agree that there is an external world. The idealist surely has the right to say that there is a *phenomenally* external world, some sort of construction out of actual and possible experience. So in order to give (3) a physical realist spin, one must emphasize that one is denying the existence of a *real* external world. But then the weight of the definition will be carried by the concept of being *really* external, and this concept, in turn, is hardly better understood than idealism itself.

Instead of (1)–(3), one might suggest the classic Berkeleyan slogan:

- (4) *Esse est percipi vel percipere.*

This thesis is often presented as the classic definition of idealism. But its meaning is unclear, because it is unclear what it takes for the *esse* of something to be to φ (e.g. to be perceived). The only interpretation that readily springs to mind is the following:

- (5) The being of x is to $\varphi \equiv \Box (x \text{ exists iff } x \varphi)$

Using (5), one can clarify (4) as

(4*) $\square$ For all x , x exist iff x is perceived or x perceives.

But (4*) fails to fit even Berkeley's own metaphysics. Berkeley sometimes suggests that the existence of physical objects reduces to conditionals about experience:

The table I write on, I say, exists, that is, I see and feel it; and if I were out of my study I should say it existed, meaning thereby that if I was in my study I might perceive it, or that some other spirit actually does perceive it. (*Principles* §3, II: 42, cf. *Principles* §58, II: 65–6 and *Third Dialogue*, II: 251)⁴

This passage implies that physical facts are reducible to facts about what is or would be perceived. But if the existence of this table is reducible to the fact that the table is perceived or *would* be perceived if someone entered the room (which might be currently empty), then (4*) is false, because it is possible that the table (which, presumably, does not perceive anything) exists when nobody perceives it.

But even if we disregard this historical point, it is hard to see why the idealist would be by definition bound to the view that brute physical objects go out of existence when they cease to be observed. But (4*) commits the idealist to that view. It also suggests that we don't exist when we are dreamlessly asleep and nobody watches us, which is implausible. Moreover, (4*) treats idealism as a necessary truth. It is hard to see why an idealist would be forced to treat dualism and physicalism as impossible. So (4*) is unduly restrictive, and since it is hard to see any other interpretation of (4) except for (4*), (4) fails as a definition of idealism.

Let's turn to suggestions by contemporary advocates of idealism. John Foster recommends a definition of idealism along these lines:

(6) The existence of minds is something over and above the obtaining of physical facts, but the existence of physical objects is nothing over and above the obtaining of facts about minds.
(based on Foster 1982: 5–7)

This formulation is problematic because the meaning of “is nothing over and above” is unclear. It sounds like a heavyweight metaphysical notion, one which (to my knowledge) isn't any better understood than idealism itself.

Howard Robinson formulates idealism in the following way:

(7) “The physical world exists only as a complex feature of experience; it exists only ‘in the minds of’ those who do or might experience it.” (Robinson 2009: 203)

⁴ See Foster (1982: 22ff) and Winkler (1989: 205f) for more on this reductive Berkeleyan doctrine.

This sentence cannot define idealism because the meaning of “existing as a complex feature of experience” is unclear. (7) seems to mean something like Berkeley’s slogan, but then we are back to the problems discussed under (4).

Timothy Sprigge offers the following formulation:

- (8) “[T]he noumenal backing [of the physical world] consists in innumerable mutually interacting centres of experience, [...] where by ‘experience’ is meant reality of the same general kind as that of which each momentary centre of experience which is our consciousness at any time is a ‘high-grade’ instance.”
(Sprigge 1983: 85)

This formula is almost wholly occult. Whatever noumenal backings are, their concept isn’t better understood than the concept of idealism.

Finally, consider the following suggestion by Robert Adams:

- (9) There are no unperceiving substances, and spatiotemporal relations are reducible to internal features of consciousness.
(based on Adams 2007: 47–9)

This definition uses the concept of substance, which is probably the single most difficult concept in metaphysics, one that certainly isn’t better understood than idealism. (If you disagree, please read *Metaphysics Z* and Berkeley’s *Dialogues* in quick succession.) Note, further, that (9) cannot be amended by replacing “substances” with “entities,” because the resulting thesis will entail that there are no unperceiving objects. But the idealist surely isn’t by definition forbidden to think that there are no (unperceiving) tables and chairs (not real ones, of course, but some sort of constructions out of experience).

It seems safe to conclude that the analytic definition of idealism is far from settled. This is bad news, because without a clear understanding of the positive claims of the idealist, arguments against materialism can backfire. Suppose, for example, that eccentric Eric finds (1)–(9) meaningless but he is convinced by contemporary idealists that (i) the physicalist conception of matter involves a vicious regress (Robinson 2009) and (ii) the physicalist conception of spacetime is contradictory (Foster 1982: 162–75). Eric’s position is not incoherent, for it might be the case that there is no intelligible conception of idealism but the physicalist conception of matter *is* viciously regressive and the physicalist conception of spacetime *is* contradictory. That would be the case, for example, if all our conceptions of reality were ultimately senseless. Eric, who finds idealism unintelligible but accepts the arguments just mentioned, has warrant for that claim. But it can hardly be the goal of the idealist to support a view like Eric’s (even if there are no real people like Eric). The idealist needs a definition that makes her positive claims very clear. And the physical realist needs that definition too. You cannot disagree with something you don’t understand.

1.2 Introducing a definition

My proposed definition of idealism says, roughly, that actual physical objects essentially coexist with subjects who can observe them. My proposal is close to the “*esse est percipi*” slogan (at least if we use (5) to interpret the latter), but instead of saying that the *esse* of physical objects is to be perceived, I’ll claim that their *esse* is to be observable. (Fine print to follow.)

I’ll frame my proposed definition in terms of a distinction between real and ideal objects. The category of the real and the category of the ideal are meant to be disjoint: no real object is ideal. Moreover, these categories are meant to be exhaustive in the sense that everything is either real or ideal or composed of both real and ideal things. I’ll also be committed to the thesis that *being real* and *being ideal* are essential properties. Nothing is contingently real or ideal.

Here are the first-pass definitions of reality and ideality, to be fine-tuned later in this chapter:

Reality:

(R) O is real $=_{df}$ O is not ideal and O has no ideal parts

Ideality:

(IS) O is ideal $=_{df}$

□ For all t , O exists at $t \supset$ Some subject S observes O at t
or S would have observed O at t if S had performed some
exploratory action prior to t .

By “exploratory actions,” I mean actions that lead directly to observation, e.g. actions that render an object visible from the subject’s vantage point or make it impact the subject’s senses in some other way. Exploratory actions include actions like moving about, looking at things, listening, touching, smelling, and tasting, and using instruments like spectacles, microscopes, cameras, detectors etc. Exploratory actions are all and only those actions the sole function of which is to lead directly to observation.

Ideality is meant to involve our *ordinary* concept of observation, the one we use in everyday life and in scientific contexts to express the fact that some concrete entity was perceived or was inferred on the basis of perception. And *Ideality* is meant to involve a *wide* concept of observation, one on which it makes sense to say things like the following: I observed my neighbor steal my morning paper today, scientists at CERN observed the Higgs boson, a blind person observed the train leave the station. (In contrast, a *narrow* concept of observation would restrict the notion to the visual modality, ruling out the third example, and/or to immediate objects of perception such as a temporal part of my neighbor’s body or a digital reconstruction of a scattering event, ruling out the first two examples.)

Ideality entails that things might exist without being real. (Ideal objects, if there are any, exist but they are not real.) This consequence of *Ideality* may offend the ears of those philosophers who take “exist” and “real” to be synonymous. If you are one of those philosophers, I ask for your indulgence on the following grounds: I propose to use “exist” in a lightweight sense that relates existence to the truth of quantified sentences—*Fs* exist iff there are *Fs*. Since the idealist surely has the right to say that there are tables (not real ones, of course, but some sort of constructions from experience), distinguishing existence from reality is justified in the present context.

With these preliminaries in place, one can define idealist worlds as worlds where all physical objects are ideal, and one can define physical realist worlds as worlds where some physical objects are real. And one can define idealism as the thesis that our world is an idealist world and one can define physical realism as the thesis that our world is a physical realist world. In short:

Idealism:

All actual physical objects are ideal.

≡ No actual physical object is real.⁵

Physical realism:

Some actual physical objects are real.

Assuming that real things are either physical or mental and that there must be real things for there to be anything, *Idealism* entails that there are real mental entities. I'll assume that these are sense data and/or immaterial minds, but otherwise I'll remain neutral about the details.

The requirement that there must be real things for there to be anything will be taken as axiomatic. The assumption that real things are either physical or mental can perhaps be doubted on the grounds that there might be entities that belong to alien sortals and are neither physical nor mental. I'll assume that worlds where some things are neither physical nor mental are remote from actuality, and hence exactly one of *Idealism* and *Physical realism* is true.

My final core assumption is that the ideality of physical objects entails that their properties are also ideal, in the following sense:

Ideal properties:

(IP) Property *P* is ideal =_{df}

- For all *t*, something is *P* at *t* $\supset$ Some subject *S* observes at *t* that *x* is *P* or *S* could have observed at *t* that *x* is *P* if *S* had performed some exploratory action prior to *t*.

⁵ To derive the first claim from the second via *Reality*, one must presuppose that if a physical object *O* is not ideal but *O* has ideal parts, then *O* has real physical parts. This premise is true if physical objects only have physical parts, parthood is transitive, and everything is either real or ideal or has both real and ideal parts. To minimize clutter, I'll tacitly take these principles for granted.

1.2.1 The basic criterion

Does *Idealism* satisfy the basic criterion that the definiens should be better understood than the definiendum? If it does, then *Idealism* is superior to (1)–(3) and (7)–(9) as a definition, because the latter violate that criterion.

Idealism appears to conform to the criterion in question, because its parts are relatively well known and clearly better known than idealism. *Idealism* involves the concept of physical objects and, via *Ideality*, the concepts of observation, exploratory action, possible worlds, and counterfactual truth. To understand the proposed definition, one only needs modal notions plus an idea of what physical objects are and what observation is. Arguably, these concepts are relatively well known and certainly better known than idealism itself.

1.2.2 Historical comparisons

The next question is whether *Idealism* is a necessary and sufficient condition of idealism. In this section, I argue that it is a necessary condition, because it is implied by historical forms of idealism.

Consider the following four historically inspired theories about the nature of physical objects:

- (I1) Physical objects are sums of immaterial mental particulars.⁶
- (I2) Physical objects are confused representations of other immaterial minds.⁷
- (I3) Physical facts are facts about potential immaterial perceptual states.⁸
- (I4) Physical fact are facts about archetypes in God's mind.⁹

These toy theories resemble physical ontologies proposed by Berkeley and Leibniz. Since Berkeley and Leibniz are the most important idealists in the history of philosophy, we have a good heuristic argument for the adequacy of *Idealism* if *Idealism* classifies (I1)–(I4) as forms of idealism.

Let's say that (I1)–(I4) are claims about *actual* physical objects, so that idealism is not assumed to be a necessary truth. *Idealism* will then classify (I1) as a form of idealism iff sums of immaterial mental particulars must coexist with immaterial minds who can observe them. Presumably, this can be taken for granted in a Berkeleian context, since Berkeley believes that ideas are mental particulars that cannot exist without being perceived.

⁶ Berkeley: *Principles* §48 (II: 61), *Third Dialogue* (II: 249).

⁷ On this Leibnizian idea, see Adams (1994: 241–53) and Hartz (1992). Note that Leibniz wavered between an idealist and a realist interpretation of the thesis that physical objects are collections of monads. (I2) is not meant to beg any hermeneutical questions about this issue. It represents a possible idealist reading of the doctrine in question.

⁸ Berkeley: *Principles* §3 (II: 42), §58, (II: 65–6), *Third Dialogue*, II: 251.

⁹ Berkeley: *Third Dialogue* (II: 248). For more on Berkeley's theory of divine archetypes, see Daniel (2001), Foster (1982: 22ff), and Winkler (1989: 205f).

Idealism classifies (I2) as a form of idealism if representations of immaterial minds must be the intentional contents of mental states. More precisely, (I2) implies *Idealism* iff the following plausible principle holds:

- (I2+) If some immaterial minds are represented in a confused manner, then there is a mind *M* such that *M* observes the representation in question.

If (I2) is conjoined with (I2+), then the existence of physical objects (= the existence of confused representations) requires the existence of minds who can observe them, which, in turn, entails *Idealism*.

Idealism classifies (I3) as a form of idealism if there cannot be facts about potential immaterial perceptual states without the existence of immaterial minds who are the potential owners of the perceptual states in question. This claim sounds plausible in an idealist context, since, intuitively, whatever facts there are in idealist worlds are facts that concern minds; specifically, facts about potential perceptual states concern minds that can have perceptual states with the kind of content in question. So (I3) entails that whenever there is a physical object (= whenever there are facts about potential experiences of specific sorts), there are subjects (the potential bearers of those experiences) who can observe the object in question, hence (I3) entails *Idealism*.

To assess (I4), let's assume that "archetypes" refers to God's conception of the physical world, the divine blueprint on the basis of which God causes us to have various experiences. With this presupposition in place, *Idealism* classifies (I4) as a form of idealism iff the existence of archetypes requires the existence of created subjects who are the intended recipients of the experiences based on the archetypes. For then it follows that the existence of physical objects entails the existence of archetypes, and the existence of archetypes, in turn, entails the existence of subjects who can have the experiences based on the archetypes, so the existence of physical objects entails the existence of subjects who can observe the objects in question.

On the other hand, if divine archetypes can exist in the absence of created subjects, then (I4) does not entail *Idealism*. For then in some possible worlds, there are physical objects (because God has certain archetypes in His mind) but there are no subjects to observe the objects in question.

To sum up, three out of four historical forms of idealism are quite easily accommodated by *Idealism*. The fourth can be taken on board if one takes on board the following alternative definition of ideality:

Ideality (theistic version):

- (IT) *O* is ideal =_{df}
- For all *t*, *O* exists at *t* $\supset$ God is causing some *S* to observe *O* at *t* or God is disposed to cause created subjects to observe *O* at *t* if created subjects perform some exploratory action prior to *t*.

This alternative definition is wholly in the spirit of the original. To see why, consider the following simplified versions of the two definitions:

Ideality (standard version):

(IS) O is ideal $\approx_{df}$

$\square O$ exists at $t \supset O$ is observed or observable by someone at t .

Ideality (theistic version):

(IT) O is ideal $\approx_{df}$

$\square O$ exists at $t \supset$ God is causing or is disposed to cause
created subjects to observe O at t .

If O only exists in worlds where there are created subjects at all times during the history of the world, then O is ideal in the sense of (IS) iff it is ideal in the sense of (IT) (assuming, for the sake of illustration, that God exists), so the two definitions coincide for worlds where history contains finite minds throughout.

The two definitions only come apart when we consider worlds without subjects and worlds where history does not always contain subjects. The theistic definition allows for the existence of ideal physical objects even in these cases, but the standard definition does not. So the theistic definition can be seen as an extension of (IS) to accommodate the hypothesis that physical objects may exist in the absence of finite minds.

To sum up, *Idealism* appears to harmonize with historical forms of idealism. (More precisely, it does if we take into consideration a kindred but more permissive theistic definition of ideality.) This corroborates the hypothesis that *Idealism* states a condition that is necessary for the truth of idealism.

1.2.3 Clash with physical realism

Idealism contradicts physical realism, so it is a sufficient condition for the truth of idealism. Or so I'll argue here.

Under (IS), *Idealism* contradicts physical realism iff physical realism entails that some actual physical object could have existed without being observable. Arguably, physical realism does entail that, since physical realism (in its typical forms, at any rate) allows for the possibility that physical objects could have failed to coexist with subjects. Take, for example, a world where something very much like our history unfolds from the Big Bang until 10 billion years ago, but then, because of a cosmic catastrophe, lethal radiation is unleashed and life never develops. Or take a world that ends shortly after the planets are formed. Such worlds are eminently conceivable if we assume that the ontology of inanimate physical objects has nothing to do with our minds. And since the existence of these worlds entails that some physical objects are not ideal in the sense of (IS), it follows that *Idealism* contradicts physical realism if (IS) is adopted as the definition of ideal objects.

Under the theistic definition of ideality, *Idealism* contradicts physical realism iff physical realism entails that some physical object could have existed without God being disposed to make subjects observe it.

Suppose that physical realism is true, and suppose that a possible world called “Hidden- α ” contains all actual physical objects. Indeed, Hidden- α is an almost perfect duplicate of the actual world. The only difference is that, at Hidden- α , the following conditional is true: If God creates subjects, then, come what may, God prevents them from observing Alpha Centauri. In Hidden- α , God is intent to deceive created subjects about the structure of the world by hiding a specific star from them. This is surely conceivable if God exists and if Alpha Centauri is a real star shining out there in space, regardless of what we do or think.

If physical realism is true, then Hidden- α exists, and if Hidden- α exists, then an actual physical object (Alpha Centauri) is not ideal in the sense of (IT). Alpha Centauri exists in Hidden- α yet God is not disposed to make subjects observe Alpha Centauri in Hidden- α . Hence, if physical realism is true, then not all physical objects are ideal in the sense of (IT), in other words, *Idealism* coupled with the theistic definition of ideality also contradicts physical realism. Hence, *Idealism* in general contradicts physical realism, which means that *Idealism* is a sufficient condition of idealism.

1.2.4 The story so far

I have introduced a definition of idealism and I have argued that (i) it satisfies the most important criterion for being a definition, and (ii) the truth of its definiens is both necessary and sufficient for the truth of idealism.

The proposed definition says, roughly, that actual physical objects necessarily coexist with subjects who can observe them. (An alternative, more permissive version says that actual physical objects only exist in worlds where God is disposed to make subjects observe them.) I claimed in 1.2.1 that this definition satisfies the basic criterion for being a definition, because it is built on modal notions, the notion of physical objects, and the notion of observation, which are, arguably, better known than idealism itself. In 1.2.2, I argued that the proposed definiens is a necessary condition of idealism because four different historical forms of idealism all entail it. In 1.2.3, I argued that the definiens is a sufficient condition of idealism, because the definiens contradicts physical realism.

I conclude that the core metaphysical commitment of idealism is probably *Idealism*, either in its standard or in its theistic formulation. I’ll remain neutral about the choice between these two until Chapter 4. For the sake of simplicity, I’ll use the standard version whenever the choice between the standard version and the theistic alternative makes no difference.

1.2.5 Too thin?

It might be objected that the proposed definiens is “too thin” in the sense that it involves mundane modal and intentional concepts and none of the exotic stuff (sense data, monads etc.) that one might expect from an idealist ontology. One might expect a definition of idealism to tell us *how* physical objects are constructed from mind-stuff. But the proposed definiens is a bare-bones modal condition about observation.

I reply that this is a feature, not a bug. My goal is to capture the minimal metaphysical condition of idealism. I have argued that *Idealism* states a necessary and sufficient condition of idealism. If my argument is sound, then *Idealism* is equivalent to idealism. And if *Idealism* is equivalent to idealism, then, thin or not, the proposed definiens is the core metaphysical commitment of idealism.

It might be objected that the definiens is too thin anyway. One might expect a definition of idealism to tell us an intricate story about immaterial minds and the way their activity gives rise to a physical world.

I reply that such stories are certainly much needed if idealism is true. But for idealism to *be* true, it is necessary and sufficient that all actual physical objects are ideal in one of the two senses of “ideal” introduced here. It seems to me that the demand for a richer definition is misguided. Consider this analogy: There are lots of interesting stories that one might want to tell about the natural numbers. But for some entities to *be* the natural numbers, it is necessary and sufficient that they satisfy a few axioms. Similarly, there are lots of interesting stories that one might want to tell about the structure of idealist worlds. But it is necessary and sufficient for the truth of idealism that *Idealism* be true.

1.2.6 Trivially false?

It might also be objected that idealism is trivially false if *Idealism* is indeed the definition of idealism. Consider a world, Barren, that is a physical duplicate of our world as far as planets, nebulae etc. are concerned but does not contain life. Various actual objects exist in Barren without being observable by anyone. Or consider Hidden- α , the world that is much like our world except that in Hidden- α , God is intent to hide Alpha Centauri from us. An actual object exists in Hidden- α without God being disposed to make us observe it. Barren and Hidden- α therefore collectively make *Idealism* false, whether we interpret the latter on the standard or on the theistic definition of ideality. It follows that idealism is easily refuted if *Idealism* is indeed the definition of idealism. Or so one might argue.

In response, the friend of *Idealism* will want to know a bit more about Barren and Hidden- α . How do we know that they contain actual physical objects?

For simplicity and without loss of generality, let's focus on Alpha Centauri, which the objector claims exists in both Barren and Hidden- α . Let's call the actual Alpha Centauri " α " and let's call the Alpha Centauri-like planets in Barren and Hidden- α " $B\alpha$ " and " $H\alpha$," respectively.

The objection at hand is predicated on the claim that α is the same planet as $B\alpha$ and $H\alpha$. Without this premise, the objector cannot conclude that an actual physical object is not ideal.

As far as I can see, there are two ways to establish that α is the same planet as $B\alpha$ and $H\alpha$: by appeal to imagination or by appeal to structure. One can imagine an object that looks like α , and mentally place it in a lifeless cosmos that looks like ours or in a world where God is a bit of a deceiver. Alternatively, one can appeal to some sort of structural description of the planets in question. One can describe α in terms of quarks, electrons, nuclear fusion etc., and attribute the same structure to $B\alpha$ and $H\alpha$.

If we construct $B\alpha$ and $H\alpha$ on the basis of imagination, we have no reason to assume that they are identical to α . All we know, on the basis of imagination, is that $B\alpha$ and $H\alpha$ look like α . But the fact that two objects look the same does not entail that they are the same object. Presumably, the ideal tree in the quad looks the same as a real tree would. If the ideal tree in the quad did not look the same as a real tree would, then the truth or falsity of idealism would be a perceptual datum. But it isn't. Or, at any rate, it would require a fairly long and strong argument to establish that it is. So imagining a real object that looks like Alpha Centauri does not prove that Alpha Centauri is a real object.

Alternatively, one can construct $B\alpha$ and $H\alpha$ through a structural description. For example, one can say that $B\alpha$ and $H\alpha$ is constituted by quarks and electrons that are arranged in the same way as the quarks and electrons in α .

In this case, the idealist will challenge the presupposition that quarks and electrons can exist without subjects. Without this implicit premise, the objection breaks down, because if quarks and electrons happen to be ideal (which cannot be ruled out at this point in the dialectic), then there are no worlds where quarks and electrons compose an unobservable object.

To sum up, the claim that *Idealism* is trivially easy to refute is based on the combination of an invalid and a question-begging move. The invalid move consists in the assumption that imagining a possible real object that looks like an actual object guarantees that an actual object is real. The question-begging move consists in the presupposition that the concepts in terms of which one gives a structural description of actual physical objects (such as the concept of quarks and the concept of electrons) are concepts of real physical objects. Since the claim that *Idealism* is trivially false can only be established through one of these moves, *Idealism* is not trivially false.

1.3 Complications

In this section, I look at various counterexamples to the proposed definition of ideality. The counterexamples will involve objects that satisfy some variant of my definition of ideality yet can coherently be conceived as real. My goal will be to refine the definition in various ways until it is immune to all relevant counterexamples.

For the sake of simplicity, I'll use the following simplified formulation of the definition as my point of departure:

Ideality-0:

O is ideal₀ =_{df}

□ O exists $\supset$ someone can observe O

The phrase “can observe” is meant to abbreviate “is observing or would observe (if some exploratory action were performed).” Since the details of this disjunction will be unimportant in the following dialectic, I suppress them.

To see why *Ideality-0* needs fine-tuning, consider the following problems:

Self-observation:

If subjects can observe themselves, then all subjects are ideal₀.

But the idealist will probably want to say that subjects are real.

Worldbound objects:

If a worldbound object O coexists with subjects who can observe it, then O is by definition ideal₀. But the notion of worldbound real objects is perfectly coherent.

Essentiality of origin:

If the causal origins of an object are essential to it, then all observable objects that were caused to exist by subjects and can be observed by subjects are ideal₀.

Psychophysical laws:

If the existence of certain organs is sufficient for the existence of subjects who can observe their own organs, then the organs in question are by definition ideal₀. But the concept of such organs does not seem to rule out it that the organs are real.

Alien observers:

If God qualifies as an observer, then all objects are trivially ideal₀.

These problems indicate that *Ideality-0* is not a precise definition of ideality. It is conceivable that some real objects are ideal₀. My goal in 1.3.1–1.3.5 will be to keep amending *Ideality-0* until the resulting definition can handle all five of these problems (and some more).

1.3.1 Self-observation

If whatever is necessarily observable is ideal_0 and subjects are essentially such that they can observe themselves, then all subjects are ideal_0 . Hence, *Ideality-0* cannot capture the notion of ideality that the idealist needs, because the idealist hardly wants to be committed to the thesis that either no subject is real or subjects are not essentially capable of self-observation.

I propose to solve these problems by making the following improvement:

Ideality-1:

O is $\text{ideal}_1 =_{df}$

$\square O \text{ exists} \supset O \text{ can be observed by someone distinct from } O$

Subjects are not distinct from themselves, so self-observation does not make them ideal_1 . Subject are only ideal_1 if they are observable by others in all worlds where they exist.

This tweak is not *ad hoc*. Intuitively, idealism concerns the relation of immaterial subjects to physical objects, and immaterial subjects are guaranteed not to be identical to physical objects. So self-observation is not relevant here. *Ideality-1* is an admissible improvement when it comes to defining the ideality of physical objects.

1.3.2 Worldbound objects

Let's take a worldbound real physical object that happens to be such that it is always observed by someone. Maybe the object in question is very important. Maybe it is the Big Red Button that can unleash the nuclear winter. Let's call it "Big." We are assuming that Big is (i) worldbound, (ii) real, and (iii) observed by someone or other during the whole of its existence.

Being worldbound, Big doesn't exist in any other world and being a very important object, Big is always observed by someone in our world. Moreover, Big is clearly not distinct from its observers (who are humans). It follows that, for all worlds W , Big exists in W only if Big is observed by beings distinct from Big, which makes Big ideal_1 . But we have assumed that Big was real, and the hypothesis that Big is real seems compatible with the rest of the scenario. Consequently, it is conceivable that some ideal_1 objects are real, so *Ideality-1* cannot be the definition of ideality.

I propose to solve this problem by the following upgrade:

Ideality-2:

O is $\text{ideal}_2 =_{df}$

$\square$ For any intrinsic duplicate dO of O , dO can be observed by someone distinct from dO

By saying that x is an intrinsic duplicate of y , I mean that x and y have the same intrinsic structure. For example, an intrinsic duplicate of Big would be a possible big red button with the same dimensions, same physical composition etc. All intrinsic duplicates of Big are big, red, and button-shaped. If Big is composed of real elementary particles, then all intrinsic duplicates of Big are composed of the same type of real elementary particles, arranged in the same configuration. But not all intrinsic duplicates of Big are in the same kind of environment and not all of them are connected to a nuclear missile control center. Generally, an intrinsic duplicate of y is a possible entity which is exactly like y in terms of its internal structure and intrinsic natural properties.

The concept of ideality_2 solves the problem of worldbound objects. Even if Big is worldbound, always observable, and hence ideal_2 , some of its intrinsic duplicates may fail to coexist with subjects if Big is real. Perhaps in a world similar to ours, an intrinsic duplicate of Big is pushed, life is destroyed forever in the nuclear winter, but Duplicate Big is spared and nobody ever looks at it again. If Big is a real object, such scenarios are certainly possible. So if Big is real, then Big will not qualify as not ideal_2 , thanks to Duplicate Big.

Alternatively, if Big is, say, a worldbound sum of red sense data, then Big will be classified as ideal_2 as long as all of its intrinsic duplicates are observed by someone, which is only to be expected if Big is composed of sense data. So *Ideality-2* solves the problem of worldbound objects.

It might be objected that *Ideality-2* is hard to apply to some of the historical ontologies discussed in 1.2.1. Suppose that the existence of Big reduces to facts about potential experience or to facts about divine archetypes. What is an intrinsic duplicate of Big in those cases? I reply that in such cases, intrinsic duplication involves the duplication of the relevant reductive base. For example, if Big is a worldbound ideal button the existence of which reduces to the existence of some archetype A in God's mind, then a duplicate of Big exists in W iff God has a duplicate of A in His mind in W .

1.3.3 Essentiality of origin

Ideality_2 can defuse another type of puzzle. Suppose that the causal origins of objects are essential to their identity and minds bring certain physical objects into existence in all nomologically accessible worlds. For example, imagine a dualist ontology where psychophysical laws guarantee that all human minds bring a real pineal gland into existence during their gestation in the womb. If origins are essential, these pineal glands will necessarily coexist with subjects. If subjects can observe their own pineal glands, then pineal glands will qualify as ideal_1 , even though they are real by hypothesis. Note, though, that the pineal glands in question will not qualify as ideal_2 if some of their intrinsic duplicates fail to coexist with minds. (E.g. because zombies are possible.)

However, *Ideality-2* gives the wrong verdict if the existence of pineal glands is sufficient for the existence of subjects. This is conceivable in dualist and physicalist ontologies where the laws of nature guarantee that each pineal gland comes with a subject attached. Pineal glands will then qualify as ideal_2 (provided that subjects can observe their own organs, which sounds OK). So some real objects can be ideal_2 after all. *Ideality-2* is not the definition of ideality either. Note, however, that this isn't a problem about the essentiality of origins any more. It is a problem about psychophysical laws that make the existence of certain organs sufficient for the existence of subjects. Let's see how ideality_2 can be tweaked to handle this phenomenon.

1.3.4 Psychophysical laws

The problem of psychophysical laws arises if the existence of certain physical organs is sufficient for the existence of subjects. Coupled with the assumption that subjects can observe their own organs, such organs will qualify as ideal_2 . But the existence of psychophysical laws is compatible with realism. So *Ideality-2* is not the definition of ideality.

I propose to solve this issue by making the following improvement:

Ideality-3:

O is ideal_3 =_{df} $\square$ For any part p_O of O : $\square$ for any intrinsic duplicate dp_O of p_O , dp_O can be observed by a subject distinct from dp_O $\equiv$ $\square$ All parts of O are ideal_2 .

Ideality-3 is meant to be neutral on the metaphysics of parthood. It can be restated under mereological nihilism without any significant change. Suppose that "Renée's pineal gland" is a plural referring term that picks out some particles. The nihilist version of *Ideality-3* will say that Renée's pineal gland is an ideal_3 plurality iff each particle that "Renée's pineal gland" picks out is such that all of its intrinsic duplicates are observable. Although *Ideality-3*, as formulated above, contains mereological notions, the idea behind it does not presuppose anything about the metaphysics of parthood. We could just as well talk about organwise arranged particles instead of organs without any change in the way of *Ideality-3* defuses the puzzles at hand. With this proviso in mind, I'll continue to use the idiom of parthood, because it is very convenient.

The gist of *Ideality-3* is that no part of an ideal object has unobservable intrinsic duplicates. Real pineal glands will not qualify as ideal_3 as long as they are composed of parts that have unobservable intrinsic duplicates. For example, if pineal glands are composed of quarks and electrons, and quarks and electrons can exist in the absence of subjects, then pineal glands are not ideal_3 .

It might be objected that *Ideality-3* breaks down if pineal glands are composed of parts the existence of which is sufficient for the existence of subjects. This problem is our next troublemaker.

1.3.5 Thoroughly psychofunctional matter

Imagine a world where physical objects are real and the existence of pineal glands is nomologically sufficient for the existence of minds. Suppose, further, that pineal glands happen to be physical simples. If minds can observe their own pineal glands, which continues to be a sensible suggestion, then *Ideality-3* delivers the verdict that pineal glands are ideal_3 . But we presupposed that they are real. So *Ideality-3* is not the definition of ideality. Or so it seems.

Let's introduce the following concept to investigate this issue:

Thoroughly Psychofunctional Matter:

A piece of matter M is thoroughly psychofunctional =_{df}

- For any part p_M of M : □ for any intrinsic duplicate dp_M of p_M , dp_M is part of the body of a subject who is distinct from dp_M and can observe dp_M .

Thoroughly psychofunctional matter is composed of parts which have no unobservable intrinsic duplicates, because all the duplicates come with a subject attached. The mereologically simple pineal glands introduced above are pieces of thoroughly psychofunctional matter. Thoroughly psychofunctional matter is matter that is organic through and through, so that the existence of any part of it is sufficient for the existence of a whole organism. The idea of thoroughly psychofunctional real matter seems to go best with exotic forms of dualism or physicalism, but it might also be compatible with strange forms of idealism. In any case, it seems coherent to assume that there could be such matter.

Thoroughly psychofunctional matter is a challenge to *Ideality-3*, because everything made of such matter qualifies as ideal_3 , even though the substance in question can be thought of as real, or, at any rate, there is nothing in the scenario to suggest that thoroughly psychofunctional matter can only exist in idealist worlds.¹⁰

In reply, a proponent of *Ideality-3* can say that thoroughly psychofunctional matter is a rather outlandish possibility. We have no reason to think that such matter actually exists. Since *Ideality-3* distinguishes real and ideal objects in all worlds where thoroughly psychofunctional matter does not exist (with actuality being one of them), and since idealism is meant to be a thesis about actuality, we can safely ignore thoroughly psychofunctional matter when we investigate idealism.

¹⁰ There is a similar problem about necessarily coexisting subjects. Suppose that type G subjects always come in groups the members of which can always observe each other: necessarily, each type G subject coexists with other type G subjects who can observe her. (This might come about if type G subjects form holistic societies that determine their identity, so that one cannot duplicate a type G subject without duplicating the rest of the society in which that subject is embedded.) All type G subjects will then qualify as ideal_3 , which sounds wrong, because the scenario does not rule out realism about subjects. This scenario formally coincides with the case of thoroughly psychofunctional matter if we take the body of type G subjects to be composed of other type G subjects.

Further, a proponent of *Ideality-3* can say that thoroughly psychofunctional real matter is not matter but *schmatter*, a substance that bears some similarity to actual matter but is, in fact, a very different metaphysical beast. There is no empirical evidence whatsoever that atoms and molecules cannot exist without constituting organisms, so it seems safe to assume that actual matter is not thoroughly psychofunctional and that thoroughly psychofunctional matter is not matter but *schmatter*, a substance that is similar to actual matter but has a very different nomological profile. Since idealism concerns the mind-dependence of matter, not the *schmind*-dependence of *schmatter*, thoroughly psychofunctional real matter can be safely disregarded when one investigates the ideality of actual physical objects. *Ideality-3* is fine for investigating idealism, because idealism is a thesis about the actual world.

1.3.6 Alien observers

One can trivialize *Ideality-3* by assuming that God can observe everything in all worlds. If this is true, then all entities are by definition ideal_3 (except for God), because all possible intrinsic duplicates of all parts of all objects are observable by a subject distinct from them. So *Ideality-3* collapses if God qualifies as an observer. (More precisely, *Ideality-3* collapses if God is an observer who is distinct from created entities.)

To construct a similar puzzle without theism, suppose that somewhere in the far reaches of our galaxy, there lurks a demon, the Grand Intuitor, who can observe any of his worldmates with his extremely powerful demonic mind. The Grand Intuitor has this capability essentially. And he is such a persistent being that he exists in all worlds that contain a duplicate of an actual physical object.

Given this setup, it follows that all duplicates of all actual physical objects are located in worlds where the Grand Intuitor, a being distinct from physical objects, can observe them. Hence, all duplicates of all parts of all actual physical objects are observable by a subject distinct from them, so all actual physical objects ideal_3 . But it seems coherent to suppose that the Grand Intuitor lives in physical realist worlds.

I propose to solve such problems by restricting the meaning of “observe” in *Ideality-3* to cognitive acts of the kind that we humans (or actual organisms in general) are capable of. Neither God nor the Grand Intuitor observes anything in that sense. Their ways of perceiving are alien.

This restriction is not *ad hoc*. Neither God nor the Grand Intuitor observes things the way we do. “Observe” is not univocal if we mix talk of alien observation with talk of human observation.

1.3.7 Summary of 1.3.1–1.3.6

Here's a quick overview of the preceding dialectic. I began with the following simplified version of the definition of ideality:

Ideality-0:

O is ideal₀ =_{df}

$\square O$ exists $\supset$ someone can observe O

There are five basic problems with *Ideality-0*. First, assuming that subjects can observe themselves, all subjects are ideal₀. Second, worldbound objects that coexist with subjects who can observe them are ideal₀. Third, if origins are essential, then all objects brought into being by subjects who can observe them throughout their existence are ideal₀. Fourth, biological organs the existence of which is sufficient for the existence of subjects are ideal₀. Fifth, if God qualifies as an observer, then everything is ideal₀. So *Ideality-0* cannot define ideality—it is conceivable that some real objects are ideal₀.

I introduced three tweaks to solve these problems. *Ideality-1* defined ideal objects as objects that are necessarily observable by a subject distinct from them (1.3.1). This takes care of the problem of self-observation. I argued that this fix is not *ad hoc* because idealism is meant to be a thesis about the relation of subjects to objects that are distinct from them and not a thesis about the relation of subjects to themselves.

Ideality-2 defined ideality in terms of intrinsic duplicates: ideal₂ objects are such that all their intrinsic duplicates can be observed by a subject distinct from them. Unlike *Ideality-1*, *Ideality-2* can defuse the puzzle about worldbound objects, because even if a real worldbound object (for example, a worldbound real button) happens to be accompanied by subjects all the time, it is to be expected then some of its intrinsic duplicates will fail to coexist with subjects (1.3.2). *Ideality-2* can also defuse a puzzle concerning the essentiality of origin, because objects that were brought into being by subjects and hence do not exist in worlds without subjects will have intrinsic duplicates that are unaccompanied by subjects, so they will not qualify as ideal₂ (1.3.3).

Ideality-3 defined ideal objects as objects composed solely of ideal₂ parts. Unlike *Ideality-2*, *Ideality-3* can defuse a puzzle concerning psychophysical laws (1.3.4).

I have argued that *Ideality-3* is fine as far as the analysis and discussion of idealism is concerned, because idealism is a thesis about the actual world and about subjects like us, but the objections to *Ideality-3* concern worlds that are not actual or subjects that are unlike us. Specifically, I argued that actual matter is not thoroughly psychofunctional (1.3.5) and that alien observers like God and the Grand Intuitor can be ruled out by restricting the meaning of “observation” in *Ideality-3* to acts that we humans are capable of (1.3.6).

1.3.8 The fine-tuned definition of ideality

The fine-tuned definition that emerges from the preceding dialectic is rather complicated, especially if one takes notice of the theistic alternative in defining ideality. The best way to state the fine-tuned definition is to break it up into smaller components:

Exists for

x exists for $y \equiv y$ can observe x

Observability

y can observe $x \equiv y$ is observing x or x would observe x if y performed some exploratory action

Distinctness

x is distinct from $y \equiv x$ and y do not overlap

In itself

x exists in itself $\equiv_{df}$

(a) (standard version)

x does not exist for any y distinct from x

(b) (theistic version)

God is not causing and is not disposed to cause created subjects to observe x

Ideality (fine-tuned version):

O is ideal $\equiv_{df}$

$\square$ For any part p_O of O : $\square$ No intrinsic duplicate of p_O exists in itself

Reality (fine-tuned version):

O is real $\equiv_{df}$

$\square$ For any part p_O of O : $\diamond$ An intrinsic duplicate of p_O exists in itself

These definitions form the basis of my official analysis of idealism as the claim that physical objects are ideal. Although the following chapters will use the simplified definition that was introduced early in this chapter, it is important to see that a counterexample-proof full definition waits ready in the wings.

The fine-tuned definitions entail that no real object is ideal. If parthood is transitive, the definitions also entail that all parts of an ideal object are ideal and all parts of a real object are real. Note, though, that if an object O has both real and ideal parts, then O is classified as neither ideal nor real. It is conceivable that the idealist finds herself committed to such objects—for example, O might be the fusion of an immaterial subject and her ideal body (the idealist version of a soul/body composite). In any case, the fine-tuned definitions entail that real objects are all and only those objects that are not ideal and have no ideal parts, as it was stated at the beginning of this chapter.

1.3.9 P.S. Idealism about abstracta

Ideality is meant to be a definition ideal *concrete* objects. It cannot be applied to abstract objects, because it sounds wrong to say that numbers or propositions are, or could be, observed. But perhaps there are relevant analogous concepts of ideality and reality for the abstract realm.

One might try to transfer *Ideality* to abstracta by replacing “observe” with something like “think of.” The claim, in its most basic form, would then be that ideal abstracta are abstracta that only exist in worlds where someone can think of them. This idea is open to an analogue of the problem of alien observers—in this case, there will be a problem about alien thinkers who necessarily coexist with, and can think of, any abstract object. For example, God will render all abstract objects ideal. There will be analogous problems about worldbound abstracta etc. Let’s pretend that we have introduced appropriate fixes for these problems by stipulating that “think of” refers to human thinking etc. For simplicity, let me use the slogan that an abstract object is ideal iff it only exists in worlds where humans can think of it.

This definition will entail that all necessary abstracta are real (provided that there are possible worlds without humans, which sounds plausible). So numbers and propositions will qualify as real. If one conceives of physical laws as propositions, then the definition will also allow the idealist to say that physical objects are not real but physical laws are real.

The definition can also ground a distinction between necessary abstracta like numbers or propositions and abstract artifacts like novels or symphonies. The former will qualify as real abstracta and the latter may qualify as ideal abstracta. Novels and symphonies will qualify as ideal abstracta if they only exist in worlds where people can read novels and listen to symphonies. Since numbers and propositions seem to be very different from novels and symphonies *qua* abstracta, this feature of the definition in question might be sensible.

I emphasize that a commitment to *Idealism* is not meant to be a commitment to idealism about abstracta. Exploring the concept of idealism about abstracta would need a whole other dissertation. I’m only mentioning this theory to set it aside. Idealism about abstracta would be the thesis that all abstract objects necessarily coexist with humans who can think of them. This theory seems to be a form of nominalism, and it is very different from idealism about physical objects. Even if it might be historically linked to idealism (e.g. by way of Berkeley’s criticism of abstract ideas), the two theories do not entail each other. *Idealism* is not applicable to abstract objects and it seems logically compatible with realism about abstracta, just as idealism about abstracta seems logically compatible with realism about physical objects.

1.4 The problem of grounding

Recently, concepts of fundamentality, dependence, and grounding have become prominent in metaphysics (Correia and Schnieder 2012, Fine 2001, Lowe 1998: Ch.6, Schaffer 2009). For simplicity, I'll use "grounding" as an umbrella term to cover the concepts at hand.¹¹ The goal of this section is to find out whether theories of grounding can challenge my definition of idealism.

Grounding is said to be a very special metaphysical relation that underlies phenomena like composition, realization, and determination. Grounding is also said to supersede supervenience in the analysis of dependence and it is said to supersede modal notions in the analysis of essence. Friends of grounding recommend that we recast familiar claims about the supervenience of the mental on the physical, or of the moral on the nonmoral, or of truth on being, or of wholes on parts, as claims like

(G1) The mental is grounded in the physical.

(G2) The moral is grounded in the nonmoral.

(G3) Truth is grounded in being.

(G4) The properties of the whole are grounded in the properties of the parts.

If friends of grounding are right, then my proposed definition of idealism may be on the wrong track, because my proposed definition is built on modal notions and does not use the idiom of grounding. But if grounding supersedes modal analyses of dependence, then idealism should probably be formulated as the following thesis:

(G5) The physical is grounded in the mental.

The goal of this section is to find out whether we have warrant for thinking that (G5), or some similar thesis that uses the concept of grounding, priority, dependence, or fundamentality, is indeed preferable to my proposed definition, *Idealism*, which does not use any of those concepts.

My strategy will be the following: I'll present the three criteria (two formal and one substantive) that are routinely invoked to characterize the phenomenon of grounding. Next, I'll check if the relation obtaining between subjects and physical objects under *Idealism* meets any of the criteria. I'll argue that it meets all three, so there is no good reason to think that *Idealism* is not the definition of a grounding relation. Consequently, the charge that *Idealism* should be rejected in favour of a claim about grounding is not very well motivated.

¹¹ One important difference between competing notions of grounding, dependence etc. concerns the categorical profile of the relation in question. Schaffer (2009) thinks that the central explanatory concept of metaphysics is grounding, which relates entities. Fine (2001) opts for a primitive notion of essence and a primitive notion of reality, the latter of which is applicable to propositions. Lowe (1998: Ch.6) argues that ontological dependence reaches across categories, e.g. events depend on substances. These issues will be ignored here.

Grounding is generally supposed to have three important features which collectively make it a special metaphysical relation in its own right:

- (C1) Asymmetry. Grounding is asymmetric: If P grounds Q , then Q doesn't ground P .
- (C2) Hyperintensionality. For some A , B , and C , A grounds B but A does not ground C even though B and C necessarily coexist (or are necessarily coinstantiated, or have the same truth value in all worlds etc.).
- (C3) Explanatory relevance. That P grounds Q explains, or helps explain, something about the metaphysics of Q . If P grounds Q , then Q exists in virtue of P and/or the nature of Q is determined by P etc.

These criteria are often invoked to dismiss modal analyses of dependence. For example, supervenience fails to satisfy (C1) and (C2), since everything supervenes on itself and supervenience is insensitive to hyperintensionality.¹²

To see whether *Idealism* satisfies (C1)–(C3), let's remind ourselves of the basic content of my proposed definition of idealism:

Idealism:

All actual physical objects are ideal $\approx_{df}$
 For any actual physical object O , $\Box O \text{ exists} \supset$
 $\supset O \text{ is observed or is observable by a subject.}$

Let's say that φ is the relation that obtains between subjects and any specific physical object O if *Idealism* is true:

- (10) subjects $\varphi O \equiv$
 $\Box O \text{ exists} \supset O \text{ is observed or is observable by a subject}$

The question is whether φ is a grounding relation. If it is, then the charge that *Idealism* should be replaced by a claim about grounding is not very well motivated, because *Idealism* then is a claim about grounding.

We know that φ isn't a grounding relation if φ does not satisfy one of the three main criteria of grounding, (C1)–(C3). Conversely, we reason to doubt that φ isn't a grounding relation if φ satisfies all of (C1)–(C3).

Apparently, φ satisfies (C1), the criterion of asymmetry. If we switch the terms in (10), substituting “physical objects” for O , we get

- (11) Physical objects φ subjects $\equiv \Box \text{ a subject } S \text{ exists} \supset$
 $S \text{ is observed or is observable by a physical object}$

Since subjects are not physical objects in idealist worlds, physical objects cannot observe anything in idealist worlds, hence (11) is false in the context of idealism. Consequently, φ is asymmetric in the context of idealism.

¹² See McLaughlin and Bennett (2011: 3.5), Fine (2001: 11, 15), Schaffer (2009: 364).

For φ to satisfy the hyperintensionality condition, (C2), there must be some x and y such that (i) x and y necessarily coexist, and (ii) subjects φx but they do not φy . Arguably, φ satisfies this condition, so φ is a hyperintensional relation. For consider the following three entities:

S_N : Socrates's nose

$\{S_N\}$: the singleton of Socrates's nose

$[S_N]$: the fact that Socrates's nose exists

These entities necessarily coexist. If idealism is true, then noses are ideal objects, so the following will be true:

(12) subjects $\varphi S_N \equiv \square$ Socrates's nose exists only if
subjects can observe it.

But the following propositions will be false:

(13) subjects $\varphi \{S_N\} \equiv \square$ The singleton of Socrates's nose exists
only if subjects can observe it.

(14) subjects $\varphi [S_N] \equiv \square$ The fact that Socrates's nose exists exists
only if subjects can observe it.

These proposition are false because sets and facts cannot be observed. They are objects of thought, not objects that impact our senses. Or, at any rate, there seems to be a coherent conception of observation on which sets, facts, numbers, propositions etc. cannot be observed, and the idealist is free to adopt that conception, making φ hyperintensional.

The third question is whether the fact that subjects φ physical objects explains what physical objects are. It seems that φ does have explanatory relevance in this sense. Specifically, (10) suggests that the existence of ideal physical objects is essentially related to actions through which subjects explore their environment. So it is reasonable to say that φ explains something about the nature of ideal physical objects.

To sum up, the relation that forms of the basis of *Idealism* satisfies the basic formal and substantive criteria of grounding.

This result can be interpreted in two ways. First, it can be interpreted as a sign that φ is a grounding relation. Then the complaint that *Idealism* should be replaced by a grounding claim is unmotivated. Alternatively, the result can be interpreted as a sign that (C1)–(C3) do not characterize grounding. Then friends of grounding must look for other criteria to distinguish true grounding claims from modal analyses like *Idealism*. Pending any persuasive suggestions about such criteria, I conclude that *Idealism* is the definition of a grounding relation, specifically, it is the definition of the grounding relation that obtains between subjects and physical objects in idealist worlds. Hence, the complaint that *Idealism* should be replaced by a grounding claim is unmotivated.

1.5 Summary of Chapter 1

Chapter 1 offered, refined, and defended a modal definition of idealism. In 1.1, I claimed that extant formulations of idealism do not meet the basic criterion that a definiens should be better understood than the definiendum. I considered seven extant formulations of idealism and I argued that they all fail this test.

Section 1.2 introduced a modal definition of idealism. The definition says that idealism is true iff actual physical objects are ideal, with ideal objects defined as objects the existence of which presupposes the existence of subjects who can observe the objects in question. I argued that this condition meets the basic criterion for being a definition (1.2.1). I also argued that if we take a theistic variant into account, the definiens is entailed by four historical forms of idealism, so we have reason to think that the definiens is a necessary condition of idealism (1.2.2). I also argued that the proposed definiens entails the falsity of physical realism, hence it is a sufficient condition of idealism (1.2.3). Finally, I defended the definition from the charge that it is too thin (1.2.4) and that it makes idealism trivially easy to refute (1.2.5).

Section 1.3 refined the definition of ideality by engaging with a series of puzzles about possible objects that can be coherently conceived as real yet are classified as ideal by one of my definitions of ideality. By the end of 1.3, I constructed a fine-tuned definition that is wholly in the spirit of the original and is immune to the counterexamples.

Section 1.4 confronted the objection that idealism should be formulated as a thesis about grounding and not in terms of modal conditions. I argued that the definiens I propose satisfies the three known features of grounding claims, so the complaint that it should be replaced by a grounding claim is not very well motivated.

The take-home message of Chapter 1 is the collection of three shorthand definitions that will be used in later chapters to argue for idealism:

Ideality (standard version):

O is ideal $\approx_{df}$

$\square O \text{ exists} \supset O \text{ is observed or is observable by someone}$

Ideality (theistic version):

O is ideal $\approx_{df}$

$\square O \text{ exists} \supset \text{God is causing or is disposed to cause subjects to observe } O.$

Ideal properties:

Physical property P is ideal $\approx_{df}$

$\square \forall x: x \text{ is } P \supset \text{Someone observes or could observe that } x \text{ is } P.$

2 The ideality of matter and space: Three contemporary arguments

This chapter reconstructs three contemporary arguments for idealism and uses the resulting dialectic to lay the foundations for the idealist argument that the dissertation ultimately defends.

Section 2.1 deals with Howard Robinson's case for idealism, which falls into two parts. The first part, called "the Power Regress," seeks to show that worlds without categorical properties are metaphysically impossible (2.1). I'll argue for a weaker version of the Power Regress, one that establishes that some nomically relevant properties are categorical. The second phase of Robinson's case for idealism is a refutation of physical realism from the conjunction of the Power Regress with the thesis that every fundamental physical property is a disposition (2.2). I'll argue that fundamental physical properties are indeed typically dispositions, or, at any rate, they lend themselves very easily to a dispositionalist interpretation (2.2.1). But Robinson's argument is vulnerable to the objection that spatiotemporal position is not a dispositional property (2.2.2). I'll also argue that positing primitive physical categorical bases that underlie physical dispositions is only warranted if one has independent reasons to prefer physical realism to idealism.

Section 2.3 reconstructs two arguments by John Foster against the reality of spacetime. Foster's *modal* argument seeks to show that the physical realist must be committed to two contradictory claims, the claim that spacetime could have sustained different laws and the claim that it could not have (2.3.1). Foster's *abductive* argument attempts to prove that the physically relevant structure of spacetime is the structure manifest in experience, hence, by inference to the best explanation, real spacetime can be thrown out of our ontology (2.3.2). I'll claim that both of these arguments are eminently defensible and physical realists can only resist them either by embracing some implausible metaphysical principle (e.g. that the laws of nature are metaphysically necessary) or by showing that we have good reasons, independent of the ontology of spacetime, to prefer physical realism to idealism.

Section 2.4 contains a summary of the whole chapter and it indicates the road ahead. My goal in Chapters 3 and 4 will be to bolster the case for idealism that emerges when one lays Robinson's and Foster's arguments side by side.

2.1 The Power Regress

Consider a billiard ball, Bill, and an empty region of space, Reggie. Suppose that Reggie has the same shape and size as Bill. What's the difference between Bill and Reggie? What properties does Bill have that Reggie fails to have?

The intuitive answer is that Bill fills space whereas Reggie *is* space. Bill has mass, which is spread out in a certain region. Bill can move to other regions. Reggie lacks mass, and Reggie cannot move to other regions because Reggie *is* a region. So having mass and having a spacetime trajectory seem to be the two main properties that set Bill apart from Reggie. Since having a trajectory is an extrinsic feature of Bill's, we must examine mass if we want to know what intrinsic feature makes Bill into a real physical object.

Mass, as it figures in classical physics, determines how an object moves and how it influences the movement of other objects. The heavier an object is, the more reluctant it is to be accelerated by outside forces and the greater its gravitational pull on other objects. (Mass is a much more complicated property in modern physics, but we may ignore this for the sake of illustration.) So, all in all, what we have is that the only intrinsic property that makes Bill into a real object is a property that tells us what Bill does under various circumstances.

The gist of Robinson's argument for idealism is that such power-like properties cannot exhaust the nature of an object, so physical realism is false. This argument involves two central claims: (i) there must be properties that are not powers and (ii) all physical properties are powers. Here is the argument for the first central claim:

We only know what [a power] *A* is if we know what kinds of thing the actualization of its potentiality gives rise to [...], what it is a power to do, what states would constitute its manifestation. Let us call the power which *A* is the power to produce "*B*." So what *A* is, is the power to produce *B*. But this is not informative unless we know the nature of *B*. *B*, being a power, is the power to produce some further power state, call it *C*. [...] It seems that we are moving in a regress. [...] And though I have stated this argument in terms of what we could know, the argument is not essentially epistemological. One could equally well say that what the nature of *A* is depends on what it is a potentiality for, for what a power is, is given by what it is a power to do. What it is a power to do is a function of what would constitute its manifestation, and if the nature of this latter can have no determinate expression, neither can the power which is determined in terms of it.

(Robinson 2009: 192, cf. Robinson 1982: 113–7)

One might try to justify this argument (“the Power Regress,” as I’ll call it) by the following thought. Dispositions are often characterized as “pointing” toward their manifestations. Fragility points toward breaking, negative charge points toward attracting positively charged objects etc. If a disposition *D* had another disposition *E* for its manifestation, *D* would point toward a pointing. It seems sensible to suggest that *D* would then ultimately point toward the state that *E* points toward, and that state would constitute *D*’s true manifestation, with *E* being a mere mediator. Even if *D* points toward another disposition *E*, what bearers of *D* do (how they can change things or can be changed) is determined by the state that *E* points to. *D* ultimately points to whatever *E* points to, in other words, dispositions only mediate the manifestations of other dispositions without constituting their manifestations. But then we get into a vicious regress if we assume that every disposition (or power—I’ll use these terms interchangeably) points toward another disposition.

To milk this intuition, consider a world *W* with a series of fundamental properties, *having type 1 vim*, *having type 2 vim* etc. And suppose that *having type N vim* is the disposition to make a neighbouring object have type *N+1 vim*. Every fundamental property in *W* is a disposition to bestow a disposition to bestow a disposition... *ad infinitum*. This is what a world without categorical features looks like. The proponent of the Power Regress claims that such worlds are metaphysically empty, because they only contain dispositions that mediate further dispositions without any of them ultimately getting a grip on reality.

In Robinson’s original version, the Power Regress is meant to show that pandispositionalist worlds (worlds without nondispositional properties) are impossible. In the following, I reconstruct the Power Regress in a version that has a weaker conclusion. Although I share the intuition that Robinson appeals to, I think that the intuition cannot secure the stronger conclusion, because many philosophers find pandispositionalism compelling¹³ and they are happy to appeal to a conflict of intuitions, dismissing the apparent regress with the help of auxiliary metaphysical principles.¹⁴

In the rest of this section, I’ll consider two versions of pandispositionalism, a necessitarian and a contingent version. On necessitarian pandispositionalism, all possible properties (of concrete particulars) are dispositions. On contingent pandispositionalism, all nomically relevant properties are dispositions, but there

¹³ Pandispositionalists include Bird (2007), Hawthorne (2001), Marmodoro (forthcoming), Mumford (2004), Shoemaker (1980), and Yates (2013).

¹⁴ Regress arguments remotely similar to Robinson’s but subtly different in detail have been proposed by Armstrong (1997: 80), Lowe (2010), and Psillos (2006), among others. For pandispositionalist responses, see Marmodoro (forthcoming) (against Armstrong), Bird (2007: 136–46) (against Lowe—note that Bird treats Robinson’s argument as a variant of Lowe’s), and Marmodoro (2009) (against Psillos).

are (or might be, for all we know) nondispositional properties in worlds where the laws of nature are different from actual laws. I'll claim that necessitarian pandispositionalism is false and contingent pandispositionalism is epistemically unstable, because it requires support from science while undermining the reliability of scientific knowledge.

The road map is as follows. In 2.1.1, I'll tease out a formal consequence of pandispositionalism, the consequence that every property instance is associated with a potentially infinite chain of nomically possible consequences. In 2.1.2, I'll argue that necessitarian pandispositionalism is false, because there are possible worlds without such chains. In 2.1.3, I'll argue that on contingent pandispositionalism, every nomically possible property instance is associated with an *unknowable* infinite chain of nomically possible consequences, and, as a result, contingent pandispositionalism undermines itself. Section 2.1.4 sums up the dialectic and ties up loose ends.

2.1.1 Infinite C-trees

The goal of this section is to show that on pandispositionalism, every property instance is associated with a potentially infinite chain of nomically possible consequences, which can be formally represented as an infinite descending chain of conditionals. I'll call such chains "infinite C-trees."

I'll argue for infinite C-trees from what I take to be the necessary and sufficient condition of being a disposition. At a first pass, I take the basic condition for a property *D*'s being a disposition to be the following:

$$(D) \quad \Box \forall x : D(x) \supset (\exists F, G : F(x) \rightarrow G(x))$$

Notation: " $\supset$ " signals material implication and " $\rightarrow$ " signals conditionals. I'll remain neutral on the relation between these two connectives.¹⁵

(D) says that, necessarily, whatever instantiates *D* is such that a conditional applies to it. Dispositions are typically associated with characteristic stimuli and manifestations—for example, whatever is fragile is such that in some cases (e.g. when struck in the absence of antidotes etc.), it breaks. Generally, whatever has a disposition is such that in certain situations, it behaves in a specific way. So (D) seems like a sensible condition on *D*'s being a disposition. (Doesn't this reduce dispositions to conditionals? I'll argue in a moment that it does not.)

One could reasonably suggest that something *more* than (D) is needed to capture our commitments about dispositions. Consider three problems.

Suppose that numbers exist and whenever the number 3 is adequately conceived of, it is found to be a prime. Then whatever has the property of being the number 3 is such that if it is adequately conceived of, it is found to be prime. Hence, if (D) is a sufficient criterion of a property's being a disposition,

¹⁵ Grice (1989) and Jackson (1987) think that the two connectives are the same. Bennett (2003), among others, disagrees. For an overview of the issues, see Edgington (2006: 2.2–2.4).

then *being the number 3* is a disposition. That sounds wrong. And even if one were prepared to treat *being the number 3* as a disposition, this strange commitment should not follow from the definition of dispositions. So (D) is not a sufficient criterion of being a disposition.

To see a similar problem outside Plato's heaven, let's restrict our quantifiers to concrete particulars and their properties, and suppose that, for any x in any world W , some conditional or other is true about x in virtue of x 's falling under a natural law. (Perhaps there are worlds with particulars that do not satisfy any conditionals, either because no laws apply to them or because the laws do not entail any conditionals about them. Let's assume, for the sake of illustration, that such worlds are not possible.) Then for any property F , all possible bearers of F are such that some conditional or other applies to them, and therefore all properties whatsoever satisfy (D). If (D) is meant to be a sufficient criterion for being a disposition, then it follows that all properties are dispositions. But this sounds wrong. Proving that there are no nodispositional properties should not be so easy.

Moreover, (D) is suspicious even if we restrict our attention to dispositions. Suppose, for the sake of illustration, that all fragile things tend to break when struck and that in some world W , all flammable things are fragile. Then flammability can satisfy (D), as far as W is concerned, simply on account of the fact that flammable things in W tend to break when struck. But this is wrong. Flammability has nothing to do with its bearers' tendency to break. So even if we restrict our attention to dispositions, (D) is suspicious because it can be satisfied for the wrong reasons.

It seems, then, that (D) is not a sufficient criterion of being a disposition. Properties can satisfy (D) for the wrong reasons—properties of abstracta can satisfy (D) on account of their extrinsic relations to thinkers, and any property of concrete particulars can satisfy (D) if all concrete particulars are subject to laws and hence make some conditional or other true. Even dispositions can satisfy (D) for the wrong reasons, on account of their bearers having some other disposition. Clearly, (D) needs to be hedged somehow to make it relevant to dispositions as such.

I propose to solve this issue by conjoining (D) with the requirement that the antecedent of the conditional in (D) involve a state that activates D and its consequent involve one of the manifestations of D :

Dispositionality:

(D*) D is a disposition iff

$$\square \forall x : D(x) \supset (\exists F, G : F(x) \rightarrow G(x))$$

where $F(x)$ is a state that activates D , and

G is a manifestation of D .

(D*) is not meant to be a *definition* of dispositionality. (D*) involves the concept of activation and the concept of manifestation, which are (arguably) parts the concept of dispositions. But, I'm not trying to define dispositionality. I'm trying to capture its immediate formal consequences.

Note that (D*) is noncommittal about the reducibility of disposition to conditionals. It is compatible with (D*) that a disposition *D* can be activated by different stimuli in different possible worlds or even across a single world. It is also compatible with (D*) that a disposition *D* has a multitude of distinct possible manifestations, depending on circumstances. So (D*) does not entail that each disposition can be paired off with a conditional.

(D*) is not vulnerable to the problems that threaten (D). Take the case of *being the number 3*. Even if, necessarily, anything that is the number three is such that it is recognized to be a prime when coherently conceived of, *being the number 3* does not satisfy the second half of (D*), because being coherently conceived of is not a state that activates the property of being the number 3.

Or suppose that every possible concrete particular makes some conditional or other true, on account of being subject to laws. It does not follow from this that (D*) is satisfied by all properties, because it does not follow that the conditional that is true about a specific particular on account of laws involves the activation condition and manifestation of all properties.

Finally, suppose that in some world *W*, all flammable things are fragile and hence flammability satisfies (D), as far as *W* is concerned, on account of the fact that flammable things in *W* break when struck. Flammability does not similarly satisfy (D*), as being struck is not a state that activates flammability and breaking is not one of the manifestations of flammability.

It appears, therefore, that (D*) is a reasonable condition on dispositions. For the purposes of the following dialectic, I'll assume that (D*) is set in stone as far as the metaphysics of dispositions is concerned. It seems to me that denying *Dispositionality* means denying a platitude about dispositions, the platitude that bearers of dispositions are such that certain activating conditions make them manifest a state that is metaphysically linked to the disposition in question.

I assume that whenever properties are mentioned, we restrict our attention to natural properties, those sparse properties that generate a full yet economical description of the world, carving nature at its joints. I also assume that natural properties are properties of concrete particulars, hence from now on, no use of "property" refers to properties of abstracta.

Pandispositionalism is the thesis that all properties are dispositions. Having accepted, for the sake of argument, that (D*) captures the necessary and sufficient criterion of being a disposition, it follows that pandispositionalists must uphold the second half of (D*) (the part which says that possessing *D*

implies some conditional that involves the activation of D and one of the manifestations of D) for all properties.

To sum up: We are assuming that (i) quantification over properties is restricted to natural properties, (ii) (D^*) is true, (iii) all properties are dispositions, in other words, pandispositionalism is true. Putting these together yields what I take to be the formal definition of pandispositionalism:

Pandispositionalism:

$$\forall F: \Box \forall x: F(x) \supset (\exists G, H: G(x) \rightarrow H(x)),$$

where $G(x)$ is a state that activates F , H is a manifestation of F , and “ F ,” “ G ,” “ H ” range over natural properties.


The rest of 2.1 explores the substantive implications of this definition. In the present section, I’ll argue that *Pandispositionalism* associates every property instance with a potentially infinite chain of nomically possible consequences. In 2.1.2 and 2.1.3, I’ll try to show that this leads pandispositionalism into grave metaphysical and epistemological trouble.

Now to infinity. Suppose, first, that x is F . By *Pandispositionalism*, there are properties G and H such that $G(x) \rightarrow H(x)$.

Pandispositionalism applies to all properties, so it applies to G and H too. Therefore, there are properties I, J, K, L such that $I(x) \rightarrow J(x)$ is implied by $G(x)$ and $K(x) \rightarrow L(x)$ is implied by $H(x)$.

Pandispositionalism applies to all properties, so it applies to I, J, K, L too. Therefore, there are properties M, N, O, P, Q, R, S, T such that... etc.

When these consequences of *Pandispositionalism* are mapped out in full, we’ll have an infinite descending chain of conditionals, linked to each other by material implication (the variable “ x ” is suppressed):



I’ll call this structure “an infinite C-tree.” An infinite C-tree is a descending chain of conditionals, linked by material implication as shown above. Note that, by *Pandispositionalism*, every conditional in an infinite C-tree is such that its antecedent is a state that activates the property next up in the tree and its consequent is one of the manifestations of the property next up in the tree.

Pandispositionalism entails that all property instances are associated with an infinite C-tree. So *Pandispositionalism* entails that every property instance is associated with a potentially infinite chain of nomically possible consequences. For suppose that for some x , $F(x)$ holds. The tree tells us that $G(x) \rightarrow H(x)$ is true in that case, so if x becomes G , then $H(x)$ will obtain, in other words, $H(x)$ is a nomically possible consequence of x 's being F . By iteration, the C-tree guarantees that x 's being F has a potentially infinite number of nomically possible consequences. Since F was picked at random, pandispositionalism guarantees that every property instance is associated with a potentially infinite chain of nomically possible consequences. The goal of 2.1.2 and 2.1.3 is to establish that this consequence of pandispositionalism is metaphysically and epistemologically unacceptable.


2.1.2 Finite nomic sequences

This section argues that there are possible worlds without infinite C-trees, and hence pandispositionalism is not metaphysically necessary.

Consider a possible world, *Lightning*, where spheres are spontaneously generated in a wide and otherwise empty space. Every newly generated sphere is red. The spheres are immovable and, apart from the phenomenon I'm about to describe, they do not interact with anything in any way.

If there are spheres around in *Lightning*, there is some chance that one of them is hit by lightning. If a freshly generated, red sphere is hit by lightning, it turns yellow. If a yellow sphere is hit by lightning, it turns green. If a green sphere is hit by lightning, it is annihilated. These changes are due to the colours of the spheres. The bolts of lightning have no other effects.

The nomologically possible behaviour of any given sphere in *Lightning* is described by the following chain of conditionals:



Since *Lightning* has no denizens except the spheres and since the spheres do not interact with anything except the lightning bolts, which have no other effect except recolouring and annihilating spheres, it follows that the nomic sequence depicted above sums up the laws of nature in *Lightning*. (Minus the law about the chancy generation of spheres, but that does not affect my argument.)

I claim that in *Lightning*, no infinite C-tree grows from redness. Hence, *Lightning* makes necessitarian pandispositionalism false.

If an infinite C-tree grew from redness in Lightning, the sequence depicted above would branch off in the leftward and the downward direction to infinity. Let's see if this is conceivable.

For the sequence to branch off downward, annihilation should imply some further conditional. But if x is annihilated, x does not exist any more, so x cannot make any conditionals true any more. Hence, annihilated spheres cannot make any further conditionals true. Hence, the nomic sequence starting from redness cannot branch off further in the downward direction.

It might be objected that annihilation could very well entail a conditional. The annihilation of a sphere might have consequences that do not concern the annihilated sphere. Normally, the annihilation of an object implies that the matter composing it is dispersed, inducing changes that can be expressed in the form of conditionals. Perhaps the annihilation of spheres also has consequences in Lightning, lending a dispositional character to annihilation.

This is a fair point about the annihilation of objects in our world. But the laws of our world do not hold in Lightning. And in Lightning, the annihilation of a sphere has no further consequences.

It might be insisted that the annihilation of a sphere does have certain consequences. For example, if there are 14 spheres around in Lightning when a green sphere called "Gollum" is hit, then the annihilation of Gollum implies that if a further green sphere is hit, there will be 12 spheres around. Hence, the chain of conditionals issuing from Gollum's redness can sprout further conditionals even beyond Gollum's annihilation.

The problem with this suggestion is that the conditional in question is not about Gollum, contrary to what (D*) requires. But even if we disregard this point, the conditional in question also violates (D*) on the grounds that its antecedent is not a state that activates the property next up in the chain and its consequent is not a manifestation of that property. A further green sphere's being hit is not a state that activates Gollum's annihilation, and there being 12 spheres left in not a manifestation of Gollum's annihilation.

It might be suggested that *being annihilated* is not subject to (D*) in the first place, because it is a special kind of property that one should lay aside when one explores pandispositionalism. (*Being generated* might be a similar special case.) Or perhaps *being annihilated* is not even a property, but the absence of properties, and hence it requires a metaphysics that is different from the metaphysics of dispositions.

But even if one accepted that point, the chain cannot continue leftward either, even though there is no sign of annihilation there. The chain could only continue leftward if one could link some conditional to the property of being struck. But there are no reasonable candidates.

It might be suggested any of the other conditionals could be linked to the property of being struck to generate infinite descent. For example, one could claim that any sphere x that is being struck by lightning is such that if x is red, then x will turn yellow. Hence, the chain can continue indefinitely in the leftward direction if one links the conditional “Red $\rightarrow$ Yellow” to *being struck*, and then appropriate further conditionals to redness and yellowness. Nothing about *Pandispositionalism* requires the infinite C-tree to be composed of infinitely many distinct conditionals, so this suggestion is formally adequate, and it makes Lightning conform to the pandispositionalist rule that a C-tree grows from every property instance. (Disregarding annihilation, perhaps.)

The problem with this suggestion is that the conditional that the interlocutor links to *being struck* violates (D*). (D*) requires that the conditional linked to a disposition D be such that its antecedent involves a state that activates D and its consequent involves one of the manifestations of D . But redness is not a state that activates *being struck*. So one cannot continue the chain leftward in the way the interlocutor suggests.

I conclude that no infinite C-tree grows from redness in Lightning. Since necessitarian pandispositionalism requires that an infinite C-tree grow from every property instance, necessitarian pandispositionalism is false.

I can see three possible pandispositionalist responses. The first is to deny that Lightning is a possible world. I see no reason to make the modal plenum impoverished in that way. Lightning is conceivable and there seems to be no reason to think that it’s impossible.

Second, the pandispositionalist can bite the bullet, agreeing that her thesis is not a necessary truth. I’ll discuss the costs of making pandispositionalism into a contingent theory in the next section.

Third, the pandispositionalist can deny that *Pandispositionalism* defines pandispositionalism. Since *Pandispositionalism* follows from *Dispositionality* plus the claim that all properties are dispositions, this move requires denying *Dispositionality*. For reasons outlined earlier, I think this is unwarranted.

2.1.3 Infinite epistemic descent

This section argues against contingent pandispositionalism. Specifically, I’ll argue that C-trees involve an infinite number of distinct properties, each one epistemically prior to the ones upstream, and I’ll go on to claim that this makes contingent pandispositionalism epistemically unstable. To kick off the dialectic, let me offer some general thoughts on contingent pandispositionalism.

As I understand it, contingent pandispositionalism is the thesis that all nomologically relevant properties are dispositions in the actual world, but there are (or might be, for all we know) possible worlds where some nomologically relevant properties are not dispositions.

Presumably, contingent pandispositionalism makes pandispositionalism into an empirical hypothesis, something to be voted for or against on the basis of our best science.

There are two problems with this view. The first is that pandispositionalism sounds more like a metaphysical thesis about what it is to be a property and less like an empirical thesis about actual properties. Making pandispositionalism into an empirical hypothesis blunts the metaphysical edge of the theory.

Second, and more importantly, contingent pandispositionalism is open to a crippling epistemological problem. Let me elaborate.


Our knowledge of dispositions is inseparable from our knowledge of stimulus conditions and manifestations. We identify, investigate, and classify dispositions by identifying, investigating, and classifying their characteristic stimuli and manifestations. Knowing a specific disposition D presupposes a concept of the stimuli that activate D and a concept of the manifestations of D . You cannot know what fragility is without knowing what it is to be struck and what it is to break. You cannot know what flammability is without knowing what fire is. You cannot know what honesty is without knowing what it means to tell the truth. And so on. If F is a typical stimulus condition of D and G an associated manifestation, then you cannot know what D is without knowing what F and G are. In other words, the following principle has considerable immediate appeal:

Known Dispositions:

- (KD) If D is a disposition that figures in actual laws, then
 $\exists F, G$ such that $F(x) \rightarrow G(x)$ is true in a typical case,
 where $F(x)$ a state that activates D , G is a manifestation of D ,
 and F and G are epistemically prior to D .

To deny *Known Dispositions*, one must deny that knowing a disposition requires a conception of the ways it typically gets activated and is manifest. But that platitude seems very hard to deny.

Known Dispositions is a threat to contingent pandispositionalism for the following reason. (KD) gives rise to C-trees with infinite *epistemic* descent. Let's revisit our paradigm C-tree:



If both (KD) and pandispositionalism are true, then this tree will involve an infinite number of distinct properties, all epistemically prior to the ones upstream. It will involve an infinite number of distinct properties, because by (KD) we know that H is epistemically prior to F , L is epistemically prior to H etc., and epistemic priority is surely irreflexive and transitive, hence $F, H, L \dots$ are distinct properties, all of which are epistemically prior to the ones upstream.

The C-trees that grow from actual property instances under the conjunction of (KD) with pandispositionalism will involve an infinite number of properties, all epistemically prior to the ones upstream. We cannot come to know such a structure in its entirety, because we cannot acquire scientific knowledge of an infinite number of properties. Hence, if contingent pandispositionalism is true, then actual laws are unknowable in their entirety.

It might be objected that we can know an infinite number of properties. For example, for all N , we have an idea of what it means to weigh N kgs. We know an infinity of properties of the form *weighing N kgs*.

This fact is irrelevant in the present context. The C-trees generated by (KD) involve infinite series of epistemically more and more fundamental properties, but weighing N kgs is not epistemically prior to weighing $N+1$ kgs, so this example is irrelevant here. Generally, a determinate P^* of a determinable P is not epistemically prior to other determinates of P . Consequently, the infinite C-trees that are generated by (KD) must involve an infinite number of properties each of which is either a determinable or neither a determinable nor a determinate.

An infinity of such properties seems unknowable in its entirety—think about a system of particles with an infinite number of properties analogous to quark colour. Or think about a biological system with an infinite number of distinct species, or with organisms having infinitely long DNA chains. It is hard to see how such systems could be known in their entirety from empirical investigation. Hence, contingent pandispositionalism entails that actual laws are unknowable in their entirety.

Now the idea that actual laws are unknowable in their entirety is, of course, perfectly tenable, and it even has a nice historical pedigree. But it jars against another idea that contingent pandispositionalism drags into the dialectic, the idea that pandispositionalism is an empirical hypothesis. If pandispositionalism is an empirical hypothesis yet pandispositionalism entails that our knowledge of reality is guaranteed to be seriously limited forever, how could science help us establish that pandispositionalism is true? If science is supposed to tell us that all properties are dispositions but science is bound to miss out on most of the properties that there are, we seem to be in an epistemically precarious position. Contingent pandispositionalism is an epistemically unstable theory, because it appeals to science while undermining the authority of science.

One might suggest that contingent pandispositionalism can be accepted on empirical grounds if science tells us that all nomically relevant properties are dispositions. The problem with this suggestion is that, by (KD), we know that contingent pandispositionalism comes with an infinite number of properties that one should know in order to have a full conception of actual laws. We cannot know such an infinite system of properties, hence we cannot attain a full conception of actual laws if contingent pandispositionalism is true. But if we cannot attain a full conception of actual laws, then the fact that science tells us that all nomically relevant properties are dispositions is not reliable evidence that all nomically relevant properties are indeed dispositions. The data is then equally well explained by the hypothesis that contingent pandispositionalism is false, some actual properties are categorical, but we cannot don't find them because we cannot attain a full conception of actual laws. Since contingent pandispositionalism comes with a compulsory unknowability clause, it renders all possible scientific evidence compatible with an alternative hypothesis, the hypothesis that our epistemic limitations preclude us from recognizing the categorical properties out there.

It might be argued that, as long as science only finds dispositional properties, the thesis that there are dispositions all the way down is justified on inductive grounds. Hence, contingent pandispositionalism can be corroborated.

The problem with this suggestion is that, by (KD), we are guaranteed to have imperfect knowledge even of those dispositions that we are aware of. Since the infinite C-trees of the contingent pandispositionalist assign infinitely many epistemically prior properties to any *F*, contingent pandispositionalism implies that our knowledge of any specific property is bound to be imperfect, since we cannot know all properties that are epistemically prior to it. As a result, contingent pandispositionalism compromises our scientific knowledge in such a way that continued confidence in the reliability of induction sounds little more than an *ad hoc* excuse to cling to pandispositionalism.

2.1.4 Objections

The story so far: 2.1.1 introduced infinite C-trees. Infinite C-trees are infinite descending chains of conditionals linked by material implication. I argued that pandispositionalism generates such a tree for every property instance. In 2.1.2, I claimed that there are worlds without infinite C-trees, so pandispositionalism is not necessarily true. In 2.1.3, I argued that contingent pandispositionalism is unconvincing and epistemically unstable. It is unconvincing because it blunts the metaphysical edge of pandispositionalism, turning it from a theory of properties into an empirical hypothesis. And it is epistemically unstable because it compromises the reliability of scientific evidence while relying on such evidence. To sum up, necessitarian pandispositionalism is false, and contingent pandispositionalism is best avoided.

I see four pandispositionalist lines of resistance. The pandispositionalist can deny the contingency of laws, rendering my argument against necessitarian pandispositionalism unsound. This is not the place to settle the question of the necessity or contingency of natural laws, so let me just state my conviction that making laws metaphysically necessary makes the modal plenum impoverished for no good reason. If that's the cost of making pandispositionalism into a necessary truth, then the cost is too high.

The pandispositionalist can resist my complaints against the contingent version of the theory in three ways: by denying (KD), by denying that infinite systems of properties cannot be known, or by denying that the unknowability of actual laws undermines contingent pandispositionalism. For reasons outlined earlier, I find the last of these suggestions implausible.

One might perhaps try to justify the suggestion that we can know an infinite system of properties by appeal to colour and shape: perhaps we are capable of distinguishing an infinite number of distinct colours and shapes, so why could we not survey a whole infinite C-tree?

The problem with this suggestion is that appeals to perceptual experience are invalid in this context, since we are talking about scientific theories. In the present context, knowing an infinite number of properties would mean, e.g., knowing a system of particles with an infinite number of properties analogous to quark colour, or knowing an infinite number of distinct species, or exploring infinite DNA chains in full detail. Such epistemic miracles are unlikely to happen in science.

We are left with the option of denying (KD), i.e. the following principle:

Known Dispositions:

- (KD) If D is a disposition that figures in actual laws, then
 $\exists F, G$ such that $F(x) \rightarrow G(x)$ is true in a typical case,
 where $F(x)$ a state that activates D , G is a manifestation of D ,
 and F and G are epistemically prior to D .

Denying (KD) requires asserting that some dispositions are epistemically prior to (or on a par with) their stimuli and manifestations. Such dispositions may help us identify their stimuli and manifestations, instead of the stimuli and manifestations helping us identify it. Or such dispositions may be knowable independently of their stimuli and manifestations. Free will and some linguistic abilities might fit this mould. In any case, I myself can't see any other potential candidates. But free will and linguistic abilities can hardly help contingent pandispositionalism, since what is at stake in this dialectic is our knowledge of the dispositions in subhuman and inorganic nature.

I conclude that there are no straightforward pandispositionalist responses to my argument, hence a weakened version of the Power Regress is defensible. Some nomically relevant properties are categorical in the actual world.

2.2 From the Power Regress to idealism

Howard Robinson originally offered the Power Regress as an argument for idealism. In the present section, I discuss how one can use the Power Regress for such purposes.

We saw at the beginning of 2.1 that paradigm physical properties, for example the property of having mass, appear to be dispositions. If it can be shown that all physical properties are dispositions, then it is to be expected that the Power Regress will cause some headache to physical realists. Let me discuss two ways this can come about, a serious headache and a milder one.

The Power Regress gives a serious headache to physical realists if the Power Regress is sound in its original Robinsonian formulation, the conclusion of which says that worlds without nondispositional features are impossible. For then one can construct the following argument:

- (1) Every fundamental physical property is a disposition.
- (2) If every fundamental physical property is a disposition, then every physical property is a disposition.
- (3) If physical realism is true, then the world is closed under physics.
- (4) If the world is closed under physics, then there is some possible world which is a physical duplicate of our world and does not contain nonphysical properties.
- (5) Therefore, by (1), (2), and (4), if the world is closed under physics, then there is a possible world where all properties are dispositions.
- (6) By the Power Regress, worlds where all properties are dispositions are impossible.
- (7) Therefore, by (3), (5), and (6), physical realism is false.

The argument has five premises, (1), (2), (3), (4), and (6).

Premise (1) is the claim that all fundamental physical properties are dispositions. This thesis will be discussed in 2.2.1 below.

Premise (2) says that all physical properties are dispositions if fundamental physical properties are dispositions. This is supported by the plausible assumption that all physical properties are reducible to, or built from, or determined by, fundamental ones. Arguably, this is what makes fundamental properties fundamental. So (2) seems like a safe assumption.

Premise (3) says that on physical realism, the world is closed under physics. This amounts to saying that nonphysical entities or properties are not causally relevant if physical realism is true. This sounds OK as far as standard forms of physical realism (physicalist monism and epiphenomenalist property dualism) are concerned.

Premise (4) says that if the world is closed under physics, then there is a possible world which is a physical duplicate of our world and does not contain nonphysical properties. This sounds quite plausible, since the intuition behind physical closure is that you can “remove” all nonphysical stuff (if there is any), leaving the rest intact.

Premise (6) is the original Robinsonian conclusion of the Power Regress.

Since the other premises, with the possible exception of (1), are fairly uncontentious, the crux of the idealist application of the Power Regress is (6).

As mentioned in 2.1, I don’t know how (6), the original strong conclusion of the Power Regress, can be upheld in the face of pandispositionalist resistance. I have defended a weaker version of the Power Regress, one which establishes that pandispositionalism is not true in the actual world, more precisely, that some nomically relevant properties are not dispositions in the actual world. But this weaker premise can also ground a fairly strong idealist-friendly argument:

- (1) Every fundamental physical property is a disposition.
- (2) If every fundamental physical property is a disposition, then every physical property is a disposition.
- (8) Therefore, if there are no nomologically relevant nonphysical properties, then every nomologically relevant property is a disposition.
- (9) By the Power Regress, some nomologically relevant properties are not dispositions.
- (10) Therefore, there are nomologically relevant nonphysical properties.

This conclusion is weaker than (7), but it still contradicts physicalism, and it harmonizes with idealism. It is bad enough for physical realists if (10) is true.

In the rest of 2.2, my goal will be to examine the soundness of these two antiphysicalist applications of the Power Regress. I have already argued that (9) is true and I noted that I find (6), the original strong conclusion of the Power Regress, plausible, so I lay these aside. Of the rest, (2), (3), and (4) seem fairly uncontroversial, so it seems that the premise that must be discussed is (1), the claim that all fundamental physical properties are dispositions. In 2.2.1 below, I argue that we have reasons to think that most fundamental physical properties are dispositions. However, as 2.2.2 points out, spatiotemporal position seems to be an exception to this rule. In 2.2.3, I argue that positing any other categorical properties that might underlie physical dispositions is unwarranted. So the upshot of the discussion will be that spatiotemporal position is the only nondispositional property in fundamental physics. In the rest of the dissertation, this fact will be the basis of my own case for idealism.

2.2.1 Are fundamental physical properties dispositions?

Ordinarily, we often think of physical objects in terms of nondispositional or not obviously dispositional features, such as shape, texture, and colour. But, as it has often been pointed out, most of these properties either fall out of the picture when we move from the manifest image to the scientific image, or they take on dispositional clothing, along with the rest of the physical world. As Simon Blackburn remarks,

Resistance is *par excellence* dispositional; extension is only of use, as Leibniz insisted, if there is some other property whose instantiating defines the boundaries; hardness goes with resistance, and mass is knowable only by its dynamical effects. Turn up the magnification and we find things like an electrical charge at a point, or rather varying over a region, but the magnitude of a field at a region is known only through its effect on other things in spatial relations to that region. A region with charge is very different from a region without: perhaps different enough to explain all we could ever know about nature. It differs precisely in its dispositions or powers. But science finds only dispositional properties, all the way down. (Blackburn 1990: 62–3)

Frank Jackson thinks that fundamental physical properties cannot help being dispositional, since physics investigates the causal role of properties and these, in turn, are cashed out in terms of dispositions:

When physicists tell us about the properties they take to be fundamental, they tell us what these properties do. This is no accident. We know about what things are like essentially through the way they impinge on us and our measuring instruments. (Jackson 1998: 23)

Similarly, Brian Ellis and Caroline Lierse assert that

With few exceptions, the most fundamental properties that we know about are all dispositional. They are of the nature of powers, capacities and propensities. (Ellis and Lierse 1994: 32)

Indeed, the fact that the sciences investigate causal roles has led a number of philosophers to embrace the thesis that *all* properties (or, at any rate, all properties instantiated in the concrete world) are dispositions.¹⁶ One purported advantage of this view is that it makes the nature of properties epistemically accessible. If there would be more to properties than their causal role, then we could not know about their real nature, at least when it comes to subatomic properties that are not directly manifest in sense experience.¹⁷

¹⁶ For a summary of this view, see Hawthorne (2001). See also the references in note 13 on p.31.

¹⁷ See e.g. Yates (2013: 95). For a critical discussion of this point by a dispositionalist, see Hawthorne (2001: 365–8).

The claim that physical properties are at bottom dispositional appears to be strongly supported by contemporary physics. Consider the following summary of the current standard model of particle physics:

Currently, the term “elementary particle” denotes some particles which are globally called *leptons*, and the *quarks*. [...] The quarks and leptons [are] currently considered to be the ultimate constituents of matter. [...] A description of the structure of matter cannot be complete without considering the interactions (forces) that “join” the particles and more generally that regulate the interactions amongst them. The *strong*, *weak*, *electromagnetic* and *gravitational* interactions were identified. Each interaction has its own force particles [*bosons*] that relay the interaction. [...] The *photons* mediate the electromagnetic interaction, the W^+ , W^- and Z^0 *vector bosons* the weak interaction, the eight *gluons* the strong interaction and the hypothetical *graviton* the gravitational interaction. (Braibant et al. 2009: 1–2)

As Braibant et al. explain, elementary particles have four fundamental properties: typical lifetime, mass, spin (which expresses the angular momentum of a particle), and the baryonic or leptonic numbers, which, through various conservation principles, control the way that particles decay into other particles (op.cit. 109–10). At least three out of these four properties (mass, spin, and baryonic/leptonic number) appear to be dispositional. (Typical lifetime is also dispositional if one interprets it as a propensity to persist.) This impression is reinforced when we consider the nature of fundamental interactions:

At present we know of four types of forces as basic interactions in Nature. The strongest of them is the *nuclear force*, attracting protons and neutrons inside the atomic nucleus, although its range is limited to distances of the order of the diameter of the atomic nucleus, i.e., 10^{-13} cm. After this, the next strongest is the *electromagnetic force*, which is exerted between electrically charged particles, and in particular which attracts protons and electrons to form the atom. Then follows the so-called *weak force*, mediating the beta decay of nuclei. This is also a short-range force. [...] Finally, the weakest is the *gravitational force*. Like the electromagnetic force, this has long range. All other interactions observed in Nature can be reduced to these four forces. (Chaichian 2014: 279)

The current standard model of fundamental physics invites us to think of the world in terms of four types of forces that influence the generation and movement of elementary particles that are in a large part individuated by dispositional properties. So contemporary physics seems very much open to the interpretation that fundamental physical properties are typically dispositions.

2.2.2 The problem of spatiotemporal position

Although the view that fundamental physical properties are mostly dispositions seems fairly uncontroversial, one can raise objections against the thesis that *all* fundamental physical properties are dispositional. One important objection is that spatiotemporal position is not, or is not easily interpreted as, a disposition. The fact that a physical object is somewhere in spacetime or that it has a certain trajectory in spacetime seems to be a structural feature about spacetime and not primarily a fact about the particle's readiness to induce or undergo changes.

One can object to this claim on the grounds that having a spatial position entails various conditionals. For example, one might argue that occupying a certain region R entails that the object that occupies R will move to a nearby region if it experiences enough force. But it is far from clear that movability is a nomologically essential consequence of occupying a specific position. It seems coherent to assume that in some world, spacetime contains an absolutely immovable object (for example, an elementary particle that has no mass and does not interact with other particles). And if this scenario is coherent, then having a certain spatiotemporal position does not, in and of itself, necessitate anything about the behaviour of the objects that occupy that position.

Some philosophers recommend a dispositionalist view of spatiotemporal properties on the basis of general relativity. Allegedly, general relativity treats spacetime as a dynamic structure that reacts to the presence of matter, affecting the spatiotemporal trajectory of objects (Bird 2007: 166, Esfeld 2009). Hence, one can argue, spatiotemporal position is a disposition according to our best current theory of spacetime.

This claim can be contested on technical grounds by claiming that certain features of spacetime are purely structural even in general relativity (Bartels 2013, Livanios 2008). Since this problematic is unintelligible for anyone but a circle of experts, I set it aside. But one can object to the suggestion at hand in a less technical fashion too, simply on the grounds that, regardless of the nature of spacetime in general relativity, the notion of a spatiotemporal arrangement is primarily a geometrical notion that does not lend itself to a purely dispositional interpretation.

Imagine that we take a snapshot of the whole universe at a specific time t (from a fixed reference frame). The result is a time-slice that contains temporal parts of actual physical objects located at various distances from each other. (Temporal parts are not essential to the example: very short-lived particles could do the job just as well.) We can imagine that the same time-slice (or a structural duplicate of it) is also found in another possible world W^N where, unlike in actuality, Newtonian laws hold and the curvature of space does not undergo dynamic changes. It seems reasonable to assume that there is such a

possible world, since we can imagine that an alternative (Newtonian) past and future is appended to the time-slice in question, as if an alternative beginning and end were spliced to a frame taken from a movie. The spatial relations that actual objects instantiated at t in the actual world are left intact in W^N , the only difference being that the evolution of their trajectories is governed by Newtonian laws before and after t in W^N . If worlds like W^N exist (and there is no reason to rule them out), then having a certain spatiotemporal position is not essentially tied to dynamic changes in the curvature of space. It is a structural, categorical feature that can be described in purely geometrical terms. Just as triangularity is not a disposition (even though it can accompany dispositions), likewise, spatiotemporal relations are not dispositions (even though they can accompany dispositions, for example, they can accompany dynamic changes in the curvature of spacetime). Or so one can argue.

I conclude that premise (1) of the antiphysicalist arguments from the Power Regress is open to the objection that spatiotemporal position is not a disposition and hence not all properties known from physics are dispositions. We saw that this claim can only be resisted by controversial technical arguments about general relativity. It follows that the soundness of Robinson's case for idealism (and of my own weaker antiphysicalist variation on it) is far from evident.¹⁸

2.2.3 Unknown categorical bases

Some philosophers argue that even if we lay aside spatiotemporal properties, we have no reason to think that the dispositions studied in physics have no categorical bases. As J. L. Mackie explains,

The old definition of mass as quantity of matter is [...] not so far from the mark after all. It is true that we do not know much about what matter is; but it is reasonable to postulate that there is a relatively permanent quantitative something-or-other intrinsic to objects [...]. In saying that an object has such-and-such a mass we may reasonably opt for the interpretation that this is to say that it is such that a certain set of conditionals holds, and that although this style of introduction is dispositional, what is introduced is an intrinsic, quantitative, but otherwise mainly unknown feature. (Mackie 1973: 151)

¹⁸ It might be argued that denying (1) on the grounds that spatiotemporal properties are categorical is a Pyrrhic victory for physical realists, because having a spatiotemporal position cannot be the only nondispositional feature of a real object. However, physical realists are free to reply that filling in space is a primitive property (one that we can understand because we understand geometry), or they can opt for the idea of primitive categorical bases discussed in 2.2.3. Once spatiotemporal position is acknowledged to be a real feature in the physical world, physical realism is guaranteed to be true one way or another. Consequently, idealists are better off attacking the reality of spacetime once they acknowledge that spatiotemporal properties are not dispositions.

George Molnar also endorses the idea of primitive categorical bases:

If the property of exerting a certain force is a definite *something* that the numbers can measure, so is being the *source* of that force. That *about the object* that makes it a source of a force is a (quantitative) power property. It is open to the dispositionalist to say that this is where the quiddity lies, this is what the numbers are numbers of. (Molnar 2003: 179)

The gist of these suggestions is that physical realists are free to assume that pieces of matter have some categorical feature in virtue of which they have the dispositions that physics investigates. For example, physical realists are free to claim that there is some quantity of matter wherever there is a disposition to move, or that there is a categorical source of electric force in charged bodies.

One problem about primitive categorical bases is that there is little reason to think that they are essentially tied to dispositions. As Robinson (1982: 121–3) points out, one can imagine that a piece of magnet is moved to another location but leaves its magnetic field behind, retaining whatever intrinsic structural features it had. If this possibility is coherent, then primitive categorical bases are only contingently associated with the dispositions that they ground, and since physics only tells us about the latter (assuming that properties known from physics are dispositional, which the friend of primitive categorical bases accepts), it follows that we cannot ever know what these primitive categorical bases are.¹⁹ Hence, primitive categorical bases fix physical realism at the price of making physical reality somewhat inscrutable.

Whether this objection to primitive categorical bases is serious depends on whether it is more reasonable for a physical realist to posit primitive categorical bases instead of rejecting physical realism when she starts worrying about the dispositional nature of physical properties. If the physical realist is fond of her ideology, she'll probably treat the hypothesis of primitive bases as an inferred truth and not as a reason to give up physical realism.

However, one can argue that positing primitive categorical bases is not warranted without independent reasons to believe in the reality of the physical. The physical realist can only treat the existence of primitive categorical bases as an inferred truth if she has some good reasons to cling to her ideology in the face of the dispositionalist worry. The physical realist cannot defuse the arguments from the Power Regress by announcing that there are primitive physical categorical bases. One needs independent reasons to believe in the reality of the physical to posit such bases. In the absence of good independent reasons, the hypothesis of real categorical bases is a gratuitous heuristic whose function is to shield physical realism from rational criticism.

¹⁹ Cf. note 17 on p. 45.

2.2.4 Summary of 2.1 and 2.2

In 2.1 and 2.2, I reconstructed Howard Robinson's argument for idealism in two stages. In 2.1, I reconstructed the Power Regress, which (in the version defended here) establishes that some nomologically relevant properties are not dispositions. My defense of this thesis proceeded in three stages. In 2.1.1, I argued that pandispositionalism assigns a potentially infinite chain of nomically possible consequences to every property instance. In 2.1.2, I claimed that there are worlds without such chains, therefore pandispositionalism is not necessarily true—if it's true, it is contingently true, so (presumably) it requires empirical support. In 2.1.3, I argued that pandispositionalism generates infinite epistemic regresses that render contingent pandispositionalism epistemically unstable, because these regresses guarantee that the actual laws of nature are unknowable in their entirety and hence that evidence from science about the stock of actual properties is not reliable.

Section 2.2 showed how the Power Regress can ground antiphysicalist arguments. I discussed a stronger argument that proves idealism and a weaker one that shows that some nomologically relevant properties are nonphysical. The crux of the antiphysicalist applications of the Power Regress was the thesis that all fundamental physical properties are dispositions. In 2.2.1, I argued that fundamental physical properties are typically dispositions. In 2.2.2, I claimed that this thesis is harder to maintain for all fundamental physical properties, because spatiotemporal position does not seem to be a disposition. I indicated that this idea might be open to objections from general relativity, but, arguably, these objections are far from decisive and they are intuitively implausible. Section 2.2.3 dealt with the claim that physical dispositions may have real categorical bases. I conceded that this is possible, but I emphasized that positing primitive categorical bases is only warranted if one has independent reasons to prefer physical realism to idealism.

All in all, the take-home message of the dialectic are the following points: (i) there are nomically relevant categorical properties, (ii) the only candidate categorical property of matter seems to be spatiotemporal position, and (iii) the hypothesis of unknown categorical bases is only warranted if one has independent reasons to prefer physical realism to idealism.

It follows that one can build an argument for idealism that is immune to the objections to Robinson's argument if one shows that (i) spatiotemporal position is not a real property and (ii) there are no good independent reasons to prefer physical realism to idealism. This move exposes the physical realist to the Power Regress without assuming that no physical property is categorical. I'll explore this strategy in a bit more detail at the end of this chapter, after John Foster's arguments against real space have been explored.

2.3 Scrambled and gappy worlds

John Foster constructed a number of ingenious thought experiments to prove that space isn't real. His dialectic, originally presented in Parts III and IV of *The Case for Idealism*, revolves around two arguments. Section 2.3.1 discusses Foster's *modal* argument and section 2.3.2 reconstructs his *abductive* argument. I'll argue that both arguments are deeply suggestive and they force the physical realist to embrace some implausible metaphysical principle or to come up with independent reasons to prefer physical realism to idealism.

2.3.1 The modal argument against spatial realism

Foster's modal argument against real space is based on the combination of two ideas, both of which are said to be entailed by physical realism. The first claim is that the laws of nature are contingent, and the second is the claim that the structure of a space is essential to its identity:

[T]here are two principles which are individually undeniable but irreconcilable from the standpoint of spatial realism. The first, which we may call the *principle of variability*, is that, for any ontologically primitive item, the physical geometry of that item, if it has one, [...] is subject to variation, through different possible worlds, with variation in the relevant laws. The second, [...] the *principle of constancy*, is that, for any genuine space, the geometrical structure of that space is essential to its identity and holds constant through all possible worlds in which the space exists, whatever the variations in nomological organization. These two principles are, clearly, jointly incompatible with the claim that physical space is ontologically primitive. For if [it is], then, by the one principle, its geometrical structure varies from world to world according to its nomological organization, while, by the other principle, its geometrical structure remains constant with the constancy of its identity. (Foster 1982: 172)²⁰

Let me introduce the concept of transworld variation to reconstruct Foster's modal argument:

Transworld Variation:

x can undergo transworld variation in terms of $y =_{df}$
 x exists in more than one possible world and x isn't
 qualitatively the same in all worlds in terms of y
 (e.g. spacetime in terms of the laws that it sustains).

²⁰ Foster defines realism about spacetime as the thesis that spacetime is ontologically primitive (or fundamental). I'll reconstruct his argument in the terminology of real vs. ideal structures instead. This move seems justified since Foster (1982: 162) identifies anti-realism about xs with the thesis that xs are not ontologically primitive.

With this concept in place, Foster's modal argument against spatial realism can be reconstructed as a simple but powerful *modus tollens*:

- (11) If spacetime is real, then it can undergo transworld variation in terms of the laws that it sustains.
- (12) Spacetime cannot undergo transworld variation in terms of the laws that it sustains.
- (13) Therefore, spacetime isn't real.

Foster appears to take (11) to be axiomatic on the grounds that laws are contingent and laws must be sustained by real objects. However, this fact in itself does not entail (11). Foster's "principle of variability" is open to the following complaint:

The laws of nature *are* contingent: our world could have had different laws. But this fact does not entail that *spacetime* could have sustained different laws any more than the modal fact that I could have worn a sandal instead of a shoe entails that my shoe could have been a sandal.

The gist of this objection is that we can make sense of the contingency of laws without assuming that the same spacetime recurs in different worlds. This is true even if we presuppose that nomic variation has to do specifically with the structure of space. Suppose, for the sake of illustration, that space is actually curved but it could have been flat because the laws of gravitation could have been different. This claim can be upheld without assuming that this very space (the one we inhabit) could have been flat instead of being curved. One can claim that space could have been flat because in some possible world, objects live in a flat spacetime.

One may venture the following reply on behalf of Foster: If spacetime is real and spatiotemporal position is a categorical property, then spatiotemporal structure can be pried away from the laws that are instantiated in spacetime and the same spatiotemporal structure can recur in worlds where the laws are different. Hence, (11) is justified if spatiotemporal position is a categorical property, which, as we saw, is a very plausible idea.²¹

Premise (12) is the claim that spacetime cannot undergo variation in terms of the laws it sustains. This idea can be established through an analysis of deceptive worlds. More specifically, it can be reduced to the thesis that a certain type of structural anomaly is metaphysically impossible. Let me elaborate.

²¹ According to Barry Dainton (2001: 232), the target of Foster's modal argument is substantivalism about spacetime. I believe that substantivalism in itself cannot give us (11), because spacetime cannot recur in nomically different worlds if spatiotemporal position is a dispositional property. On the other hand, if spatiotemporal position is a categorical property, then (11) seems justified under relationalism as well.

Imagine that, unbeknownst to you, your bathroom is on the Moon but, due to a nomic irregularity, you never notice this, because the moment you cross the threshold of your bathroom, you are imperceptibly teleported to the Moon, where the laws of gravity are locally overruled and everything feels like home. The illusion is so perfect that electromagnetic rays traveling to your bathroom also teleport to the Moon, so visual experience suggests that you are still in your apartment on Earth. Similar irregularities guarantee that astronomers never catch you in the shower while they are scanning the heavens. All relevant laws of nature are mucked up in such a way that the illusion is simply perfect. Let's say that the world is *scrambled* iff such deceptions obtain. Scrambled worlds are easily conceivable if we assume that space is real and that the descrambling takes place in our representation of the world.

An even more radical deviance would occur if the descrambling did not belong to our conception of spacetime but were somehow part of physical reality. This would be the case if the primitive geometry of spacetime (the primitive metrical relations between objects, e.g. the primitive distance between your bathroom and the Moon) did not match the kinematics of the world (e.g. because some inertial trajectories would qualify as discontinuous according to the primitive metric). If such deviance obtained, it could be the case that your bathroom is geometrically on the Moon (because the primitive distance between it and portions of the Moon is small), but, kinematically, it is on Earth (because, say, it takes relatively little energy to move from your bathroom to your living room but it takes a lot of energy to move from your living room to the Moon).

Let's call worlds with such a split between the primitive geometry and the kinematic structure of space "*Objectively Scrambled Worlds*." Objectively Scrambled Worlds are scrambled worlds where the descrambling is done at the level of kinematic laws and not at the level of our representation of space.

If premise (12) is false, then Objectively Scrambled Worlds are possible. To see why, suppose that spacetime can undergo transworld variation in terms of the laws that it sustains ($= \sim(12)$). Let world *A* and *B* have the same spacetime but different kinematic laws. Arguably, the identity of a spacetime depends on its structure, hence the *A*-spacetime and the *B*-spacetime must have the same underlying structure to be the same spacetime. This underlying structure cannot encode the kinematic laws, since those differ between *A* and *B*. So there is a difference between the primitive geometry of *A* and *B* (which remains constant as we move from *A* to *B*) and their nomologically relevant structure (which varies from *A* to *B*). So the negation of (12) entails that the primitive geometry and the kinematic structure of a spacetime can come apart. Objectively Scrambled Worlds are worlds where these two structures come apart, so if (12) is false, then Objectively Scrambled Worlds are possible. Hence, if Objectively Scrambled Worlds are impossible, then (12) is true.

Objectively Scrambled Worlds are impossible. If they were possible, then there would be some metaphysically relevant difference between worlds where your bathroom, or this building, or Africa etc. is on Earth and worlds where they are on the Moon but everything behaves the same way. But no difference is discernible between such cases. Suppose that in world *A*, Africa is near Europe, and in world *B*, Africa is on the Moon but kinematically, things behave as if Africa were near Europe. Suppose that a meteor hits the Mediterranean in both worlds, causing a tsunami. Given the identity of kinematic laws, the tsunami will have the same effects on Africa and Europe in *A* and *B*. So the fact that Africa is on the Moon in *B* makes no physical difference whatsoever.²²

The metaphysical absurdity of Objectively Scrambled Worlds can be demonstrated by a different route as well. Suppose we have two representations of our (unscrambled) spacetime, $R\alpha$ and $R\beta$. Both representations assign three numbers to all spatial points in the universe. $R\alpha$ does this in a natural way, using an arbitrary point (e.g. the center of the Sun) as the origin. We generate $R\beta$ by setting up an arbitrary bijection between the points in Africa and the points in some region of the Moon, swapping their $R\alpha$ -coordinates and leaving $R\alpha$ the same otherwise. The result, $R\beta$, is a different way to label the points of the universe with triples of numbers.

If Objectively Scrambled Worlds were possible, then the difference between $R\alpha$ and $R\beta$ could correspond to a real difference between two possible worlds that are otherwise exactly alike. Specifically, it would correspond to the worlds *A* and *B* from the earlier example, where *A* is an unscrambled world and *B* is a world where Africa is on the Moon but everything behaves as in *A*. But when we move from $R\alpha$ to $R\beta$, we just change our system of coordinates: we switch from a neat Sun-centered system to a gerrymandered system. So primitive topological differences between worlds that are otherwise alike make no sense.

To sum up, premise (12), the principle of constancy, is supported by the following reasoning:

- (14) If spacetime can undergo transworld variation, then the primitive geometry and the kinematic structure of spacetime can come apart.
- (15) If it the primitive geometry and the kinematic structure of spacetime can come apart, then Objectively Scrambled Worlds are possible.
- (16) Objectively Scrambled Worlds are impossible.
- (12) Therefore, spacetime cannot undergo transworld variation in terms of the laws that it sustains.

²² This point was made by Howard Robinson in conversation.

Note, further, that Foster's modal argument would lose very little of its force if one insisted on the possibility of primitive geometric differences between kinematically equivalent spacetimes. As Foster (1982: 152–60) points out, the physically relevant structure of spacetime will not include the primitive geometry anyway, because the primitive geometry is physically inert. And this fact, in turn, can ground a modified version of the modal argument:

- (11*) If the physically relevant structure of spacetime is real, then it can undergo transworld variation in terms of the laws that it sustains.
- (12*) The physically relevant structure of spacetime cannot undergo transworld variation in terms of the laws that it sustains.
- (13*) Therefore, the physically relevant structure of spacetime is not real.

This weakened conclusion is still a nice blow to physical realism, because it entails that, as far as the physically relevant structure of spacetime is concerned, idealism is more sensible than physical realism.

To sum up: Foster's modal argument against spatial realism is based on two premises, the principle of variability and the principle of constancy, which together entail both the transworld variability and the transworld constancy of the structure of space. I argued that the principle of variability requires the assumption that spatiotemporal position is a categorical property. I also argued, closely following Foster, that the principle of constancy can be established by considering deviant worlds where space is "objectively scrambled" in the sense that the laws about inertia and acceleration do not match the (physically inert) primitive geometry.

All in all, my conclusion is that Foster's modal argument is sound if spatiotemporal position is a categorical property. We saw earlier that denying that assumption exposes the physical realist to Robinson's case for idealism. So treating spatiotemporal position as a categorical property is not a very good move against Foster's modal argument. As a result, Foster's modal argument seems eminently defensible.

The only straightforward physical realists response to Foster's argument is to deny that the laws of nature are contingent. This move makes the principle of variability false, rendering the argument unsound. The situation here is similar to the case of unknown categorical bases (2.2.3): if the physical realist is fond of her ideology, she'll treat the modal argument as proof that the laws of nature are necessary and not as proof that physical realism is false. However, as in the case of primitive categorical bases, this move has no rational appeal unless one has strong independent reasons to prefer physical realism to idealism.

2.3.2 The abductive argument against spatial realism

Foster's second argument against spatial realism is based on an important principle that Foster calls "the principle of representation":

[T]he physically relevant structure of the underlying external reality is that structure under which the reality is empirically represented at the human viewpoint—that structure under which the reality is disposed to reveal itself [...] through the constraints on human experience. I shall call this the *principle of representation*. (Foster 1982: 212)

Foster argues, on the basis of this principle, that real spacetime is an irrelevant posit which, by inference to the best explanation, can be discarded:

[T]he constraints [on experience], on their own, suffice for the creation of the physical world, irrespective of what (if anything) lies behind them [...]. [I]f there is an external reality, it contributes nothing to the existence of the physical world save what it contributes to the obtaining of the constraints [...]. [A]ny alteration of the external reality which makes no difference to the constraints makes no difference to the physical world, or, indeed, to the total physical reality comprising both the physical world and the psychophysical laws. (op.cit. 208)

As I understand it, Foster's abductive argument has the following shape:

- (17) It is possible that the physically relevant structure of spacetime is the structure manifest in experience and not the structure that is really there (if there is one).
- (18) *Being physically relevant* is not a contingent property.
- (19) Therefore, by (17) and (18), the real structure of spacetime (if it has one) is physically irrelevant.
- (20) If something is physically irrelevant, it has no explanatory value in physical ontology.
- (21) Therefore, by (19) and (20), the real structure of spacetime (if it has one) has no explanatory value in physical ontology.
- (22) If a possible physical structure has no explanatory value in physical ontology, we have warrant for thinking that it does not exist.
- (23) If spacetime has no real structure, then spacetime isn't real.
- (24) Therefore, by (21)–(23), we have warrant for thinking that spacetime isn't real.

The argument has five premises, (17), (18), (20), (22), and (23).

Premise (17) says that the real structure of spacetime can be different from the structure manifest in experience. This is supported by a type of deceptive possible worlds one might call “gappy worlds” (Foster 1982: 217–8). Suppose that the region where this room seems to be does not exist. Due to an aberration in the structure of spacetime, the walls of the room touch one another from the inside, so that there is nothing inside. The illusion that there is a room here is due to an irregularity in the way we represent reality. When we approach the doorway, we have the experience as of entering a room, but in reality, we slow down and stop, entering a hallucinatory state. Gappy worlds are conceivable, and there is no reason to deny that they are impossible. So (17) is justified: it is possible for the real and apparent structure of spacetime to come apart.

Premise (18) is the thesis that *being physically relevant* is not a contingent property. In the context of (17)–(24), it is enough to establish this for spacetime structures: we need to establish that if a certain real spacetime structure S_R is physically irrelevant in some world, then S_R cannot be physically relevant in any other world. Coupled with (17), this thesis entails that the real structure of spacetime is physically irrelevant (= (19)).

Foster (1982: 220–23) supports (18) by asking us to consider a series of gappy worlds where the gap gets bigger and bigger. Suppose that W_0 looks like our world and is gap-free, W_1 is a world where there is a small gap (e.g. the region where this room seems to be does not exist), W_2 is a world where there is a slightly bigger gap (e.g. the region where this building seems to be does not exist)... and so on. We may imagine, as a limiting case, a world W_∞ where there is no real space left, and the structure we take to be the real structure of spacetime is wholly ideal.

It is obvious that the real structure of spacetime is physically irrelevant in W_∞ , because spacetime has no real structure there. Moreover, the real structure seems physically irrelevant in the higher regions of the series from W_1 to W_∞ , because the real structure is so far removed from the apparent world there that it would not enter physical explanations even people knew about it.

Foster invites us to determine at which point between W_∞ and W_0 the real structure of spacetime becomes physically relevant. If the real structure is ever physically relevant, it surely is in W_0 , where the real and apparent structures coincide. And since the real structure is physically irrelevant in W_∞ , if the real structure is ever physically relevant, there must be a point between W_∞ and W_0 where it becomes physically relevant. Where is that point?

Foster argues that there is no good answer to this question, because there is no non-arbitrary criterion for claiming that, say, W_{297} (or a vague area around it) is the place in the series where the real structure of space becomes physically relevant.

It might be objected that this reasoning is not different from standard sorites paradoxes, so there is no reason to believe its conclusion. One can imagine a row of heads, H_∞ to H_0 , such that H_0 is completely bald and H_∞ has as much hair as metaphysically possible. We can't say where baldness first enters the series between H_∞ to H_0 , but it would be rash to conclude that there are no bald people. Likewise, the fact that the physical relevance of real spacetime cannot be caught in the act somewhere between W_∞ and W_0 does not mean that real spacetime structure is physically irrelevant. Or so one might argue.

In response, a proponent of Foster's abductive argument can point out that there is an important disanalogy between standard sorites series and the series running from W_∞ to W_0 . In the series from H_∞ to H_0 , we have one specific (kind of) entity, a head, which loses one property (the property of being hirsute) and comes to acquire another property (the property of being bald). In the series running from W_∞ to W_0 , we have two distinct structures, the apparent and the real structure of spacetime. The apparent structure stays physically relevant all through the series, and the puzzle concerns whether the real structure can come to acquire the same property at some point in the series. Foster seems to imply that once the apparent structure of spacetime is acknowledged as physically relevant throughout, there is no no-arbitrary way of picking a part of the series where the real structure kicks in and suddenly becomes important.

One might suggest that the real structure of spacetime is physically relevant in the later part of the series from running from W_∞ to W_0 , the part where real structure more or less matches the apparent structure. As long as there are only a few spots where appearance and reality mismatch, or as long as the missing regions are relatively rarely visited by humans, one is justified to say that the real structure of spacetime is physically relevant.

The problem with this suggestion is that once the apparent structure is agreed to be physically relevant throughout the series, it is hard to see why the fact of approximate isomorphism would make the real structure physically relevant at a certain point. The apparent structure performs the same role, as far as subjects are concerned, throughout the series, regardless of whether there is a real structure that roughly matches it. This suggests that the existence of the apparent structure is sufficient for the existence of a physical world. A roughly matching real structure that is tacked on at some point does not seem to do any extra work. So Foster's contention that there is no nonarbitrary criterion for saying that the real structure of spacetime becomes physically relevant at some point in the series is not entirely unreasonable.

An interlocutor might suggest that real spacetime is physically relevant because, together with its occupants, it acts as a truthmaker for truths about the physical world, including truths about appearances. Hence, real spacetime is physically relevant up until (but excluding) W_∞ .

A friend of Foster's abductive argument can reply that the real *structure* of spacetime does not seem to do any truthmaking as far as the apparent world is concerned. Presumably, the same real structure can underlie different apparent structures if psychophysical laws vary. Hence, we don't have a case here for the thesis that the real structure of spacetime is physically relevant.

One might urge that real structure is *causally* relevant because it makes real causation (and hence experience) possible. Causation is a paradigm physical phenomenon, so this makes real spacetime structure physically relevant.

But this idea is open to the complaint that in extreme far regions of the series from W_0 to W_∞ , that is, in worlds where there is almost no real space left, the causal relevance of physical reality seems minimal. What inhabitants of those worlds call "causation" has almost nothing to do with what goes on in real space. And if one suggests that the causal (and hence physical) relevance of real space is limited to the lower regions of the series, one is back to the problem we started with: As far as subjects are concerned, nothing about the physical world changes when real spacetime kicks in causally. The only difference is that what subjects take to be the physical world is implemented in a slightly different way from then on. Hence, as far as our paradigmatic human uses of "the physical world" are concerned, real spacetime structure seems physically irrelevant even if some real causal structure underlies the appearances.

To sum up, the claim that *being physically relevant* is not a contingent property of spacetime structures seems plausible, so premise (18) is acceptable for present purposes. Premises (17) and (18) together establish the principle of representation, according to which the physically relevant spacetime structure is the structure manifest in experience.

Premise (20) says that physically irrelevant posits have no explanatory value in physical ontology. I take it that this is analytic.

Premise (22) says that possible physical structures that have no explanatory value in physical ontology can be left out of physical ontology. This premise is true if the following general abductive principle is true: One can discard possible posits that have no explanatory value in the ontology of the actual world. Conversely, there seems little motivation to uphold (22) without a commitment to such an abductive principle. So denying premise (22) requires the thesis that inference to the best explanation is not valid in metaphysics. I see no good reason to defend that thesis. Inference to the best explanation is as good a principle in metaphysics as anywhere else.

Premise (23) is the claim that spacetime cannot be real without having a real structure. This idea is unlikely to be contentious. There is little more to a spacetime than its structure, so a real spacetime must have a real structure.

Since three out of its five premises are uncontroversial and the other two are plausible, we have reason to think that Foster's abductive argument is sound.

2.3.3 Summary of 2.3

I have reconstructed two arguments by John Foster against spatial realism. Section 2.3.1 focused on Foster's *modal* argument, the gist of which is that the physical realist must be committed to two contradictory principles, the principle of variability, which says that spacetime could have sustained different laws, and the principle of constancy, which says that spacetime could not have sustained different laws. I argued that the principle of variability is entailed by the conjunction of two claims, the independently plausible hypothesis that laws are contingent and the defensible premise that spatiotemporal position is a categorical property. Following Foster, I argued that the principle of constancy is sound because there are no worlds where the primitive geometry of spacetime and the kinematic structure of spacetime come apart.

Section 2.3.2 reconstructed Foster's *abductive* argument, the gist of which is that real spacetime structure is an idle posit which, by inference to the best explanation, can be thrown out of physical ontology. The crux of this argument is the principle of representation, according to which the physically relevant structure of spacetime is the structure manifest in experience. We saw that this principle can be established by considering a series of worlds where the mismatch between the real and the apparent structure of spacetime becomes more and more severe, so that, high up in series, real spacetime structure is physically wholly irrelevant. I supported Foster's contention that there is no nonarbitrary criterion for saying that real structure becomes physically relevant somewhere in the lower regions of the series.

All in all, both of Foster's arguments against spatial realism appear to be eminently defensible. We saw that the modal argument can be resisted in two ways, by denying that spatiotemporal position is a categorical property or by denying that the laws of nature are contingent. In 2.2, we saw that treating spatiotemporal position as a dispositional property exposes the physical realist to the arguments from the Power Regress, so resisting the modal argument by denying the categorical nature of spacetime position is ill-advised. It follows that the physical realist can only resist the modal argument by treating laws as necessary. I argued in 2.3.1 that denying nomic contingency is only warranted if one has independent reasons to prefer physical realism to idealism. Such reasons would allow the physical realist to resist Foster's abductive argument too, since inference to the best explanation cannot lead to idealism if idealism is known to be an inferior theory.

The upshot of the dialectic in 2.3 is that spatial realism is not motivated by the metaphysics of spacetime itself but must be a consequence of independent premises that go beyond the ontology of space. These reasons must ground the verdict that physical realism is in general preferable to idealism.

2.3 Summary of Chapter 2

I have reconstructed three contemporary arguments for idealism. In 2.1 and 2.2, I focused on Howard Robinson's case for idealism, which consists of two parts, the Power Regress (an argument against worlds without categorical features), and a refutation of physical realism from the conjunction of the Power Regress with the thesis that all known physical properties are dispositional. I claimed that a slightly weakened version of Robinson's argument is sound, one that proves the existence of nomologically relevant nonphysical properties. I pointed out, however, that antiphysicalist applications of the Power Regress are open to the objection that spatiotemporal position is a categorical property, hence not all fundamental physical properties are dispositional. I also argued that positing unknown categorical bases that underlie physical dispositions is unwarranted without independent reasons to prefer physical realism to idealism.

Section 2.2 reconstructed two arguments against real space by John Foster. Foster's modal argument is an attempt to show that physical realism entails that spacetime could have and could not have sustained different laws. Foster's abductive argument establishes that real space is physically irrelevant, so, by inference to the best explanation, we have reason to think that it does not exist. I claimed that the modal argument can be resisted by treating spatiotemporal position as a dispositional property (a move that exposes the physical realist to Robinson's case for idealism) or by denying the contingency of laws. Both of Foster's arguments can be resisted by citing strong independent reasons to prefer physical realism to idealism.

All in all, Chapter 2 corroborates the following hypothesis: Physical realism (especially in its antidualist form) is not strongly supported by our current best knowledge about the nature of the physical. That knowledge lends itself easily to an idealist interpretation, and one needs weighty extra premises (such as the hypothesis of unknown categorical bases or the denial of the contingency of laws) to make the scientific image of the material world fully in line with physical realism. If idealism is false, its falsity is not shown by what we know from physics but is established by independent reasons which override the metaphysical principles that lead from physics to idealism in a rather natural way. Hence, the idealist arguments discussed in this chapter go a long way toward a sound defense of idealism.

The goal of the rest of the dissertation is threefold. First, I'd like to offer a new argument for the ideality of spacetime. Chapter 3 will argue that idealism can, and physical realism cannot, solve a tough puzzle in the philosophy of relativity. My hope is that Chapter 3 lends independent support to Foster's abductive argument by showing that idealism has explanatory advantage over physical realism in the context of actual physics.

Second, I'd like to see if there are good independent metaphysical reasons to prefer physical realism to idealism. We saw at various points in the dialectic that the last-resort defense for physical realists is the suggestion that, regardless of issues about fundamental physical properties, physical realism is preferable to idealism for independent reasons. Chapter 4 looks at the two candidate reasons that I know of, the claim that idealism cannot supply truthmakers for physical truths and the claim idealism cannot explain natural laws.

Third, I'd like to start building my own case for idealism. My own case for idealism consists of a defense of the following argument:

- (I) If physical realism is true, then the particles and fields studied by fundamental physics are real.
- (II) If physical objects are real, they have real categorical properties.
- (III) The only categorical property of the particles and fields studied by fundamental physics is spatiotemporal position.
- (IV) Spatiotemporal position is not a real property.
- (V) Therefore, physical realism is false.

The present chapter has argued for (III) in 2.2.1–2.2.3, and it marshalled some evidence for (IV) in the form of Foster's arguments. A further argument for (IV) is coming up in Chapter 3.

Premise (II) can be justified by appeal to the Power Regress. The gist of the Power Regress, both in its original form and in the weaker form defended here, is that there should be nomologically relevant categorical properties, otherwise an unpleasant regress develops. We know from the definition of reality that real physical objects are modally separable from subjects in the sense that they can exist without being disposed to be observed by subjects and hence without being disposed to interact with anything nonphysical (assuming, as I did for the purposes of this dissertation, that anything that is nonphysical is mental).²³ Consequently, it is not of the essence of real physical objects to interact with the nonphysical, therefore real physical objects can be considered in nomic isolation from anything nonphysical. More technically, one may assume that if physical objects are real, then some world that is a physical duplicate of our world has laws that only involve physical properties. If physicalism is true, then that duplicate is the actual world; if dualism is true, then the duplicate in question is the physical part of our world, with the nonphysical components stripped away. (A similar conclusion would follow from the standard physical realist assumption that the world is closed under physics.)

²³ On my official definition of reality, this sort of separability works for inanimate objects only, since a human organism might be real without being modally separable from subjects. (The official definition appeals to the separability of parts to get around this issue.) But since (I)–(V) concerns fundamental particles and fields, which are inanimate, this limitation is not important for the point I'm making.

If physical objects have no real categorical properties, then such physical duplicate worlds will only involve dispositional properties, making our physical system (which is identical to the duplicate in question) open to the complaints from the Power Regress. Hence, in order to avoid the problems that the Power Regress draws attention to, real physical objects need real nondispositional properties. So premise (II) seems relatively straightforward once some version of the Power Regress is accepted.


The only premise unaddressed so far is (I), the thesis that fundamental particles and fields are real if physical realism is true. I'll postpone the discussion of this premise until the conclusion, where I'll argue that denying (I) wrecks contemporary mainstream physicalism, so the denial of (I) is not a good move in a present dialectic, the point of which is to see if contemporary mainstream physicalism can find a solid metaphysical foundation in real matter.

3 Real spacetime as excess structure


This chapter argues against real spacetime. In 3.1, I use a toy example to show how idealism can solve underdetermination issues about spatiotemporal structure and how it can motivate the view that real spacetime is a posit that makes one committed to metaphysically useless excess structure. In 3.2–3.4, I deploy this strategy in the context of a puzzle about relativity. Specifically, I'll argue that the puzzle of the conventionality of distant simultaneity is easily solved by idealism along the lines sketched in 3.1, but it has no straightforward and uncontroversial physical realist solution. I'll also indicate that idealism promises to deliver further similar results, hence, by inference to the best explanation, we may conclude that spacetime is not real but ideal.

3.1 Albert the ant

Albert is an extensionless ant who lives on a line. He can see and touch things on the line, he cannot see or touch anything above or below the line, and he cannot perceive changes in his vertical position. He has a meter stick that he drags around with himself, measuring lengths and distances:



At certain places, there are humps on the line where Albert lives:



Albert, being confined to his own one-dimensional perspective and unable to perceive vertical changes in his position, cannot be directly aware of these humps. (It is consistent with this assumption that Albert can entertain the *idea* of unperceivable humps. He just happens to be unable to perceive the humps. As we'll see, the question is precisely whether he has any reason to think about his world in terms of unperceivable humps.)


Suppose, further, that the physical laws in Albert's world are such that whenever Albert reaches a hump, his meter stick starts oozing upward, its far end disappearing from Albert's sight:

$t = 1$



Then, as Albert mounts the hump, the stick forms a tangent, touching the hump at a single point only:

$t = 2$



As Albert travels along the hump, the stick continues to form a tangent to it:


$t = 3$



$t = 4$



$t = 5$




Finally, when Albert is near the far end of a hump, the stick oozes down on the far side, regaining its original shape and orientation:

$t = 6$




$t = 7$




Now let's see what this process looks like from Albert's perspective. Since Albert cannot experience elevation and curvature (he is unaware of the second, vertical spatial dimension), he has the experience as of moving in a straight line even when he is on the hump. But since parts of his meter stick disappear into the spatial dimension that Albert cannot perceive, Albert will none the less notice some change, namely, he'll see the meter stick first shrink to a point, then grow back to its original size:

Albert's perception of the meter stick:



As far as Albert's experience as a one-dimensional surveyor is concerned, the humps in his world appear to be areas where meter stick shrinks to a point, then gradually regains its original size.

Let's call the world we've discussed so far "world A." And let's consider a slightly different world *B*, which is much like *A*. *B* includes Albert, his meter stick, and a line where Albert lives. The only difference is that at the places where *A* contains humps, *B* contains *contraction zones*:



A contraction zone is a segment of the line where meter sticks get distorted because of a force field. When Albert pushes a meter stick into a contraction zone, the stick starts shrinking, hovering at the border of the zone:

$t = 1$



When the stick has shrunk to a point, it finally enters the contraction zone, and it stays contracted to a point throughout the zone:

$t = 2$



$t = 3$



$t = 4$




$t = 5$




Finally, when Albert is about to leave the contraction zone, the stick starts growing again, regaining its original size the moment Albert leaves the zone:

$t = 6$




$t = 7$




Let's compare Albert's experiences in the two worlds:

Albert's perception of the meter stick:

World A (containing humps)



World B (containing contraction zones)



As far as Albert's visual experience of the meter stick is concerned, he could just as well be in world *A* as in world *B*. Now consider two cases:

Case 1

In Albert's world, space is real.

Case 2

In Albert's world, space is ideal.

If Case 1 obtains, then it is to be expected that world *A* and world *B* are two different possible worlds, both of which could be Albert's world, as far as he can tell on the basis of experience. There seems to be a clear difference between worlds that contain humps like *A* and worlds that do not, so *A* and *B* must be different worlds. And since Albert cannot tell the difference between *A* and *B* on the basis of experience, he could be in either of the two worlds for all he knows, no matter which world he is actually in.

If Case 2 obtains, the humps in the first series of pictures do not correspond to anything real. Moreover, worlds *A* and *B* contain the same concrete facts in this case, since they contain experiences of the same phenomenal character. Those features that are supposed to distinguish *A* from *B* (e.g. that some line segments have nonzero curvature in *A* and zero curvature in *B*) are not observable by Albert, and since physical facts entail observability in idealism, there will be no facts of the matter about the features in question in idealism. It follows that in Case 2, world *A* and world *B* are the same world. The pictures depicting *A* and *B* (and the corresponding mathematical models that Albert

might use to predict future experience) are different representations of the same underlying (phenomenal) reality. Both representations correctly predict future experience, and neither is an actual *picture* of reality. The humps and contraction zones in the pictures are mere mental aids. Albert is free to choose the representation he likes best to formulate the laws of nature.

It follows that idealism can get Albert out of an epistemological quandary. If Albert's space is ideal, then Albert will know what the real structure of his world is (since, in that case, real structure is the structure of experience). In contrast, if Albert's space is real, Albert is faced with an underdetermination problem that he can only solve by appeal to epistemic luck or by reverting to pragmatic considerations (e.g. by choosing the model that seems more simple or convenient or aesthetically pleasing).

This toy example is meant to illustrate two things. First, it is meant to show how an idealist can turn puzzles about empirically equivalent theories into arguments for idealism. Whenever one is faced with a choice between two empirically equivalent physical theories, the idealist might have a chance to argue that, on her view, the alternatives in questions are just ways to represent the same facts about experience, and, as such, they do not constitute an epistemological dilemma any more than the choice between orthogonal and spherical coordinates constitutes an epistemological dilemma.

Second, by the same token, the idealist can argue that physical realism makes one committed to excess structure that creates epistemological trouble without explaining anything. For example, realism about Albert's space makes one committed to real facts about the curvature of Albert's line. This excess structure is necessary to secure realism, but it introduces an inscrutable component into Albert's world without doing any other work.

To sum up, if Albert's case can be treated as a paradigm, then idealists can argue that physical realism is unfavourable on two counts, first because it makes the structure of reality partly inscrutable, and second because physical realism clutters physical ontology with useless excess structure.

The toy example also indicates how physical realists can resist this type of reasoning. The best counter-move for physical realists is to insist that there is no physical difference between the rival models in question. In Albert's case, the physical realist should argue that there is no difference between forces that distort measuring instruments and variations in the curvature of space.


Needless to say, none of this is meant to be a real *argument* for idealism. The goal of the following sections is to explore a puzzle about spacetime from the philosophy of relativity and to argue that it can motivate a similar case for idealism in the context of actual physics.

3.2 The conventionality of simultaneity


The following sections explore a strange and recalcitrant puzzle from the philosophy of relativity. The puzzle is called “the conventionality of (distant) simultaneity.” Briefly, the problem is that special relativity seems compatible with observation-transcendent hypotheses about variations in the speed of light. Specifically, it seems compatible with the claim that the speed of light varies in different directions but the discrepancies even out. The present section explains the details of this puzzle. Section 3.3 surveys physical realist solutions. I’ll argue that the puzzle has no straightforward and uncontroversial physical realist solution. Section 3.4 outlines the idealist solution. I’ll argue that it is much more satisfying than physical realist suggestions, hence the conventionality of simultaneity can motivate the view that spacetime is ideal.

3.2.1 Measuring the one-way velocity of light

Suppose that the space of an observer is single line, with the observer standing at point P and two mirrors positioned on both sides at unit distance:




Now let’s add a time axis to generate a two-dimensional spacetime:




A line in this coordinate system represents both the spatial trajectory of an object and the times at which it visited a particular place. For example, the spatiotemporal trajectory of the observer stationed at P is represented by the t axis. Points in the diagram represent momentary events at a point in space. For example, each point on the t axis “is” the place P , but at different times. We can think of these points as momentary events at P .

Suppose we choose our unit distance to be c meters, so that the speed of light (on the standard conception) is 1 unit distance per second. With this convention in place, the spacetime trajectory of two light rays, emitted from P at $t = 0$ toward the mirrors on each side, will look like this:



Generally, the path of light rays traveling to or from P will be represented as lines tilted at 45 degrees to the x axis:



Rays 1 and 2 travel towards P from a distant source, reaching P at $t = 0$. Rays 3 and 4 travel from P toward some distant target, leaving P at $t = 0$. (Perhaps Rays 3 and 4 are Rays 1 and 2, crossing at the origin or reflected back toward their respective sources.)

The shaded area enclosed by the four rays is called the *lightcone* of the origin (the origin itself being the momentary event occurring at P at $t = 0$). Every point in the spacetime diagram has its own lightcone. A lightcone is the collection of those point-events that can be causally connected to the point whose lightcone they constitute. In the previous picture, Rays 1 and 2 represent the “outermost” points that can causally influence the point-event at the origin, and Rays 3 and 4 represent the “outermost” points that can be causally influenced by the point-event at the origin. If a point from outside the shaded area were causally connected to the origin, then the causal influence would have to propagate faster than light, which is impossible in special relativity.


An important premise of relativity is that the velocity of light is the same for all observers, no matter their relative motion (Einstein 1923: 41, 2005: 78).²⁴ This idea marks a radical departure from Newtonian physics, where the apparent velocity of objects varies with the velocity of the observer. If Superman is flying at 599 mph, chasing a plane that is flying at 600 mph from the standpoint of a stationary observer, then, in Superman’s frame of reference, the plane will crawl at 1 mile per hour. This logic applies to relative velocity across the board in classical physics. But in relativity, a light beam traveling at c from the point of view of the stationary observer will also travel at c from the point of view of Superman (and everyone else). The invariance of the speed of light results in well-known oddities like time dilation and length contraction. It also entails that all observers will agree on the shape and position of the light cones, hence the basic structure of the causal order will be objective.

Now consider the following problem. The observer at P wants to measure the speed of light. A seemingly straightforward way to do this is to send out a light beam from P at $t = 0$ toward one of the mirrors, and record the time, T , when the beam, reflected back from the mirror, returns to P . Dividing 2 by T then gives us the velocity of light, since the light beam traveled 2 units of space in T units of time. Given that we chose our metric unit to be c meters, we expect $2 / T$ to equal 1, that is, we expect the round-trip to take 2 seconds.

Notice, however, that if the observer is confined to P and this experiment is her only means of measuring the speed of light, then she cannot be sure that light travels with the same velocity to the mirror and back. The observation that it takes light 2 seconds to make the round-trip is compatible with the hypothesis that, say, the speed of light toward the mirror is $0.625c$ and its speed back is $2.5c$. In such a case, it will take the beam 1.6 seconds to reach the mirror and 0.4 seconds to come back, so the round-trip will take 2 seconds, just as in the nondeviant case where the speed is the same in all directions.

²⁴ The second basic principle of relativity theory is that observation and experiment leads to the same formulation of the laws of nature in all frames of reference. This principle will not be important in the present dialectic.

If, in fact, light travels at $0.625c$ to the right and at $2.5c$ to the left, then the spacetime diagram of the experiment will look like this:



A nondenumerable infinity of such deviant scenarios are compatible with the observed fact that it takes light 2 seconds to bounce back from the mirror. In principle, each point on the dotted line could be the event of the beam's hitting the mirror. Hence, a lonely observer who is confined to P and can only use mirrors (or other distant objects that “signal back”) cannot measure the one-way velocity of light. She can only tell that the average round-trip velocity is c m/s (in vacuum, at any rate—let's ignore this).

What is interesting about this problem is that the deviant scenarios sound trivially easy to disconfirm once we discard the unreasonable hypothesis that we have a single observer who is confined to P . But, surprisingly, the problem at hand has no easy solution. It is one of the most recalcitrant puzzles in the philosophy of relativity, one that continues to generate controversy, more than a hundred years since it was discovered by Einstein.

To see why the one-way velocity is not easy to measure, suppose we discard the unreasonable assumption that the observer is confined to P and can only use mirrors. Suppose we have two clocks instead, one at P and one at $x = 1$. The observer at P sends out a light beam at $t = 0$, and the clock at $x = 1$ (or a second observer who has access to the distant clock) records the time, T , when the beam arrived. Dividing 1 by T , we have the one-way velocity of light.

To make this experiment work, the two clocks must be synchronized: we must have good grounds for saying that the second observer's clock read “0” (or some other known value) when the light beam left P . Otherwise the readout on the second clock won't tell us how long it took light to traverse unit distance, because we will have no grounds for saying that when we released the light beam at $t = 0$ from P , the second clock read “0” (or some other known value).

Let clock A be at P and clock B at $x = 1$. We connect the two clocks by a wire. Before launching the light beam, we send an electric signal from A to B , setting B to a known value that guarantees that the two clocks both read “0” when the light beam takes off from P . Then the readout on B will give us the time it took light to traverse unit distance.

To guarantee that clock B reads “0” when the light beam leaves P at $t = 0$ (on A ’s time), we have to know exactly how long it takes for the synchronizer signal, sent from A to B , to reach clock B . Instantaneous signaling is impossible in relativity (because nothing can travel faster than light), so the synchronizer signal will take some time to reach B . Suppose it takes S seconds. Then we have to send the synchronizer signal at $t = -S$ (on A ’s time), telling B to set itself to 0, if we want the two clocks to read “0” at the same moment.

The trouble is that if we use electric signaling, then the magnitude of S will depend on the speed of light, since the speed of light is involved in electric phenomena.²⁵ So we cannot use this method to synchronize two distant clocks to measure the velocity of light from A to B , because we must already know the one-way speed of light to synchronize the two clocks. We’ll have to know the constant we want to measure prior to measuring it.

Note, further, that if we arbitrarily choose an admissible value for the one-way velocity, then our measurements will not detect an anomaly even if there is one. For simplicity, let’s assume that electric signals travel at the speed of light, and suppose that in reality, light travels at $0.625c$ to the right and at $2.5c$ to the left. Our observer assumes uniform velocity, so she believes that it will take light 1 second to travel from clock A to clock B . At $t = -1$ (on A ’s time), the observer sends out a signal that tells clock B to set itself to 0. The observer reckons that the signal will hit B the moment A reads 0, so both clocks will read “0” when the light beam leaves A .

Let’s see what happens in reality. In reality, the synchronizer signal travels at $0.625c$ from A to B , so it reaches B in 1.6 seconds. Since it is sent at $t = -1$ (on A ’s time), B will read “0” when A reads “0.6.”

Meanwhile, the light beam whose one-way velocity is being measured is emitted when A reads “0.” This beam also takes 1.6 seconds to arrive to B , so it hits B when A reads “1.6.” But B is running 0.6 seconds late in comparison to A , so B will read “1” when the light beam hits it. Hence, when the observer, or her buddy, checks the second clock, she’ll think that it took light 1 second to travel unit distance, and she’ll figure that the one-way velocity of light is c m/sec. The difference between this value and the real one will be undetectable, because an incorrect assumption about the result is already built into the synchronization method and the mistakes even out.

²⁵ See Torretti (1979: 303) for a note on how deviances in the one-way velocity affect the laws of electrodynamics.

It follows that remotely synchronized clocks cannot help us determine the one-way velocity of light.²⁶

Suppose we do the following instead: We set two clocks side by side at P , and we synchronize them locally. Then we transport clock B to $x = 1$, and, when A reads “0,” we send a light beam to B . Since the two clocks are already synchronized, the readout on B will tell us the one-way velocity.

The problem with this idea is that special relativity predicts that moving clocks slow down. Hence, when we move clock B from P to the measurement post, clock B will show a different time than A upon arrival. The slowdown depends on the velocity of light, so we’ll have to make an assumption about the one-way velocity to correct for the slowdown. And if we make an incorrect assumption, our setup will be mucked up in such a way that we’ll end up measuring c for the one-way velocity even if it isn’t. (See Appendix.)

To sum up, there is no easy way to find out whether light travels with the same speed in all directions. In his seminal paper on special relativity, Einstein claimed that we simply *stipulate* that the one-way speeds are uniform:

If at the point A of space there is a clock, an observer at A can determine the time values of events in the immediate proximity of A by finding the positions of the hands which are simultaneous with these events. If there is at the point B of space another clock in all respects resembling the one at A , it is possible for an observer at B to determine the time values of events in the immediate neighbourhood of B . But it is not possible without further assumption to compare, in respect of time, an event at A with an event at B . We have not defined a common “time” for A and B , for the latter cannot be defined at all unless we establish by definition that the “time” required by light to travel from A to B equals the “time” it requires to travel from B to A . (Einstein 1923: 39–40)

The thesis that the one-way velocity of light is a matter of convention entered the philosophy of science through Hans Reichenbach’s (1957 [1927]) landmark interpretation of relativity theory, and it entered English-language philosophy of science through Adolf Grünbaum’s (1955, 1963) endorsement of Reichenbach’s claims.²⁷ The thesis is usually known as “the conventionality of (distant) simultaneity.” It has been a subject of debate ever since.


²⁶ A similar masking effect occurs if we try to use two-way synchronization, letting clock B tell clock A when it received the synchronizing signal. As long as the round-trip average velocity is c , discrepancies in the one-way velocities will be undetectable.—Notice that, properly speaking, the length of time units must also be synchronized in order for two clocks to beat the same time (see Einstein 1993 and Jammer 2006: 124–5). Factoring this in would also involve clock B signaling back to A .

²⁷ See Reichenbach (1957: §19, 123–7) and Grünbaum (1955, 1963: 341–68) for two classic expositions, and see Jammer (2006: Chs. 9–10) on the history of the puzzle.

3.2.2 Consequences of deviance

In the last hundred years or so, various deviant models of special relativity have been proposed, models where the speed of light is not uniform. One of the most famous is John A. Winnie's (1970), which applies to the kind of 2D spacetimes that we are considering. There are more general models as well, all of which are claimed to generate the same observable predictions as models where the speed of light is uniform.²⁸ The goal of this section is to present some basic non-technical consequences of the assumption that the one-way speed of light is not uniform.

First of all, to see why the problem at hand is called “the conventionality of simultaneity,” consider three possibilities about the one-way speed of light:




The three scenarios agree that the average round-trip velocity of light is c , i.e., that it takes 2 seconds for the beam to come back. (Remember that the diagrams use c meters as the unit of distance.) But the three scenarios disagree on the one-way velocity.

Scenario 2 represents the standard (nondeviant) assumption that it takes light the same amount of time to reach the mirror and come back. Graphically, this means that the moment when the beam is reflected back is halfway on the dotted line, coinciding with $t = 1$ on the P -observer's clock.

Scenarios 1 and 3 are deviant. In Scenario 1, light goes much faster than c to the right and much slower than c to the left, so that the moment of reflection occurs comparatively early (earlier than $t = 1$) on P 's clock. In Scenario 3, it is the other way around: light travels slower than c to the right and faster than c to the left, so the moment of reflection occurs later than $t = 1$ on P 's clock.

²⁸ For a sample of nonstandard models of special relativity where the one-way speed of light is not uniform, see Scott-Iversen (1944), Edwards (1963), Winnie (1970), Abraham (1986), Minguzzi (2002), and Ben-Yami (ms).

The three scenarios, therefore, disagree not only on the one-way velocities, they also disagree on which moment in P 's local history (i.e. which point on the t axis) is simultaneous with the light beam's hitting the mirror:




If Scenario 1 obtains, then the moment when the light beam hits the mirror (the moment of reflection) is simultaneous with E_1 (which is a point-event in the local history of the observer at P). If Scenario 2 obtains, the moment of reflection is simultaneous with E_2 (which is a different point-event in the local history of the observer at P). If Scenario 3 obtains, the moment of reflection is simultaneous with E_3 (which is, again a different event).

Now suppose that, as Einstein implied and as Reichenbach and Grünbaum explicitly argued, it is a matter of convention whether we assume the one-way speeds of light to be uniform or deviant. Then it is a matter of convention which one of the these three scenarios obtains. And, consequently, it is a matter of convention which one of E_1 – E_3 is simultaneous with the moment of reflection. The observer at P can simply stipulate which moment in her local history is the moment when the light beam hits the mirror at a faraway point. (Within certain limits, of course. But as the diagram shows, these limits are considerably wide.)

This is why the problem at hand is called “the conventionality of (distant) simultaneity.” If the one-way velocity of light is conventional, then it is a matter of convention when distant events occur according to my local time. And, conversely, if this is a matter of convention, then the one-way velocity of light is also a matter of convention.


It seems profoundly unintuitive to think that there is no fact of the matter about these relations and that it is a matter of definition (or convenience, or theoretical elegance) whether light moves with the same speed in all directions. The existence of deviant models has therefore spurred various responses that seek to undermine or downplay their importance. Before looking into these responses, let's consider a further important consequence of deviance.

Variations in the one-way velocities of light will also affect the shape of light cones. The light cone of the origin (that is, the collection of points that can be causally connected to the origin) looks like this in the standard case when the one-way velocity of light is uniform:




Uniform one-way velocity (Scenario 2)

If the one-way velocities are not uniform, the light cone will take a different shape. For example, if light goes faster to the left than to the right (with the two velocities adding up to the theoretically required round-trip speed of c), then the light cone of the origin will look like this:



Deviant case: different one-way velocities
(Scenario 3: Light travels at $\sim 0.625c$ to the right
and at $\sim 2.5c$ to the left)

Hanoch Ben-Yami (ms.) constructed a deviant model of special relativity where the difference from the standard case is even more striking, and the light cone of the origin looks like this:



The conventionality of simultaneity entails that the shape of light cones is also conventional. (In fact, they aren't necessarily cones any longer.)

The rest of this chapter falls into two parts. In 3.3, I look at various physical realist responses to the puzzle of conventionality and I'll argue that none of them are straightforward and uncontroversial. In 3.4, I'll outline the idealist solution, which, I'll claim, is straightforward and uncontroversial, giving a neat and satisfying explanation of the puzzle. I'll conclude that the conventionality of simultaneity can motivate the view that spacetime is ideal.


3.3 Physical realist solutions

This section surveys possible physical realist responses to the conventionality of distant simultaneity. Not all of the theories below entail physical realism, but, as far as I can see, they exhaust the spectrum of dialectical strategies that are compatible with physical realism.

3.3.1 Experimental tests

There is a venerable history of suggestions about experiments to measure the one-way speed of light. As Max Jammer's (2006: Ch.12) extremely detailed survey indicates, an assessment of all these (apparently, failed) attempts is more of a physicist's job than a philosopher's. But it is useful to look at an example in order to appreciate the complexities involved.

Tim Maudlin (2012: 121–4) asserts that the one-way velocity of light can be (and has been) measured (and has been found uniform) using the following contraption:



Two disks, each with a slender slit on it, is fixed to the ends of a rod that can be rotated with a speed that we can control. A light beam is flashed from one side at the first disc. If the rod is rotating with the right speed, then the light beam getting through the first slit will reach the second the moment the second slit is at the position where the first was when light beam entered the area between the two discs, so light will get through the contraption and leave a mark on the screen on the other side. Otherwise, if the rod is set to rotate at a speed so that the beam entering through the first slit will fail to hit the slit on the second disc, no mark will appear. If we know the distance between the two discs, the speed at which they rotate, and the relative position of the two slits, then we can calculate what speed a light beam must have in order to leave a mark on the screen. Next, all we have to do is (i) set the rod to rotate at a speed such that light traveling at c will get through the second slit and (ii) turn the whole contraption in various directions. If light leaves a mark on the screen in all directions, then the one-way velocity of light is c in all directions.

As Max Jammer (2006: 224) explains, this test (which is a variation on an experiment performed by H. L. Fizeau in 1849) presupposes that the rod that moves the two discs is perfectly rigid. But that principle is false in relativity:

The denial of actions at a distance, which were admitted in Newtonian physics, also denies the existence of perfectly rigid rods. The definition of perfect rigidity, as preservation of geometrical shape, implies that if one end of such a rod would be set into motion, the other end would instantaneously start moving as well, so that the rod could serve as a generator of actions at a distance. If perfectly rigid rods existed, the problem of distant synchronization could, of course, be solved simply by coupling clock mechanisms by such rods. In fact, numerous synchronization procedures, proposed to disprove the conventionality thesis, are but more or less disguised versions of such coupling proposals. (Jammer 2006: 222)

To measure the one-way velocity the way Maudlin suggests, one must suppose that if a motor starts turning one end of the rod, giving it a certain angular velocity, then the far end of the rod will start moving at the same instant, so that the two slits will stay in the same relative position during rotation that they occupied at rest. Knowing the relative position then lets us calculate the exact speed that is required for light to escape through the second slit. However, if the rod that moves the two discs is not perfectly rigid, then the discs do not start spinning at the same time, so the relative position of the slits changes. One must factor in these changes in relative position to know what speed light needs to escape through the second slit once the contraption has started operating. But material forces (including the force that moves the rod) propagate at speeds that ultimately depend on the speed of light, hence in order to calculate the time lag between the moments the two discs start rotating, one must know the one-way velocity of light. So this contraption does not allow us to calculate the one-way velocity without making assumptions about it and thereby begging the experimental question.

Numerous other experimental setups have been proposed, a thorough survey of which would require a separate book filled with technical discussions in physics.²⁹ As the surveys by Jammer (2006) and Janis (2010) indicate, there is no clear sign of an emerging consensus. In the words of John Norton,

The quest for the one way velocity of light is beginning to look like the quest for a perpetual motion machine, for in both cases the fruitlessness of the quest can be demonstrated by quite elementary means. [...] [I]t reduces to the simple question of whether special

²⁹ Here is a sample of proposed methods and conventionalist rejoinders, from Jammer (2006). Experimental setups using mechanical devices like the one mentioned by Maudlin have been proposed by Eagle (1938), Feenberg (1974), and Jackson and Pargetter (1977), among others. Øhrstrøm (1980), Torretti (1979), and Townsend (1980) argue that these methods are not convention-free. Stolakis (1986) discusses the possibility of measuring the one-way speed of light using refraction, but Clifton (1989) claims that this method isn't convention-free either. A number of authors (e.g. Burniston Brown 1967, Froome and Essen 1969: 3) say that the method that Ole Roemer used in the 17th century to measure the speed of light can give the one-way speed. Karlov (1970), Babovic et al. (1991), Shea (1998), and Jammer (2006: 230) argue that this method tacitly relies on the assumption of uniform speed. Essen and Gordon-Smith (1948), Bol (1950), and Liebowitz (1956) have proposed measurements using microwave resonance, but Grünbaum (1956) and Salmon (1977) argue that these methods are not convention-free. Ruderfer (1960) and Møller (1962) proposed to test the one-way velocity using laser or maser, but Sjödin (1979) and Podlaha (1980) claim to have showed that these methods cannot measure the one-way velocity. (Note that Ruebenbauer (1980) claimed to have found evidence for uniform one-way velocity using such a method, while Marinov (1974) claimed to have found evidence for non-uniform velocity.) Ellis and Bowman (1967) and Prokhovnik (1973) proposed transporting clocks at very slow speeds to get around the time dilation problem and create perfectly synchronized distant clocks. Grünbaum, Salmon, van Fraassen and Janis (1969), Friedman (1977), and Winnie (1970: 223–9), among others, have argued that this method either cannot detect deviances or it solves the problem by *fiat*, swapping one kind of convention for another. This list of proposed tests and conventionalist rejoinders is far from complete and most of the debates belong to physics proper. Nonetheless, this brief survey indicates that the testability of the one-way velocity is contentious at best.

relativity can be formulated in certain ' ϵ -Lorentz coordinate systems' rather than just the 'Lorentz coordinate systems' used in the familiar standard formulation of the theory. That this is possible has been known in principle since as early as 1913, when Einstein introduced techniques which would enable special relativity to be formulated in arbitrary spacetime coordinate systems. The quest for the 'true' value of ϵ [= the ratio between c and twice the 'real' velocity] and the (coordinate dependent) one way velocity of light which it determines, is as fruitless as the quest for the subset of 'true' coordinate systems in which special relativity can be formulated. For this task, all coordinate systems are equally viable. (Norton 1986: 119)

Given the long history of unsuccessful attempts to devise empirical tests for the one-way velocity, and given the opinion of notable philosophers of science that such tests can't exist, the thesis that the one-way velocity of light is not an empirical issue seems defensible. I'll assume, for the purposes of this dialectic, that the one-way velocity is either observation-transcendent, or it does not exist. The question for the rest of 3.3 is to find out how physical realists can cope with this situation.

3.3.2 Appeals to symmetry

Michael Friedman (1984: 156–76) thinks that the hypothesis of non-uniform one-way speed commits one to a physically irrelevant asymmetry that inference to the best explanation gets rid of easily. To see Friedman's idea in a bit more detail, suppose that light goes faster to the right than to the left in a 2D spacetime, so that light cones are tilted to the right. Friedman argues that in such a world, a spatial direction is singled out as somehow unique, but this asymmetry (or, in technical language, anisotropy) has no physical significance whatsoever and does not do any explanatory work. Hence, those who worry about the one-way speed worry about a possible asymmetry that adds nothing substantial to our conception of the physical world except for making it more complex. Hence, the puzzle about the conventionality of simultaneity can be solved by a simple abductive inference.

Friedman's argument has the following basic structure:

- (Exp) Deviant models of relativity do not explain anything.
- (Abd) Hypotheses that don't explain anything can be disregarded.
- $\therefore$ Hence, deviant models of relativity can be disregarded.

Now if we move in the context of *physics*, (Exp) is certainly true. Physics is concerned with prediction, and since deviant models are empirically equivalent to the standard one, variations in the one-way speed of light can have no explanatory role in physics.

However, if (Exp) is interpreted in this way, then it is a pragmatic point about physics and not a point about the metaphysics of spacetime. If a deviant model of relativity suddenly turned out to be more workable for certain calculations, (Exp) would cease to be true and physicists would be justified to switch to that model. But, presumably, Friedman does not want to say that the structure of spacetime would change as a result. Hence, the fact that variations in the one-way speed of light have no explanatory relevance in physics does not entail that the one-way speed of light is uniform.

In response, the Friedmannian can insist that her point is not pragmatic but ontological. We have reason to think that there are no arbitrary asymmetries in nature. Nature seems highly uniform, and physical laws are typically invariant with respect to time and place. Hence, we have reason to think that no arbitrary asymmetry affects the one-way speed of light, and, as a result, the argument from simplicity does have ontological consequences.

The gist of this reply is that deviant models of relativity contradict our best inductive principles. The Friedmannian, in effect, says that the deviant models represent sceptical scenarios that can be disregarded even if they are empirically equivalent to the normal case. The hypothesis that we are brains in vats is empirically equivalent to the hypothesis that we are normal human beings, but it does not follow that BIV worlds must be taken into account when we look for the most plausible ontology of the actual world. Or, to take an example from physics, consider epicycles in Ptolemaian cosmology. Ptolemaian cosmology can be made to be empirically equivalent with the heliocentric theory of Kepler at the price of adding more and more complexity. But it does not follow that Ptolemaian worlds are relevant when we consider the structure of the actual physical world. Similarly, worlds where the one-way speed of light is not uniform can be disregarded when we consider the ontology of the actual world, even if the former are empirically equivalent to our preferred standard theory.

In response, the anti-Friedmannian can say, first, that the way natural laws are distorted in deviant models is different from the way reality differs from our preferred simple theories in BIV worlds and in worlds where Ptolemaian cosmology is true. Deviant models of special relativity differ from our preferred simple theory in terms of the observation-transcendent value of a constant that is found in both standard and deviant models, playing the same theoretical role. So the difference between worlds where the one-way speed of light is uniform and worlds where it is not uniform is much smaller than the difference between BIV worlds and normal worlds or the difference between heliocentric worlds and observationally equivalent Ptolemaic ones. As a result, the fact that BIV worlds and Ptolemaic worlds can be safely disregarded when we investigate actual physical ontology does not entail that worlds where the one-way speed is not uniform can also be disregarded.

Moreover, the anti-Friedmannian can point out that even if the charge of skepticism is sound, it only secures the conclusion that the one-way speed of light is *actually* uniform. The Friedmannian admits that deviant worlds are metaphysically *possible*. But this fact is quite puzzling in itself, even if our actual hypotheses about the one-way speed are correct because of a piece of epistemic luck. Intuitively, it would much better to solve the puzzle by showing either that the deviant models are not consistent with actual laws or by showing that the difference between deviant and non-deviant models do not correspond to a real difference. Friedman's solution entails that, possibly anthropocentric appeals to symmetry aside, there is no reason why nature should care about the one-way speed of light, which seems to be a peculiar feature of physical reality. Friedman's solution only dispels the puzzle at an epistemic level (as a challenge to our knowledge of actuality), hence it is inferior to those solutions that dispel the puzzle at a metaphysical level as well.³⁰

3.3.3 Mathematical arguments

Some philosophers and physicist have proposed mathematical proofs against deviant models. The most famous such proof is by David Malament (1977), who showed that the hypothesis of uniform one-way velocity is the only one that can satisfy certain symmetry principles. Briefly, and disregarding technical details, Malament proved the following: If we have an unaccelerated path UP in spacetime and we perform a reflection, rotation, or translation of the whole content of spacetime in such a way that UP is mapped onto itself, then the operation will only leave the rest of the spacetime intact if the one-way speed of light is uniform.³¹

The intuitive content of this proof can be illuminated by a simple example. Suppose that, in Case 1, a stone flies from spatial point A to spatial point B in T seconds. In Case 2, the same process unfolds in $2T$ seconds, with no change in the stone's trajectory except for the multiplication of the time coordinates by 2. Intuitively, if this transformation is applied to the whole content of spacetime (that is, to all point-events), then all other physical processes should look the same (apart from the time dilation). Generally, one can expect the content of space to remain invariant under transformations that map a specific physical process onto itself. Malament proves that this requirement entails that the speed of light is the same in all directions.

³⁰ For a more technical exploration of a similar point, see the exchange between Ohanian (2004) and Macdonald (2005). Ohanian argues that the kinematic laws will be much more complex in deviant models, but Macdonald points out that this complexity in itself does not entail that no deviant model is true; it only entails that we naturally prefer the convention that laws are represented in a relatively simple way. See Salmon (1977: 273) for a similar point about the shape of conservation laws in deviant models.

³¹ For summaries of the argument, see Jammer (2006: 257) and Ben-Yami (2006: 464–6).

There are two problems with Malament's proof. As it was pointed out by Ben-Yami (2006: 466–71) and Sarkar and Stachel (1999), among others, the proof must allow for reflections that reverse the direction of time, in other words, Malament's admissible transformations include scenarios when history unfolds backward with respect to the original path *UP*. It is far from intuitive that such symmetries must obtain for actual systems, so Malament's proof can be criticized on the grounds that its premises are not self-evident.

Second, the requirement that spacetime be symmetrical in a way that happens to entail uniform one-way velocity appears to beg the question against the conventionalist (Grünbaum 2010, Janis 1983: 107–9). This stipulation carries no more weight in itself than Friedman-style appeals to simplicity.³²

3.3.4 Eternalism

Some philosophers believe that the conventionality of simultaneity can be solved by adopting eternalism. Here's a brief statement of the idea:

[T]he message, which the vicious circle [about the measurement of the one-way velocity] has been trying to convey to us is truly amazing—reality is not a three-dimensional world, because if it were, what exists would depend on our choice and would also be a matter of convention. But what exists will not be a matter of convention if reality is a four-dimensional world with one temporal and three spatial dimensions. [...] The vicious circle in determining the one-way velocity of light finds a natural explanation—light does not move at all since it is a forever-given worldline in spacetime. We arrived at the vicious circle because we asked an incorrect question about the real magnitude of the one-way velocity of light, whereas light does not propagate in spacetime (in the absolute reality according to Minkowski), and therefore does not possess such a property as velocity. (Petkov 2009: 159–60)³³

This argument has the following shape:

- (E1) If eternalism is true, then there is no movement (since the whole of spacetime exists eternally as a 'block').
- (E2) If there is no movement, then there are no facts of the matter about the movement of light.
- (E3) Therefore, if eternalism is true, then nothing corresponds in reality to the (purported) one-way velocity of light.

³² See Janis (2010: section 5) about other mathematical proofs. Feinberg (1967) argues against conventionality from tachyons (rejoinder: Pirani 1970), and Zangari (1994) uses considerations about spin-1/2 particles (rejoinder: Gunn and Vetharaniam 1995).

³³ For earlier versions of this argument, see Wingard (1972) and Petkov (1989). Note that this argument is different from Putnam's (1967) contention that presentism is false because of the *relativity* of simultaneity, i.e. because two observers can disagree on temporal relations.

If this reasoning is sound, it constitutes its own *reductio*, because it entails an absurd consequence. If (E1)–(E3) is sound, then the following argument must also be valid:

- (E1) If eternalism is true, then there is no movement (since the whole of spacetime exists eternally as a ‘block’).
- (E2) If there is no movement, then there are no facts of the matter about the movement of light.
- (E4) Therefore, if eternalism is true, there are no facts of the matter about the speed of light in general; specifically, there are no facts of the matter about the average round-trip speed of light.

But this conclusion is obviously absurd, since the average round-trip speed of light is a physical constant with a known value. Everyone in this debate agrees that the average round-trip speed is c (in vacuum, at any rate). If the argument from eternalism proves that there are no facts of the matter about this, then the argument proves too much.

In response, a friend of the eternalist solution must argue that the concept of one-way speed is empty in eternalist universes, but the concept of *round-trip* speed is not. But it is hard to see any non-ad hoc way to establish this. Petkov’s point is that there is no movement in an eternalist universe, hence light does not move at all in any specific direction. But if light does not move at all in any specific direction, then, *a fortiori*, it does not move in any specific direction *and back* either. Hence, the eternalist cannot block (E4), and (E4) constitutes a *reductio* of the eternalist argument.

3.3.5 Gauge freedom

Some philosophers think that the conventionality of simultaneity is due to the fact that the one-way velocity of light is not a real feature of the physical world even though the physical world itself is real. Adolf Grünbaum is one of the champions of this view:

[It is a mistake to think that] Einstein’s repudiation of Newton’s absolute simultaneity rests on a mere *epistemic* limitation on the ascertainability of the existence of relations of absolute simultaneity. To be sure, human operations of measurement are indispensable for *discovering* or *knowing* the physical relations and thereby the time relations sustained by particular events. But these relations are or are not sustained by physical macro-events quite apart from *our* actual or hypothetical measuring operations [...]. [It] is because no relations of absolute simultaneity *exist* to be measured that measurement cannot disclose them. (Grünbaum 1963: 368)

Grünbaum's point is that there is nothing in reality corresponding to the hypothetical relation of distant simultaneity: in reality, there are no facts of the matter about which moment in my history is simultaneous with the light beam hitting a distant mirror or detector.³⁴ As we saw, the problem of distant simultaneity is equivalent to the problem of the one-way velocity of light. So Grünbaum's solution boils down to the claim that there are no facts of the matter in physical reality about the one-way velocity of light.

The gist of this solution is that the way we represent the one-way velocity of light is a matter of choosing a certain set of coordinates instead of another, akin to the choice between measuring a certain length in feet or meters.³⁵ More technically, one can say that our representations of light are subject to a certain *gauge freedom*—just as physical reality doesn't care whether we count distances in feet or meters, it doesn't care whether we represent light as having a uniform or a non-uniform one-way speed.

Gauge freedom is a prevalent phenomenon in spacetime physics. As Tim Maudlin explains:


[T]he geometrical structure of a spacetime diagram is not the same as the geometrical structure of the space-time being represented. What we are doing when we draw a diagram is using one kind of geometrical object to represent another. Since the geometries of the two objects are not the same, we must pay close attention to which aspects of the diagram correspond to real physical facts and which are merely conventions. In the case of Newton's own account of space and time, [...] the *angle* of a trajectory on the diagram has a physical significance: it represents the *absolute velocity* of a body, with objects at absolute rest occupying vertical trajectories. But in Galilean space-time [= a relationalist spacetime for Newtonian laws], there are no absolute velocities. It is a matter of arbitrary choice which straight trajectories are depicted as vertical and which as "tilted" in the diagram. (Maudlin 2012: 63)

Just as there are no absolute velocities in relationalist spacetimes, likewise, there is no fact of the matter about the one-way speed of light. Just as we are free to shift the contents of a Galilean spacetime diagram 1 meter to the right, without any change to the real content of the resulting representation, likewise, we are free to assume either uniform or non-uniform one-way speed in special relativity, without any change in the real content of the resulting representation. Or so the friend of Grünbaum's solution claims.


³⁴ It is interesting to note that Reichenbach, the other great 20th century champion of the conventionality of simultaneity, believed the exact opposite: he thought there might be facts about the one-way velocity, but they are unknowable (Jammer 2006: 200, Friedman 1977: 426).

³⁵ Apart from Grünbaum, Mittelstaedt (1977) and Weingard (1985) also espouse this view.


One way to clarify Grünbaum's solution is to say that physical reality only contains causal relations, and these relations can be metrized in different ways, corresponding to different (conventional) assumptions about the one-way velocity of light. Suppose, for example, that we have two point-events, E and F , such that F is causally connected to E via a light signal and E is causally connected to the origin via a light signal. Assuming that the one-way velocity of light is uniform, one will represent the relative spatiotemporal positions of the three events the following way:



If one adopts a deviant model where light goes more slowly to the right than to the left, then one will adopt the following alternative representation of the relative positions of E , F , and the origin:



If one adopts Ben-Yami's deviant model (see p.79), then the relative position of the three events will be represented like this:



And so on. Every specific deviant model will give rise to its unique way to represent the relative positions of E , F , and the origin. The differences between these ways to represent the events in question will depend on differences between the shape of light cones in different models.

Grünbaum's point can now be reformulated in the following way: There are no facts of the matter about the 'real' shape of the light cone; the real facts are confined to the relations of causal precedence between various events. In our example, physical reality contains facts about the causal precedence of F over E and about the causal precedence of E over the origin. But when we start representing these three events in a coordinate system, the rest of the story is up to gauge freedom. We are free to assign numbers to these three events in any way we like (within certain limits), just as we are free to represent the length of the meter stick as 1 or (counting in feet) as $1 / 0.3048$. Reality does not force us to count in meters as opposed to feet. Similarly, reality does not force us to assign one specific one-way velocity to light instead of another (as long as the round-trip velocity constraints are respected).

The trouble with Grünbaum's solution is that it seems hard to accept that there are no metrical facts about the physical world, only relations of causal precedence. As Michael Friedman remarks:

Grünbaum has given us no reason to accept the view that the only objective temporal relations are constituted by causal relations. Indeed, how could one possibly support such a view? Our only grip on which properties and relations are objective constituents of the physical world is via our best theories of the physical world. The properties and

relations that we hold to exist objectively are those that our best physical theories postulate. And since our best theories do not merely postulate the kind of *ordinal* (causal) temporal relations favored by Grünbaum—they postulate metrical relations as well—we have no reason to grant such ordinal (causal) relations the privileged ontological status that Grünbaum wants to ascribe to them. (Friedman 1977: 430)

Notice that Friedman's point does not presuppose absolutism about physical quantities. There can be facts of the matter about the one-way velocity even if physical quantities must be understood in a relationalist way, since there will be an (uneven) relation between the speed of light in one direction and its speed in another direction. So one cannot dismiss Friedman's point by denying that physical quantities are absolute.

The basic problem about Grünbaum's solution is similar to the problem about the eternalist solution. The gist of the eternalist solution is that there is no movement in an eternalist cosmos, hence there are no facts of the matter about the one-way speed of light. This suggestion is shipwrecked by the fact that the average round-trip speed of light is a physical constant with a known value. If the eternalist solution were sound, then there would be no facts of the matter about the round-trip speed either. Similarly, if Grünbaum is right in claiming that the one-way speed of light is subject to considerable gauge freedom, then, intuitively, the average round-trip speed of light should be subject to the same kind of gauge freedom. But it isn't. Granted, the average round-trip speed of light is subject to some gauge freedom in the sense that one is free to express it in meters per second, feet per lunar month, or any other speed unit one likes. But this is not the same kind of gauge freedom that affects the *one-way* speed of light, which can be freely chosen within certain limits even after we fix a speed unit. Since there seem to be metrical facts about the physical world (e.g. the round-trip speed of light), the claim that the only real physical facts are facts about causal precedence seems to be false. Hence, it is far from clear that Grünbaum's solution really works.

3.3.6 Summary of 3.3

The previous five sections surveyed those responses to the conventionality of simultaneity that can be adopted by physical realists. In 3.3.1, I discussed attempts to measure the one-way velocity of speed. If any of these attempts had some chance of succeeding, the puzzle would find an empirical resolution. I presented a seemingly straightforward way to measure the one-way velocity, then I pointed out that it contradicts one of the basic assumptions of relativity. I also indicated that we have no reason to think that any other proposed method works.

Section 3.3.2 looked at arguments that appeal to symmetry. The gist of these arguments, championed by Michael Friedman among others, is that positing non-uniform one-way speed introduces an asymmetry into spacetime without explaining anything, hence deviant models can be disregarded. I argued that Friedmann's point concerns pragmatics rather than ontology: if physicists suddenly found it more convenient to use deviant models for their calculations, the real one-way speed would not change as a result. But even if we disregard this point, appeals to symmetry can only provide an epistemic, as opposed to metaphysical, solution to the dilemma, since these appeals presuppose that the one-way speed could have been different. Arguably, it would much be better to solve the puzzle by showing that deviant models are impossible or by showing that they do not represent a real difference from the standard case.

Section 3.3.3 briefly discussed David Malament's celebrated proof that only the hypothesis of uniform one-way speed satisfies certain symmetry constraints. Drawing on the work of Ben-Yami, Grünbaum, and others, I claimed that (i) Malament's proof presupposes that worlds where our history unfolds backward are relevant in terms of actual symmetries, which is a questionable premise, and (ii) the assumption that spacetime must conform to those specific symmetries that Malament favours appears to beg the question against the conventionalist. Hence, Malament's proof is not substantially different from Friedman-style appeals to simplicity.

Section 3.3.4 addressed the idea that the puzzle of the conventionality of simultaneity is dispelled by eternalism. The gist of the eternalist solution was that there is no movement in an eternalist cosmos, hence there are no facts of the matter about the one-way speed of light. I claimed that if this argument is sound, it entails its own *reductio*, because it establishes that there are no facts of the matter about the average round-trip velocity of light, which is a physical constant with a known value.

Finally, 3.3.5 addressed Grünbaum's contention that the one-way velocity of light is subject to gauge freedom. According to Grünbaum, physical reality involves relations of causal precedence only, and those relations are left intact when we move from deviant to non-deviant models or from one deviant model to another. Drawing on a suggestion by Michael Friedman, I argued that Grünbaum's theory is implausible because it entails that there are no metrical facts about physical reality whatsoever. Specifically, I pointed out that this idea contradicts the uncontroversial fact that the average round-trip speed of light is a physical constant with a known value.

To sum up, it seems reasonable to conclude that the puzzle about the conventionality of simultaneity has no straightforward and uncontroversial physical realist solution. Even if there is a sound solution somewhere, its details are far from evident and it doesn't seem to be much publicized.

3.4 The idealist solution

3.4.1 Sketch of the solution

The idealist solution to the puzzle of the conventionality of simultaneity is quite simple. We saw in Chapter 1 that the definition of ideal physical properties is roughly the following:

Ideal physical properties:

- (IP) Physical property P is ideal $\approx_{df}$
 $\square \forall x: x \text{ is } P \supset \text{Someone observes or could observe that } x \text{ is } P.$

(The parts of the definition that are missing, namely time indices, the condition that observations are preceded by exploratory actions, and the parts that are needed to defuse the counterexamples from 1.3, are not relevant here.)

To find out how idealism handles the puzzle of conventionality, let's apply (IP) to velocity:

Ideal velocities:

- (IV) Velocity is an ideal property $\approx_{df}$
 $\square \text{Some } x \text{ has velocity } V \supset \text{Someone observes or could observe that } x \text{ has velocity } V.$

As indicated in 1.2, it is natural to assume that all physical properties are ideal if idealism is true. It follows that

- (i) If idealism is true, then light does not have a real velocity (because all physical properties are ideal).
- (ii) By (IV) and (i), if idealism is true, there are no unobservable velocities.

Since the one-way velocity of light is unobservable, it follows from (ii) that light does not have a one-way velocity under idealism. In idealist worlds that have the kind of laws we do, there is no fact of the matter about the one-way velocity of light. Hence, if idealism is true, we don't have to ascertain that light moves with the same one-way speed in all directions, because there are no facts of the matter about those speeds. The puzzle about the conventionality of simultaneity is smoothly dissolved by idealism.

3.4.2 Elaborating the solution

It is worth investigating the idealist solution in a bit more detail, because even if the solution is formally straightforward, one might complain that it does not render the puzzle any less puzzling. How come that light has no speed going from here to the mirror?

To get a feel for the idealist answer, let's investigate a toy idealist world containing two immaterial subjects, Alice and George. We also assume that the laws of relativity (or their idealist equivalents, at any rate) hold in this world.

Facts about the world of Alice and George are facts about sensory episodes that are structurally and temporally harmonized. One can build an intuitive picture of this by invoking the concept of points of view. We all have an intuitive notion of what it means to have a point of view. It means, typically, that you have the visual experience of a spatially extended environment, perhaps of a quad like this:



Fig. 1 *A visual field*

The visual field can undergo various changes. The shapes that populate it at one moment give way to different shapes at later moments, giving rise to the experience of movement and change. Let's call the sum total of changes in the visual field of S "the visual stream of S ." (To be fully precise, one should also consider other sensory media like hearing and smell, plus proprioception and other forms of inner sense. I'll ignore these important aspects of experience, because they are not directly relevant to the puzzle at hand.)

Our toy idealist world contains two visual streams. Their owners have the experience of standing in a quad like the one in Fig.1, facing each other. The owner of the first stream is Alice. This is what Alice sees right now:



Fig. 2 *Alice's visual field at t_0
(on Alice's time)*

The owner of the second stream is George. This is what George sees now:



Fig. 3 *George's visual field at t_0
(on George's time)*

Alice and George can initiate changes in their visual streams. For example, George can perform an action he calls "raising my left hand." This action results in the following change in his visual field:



Fig. 4 *George's visual field at t_1 (on George's time)*

When George raises his left hand, Alice's visual field changes into this:



Fig. 5 *Alice's visual field at t_1 (on Alice's time)*

Such rules about the relationship between action and experience make it possible for Alice and George to interact. Alice and George live in the same world in virtue of the fact that their sensory streams are harmonized.

The visual streams of Alice and George are harmonized (i) *structurally*, so that if George raises his hand (from his perspective), Alice sees a hand raised (from her perspective), and (ii) *temporally*, so that if Alice sees George raise his hand, and raises her own, this fact registers in George's visual stream later than his own hand-raising. We may assume that these rules cover a large range of possible interactions (e.g. walking around the quad, handing each other objects, communicating etc.).

Suppose that Alice and George start testing various consequences of special relativity. They are faced with a choice between the standard theory where the one-way speed of light is uniform in all directions (let's call this theory "*U*") and various deviant models where the one-way speed of light is direction-dependent (let's call a representative deviant model "*NU*"). Now consider two scenarios.

In the first scenario, Alice and George synchronize two clocks (remotely or locally). They proceed to test some predictions of special relativity derived from *U*. Their measurements confirm all predictions.

In the second scenario, Alice and George synchronize their clocks and they test special relativity using *NU*. Their calculations take a bit more time but their predictions are again confirmed.

If Alice and George were living in a real spacetime, there would be four relevant possibilities concerning the choice between *U* and *NU*:

Realist Case 1

Alice and George use the *U*-theory AND
The one-way velocity of light is uniform AND
The predictions of *U* are corroborated.

Realist Case 2

Alice and George use the *U*-theory AND
The one-way velocity of light is not uniform AND
The predictions of *U* are corroborated.

Realist Case 3

Alice and George use the *NU*-theory AND
The one-way velocity of light is not uniform AND
The predictions of *NU* are corroborated.

Realist Case 4

Alice and George use the *NU*-theory AND
The one-way velocity of light is uniform AND
The predictions of *NU* are corroborated.

In two of these cases (2 and 4), Alice and George get the structure of the physical world wrong, even though their predictions check out and they have no way to tell that they got something wrong.

However, given that Alice and George live in an idealist spacetime, there are only *two* relevant possibilities:

Idealist Case 1


Alice and George use the *U*-hypothesis AND
The predictions of *U* are corroborated.

Idealist Case 2

Alice and George use the *NU*-hypothesis AND
The predictions of *NU* are corroborated.

In idealist worlds, Alice and George cannot get the structure of the physical world wrong by choosing *U* instead of *NU* (or vice versa). The choice between those two models will concern pragmatic and aesthetic issues only. Perhaps one of those models (presumably *U*) is easier to work with or it is more pleasing to the intellect.

To see how this state of affairs is underwritten by the ontology of idealist worlds, let's consider three hypotheses about the one-way speed of light:



In Scenario 1, light goes faster than c to the right and slower than c to the left. In Scenario 3, it's the other way around. Scenario 2 represents the standard case when the one-way velocities are uniform.

If the spacetime of Alice and George were real, then at least two of these trajectories would represent false hypotheses about what happens in real space between Alice and George.


However, given that the spacetime of Alice and George is ideal, nothing in their world corresponds to the lines on diagram. Real concrete facts are facts about the visual streams of Alice and George, so real concrete facts about the measurement process are exhausted by the fact that Alice has the experience of sending a light beam at a certain moment, then sees it return, and George has the experience of reflecting a beam back at Alice. The lines on the spacetime diagram do not register in the visual streams of Alice and George. The only role of these lines is to help Alice and George anticipate the results of experiments. The question about which line represents reality is replaced by considerations about the most convenient way to anticipate future experience.

3.4.3 Two objections

I can anticipate two objections to this solution. The first is this:


You claim that there are no facts of the matter about the one-way velocity of light in idealist worlds because there is nothing in reality corresponding to the lines that represent the trajectory of the light beam in spacetime. But suppose that, apart from Alice and George, there is a third subject *S* in the toy world, and *S* is looking at Alice and George from the side. *S* will see a light beam pass from Alice to George and back. There will then be phenomenal facts that go proxy for facts about real light-beams. Hence, the idealist cannot solve the puzzle of the conventionality of simultaneity by claiming that nothing in reality corresponds to the lines in the spacetime diagram in idealist worlds.

The objector is describing the following situation. Suppose that a spectator, *S*, is situated in such a way that the light beam traveling from Alice to George travels horizontally in *S*'s visual field:




The objector claims that, given this setup, there will be phenomenal facts that play the same theoretical role as the trajectory of a real light beam. Hence, the idealist cannot solve the conventionality issue by saying that nothing real corresponds to the lines in the spacetime diagram.

To assess this objection, let's see how the event of the light beam's reaching George registers for *S*. Suppose that *S* is looking toward George, anticipating the moment when the light beam reaches George:




The objector claims that there will be phenomenal facts of the matter about which of the two scenarios on the diagram obtain (if any), hence there will be facts of the matter about the one-way velocity even in idealist worlds.

To see why this worry is mistaken, let's take a cross-section of spacetime along the line SG :



The goal of S is to find out which point on the dotted line is the event of the light beam's reaching George.

In order for S to see that the light beam has reached George, light must travel from G to S . And in order for S to know the exact moment at which the light signal reached George, S must know how long it took light to travel from G to S . And that duration, in turn, will depend on the one-way velocity of light from G to S . For example, the fact that S registers at $t = 2$ (on S 's time) that the light beam reached George will be compatible with both standard and deviant assumptions about the one-way velocity, leading to competing hypotheses about the moment of reflection:



It follows that facts about the visual stream of a third spectator cannot go proxy for facts about the trajectory of a real light beam. The phenomenal facts of a third spectator will be subject to the same kind of underdetermination as the experience of Alice and George.

A second, less formal, objection to the idealist solution is that it does not dispel the mystery of conventionality but locates it at a different level. For even if there is no real light beam in idealist worlds, and hence nothing to correspond to the lines between Alice and George in the spacetime diagram, there will be a real mental event corresponding to the arrival of the light beam to George. The moment when George sees a flash of light coming from Alice will come at some point between the moment when Alice sends the light beam and the moment when the beam comes back to her. If the one-way speed of light is a matter of convention, it will also be a matter of convention when, on Alice's timeline, that event occurs. And it is no less mysterious how such a relation can be purely conventional. More to the point, the same kind of underdetermination that we find in the realist case is reproduced in idealism at the level of distant mental events.

In response, the idealist can say two things. The first is that as long as Alice and George are less than 100 miles apart, competing admissible scenarios about the one-way speed will disagree about the moment of the light beam's arrival by less than 0.0005 seconds. Such minuscule periods cannot be consciously apprehended, therefore facts about the distant simultaneity of mental events are not underdetermined in idealism as long as subjects are not far removed from each other.

Granted, the underdetermination can be cranked up to a level that is on the scale of conscious apprehension if subjects are very far from each other. Suppose, for example, that Alice is on Venus, George is on Earth, and they communicate via radio. The round-trip time of the radio signal is 5 minutes. The idealist must say that there is no fact of the matter about when, on Alice's timeline, George receives the signal during the 5 minutes that pass while she is waiting for an answer. But the idealist can bite the bullet here and say that a common 'now' requires spatial proximity. Subjects at cosmic distances are not part of the kind of shared time order that we know from ordinary experience.

3.4.4 The superiority of the idealist solution

It is clear that idealism solves the puzzle of the conventionality of simultaneity in a very straightforward way. Moreover, the solution is also uncontroversial in the sense that once the thesis of idealism is accepted, the explanation about the source and nature of the underdetermination that affects the one-way speed of light is not up to dispute. In this sense, the idealist solution to the puzzle is superior to physical realist ones. Possible physical realist solutions are far from uncontroversial in the same sense.

Since the idealist has a straightforward and uncontroversial solution to the apparent conventionality of distant simultaneity, idealism explains at least one important feature of spacetime better than physical realism. Hence, inference to the best explanation can motivate the view that spacetime is not real.³⁶

Ideally, the discussion should be carried further into an examination of general relativity. I have neither the space nor the expertise to do so. However, it seems to me that idealism might be able to solve the debate between field and geometrical interpretations of general relativity (see e.g. Ben-Menahem 2006: Ch.3 for an overview of the problem). Idealism might be able solve this puzzle in a way remotely analogous to the example in 3.1, that is by denying, on phenomenological grounds, that there is a real difference between curved space and forces acting in space. Of course, this is merely speculation at this point, and the attendant problems are extremely complex.

³⁶ One can interpret the puzzle in an antirealist manner as well, denying that our models are meant to fit something called "reality." As indicated in the Introduction, this dissertation is concerned with the dispute between idealism and physical realism only.

Finally, let me note that the idealist solution to the conventionalist puzzle follows the schema that was introduced in 3.1. The one-dimensional world of Albert the ant illustrated that idealism transforms underdetermination problems into problems about the most convenient way to predict future experience, while realism, when faced with the same problems, postulates excess structure that makes the physical world partly inscrutable. In 3.2–3.4, the same contrast emerged in the context of a puzzle about actual physics. We saw that if idealism is true, then the way we represent the one-way velocity of light is a pragmatic-aesthetic issue about the best tool for anticipating future experience. Physical realism, on the other hand, when faced with the same problem, introduces an inscrutable extra feature into reality.

All in all, this chapter corroborates Foster's abductive thesis that real spacetime is an explanatorily idle posit that should be thrown away in the interpretation of physics.³⁷

³⁷ It is interesting to note that Einstein seems to have been aware of the idealist solution. F.S.C. Northrop (1941) recalls that, in conversation with Einstein, he interpreted Whitehead as claiming that reality ultimately consists of mental events and relativity theory is to be understood in terms of facts about experience. "Oh! Is that what he means?," Einstein is reported to have replied. "That would be wonderful! So many problems would be solved were it true! Unfortunately, it is a fairy tale. Our world is not as simple as that." And after a moment's silent reflection he added: "On that theory there would be no meaning to two observers speaking about the same event" (Northrop 1941: 204, quoted by Jammer 2006: 162). This anecdote contains the seeds of a further possible objection to the idealist solution, namely the objection that idealism cannot account for the publicity of the physical world. Addressing this complaint would require a separate essay on singular terms and intersubjectivity that I cannot hope to include here. For suggestions, see the *Third Dialogue* (II: 247–8) on the identity of physical objects in idealism, and Foster (1982: Part V) on intersubjective time.

4 Truthmakers and laws in idealism

The goal of this chapter is to face two important metaphysical worries about idealism, the truthmaker objection and the nomic objection. According to the truthmaker objection, the idealist cannot supply truthmakers for physical truths. According to the nomic objection, the idealist cannot explain the existence of natural laws. I'll argue that these objections collectively debilitate non-theistic forms of idealism but can be defused quite easily if idealism is upheld in conjunction with theism.

Section 4.1 addresses the truthmaker objection. The gist of this worry is that the idealist cannot explain what makes truths about unobserved physical objects (or, more generally, truths about persisting physical objects) true. I'll examine this worry in two steps. Section 4.1.1 looks at truths about prehistory. This area seems especially fertile ground for developing the truthmaker worry, since the claim that physical objects are not real seems to entail that most of the alleged history of the cosmos did not really happen. I'll argue that this worry cripples non-theistic idealism but can be solved by positing God (or, more precisely, God's conception of the physical worlds as a seemingly autonomous system) as the truthmaker for prehistory. Section 4.1.2 looks at the case of contemporary unobserved objects, reaching a similar conclusion.

Section 4.2 addresses the nomic objection, according to which the existence of natural laws is a mystery for idealists. I'll address this worry in three steps, by looking at three markedly different, influential accounts of laws, and examining whether the idealist can adapt them to her own ontology. Section 4.2.1 deals with contemporary Humeanism about laws, section 4.2.2 looks at the theory of laws as second-order universals, and section 4.2.3 focuses on laws derived from intrinsic powers. I'll argue that the idealist can adapt each of these theories to her own ontology, the only serious requirement being that idealism must be conjoined with theism if nonHumean theories are adopted. Since the idealist can explain laws just as well as the physical realist under at least three different and influential metaphysical accounts of laws, the force of the nomic objection is questionable at best. Or so I'll argue.

The upshot of the discussion will be that idealism can easily ground the structure of the physical world in real facts if idealism is conjoined with theism, and idealism has a hard time doing so if it is not conjoined with theism. This suggests that idealism intrinsically harmonizes with theism.

4.1 Truthmakers

Intuitively, it is not clear how the idealist can explain truths about the physical world. If real physical things do not account for the presence of persisting objects of perception, then what does? David Armstrong (2004: 1n1) recalls that truthmaker theory was developed partly to make this worry more explicit. Although the worry can be spelled out in other ways as well, I'll use the idiom of truthmaking, because it is very convenient for present purposes. More specifically, I'll adopt the following principle:

Truths Need Truthmakers:

- (T) If P is a true proposition about the actual world, then there are some actual entities, the x s, such that the x s make P true in the sense that the existence of the x s entails that P is true.

For example, Fido the dog is a truthmaker for the proposition that dogs exist. Nuclear missiles are truthmakers for the proposition that nuclear disarmament has not occurred. And so on.

Suppose that there is a table in room 412 and there are no observers in 412 at the moment. Conditions are normal, so the following proposition is true:

- (P) If someone enters room 412, they will see a table.

The truthmaker objection can then be formulated as the following argument:

- (1) (P) is a true proposition about the actual world.
- (2) Therefore, by (T) and (1), some x s make it true that those who enter room 412 will see a table.
- (3) If idealism is true, then the only candidate truthmakers for (P) are immaterial minds and their mental states.
- (4) Immaterial minds and their mental states cannot make it true that those who enter 412 will see a table.
- (5) Therefore, by (2)–(4), (P) cannot be true if idealism is true.

The idealist obviously needs truths like (P) to build a sane physical ontology, so idealism is in trouble if (1)–(5) is sound.

The argument has three premises, (1), (3), and (4). Premise (1) is true by hypothesis. Premise (3) sounds plausible, since there are no real concrete entities apart from minds in idealist worlds. So the soundness of the argument depends on (4), which says that immaterial minds cannot make claims about unobserved physical objects true.

In the following, I examine the possible idealist responses to this problem in two stages. Section 4.1.1 deals with truths about prehistory (where the problem is especially acute), and 4.1.2 looks at truths concerning the present. My claim will be that defusing (1)–(5) requires the adoption of theism.

4.1.1 Truthmakers for prehistory

Intuitively, the prehistoric past is a domain where the truthmaker problem is particularly vexing for the idealist. If physical objects are not real, then it is hard to see how there could have been physical facts prior to human history.

The doctrine that there were no physical facts prior to human history was famously expounded, independently of idealism, by a Christian fundamentalist, P. H. Gosse, in *Omphalos* (1857). The title (which means “navel” in Greek) hints at Gosse’s claim that the biblical Adam had a navel not because he had a real mother but because God wanted to make Adam’s body conform to His conception of the natural order. More generally, we can take Omphalism to be the thesis that history is roughly as long as human history and the world contains apparent traces of an unreal past.

Idealism is often charged with Omphalist leanings (see BonJour 2011: 2.1, Price 1950: 298, Sellars 1963: 84). The thought behind the charge is roughly the following: If the physical world is a construction from experience, then no physical events predated experience. So most of our alleged prehistory (the Big Bang, the formation of the planets, the continental drift, the Jurassic period etc.) did not happen. More technically, the claim is that the idealist cannot supply truthmakers for prehistory. This sounds quite problematic.

To assess the objection from Omphalism, which is a variant of the general truthmaker objection, let’s invoke the (simplified) standard definition of idealism:

Ideality (standard version):

(IS) O is ideal $\approx_{df}$

$\square O \text{ exists} \supset O \text{ is observed or observable by someone}$

Now suppose that the idealist wants to ground the following (hypothetical) truth about prehistory:

(6) At the place where the Big Ben stands today, a Brontosaurus grazed 150 million years ago.

Suppose, further, that dinosaurs were unminded physical objects. (If this idea is contested on the grounds that dinosaurs had minds, the argument can be restated in terms of prehistoric volcanic eruptions or other brute physical events.) Let t_{-150M} be a moment 150 million years ago and let p_{BB} be the place where the Big Ben stands today. Then (6) and (IS) entail that

(7) At p_{BB} at t_{-150M} , a Brontosaurus was observed or observable by someone.

Ex hypothesi, there were no subjects 150 million years ago, so (7) reduces to

(8) At p_{BB} at t_{-150M} , a Brontosaurus was observable by someone.

If the idealist can supply truthmakers for (8), the charge of Omphalism is deflected. And if the idealist cannot supply truthmakers for (8), then the charge of Omphalism is formally established.

To find truthmakers for (8), the idealist must find truthmakers for counterfactuals about potential prehistoric experience. But it is hard to see how something like (8) can be true on idealism. If there are no subjects around, then nothing is observable by subjects, one would think. Worse, it seems clear that in idealism, nothing exists if subjects don't exist. It follows that nothing existed 150 million years ago if idealism is true, and therefore nothing could have made (8) true if idealism is true.³⁸

In reply, that the idealist might come up with the following story:

Backward Projection:

(8) is true because the laws of nature, when projected backward through a 150-million-year-period from the start of human history, yield the conclusion that, if someone had been around at p_{BB} at t_{-150M} , they would have observed a Brontosaurus.

Backward Projection is a coherentist theory of prehistoric truth. Its upshot is that the hypothesis that (8) is true coheres well with human history plus laws.³⁹

Backward Projection passes the buck to the idealist account of natural laws. The idealist is not much better off with this move, because it isn't immediately obvious how immaterial minds ground laws. (This issue will be explored in 4.2 below.) But suppose we disregard this problem. Then *Backward Projection* is still open to two serious objections.

The first objection is that *Backward Projection* is parasitic on the solution to a different truthmaker problem that cannot itself be solved in a way analogous to *Backward Projection*. Suppose we want to ground the existence of English dinosaurs using *Backward Projection*. We'll have to say that there are dinosaur bones and other apparent traces of the prehistoric past around, and these traces, when assembled into a coherent picture according to the laws of nature, entail

³⁸ In principle, there are two ways for the idealist to argue that human prehistory also involved finite minds. First, the idealist can assume that all physical objects are parts of the body of some mind. Leibniz seems to have believed this (Adams 1994: 241–53). This theory entails that dinosaurs, prehistoric volcanoes etc. also had minds in some sense, therefore (6) is a complex fact about experience, with minds and occurrent mental states as truthmakers. Alternatively, the idealist can say that facts like (6) are facts about the experiences of supernatural minds (e.g. angels) who witnessed the whole of prehistory. I ignore these appeals to exotic other minds, because it seems to me that they cannot ultimately escape theism, and theism can defuse the truthmaker objection all by itself.

³⁹ Descartes seems to have endorsed this view. In Discourse 5, after the outline of the formation of planets, he says: "I did not wish to infer from all these things that this world had been created in the way I described, for it is very much more likely that, from the beginning, God made it as it was to be. But [...] provided that he had established the laws of nature and lent it his preserving action to allow it to act as it does customarily, one can believe, without discrediting the miracle of creation, that in this way alone, all things which are purely material could in time have made themselves such as we see them today" (Descartes 1968: 64).

that a specific number of dinosaurs visited various places during the Jurassic. But note that this explanation crucially depends on the assumption that there are buried dinosaur bones and other traces around. Presumably, most of these traces have not been observed (we haven't dug up all dinosaur bones). As a result, *Backward Projection* is parasitic on the solution of a different truthmaker problem, the problem of truthmakers for buried dinosaur bones and other contemporary traces of the unreal past.

The second objection to *Backward Projection* is that it secures a determinate past only if the initial state of human history, together with laws, is compatible with only one possible prehistory. But this is unlikely if the laws of nature are indeterministic. Indeterministic laws can generate the same state from different prehistories. If our world is indeterministic, then the fact that a tribe of immaterial subjects woke up one day somewhere in Africa is compatible with different projected prehistories. And in that case, backward projection will only tell us that something like (8) was *probable*. But it is clear that facts about the past are not meant to be probabilistic.⁴⁰

All in all, *Backward Projection* is not a good solution to the problem of prehistoric truths. And since it is hard to see any other way to solve this problem under the standard definition of idealism, I conclude that the standard definition of idealism, without any further commitments, leads to Omphalism.

As discussed in Chapter 1, idealism can also be defined in an explicitly theistic fashion:

Ideality (theistic version):

(IT) O is ideal $\approx_{df}$

$\square O \text{ exists} \supset \text{God is causing or is disposed to cause subjects to observe } O.$

(IT) generates the following alternative to (8):

(9) God was disposed to cause subjects to observe a Brontosaurus at p_{BB} at t_{-150M} .

To secure truthmakers for (9), the theist idealist might suggest the following:

Divine Intentions:

(9) is true because God was disposed to cause humans specific types of experiences in case humans came into being (that is, in case God had decided to create them) earlier.

⁴⁰ In principle, the idealist could also claim that (8) is a conditional about the potential experiences of time travelers (with some daleks thrown in, perhaps). This theory is subject to a worry similar to the first objection to *Backward Projection*. Something must make it true now that when I travel back 150 million years, I'll see a Brontosaurus at the place where the Big Ben stands today. One will either need contemporary truthmakers for the potential experiences of time travellers, or, if one locates the truthmakers in the past, one has to postulate a time traveling witness for every prehistoric event.

The upshot of *Divine Intentions* is that God has a certain story about the physical world as a seemingly autonomous system. Facts about prehistory are facts about this divine conception, entailing dispositions to cause subjects experiences of specific types. Hence, the truthmaker for prehistory is God.⁴¹

It might be objected that *Divine Intentions* cannot secure the kind of prehistory that a sane physical ontology needs. Intuitively, when we claim that dinosaurs roamed the Earth, we are not talking about God's unmanifested dispositions to cause immaterial mental states. Neither are we talking about the divine conception of the physical world as a seemingly autonomous system. When we say that dinosaurs roamed the Earth, we mean that dinosaurs were *here*, breathing the same air we breathe, and leaving their bones behind when they died. *Divine Intentions* cannot secure such facts about prehistory; it can only secure a ghostly, unreal past.

This complaint can be understood in two ways. On one reading, the complaint asserts that dinosaurs were real. This boils down to the assertion that our intuitions demand a physical realist ontology. Idealists lack such intuitions and they deny that such intuitions are reliable.

On a different reading, the complaint is that the content of our intentional states about dinosaurs does not match the kind of content that the theist idealist attributes to them. When we think about dinosaurs, we think about big-boned prehistoric lizards that used to inhabit this planet and are causally connected to us through a long and ramified chain of influence. In contrast, the theist idealist asserts that our intentional states about dinosaurs refer to the contents of the Divine Mind. This sounds like a gross misconstrual of the contents of our thoughts. Or so the physical realist can argue.

The theist idealist can reply the following: When we think about dinosaurs, we imagine what dinosaurs looked like and we try to form a conception of their biology and evolutionary history. These are precisely the kind of contents that our dinosaur-directed intentional states would have had if we had come into being during the Jurassic and had the chance to explore the environment. Assuming that God is the cause of our experiences, it follows that the contents of our (true) dinosaur-directed thoughts are very similar to the contents of the experiences that God would have caused us if we had been around in the Jurassic. So *Divine Intentions* does not misconstrue the content of our dinosaur-directed thoughts. It only violates the physical realist intuitions that typically accompany these thoughts. But the idealist lacks such intuitions and she denies that such intuitions are reliable.

⁴¹ Or, rather, the truthmaker is the state of God's having a certain conception of our prehistory. God's existence only necessitates our prehistory if God necessarily creates a world like ours, which is implausible. Generally, truthmakers cannot be concrete objects but only states thereof (cf. Merricks 2007: 17–22). Let me ignore this subtle point.

4.1.2 Truthmakers for the present

The existence of unobserved contemporary objects gives rise to a similar but subtly different dialectic. Suppose that nobody is in room 412 right now and the following proposition is true:

(10) There is a table in room 412 now.

On the standard definition of ideality, (10) entails that

(11) A table in room 412 is now observed or observable by someone.

If nobody is in room 412 now, (11) reduces to:

(12) A table in room 412 is now observable by someone.

Generally, the existence of unobserved contemporary objects will require truths like (12) under the standard definition of ideality. As before, it isn't easy to see how the idealist can supply the requisite truthmakers. The idealist might suggest the following:

Forward Projection:

(12) is true because the laws of nature, when projected forward from an earlier moment when someone saw a table in room 412, yield the conclusion that those who now enter 412 will see a table.

Forward Projection is subject to the same initial complaint as *Backward Projection*: it requires an idealist account of laws. But suppose we bracket this issue. Then *Forward Projection* is still open to a decisive objection.

Forward Projection is predicated on the idea that earlier phenomenal facts can ground later physical facts, because, together with the laws, they entail determinate truths about potential experience. For example, some subjects may have had the experience of buying a table, taking it to 412, assembling it there, other subjects may have recently visited room 412, observing a table there etc. *Forward Projection* implies that these sensory episodes, when assembled into a coherent picture according to the laws, entail (12). But this theory can only generate a determinate physical world if all physical objects that exist now have been observed at least once.⁴² It is unlikely that our past is so rich in phenomenal detail. For example, we have no reason to think that we have observed all planets, grounding later facts about the structure of the cosmos. Consequently, *Forward Projection* must be complemented with an implausible plenitude principle that isn't independently motivated.⁴³

⁴² Perhaps it is enough to say that only traces have been observed (e.g. seeing smoke in the distance may be enough to ground the existence of a fire far away). But those traces must be so numerous that they are unlikely to invalidate the objection at hand.

⁴³ Note that *Forward Projection* is not subject to the complaint about indeterministic laws. If laws are indeterministic, then (10) will only entail that (12) is probable, and *Forward Projection* can generate this result.

As before, theistic idealism has a different, and more coherent, story to tell. On the definition of theistic idealism, (10) entails that

- (13) God is now causing or is now disposed to cause subjects to observe a table in room 412.

Coupled with the assumption that there are no observers in 412, (13) entails the following condition for the existence of the table:

- (14) God is now disposed to cause subjects to observe a table in room 412.

The upshot of (14) is very similar to the upshot of *Divine Intentions*. It implies that God has a detailed conception of our world, and, on the basis of that conception, He is disposed to cause us experiences of specific types. This theory allows the idealist to claim that the truthmaker for (10) is God.

As before, the physical realist can complain that God is not the kind of truthmaker that we intuitively associate with physical truths. And as before, the idealist can explore two readings of this complaint.

On the first reading, the complaint appeals to physical realist intuitions. Since the idealist lacks such intuitions and she thinks that such intuitions are unreliable, this version of the complaint is ineffective.

Alternatively, the complaint can be seen to highlight a mismatch between the content of our world-directed thoughts and the idealist construals of those thoughts. When we think about tables, we think about hunks of matter located in space. We do not think about the divine conception of the world. Hence, treating God as the truthmaker for truths about tables etc. misconstrues the contents of our world-directed thoughts.

The theist idealist can reply the following: When we think about unobserved physical objects, we imagine what they look like and we try to form a coherent conception of their structure. These are precisely the kind of contents that would characterize our phenomenal states, and the inferences we make on the basis of them, if we were actually observing the objects in question. Assuming that God is the cause of our experiences, it follows that the contents of our (true) thoughts about unobserved physical objects are similar to the contents of the experiences that God would cause us if we were observing the objects in question. So treating God as a truthmaker for (10) does not misconstrue the content of our world-directed thoughts. It only violates the physical realist intuitions that typically accompany those thoughts. But the idealist lacks those intuitions and she has no reason to be intimidated by them. So neither of the two interpretations of the complaint at hand has a purchase on the contention that God is the truthmaker for truths about unobserved physical objects.

4.1.3 Summary of 4.1

I have argued that theistic idealism can supply truthmakers for physical facts while non-theistic forms of idealism have a hard time doing so. In 4.1.1, I examined truths about prehistory. I argued that *Backward Projection*, the only plausible story that the non-theist idealist can come up with, cannot cope with indeterministic laws and presupposes truthmakers for the present. I also claimed that the theist idealist can coherently maintain that the truthmaker for prehistoric truths is God. I defended this thesis against two complaints, one that appealed to physical realist intuitions and one that charged the idealist with misconstruing the content of our past-directed thoughts.

In 4.1.2, a similar dialectic unfolded about currently unobserved physical objects. I claimed that the only non-theistic idealist solution in sight is *Forward Projection*, which is predicated on the idea that one can explain facts about potential experience by extrapolating from past experience via laws. I pointed out that this theory requires an implausibly rich stock of past experiences. Finally, I argued that theistic idealism can supply a truthmaker for truths about currently unobserved objects, and I defended this thesis against the same objections that were raised against its counterpart in 4.1.1.

4.2 Laws

One common objection to idealism is that it cannot explain the laws of nature. The basic intuition seems to be that once everything is reduced to experience, there is no reason why our experiences should conform to one kind of order instead of another—there is no external reality that constrains what kind of experiences can succeed one another. Here's a detailed exposition of this complaint by Laurence Bonjour:

Why, according to the [idealist], are the orderly sense-data in question obtainable [...] ? What is the *explanation* for the quite complicated pattern of actual and obtainable sense experiences that, according to [idealism], constitutes the existence of a material object or of the material world as a whole, if this is not to be explained by appeal to genuinely external objects? The only possible [idealist] response to this question is to say that the fact that sensory experience reflects this sort of order is simply the most fundamental fact about reality, not further explainable in terms of anything else. (Bonjour 2011: 2.1)

In the original text, Bonjour refers to phenomenalism instead of idealism, but his complaint can be applied to idealism without any further ado.

Note that Bonjour's complaint only has bite if the physical realist is not faced with the same problem. The physical realist must establish that laws of nature *are* explained by the existence of real physical objects.

J. L. Mackie makes the following suggestion about the advantages of physical realism in explaining laws:

The postulation that each material object has some intrinsic quantitative feature which reacts contingently but lawfully with imposed forces is just what we need to explain the otherwise remarkable coincidence that [specific instances of the same Newtonian law] hold for all objects. (Mackie 1973: 151)

Suppose, for illustration, that Newton's law of gravitation is indeed a law, and consider two balls, ball *A* (mass: 1 kg) and ball *B* (mass: *m* kgs), located 1 meter apart. Newton's law of gravitation tells us that *A* will attract *B* with a force of Gm , where G is a gravitational constant. No matter the magnitude of *m*, the force will be Gm . It follows that Newton's law sums up a nondenumerable infinity of specific rules about the relationship between distance, mass, and gravitational force. Mackie suggests that the hypothesis that *A* has a real intrinsic property (e.g. some quantity of mass) can explain why *A* obeys this nondenumerable infinity of specific rules. And this remarkable coincidence seems unexplained if there is no real matter at the place where *A* is and the rules in question just happen to co-obtain. Or so the physical realist can argue, and her contention certainly merits discussion.

To sum up, we have two distinct but complementary theses:

(15) Laws of nature are brute facts if idealism is true.

(16) Laws of nature can be explained by the presence of real matter.

If both (15) and (16) are true, then we have reason to prefer physical realism to idealism. (More precisely, we have reason to do so if we value the explanatory power of metaphysical theories. We may assume this for the sake of argument.) On the other hand, if one of (15) and (16) is false, then we have no reason to prefer physical realism to idealism, as far as laws are concerned.

To clear the ground, let me note, and push aside, two easy idealist rejoinders to (15). An easy theist idealist rejoinder would be that our world has the laws it does because of God's providential concerns. Laws of nature are not brute facts because God had some purpose in mind when He instituted them.⁴⁴

⁴⁴ This suggestion is clearly present in the history of idealism, even though idealists have disagreed on the details of the story. Berkeley (*New Theory of Vision* §147, I: 231) thinks that the laws of nature serve the aid and comfort of mankind. Leibniz, prefiguring the "best system" account of laws advocated by Ramsey and Lewis (see p.113 below), claimed that God chose those specific laws that He did because "simplicity of the ways [of nature] is in balance with the richness of the effects" as a result (*Discourse on Metaphysics* §5, Leibniz 1989: 38–9), in other words, the actual laws strike an optimal balance between simplicity and strength. It is interesting to note that even non-theistic forms of idealism can derive laws of nature from cosmic teleology. The strongly idealistic Yogācāra school of Buddhism explained laws by reference to our karma and our collective recollection of past lives (see section XXXII of Vasubandhu's *Vimśatikā* in Tola and Dragonetti 2004: 113–4, 147–8).

In response, the physical realist can agree that, if theism is assumed, then laws were designed by God on the basis of specific providential plans. But this fact does not help the idealist solve the nomic objection, because the fact that God had certain reasons for instituting these specific laws does not constitute a *metaphysical* explanation of laws. It does not tell us what it is about physical objects that makes them conform to a highly coherent and extremely complex set of laws. The physical realist thinks she has a good story to tell about this. For example, she can join Mackie and claim that quantities of real mass glue specific instances of Newton's third law together. But the idealist does not seem to have a similar story to tell about the metaphysics of laws. To make matters more difficult for the idealist, I'll assume, therefore, that (15) cannot be denied by reference to God's providential plans.

Another easy rejoinder to the nomic objection would be the comment that, on a number of respectable theories of the metaphysics of laws, laws *are* brute facts, regardless of the truth or falsity of idealism.⁴⁵ If these theories are sound, then (16) is false, and the nomic objection against idealism is undermined. But, as before, I will not assume that this move settles the debate, because the nomic objection can be reformulated by conjoining it with the denial of some premise or intuition that motivates the theory of brute laws. I'll assume that laws of nature are not brute facts but need some sort of explanation, and I'll grant the physical realist (16), the claim that physical realism can explain laws.

With these preliminaries in place, let me outline my dialectic against premise (15) of the nomic objection, which is the thesis that the idealist cannot explain laws. In 4.2.1–4.2.3, I examine three different metaphysical accounts of laws. In 4.2.1, I focus on Humeanism about laws. In 4.2.2, I look at the theory of laws as relations between universals, and in 4.2.3, I discuss laws derived from intrinsic powers. In each case, I will argue that the idealist can adapt the theory in question to her own ontology, generating a metaphysical explanation of laws that is at least as good as its physical realist counterpart. Consequently, we have reason to think that (15) is false.

Needless to say, I don't expect the following dialectic to settle all problems about the nature of laws in idealism. A thorough treatment of this topic would require a separate dissertation. But, arguably, the force of the nomic objection is severely diminished if (15) turns out to be false under three very different and influential conceptions of laws. Minimally, this indicates that the nomic objection presupposes some special theory of lawhood or some other nontrivial premise hidden, so it cannot constitute a straightforward objection to idealism.

⁴⁵ Brutal lawhood is endorsed by Carroll (1994), Lange (2009), and Maudlin (2007), among others. Note also Foster's (2003) argument that all non-theistic ontologies must treat laws (and the validity of induction) as ultimately unexplained.

4.2.1 Humean laws

On Humean conceptions, laws of nature reduce to global regularities. However, sane Humeans do not straightforwardly identify laws with regularities, because some regularities are accidental. There are no mountains of gold and no mountains of uranium, but the nonexistence of the latter is a law while the nonexistence of the former isn't. Moreover, not any old regularity is fit to be a law, even if we rule out accidental regularities. For example, the conjunction of Einstein's field equations with the axioms of set theory is not a law, even though it entails a huge number of non-accidental exceptionless regularities.

For the purposes of the dialectic about Humean laws, I propose to adopt the most detailed and comprehensive contemporary Humean theory of laws, that of David Lewis. Lewis himself credited F. P. Ramsey with the core idea:

I adopt as a working hypothesis a theory of lawhood held by F. P. Ramsey in 1928: that laws are "consequences of those propositions which we should take as axioms if we knew everything and organized it as simply as possible in a deductive system." [...] What we value in a deductive system is a properly balanced combination of simplicity and strength—as much of both as truth and our way of balancing will permit. We can restate Ramsey's 1928 theory of lawhood as follows: a contingent generalization is a *law of nature* iff it appears as a theorem (or axiom) in each of the true deductive systems that achieves a best combination of simplicity and strength. (Lewis 1973: 73)

This theory implies the following protocol for determining the laws of a possible world W :

- (i) We write down all the (concrete) facts of W . Let's call the resulting class of propositions " Φ ."
- (ii) We write down all systems of general propositions that entail only propositions in Φ . Let's call the resulting class of systems " Σ ."
- (iii) We take the intersection of those members of Σ that entail the greatest part of Φ at the price of the least amount of complexity. Let's call their intersection " S ."
- (iv) S is the system of Humean laws (H-laws) in W .

There are a number of well-known complications about this protocol. The first is that it is unclear what guarantees that step (iii), the step where we pick a specific set of laws, can be carried through. What guarantees that any system strikes the optimal balance between simplicity and strength, and that those systems that do have a nonempty intersection? Nothing, it seems. Hence, nothing guarantees that W has H-laws.

A second, related problem is that it is unclear whether the demands of unity and simplicity are free from anthropocentric bias. Perhaps a system *S* that looks very simple to us looks very complicated to Martians. If the property of being simple is subject to such perspectival bias, then a world *W* might have different H-laws depending on who runs the protocol. And even if one admits a notion of objective simplicity, it is far from clear that what is objectively simple is also simple for us.

A third problem about H-laws is that worlds that contain relatively few facts will contain relatively few H-laws. To see why this is trouble, take a world *W* that contains nothing but a single electron, motionless in a space that lacks any electric fields. Intuitively, this scenario seems coherent. However, this scenario entails that some laws in *W* are uninstantiated—for example, the law that electrons are attracted by protons is uninstantiated in *W*. Once we assume that the object that inhabits *W* is an electron, it seems hard to deny that it is a law in *W* that electrons are attracted by protons. In any case, it does not seem logically incoherent to make that assumption. But it is clear that the H-laws of *W* will not include the law that electrons are attracted by protons.

Whether these objections wreck Humeanism about laws is an open question.⁴⁶ I will not argue either way. I only mention these issues in order to indicate the kind of problems that the idealist does *not* have to worry about once we assume that the laws of nature are H-laws. If the laws are H-laws, then the problems I've just outlined are problems for physical realists and idealists alike, hence they are irrelevant for evaluating the comparative merits of idealism and physical realism in the context of Humeanism about laws.

Let's assume, then, that rules (i)–(iv) are set in stone as far as lawhood is concerned and all general objection to (i)–(iv) have been answered or laid aside. The question now is whether the physical realist has an easier time construing laws as H-laws than the idealist.

The gist of Humeanism about laws is that laws supervene on the global distribution of local occurrences. If pieces of radium typically turn into radon over time throughout the history of a physical realist world, then it will be a law, or a consequence of some laws, that radium decays into radon. Let's see whether the idealist can hijack this idea.

⁴⁶ On the mind-independence of H-laws, see Lewis (1994: 479) and Loewer (2004: 185–7). On the problem of simple worlds, see Loewer (2004: 192–4). I am not aware of Humean responses to the objection that there may not be a unique best system. Perhaps Humeans are willing to live with this consequence but they assume, by abduction from actual physics, that our world does have a unique best system. Note also that even if the three objections are impossible to rebut, the Humean can argue that laws should be identified with H-laws. Lewis himself used H-laws to build an elegant account of counterfactuals (H-laws ground closeness of possible worlds, cf. Lewis 1973: 74–5), which, in turn, yielded a powerful account of causation (Lewis 1986a: 23, 1986b: 164–5, 2000). Problems about H-laws may be outweighed by the theoretical utility of H-laws.

Here's a straightforward idealist way to construct H-laws about radium: Observations of radium are typically followed by observations of radon, with subtle and stable correlations between these two types of observations (e.g. there is a stable regularity about changes in mass). Hence, assuming that our world is idealistic, it will be an H-law, or a consequence of some H-laws, that radium (conceived as an ideal kind) decays into radon over time.

Generally, given a rich stock of experiences, H-laws can be generated in idealist worlds just as well as in physical realist worlds. So physical realism does not have an explanatory advantage over idealism in terms of explaining laws if the laws are the H-laws.

It might be objected that there may not be enough experiences to ground some H-laws, because some laws might lack observed instances. If this complaint is sound, then the idealist can only generate an impoverished stock of H-laws. This complaint, in turn, can be interpreted in two different ways. On the first interpretation, the complaint is that in small worlds (e.g. in a world containing only a few subjects), the stock of experiences will be so meagre that idealist H-laws will not cover all physical laws.

In response, the Humean idealist can reply that we buy into this problem once we buy into H-laws. We saw in connection with the lonely electron that the physical realist Humean has a similar problem about uninstantiated laws. So the complaint at hand does not seem to be a complaint about idealism.

Alternatively, it might be objected that the H-laws of our own world are much harder to generate under idealism, because there will be much fewer concrete facts. Suppose that the world ends tomorrow. Then the concrete facts of our world are facts about the phenomenal history of a few billion humans. By contrast, if physical realism is true, then history includes facts about trillions and trillions of elementary particles. It might be argued that the relative paucity of concrete facts under idealism makes it hard for idealists to generate all actual H-laws. Nothing guarantees that every physical law is instantiated at least once in the phenomenology of some subject.

In reply, the Humean idealist can argue that she can generate all the H-laws that are experimentally confirmed during the history of the world, so she can generate all H-laws that will ever be known. The existence of all the H-laws we'll ever know is guaranteed even if our world is idealistic, because laws can only be known on the basis of a sufficient stock of observations.

It might be objected that phenomena that are exceedingly hard to observe will not have enough observed instances in idealist worlds, so the set of idealist H-laws will be impoverished anyway. Suppose that ξ -bosons are observed only a few times, so observation cannot ground idealist H-laws about ξ -bosons. Nonetheless, there might be laws about ξ -bosons.

In reply, the Humean idealist can say that statements can qualify as H-laws simply in virtue of cohering well with the other laws. Hence, even if ξ -bosons are very rarely seen, a general principle that mentions ξ -bosons may be a law in idealism by helping to achieve the optimum trade-off between simplicity and strength in the context of a specific deductive system.

Granted, the Humean idealist cannot admit laws that have never been subject to any kind of confirmation. For example, the Humean idealist cannot generate H-laws laws about Unobtanium₁₄₀, a substance that is found in remote parts of the galaxy and cannot interact with anything we'll ever know of. But it is unclear whether this is a bug or a feature. After all, the Humean intuition is that laws depend on widespread regularities. Hence, laws of idealist worlds will depend on widespread phenomenal regularities. If there are no phenomenal regularities concerning Unobtanium₁₄₀, then a self-respecting Humean idealist will conclude that there are no laws about such a substance, just as a Humean physical realist will say that there are no laws about Unobtanium₁₄₀ in worlds that do not contain Unobtanium₁₄₀.

Moreover, the Humean idealist can press the point that her theory gives us all the laws that science will ever discover. Hence, the theory fulfills the minimal requirement on any decent metaphysics of laws: it accounts for all the results of science.

Finally, it might be objected that the deductive systems that we are inclined to identify with the laws of our world (e.g. Einstein's field equations, or the laws of quantum mechanics) are not statements about experience. They are statements about physical structures, expressed in a mathematical form. Hence, idealism cannot give us actual H-laws, because actual H-laws are mathematical statements about physical structures, while the H-laws of idealist worlds are statements about experience.

This complaint rests on the misconception that the idealist is obliged express every physical fact in terms of facts about phenomenal states. But it is hard to see why she should be. The idealist can agree that our candidate laws are mathematical statements about physical structures while maintaining that those structures are ideal. It is certainly not obvious that they are ideal, because it is all too easy to reify our notion of the physical into something purportedly real. But this fact in itself does not invalidate the idealist take on H-laws.

To sum up, Humean idealism is not vulnerable to the complaint that the stock of phenomenal facts is too sparse. If laws are H-laws, then idealist worlds will contain all the laws that scientists in those worlds will ever discover. Demanding more betrays an unHumean bias that has nothing to do with the relative merits of idealism and physical realism in the context of Humeanism.

4.2.2 Laws as second-order universals

An influential, and markedly unHumean, account of natural laws identifies laws with relations between universals.⁴⁷ For an illustration, imagine an unripe plum. One can think of *being an unripe plum* as a complex universal composed of further universals like *being green*, *being sour* etc. When the plum ripens, it ceases to instantiate some of these universals and starts to instantiate others. *Being green* is replaced by *being dark blue*, *being sour* is replaced by *being sweet* etc. Those who conceive of laws as relations between universals invite us to think of this process by thinking of *being an unripe plum* and *being a ripe plum* as two (complex) universals that are linked by a nomic necessitation relation (which is also a universal, but a second-order one). Whenever *being an unripe plum* is instantiated (in normal circumstances), it drags its nomic complement, *being a ripe plum*, into existence via the nomic necessitation relation. Generally, the theory of laws as second-order universals says that lawlike regularity arises because universals are connected to each other through nomic necessitation relations, forcing the instantiation of further universals.

There are a number of well-known objections to this theory. David Lewis famously complained that calling a relation “nomic necessitation” does not give us a grip on what laws are any more than calling someone “Armstrong” can guarantee that the person in question has mighty biceps (Lewis 1983: 366). There are other, more technical objections as well.⁴⁸ I bracket these problems, because my goal is not to amend or criticize the theory at hand. The question is whether it puts the physical realist at an advantage.

One might try to adapt the theory of second-order universals to idealism by telling a story about links between universals that characterize the experience of immaterial subjects. One might suggest that whenever a subject looks at an unripe plum, her mental states involve various universals that qualify her visual experience as the experience of a small elliptical greenish object. Let’s call the universal that characterizes this experience as a whole “ E_1 .” The idealist might suggest that E_1 is linked via a nomic necessitation relation to another universal, E_2 , which characterizes the visual experience of a ripe plum (of the same size etc.). When ideal plums ripen, E_1 is replaced by E_2 in the phenomenology of observers. The fact that ideal plums ripen is explained by the second-order universal that links E_1 to E_2 .⁴⁹ This is the first-pass idealist version of the theory of laws as second-order universals.

⁴⁷ The classic sources are Dretske (1977), Tooley (1977), and Armstrong (1983, 1997).

⁴⁸ E.g. van Fraassen (1987) argues that the theory cannot account for probabilistic laws.

⁴⁹ To be absolutely precise, one should emphasize that the ripening depends on environmental factors. Strictly speaking, the universal corresponding to the experience of an unripe plum is linked to the universal corresponding to the experience of a ripe plum together with various other universals that guarantee that the environment is conducive to ripening. Let me ignore this complication.

There are two problems with this story. The first is that it does not explain how unobserved plums come to ripen. The story is hard to adapt to the case when an unripe plum hangs from an unobserved tree and becomes ripe without anyone looking.⁵⁰ The story just outlined implies that all physical objects that are subject to lawlike change are observed all the time, which is incredible.

Another, more subtle problem concerns phenomenally indiscriminable states. Suppose that a certain greenish piece of wax G , when looked at from a certain angle, looks exactly like an unripe plum. Subjects who look at G are in the same phenomenal state as subjects looking at an unripe plum. If E_1 is the universal that characterizes the visual experience of an unripe plum, then, by hypothesis, E_1 is also the universal characterizing the visual experience of G . If a nomic necessitation relation links E_1 to E_2 (the universal corresponding to the visual experience of a ripe plum), then it follows that those who look at G long enough will see it turn into a ripe plum. This is wrong. So the first-pass attempt at an idealist reading of the theory of laws as second-order universals fails.

Consider the following, alternative, theory: We conjoin idealism with theism and we assume that subjects instantiate various phenomenal states in virtue of standing in specific (causal) relations to God, with these relations, in turn, conceived as universals. Suppose, further, that the causal relations between God and finite subjects are underwritten by God's conception of the sensible world. If a subject instantiates E_1 while having the visual experience of an unripe plum, she stands in an appropriate causal relation R_1 to God, and there is a corresponding intention or structural conception S_1 in God's mind such that the existence of S_1 entails that God will bear R_1 to subjects under specific circumstances (e.g. when someone looks at a plum tree in spring). And there is another intention or structural conception S_2 in God's mind, linked to S_1 in such a way that once God enters into R_1 with subjects, God is then disposed to enter into another relation R_2 with subjects, with R_2 entailing that subjects will later instantiate E_2 and have the visual experience of a ripe plum.

This theory construes laws as second-order universals in the sense that God is assumed to have the same intention or structural conception about numerically distinct but qualitatively identical plums, so the relation between S_1 and S_2 crops up repeatedly in God's conception of the world, and laws about ripening are indeed relations between universals. The theist idealist may even assume that S_1 and S_2 involve the relevant phenomenal universals as well, so that E_1 to E_2 are also linked through a lawhood-conferring relation.

⁵⁰ Note that, on the present definition of idealism, there can very well be unobserved plums. All that one needs for the existence of unobserved plums is that they be observable (or that God be disposed to cause subjects to observe it). Idealism itself is not threatened by unobserved plums. What is problematic is the idealist explanation of ripening in terms of relations between universals that characterize sensory experience.

The theist idealist take on the theory of laws as second-order universals is not open to the objections that wreck the first-pass attempt. Let me return to those objections in turn.

Physical objects that are subject to lawlike change need not be observed all the time in the theist idealist story. The only requirement is that God have a determinate conception of all the physical objects that there are.

The objection from phenomenally indistinguishable states has no bite either. Suppose that Alice is looking at a greenish piece of wax G . Her phenomenology is characterized by the same universal, E_1 , that she would instantiate if she were looking at an unripe plum. On the theist idealist story of laws as second-order universals, what happens when Alice is looking at G is that she instantiates E_1 in virtue of bearing a certain relation R_G to God, and this relation, in turn, obtains because God has a certain intention or structural conception S_G in mind. The objection is blocked because the theist idealist has no reason to suppose that S_G is the same intention or conception that underlies our experience of unripe plums. Specifically, the theist idealist can say that S_G is linked to intentions or conceptions that entail that God will cause us waxlike experiences later. Therefore, even though the same phenomenal state, E_1 , figures in both God's conception of plums and in His conception of plumlike pieces of wax, these two conceptions are distinct. God's conception of G does not entail the later visual experience of ripe plums, because S_G is not nomically connected to God's conception of ripe plums. The theist idealist can rebut the objection from phenomenally indistinguishable states because, unlike the atheist idealist, she does not have to assume that the second-order universals that underlie laws link phenomenal qualities directly.

4.2.3 Laws from powers

Finally, consider a second unHumean theory of laws, the theory that laws derive from the intrinsic powers possessed by the objects falling under laws.⁵¹ For example, one can conceive of electrons as having an intrinsic power, the power of attracting positively charged particles. The laws of electricity will then be the consequences of such powers. Or, to return to the previous example, one can conceive of unripe plums as having an intrinsic power, the power to ripen. This power is naturally activated when circumstances are normal, resulting in the observed regularity that unripe plums tend to ripen. Generally, one can think of laws as grounded in the powers of objects, with powers conceived of as intrinsic properties that entail, but are not reducible to, conditionals. (One can think of powers either as universals or as tropes. This issue is unimportant in the present context.)

⁵¹ See e.g. Bird (2007) and Ellis (2001).

At a first pass, the idealist might try to adopt power-based laws to her own ontology by arguing that immaterial subjects have certain primitive perceptual powers (for example, the power to have visual experiences various kinds, the power to have auditory experiences etc.), and laws arise when these powers are activated. The idealist might argue that our perceptual powers have a very similar overall structure, hence their activation gives rise to the experience of a shared orderly world.

There are two problems with this story. The first is that it is unclear what activates primitive perceptual powers in this conception. Presumably, our primitive perceptual powers don't get activated just by themselves, and it is also implausible to suggest that we activate each other's perceptual powers, since even though we can initiate changes in each other's phenomenology, we are not responsible for the whole content of each other's experiences, or, at any rate, it is hard to see how we could be.

But even if we disregard this issue, the existence of primitive perceptual powers cannot explain the existence of laws. To see why, suppose that Adam and Eve are the only inhabitants of an (ideal) island. In the middle of the island, there is an object *O* whose existence reduces to facts about the primitive powers of Adam and Eve. Eve thinks that *O* is a tree. Her belief that *O* is a tree is warranted by appropriate tactile and audiovisual experiences that lawfully accompany her attempts to approach and examine *O*. Adam, on the other hand, thinks that *O* is a stationary rhinoceros. His belief that *O* is a stationary rhinoceros is likewise supported by appropriate phenomenal regularities that are lawfully correlated with his actions. Finally, let's suppose that this incoherence is undetectable. When Eve climbs the tree, Adam sees her climbing a rhino, when Adam remarks that the rhino is very friendly, Eve hears a remark about a tree or does not hear anything (without this lapse in the conversation registering for either of them) etc.

This scenario is compatible with the existence of perceptual powers, but in this scenario, *O* is not a public tree and not public rhino; it is a tree for Eve and a rhino for Adam. As a result, there cannot be physical laws about *O*, because *O* is not a determinate physical object. This scenario is easy to generalize, so it follows that the existence of laws requires something extra over and above the existence of primitive perceptual powers in idealism. It requires that the powers of subjects be activated in a harmonious way. This extra requirement, in turn, wrecks the proposed explanation. If laws ultimately depend on the harmony between the way powers are activated, then powers play a subordinate role: they are the items that need to be in harmony in order for there to be laws. The nomological heavy lifting is done by the harmony and not by the powers. And since harmony does not arise from the existence of primitive perceptual powers, the story at hand is a brute-fact theory of laws in disguise.

The theist idealist is, again, in a better position to deflect this objection. The theist idealist can deflect this objection by arguing that, given God's decision to create a physical world, the perceptual powers of subjects are guaranteed to be activated in a harmonious way. Let me elaborate.

If the perceptual powers of subjects are activated in a non-uniform fashion in an idealist world, then subjects cease to inhabit a common physical world and start inhabiting a number of loosely overlapping individual environments. To see why, suppose that the representational incoherence that affects Adam and Eve is a bit more pervasive. Adam thinks that the island is filled with coconut trees and Eve thinks it is filled with banana trees, without this discrepancy registering for them. Whenever Adam picks a coconut, Eve sees him pick a banana, whenever Eve hands Adam a banana, Adam has the experience of receiving a coconut, whenever Adam makes a remark about coconut trees, Eve hears a remark about banana trees etc. And suppose that this kind of hidden incoherence is pervasive, so that, in fact, there is nothing common in Adam's and Eve's conceptions of their environment beyond the idea that they live on an island. Clearly, Adam and Eve do not inhabit a common physical world in that case—they live in their own individual environments that are loosely correlated but fail to mesh into a common world. Since the physical is public, it follows that there is no physical world in the world of Adam and Eve in this scenario. And since this example is easy to generalize, it follows that there is no physical world in idealist worlds where the perceptual powers of subjects are activated in a non-uniform fashion.

Consequently, the theist idealist can argue that, given God's decision to create a physical world, it is *a priori* guaranteed that the perceptual powers of subjects are activated in a harmonious way and hence there are laws. (Note that physical realism cannot generate the same result. There is no reason why a real physical world could not be totally chaotic.)

One might object that this story ultimately appeals to teleology, specifically, it appeals to God's providential plans, and therefore it does not explain laws from powers. Powers are the metaphysical matter for an order that is imposed by God on the basis of providential concerns.

In response, the idealist may argue that the existence of laws must ultimately be explained by reference to God's providential plans anyway.⁵² But even if this bold response is avoided, the idealist can reply that the harmony of powers is not automatically guaranteed on physical realism either. Consequently, the problem at hand is a general problem about the theory of power-based laws, and not a problem about the relative merits of physical realism and idealism in the context of power-based laws.

⁵² See Foster (2003) for an argument that explaining laws requires theism.

4.2.4 Summary of 4.2

Section 4.2 addressed the nomic objection to idealism. The gist of this objection was that idealism is inferior to physical realism because the idealist must treat laws as brute facts. I have examined three different accounts of laws, a Humean one and two unHumean ones. In 4.2.1, I argued that the Humean idealist can ground laws just as well as the physical realist. In 4.2.2 and 4.2.3, I reached a similar conclusion about laws conceived as second-order universals and laws derived from intrinsic powers.

I have not offered a comprehensive treatise on idealist laws. But I hope to have shown that the case for the nomic objection is not very strong. Humean and unHumean conceptions of lawhood can alike be adapted to idealism. The only requirement seems to be that idealism must be conjoined with theism if the idealist goes unHumean.

4.3 Summary of Chapter 4

This chapter discussed two metaphysical objections to idealism. The truthmaker objection says that the idealist cannot supply truthmakers for truths about the physical world. In 4.1, I argued that this objection is justified if idealism is upheld in a non-theistic version and it is unsound under theistic idealism. The nomic objection to idealism said that the idealist cannot explain the existence of natural laws. In 4.2, I argued that the nomic objection is weak because three very different and influential accounts of the metaphysics of laws seem equally well suited to serve an idealist ontology.

All in all, the dialectic of this chapter warrants the following conclusion: Idealist ontologies can explain the structure of the physical world just as well as physical realist ontologies, provided that idealism is upheld in conjunction with theism. The widespread belief that idealism is at a serious disadvantage here might be due to the fact that phenomenalism was typically championed without theism in the 20th century.

Conclusion

This preceding chapters were offered in defense of the following argument:

- (I) If physical realism is true, then the particles and fields studied by fundamental physics are real.
- (II) If physical objects are real, they have real categorical properties.
- (III) The only categorical property of the particles and fields studied by fundamental physics is spatiotemporal position.
- (IV) Spatiotemporal position is not a real property.
- (V) Therefore, physical realism is false.

Chapter 1 clarified the concept of real and ideal objects. I defined ideal objects as objects whose existence presupposes that someone can observe them. I also introduced a theistic variant of this definition, and I showed how the definitions can be tweaked to make them immune to counterexamples.

Chapter 2 defended premise (II) by defending Howard Robinson's Power Regress argument, the gist of which is that worlds without categorical features give rise to unpleasant regresses. I looked at a stronger and a weaker version of this claim, and I gave a detailed defense of the weaker one.

Chapter 2 also defended premise (III) in the context of a critique of Howard Robinson's case for idealism. Robinson establishes the falsity of physical realism from the conjunction of the Power Regress with the claims that (i) all physical properties are dispositional and (ii) the physical world is closed under physics if physical realism is true. I claimed that (ii) is plausible but one can object to (i) on the grounds that spatiotemporal position does not seem to be a dispositional property. I pointed out, however, that premise (III), the thesis that all fundamental physical properties apart from spatiotemporal position are dispositional, is defensible. I also argued that even though physical realists may posit unknown categorical bases that underlie the dispositions known from physics, positing such bases is only warranted if one has independent reasons to prefer physical realism to idealism.

Chapter 2 also defended premise (IV) by reconstructing John Foster's two arguments against spatial realism. The upshot of Foster's modal argument was that physical realism entails that spacetime could have and could not have sustained different laws. The upshot of Foster's abductive argument was that the physically relevant structure of spacetime is the structure manifest in experience, hence, by inference to the best explanation, real spacetime can be thrown out of physical ontology. I argued that these arguments can only be effectively resisted if the physical realist has independent reasons to prefer physical realism to idealism.

Chapter 3 presented a new abductive argument for (IV). The argument was based on the puzzle of the conventionality of simultaneity, the gist of which is that special relativity seems compatible with the hypothesis that light does not travel with the same speed in all directions. I argued that this puzzle has no straightforward and uncontroversial physical realist solution but it is easily solved by idealism, hence the hypothesis that spacetime is ideal has explanatory value in physical ontology, and, unless one has strong independent reasons not to believe in idealism, inference to the best explanation can motivate the view that spacetime is ideal.

Finally, Chapter 4 considered two metaphysical reasons to deny idealism, the truthmaker objection, which says that the idealist cannot supply truthmakers for truths about the physical world, and the nomic objection, which says that the idealist cannot explain the existence of laws. I argued that the truthmaker objection is very hard to solve under non-theistic forms of idealism but can be solved straightforwardly if idealism is conjoined with theism. I also argued that the force of the nomic objection diminishes if we look at existing accounts of the metaphysics of laws, because the idealist can adapt them fairly easily to her own ontology. The only serious requirement seems to be to conjoin idealism with theism if nonHumean conceptions are adopted.

The upshot of the whole dialectic is that premises (II), (III), and (IV) are defensible if the idealist goes theist when defending (III) and (IV). Hence, if premise (I) is true, then my defense of idealism is known to be wrong iff it is known that God does not exist. Since nobody has established that God does not exist, I conclude that either premise (I) is false or my defense of idealism is not known to be wrong.

Premise (I) is false iff physical realism is true but the particles and fields studied in fundamental physics are not real. A proponent of this view claims that there are real physical objects but the particles and fields of fundamental physics are not among them. Presumably, a proponent of this view treats fundamental particles and fields as useful fictions that help us understand and manipulate the real physical objects out there.

How much of the scientific image can be salvaged for physical realism once this road is taken? Not much, it seems. If fundamental particles and fields are not real, then, presumably, neither are atoms and molecules, since atoms and molecules are composed of fundamental particles and are governed by fundamental fields. And if atoms and molecules are not real, neither are chemical processes. And if chemical processes are not real, neither are the neural processes in the brain and throughout the body. Once the physical realist denies (I), the scientific image starts to unravel and cannot be reified into the purported real structure of the world.

It follows that physical realists can only deny (I) at the price of giving up the right to build the ontology of the mental from the scientific image of the brain and the body. Consequently, resisting (I)–(V) by denying (I) forces the physical realist to abandon contemporary mainstream physicalist views of the mental.

The goal of this dissertation was twofold: (i) to formulate idealism in the context of analytic metaphysics, and (ii) to see if idealism can challenge contemporary mainstream physicalism about the mental. Chapter 1 was meant to achieve goal (i), and the rest of the dissertation defended (I)–(V) to meet (ii). I hope to have amassed some evidence that idealism can be reformulated in contemporary terms and that the scientific image of man as a physical machine lacks a solid metaphysical foundation.

Appendix

Clock transport synchrony

Suppose we want to test the one-way speed of light using a transported clock. We set clocks *A* and *B* side by side at *P*, and we synchronize them locally. Then we transport clock *B* to $x = 1$, and, when *A* reads “0,” we send a light beam to *B*. Since the two clocks are synchronized, the readout on *B* when the beam arrives tells us the one-way velocity.

One important issue that one must take into account when thinking about this setup is that special relativity predicts that moving clocks will slow down. According to the predictions of special relativity (assuming uniform speed for light), if a clock is moved at v m/sec for t_A seconds (measured on a clock at relative rest), the time on the moving clock is given by

$$(1) \quad t_B = t_A \cdot \sqrt{1 - v^2/c^2}$$

For example, if clock *B* is locally synchronized with *A* and then transported at a velocity of $0.4c$ to its remote location, then *B* will show roughly 0.811 times the time on *A* upon arrival.

We must correct for this discrepancy before we begin the experiment, otherwise the two clocks will not be in synchrony once *B* is transported to its remote location. For example, if *B* is transported to $x = 1$ (taking 1 unit distance to be c meters) at $0.4c$ after the clocks are set to 0 locally at $x = 0$, then *A* will read 2.5 seconds and *B* will read roughly 2.0275 seconds when *B* arrives (assuming standard one-way speed). In order for the two clocks to stay in synchrony, *B* must be reset by adding 0.4725 seconds after arrival.

It is easy to anticipate how the dialectic develops from this point. Since the time on *B* must be corrected after arrival for the two clocks to stay in synchrony and the correction factor involves the velocity of light, we cannot re-synchronize *B* with *A* after *B* is transported without assuming something about the one-way velocity of light and thus begging the experimental question. Let’s see how this problem emerges in a simple 2D spacetime.

Suppose we synchronize A and B locally, we transport B at velocity v to a remote location, then we re-synchronize B with A using the standard formula mentioned above. As Winnie (1970) explains in his seminal paper on deviant models of special relativity, two things go wrong in this process if the speed of light is not uniform. First, the slowdown of B will given by

$$(2) \quad t_B = t_A \cdot \sqrt{1 - v^2/c^2} \cdot \frac{c}{c + v(2\varepsilon - 1)}$$

, where 2ε is the ratio between the average round-trip speed of light, c , and its speed in the direction in which B moves (Winnie 1970: 85, 89–91).

To simplify, let's express (2) as

$$(3) \quad t_B = t_A \cdot \sqrt{1 - v^2/c^2} \cdot \text{fudge factor}$$

This formula highlights the fact that in deviant models, the standard time dilation formula, (1), is modified by a factor that depends on the speed of light in the direction of B 's movement.

Deviant models also involve a difference between the real velocity of an object, v_R , relative to an observer O , and its velocity v relative to O measured under the assumption that the one-way speed of light is uniformly c . As Winnie (1970: 84–86) shows, the relation between these magnitudes will be

$$(4) \quad v_R = v \cdot \text{fudge factor}$$

, where the fudge factor is the same as in (3).

The reason why the real velocity relative to O will be different from the velocity measured under the standard assumption is that we can only measure velocities by sending and receiving signals that indicate locomotion. Consider two ways that an observer O can measure the speed of an object.

If O sets her clock to 0 when an object X passes her, and then she notes that at T , a signal was received from X indicating that X was at a distance of n meters, then O can only use these data to determine X 's relative speed if she knows how long it took the signal from X to reach her. The value that O calculates for X 's speed will be different under different assumptions about the one-way speed of light.

Alternatively, if O tries to measure the relative speed of X by attaching a speedometer to X , then the speed of X will be measured by a moving clock, and in order to get the speed of X relative to O 's frame from the speedometer reading on X , O will have to relate the readings on the moving clock to the readings on her own stationary clock. And we know from (3) that moving clocks will be asynchronous in a slightly different way in deviant models than in the standard one. Consequently, relative speeds will also be different in deviant models than in the standard one, because the time on moving clocks will change in a slightly different way.

Now suppose that we live in a deviant world where the speed of light is not uniform but we are assuming that it is. We synchronize two clocks, A and B locally at P , and then we transport B to 1 distance unit away at what we take to be velocity v . We assume that the time on A , t_A , will equal $1/v$ seconds when B reaches its destination. We assume uniform one-way speed for light, so we use the standard time dilation formula (1), $t_A \cdot \sqrt{1 - v^2/c^2}$, to re-synchronize clock B after it arrives at its remote location.

Let's see what happens in reality. In reality, the velocity of B is different than we think it is. Its real speed, v_R , is given by (4). Since the real speed of B is v_R , clock B reaches its destination not when clock A reads $1/v$ (as we think) but when clock A reads $1/v_R$, which is a different number than $1/v$. Let's distinguish these two values:

$$(5) \quad t_A^O = \frac{1}{v} \quad \text{(what the observer thinks the time on clock } A \text{ is when } B \text{ arrives at its remote location)}$$

$$(6) \quad t_A^R = \frac{1}{v_R} \quad \text{(the real time on clock } A \text{ when } B \text{ arrives at its remote location)}$$

We also know from (3) that in reality, the slowdown on B will be different than in the standard case. The readout on clock B upon arrival will be

$$\begin{aligned} (7) \quad t_B^R &= t_A^R \cdot \sqrt{1 - v^2/c^2} \cdot \text{fudge factor} = && \text{from (2)} \\ &= \frac{1}{v_R} \cdot \sqrt{1 - v^2/c^2} \cdot \text{fudge factor} = && \text{from (6)} \\ &= \frac{1}{v \cdot \text{fudge factor}} \cdot \sqrt{1 - v^2/c^2} \cdot \text{fudge factor} = && \text{from (4)} \\ &= \frac{1}{v} \cdot \sqrt{1 - v^2/c^2} = \\ &= t_A^O \cdot \sqrt{1 - v^2/c^2} && \text{from (5)} \end{aligned}$$

It follows that the way clock B becomes asynchronous in deviant scenarios will not tell the observers that the assumption of uniform velocity is wrong.

Suppose that we re-synchronize B with A after B has arrived at its remote location. We do this by setting B to t_A^O immediately after arrival, since we think that clock A reads t_A^O at that moment. In reality, A reads t_A^R at the moment in question, so B will run late or early compared to A , by the following amount:

$$\begin{aligned} (8) \quad t_A^O - t_A^R &= \frac{1}{v} - \frac{1}{v_R} = \frac{1}{v} - \frac{1}{v \cdot \text{fudge factor}} = \\ &= \frac{1}{v} - \frac{1}{v} \frac{c + v(2\varepsilon - 1)}{c} = \frac{2\varepsilon - 1}{c} \end{aligned}$$

As indicated under (2), 2ε is the ratio between the average round-trip speed of light, c , and its real speed in the direction in which B was transported. It follows that in reality, light beams travel at a speed of $c / 2\varepsilon$ from A to B . Since B is at unit distance from A , it takes light $2\varepsilon / c$ seconds to go from A to B .

Suppose that, once we re-synchronized B with A , a light beam is emitted from A at $t_A = 0$. We know from (8) that when this happens, B reads $(1 - 2\varepsilon) / c$. It takes light $2\varepsilon / c$ seconds to go from A to B , so the readout on B at the moment of the light beam's arrival will be

$$(9) \quad \frac{1 - 2\varepsilon}{c} + \frac{2\varepsilon}{c} = \frac{1}{c}$$

It follows that the readout on B will confirm the erroneous hypothesis that the one-way speed is c . Consequently, locally synchronizing and transporting clocks cannot help us measure the one-way speed of light.

References

Works by Berkeley

Quoted by volume and page number in the Luce/Jessop edition (9 vols., Thomas Nelson, 1948–57). Section numbers are also given if available.

Adams, Robert M. (1994): *Leibniz: Determinist, Theist, Idealist*. Oxford University Press.

—— (2007): Idealism vindicated. In P. van Inwagen and D. Zimmerman (eds.), *Persons: Human and Divine*, Clarendon Press, 35–54.

Ameriks, Karl (1991): Hegel and idealism. *The Monist* 74, 386–402.

Armstrong, David M. (1983): *What is a Law of Nature?* Cambridge University Press.

—— (1997): *A World of States of Affairs*. Cambridge University Press.

—— (2004): *Truth and Truthmakers*. Cambridge University Press.

Babovic, V. M., D. M. Davidovic, and B. A. Anicin (1991): The Doppler interpretation of Roemer's method. *American Journal of Physics* 59, 515–19.

Bartels, Andreas (2013): Why metrical properties are not powers. *Synthese* 190, 2001–13.

Beiser, Frederick C. (2002): *German Idealism*. Harvard University Press.

Ben-Menahem, Yemima (2006): *Conventionalism*. Cambridge University Press.

Bennett, Jonathan (2003): *A Philosophical Guide to Conditionals*. Clarendon Press.

Ben-Yami, Hanoch (2006): Causality and temporal order in special relativity. *The British Journal for the Philosophy of Science* 57:3, 459–79.

—— (ms): Observed simultaneity.

Bird, Alexander (1998): Dispositions and antidotes. *Philosophical Quarterly* 48, 227–34.

—— (2007): *Nature's Metaphysics*. Clarendon Press.

Blackburn, Simon (1990): Filling in space. *Analysis* 50: 2, 62–5.

Bol, K. (1950): A determination of the speed of light by resonant cavity method. *Physical Review* 80, 298.

BonJour, Laurence (2011): Epistemological problems of perception. *Stanford Encyclopedia of Philosophy*.

Braibant, S., G. Giacomelli, and M. Spurio (2009): *Particles and Fundamental Interactions*. Springer.

- Brandom, Robert (2009): Animating ideas of idealism: a semantic sonata in Kant and Hegel. In his *Reason in Philosophy* (Belknap Press), 27–110.
- Burniston Brown, G. (1967): What is wrong with relativity? *Bulletin of the Institute of Physics* 18, 71–7.
- Burnyeat, Myles F. (1982): Idealism and Greek philosophy. *The Philosophical Review* 91: 1, 3–40.
- Carroll, John (1994): *Laws of Nature*. Cambridge University Press.
- Chaichian, M., H. Perez Rojas, and A. Tureanu (2014): *Basic Concepts in Physics*. Springer.
- Choi, Sungho (2008): Dispositional properties and counterfactual conditionals. *Mind* 117, 795–841.
- Clifton, Robert K. (1989): Some recent controversy over the possibility of experimentally determining isotropy in the speed of light. *Philosophy of Science* 56, 688–96.
- Correia, Fabrice and Benjamin Schnieder (2012): Grounding: an opinionated introduction. In their (ed.) *Metaphysical Grounding*, Cambridge University Press, 1–36.
- Dainton, Barry (2001): *Time and Space*. Acumen.
- Daniel, Stephen H. (2001): Berkeley's Christian Neoplatonism, archetypes, and divine ideas. *Journal of the History of Philosophy*, 39: 2, 239–258.
- Descartes, René (1968): *Discourse on the Method and Meditations*. Penguin Books.
- Dretske, Fred (1977): Laws of nature. *Philosophy of Science* 44, 248–68.
- Eagle, Albert (1938): A criticism of the special theory of relativity. *Philosophical Magazine* 26, 410–14.
- Edgington, Dorothy (2006): Conditionals. *Stanford Encyclopedia of Philosophy*.
- Edwards, W. F. (1963): Special relativity in anisotropic space. *American Journal of Physics* 31, 482–9.
- Einstein, Albert (1923 [1905]): On the electrodynamics of moving bodies. In H.A. Lorentz, A. Einstein, H. Minkowski, *The Principle of Relativity* (Dover Publications), 35–65.
- (1993): The principle of relativity and its consequences in modern physics. In his *Collected Papers*, vol. 3 (Princeton University Press), 131–74.
- (2005 [1917]): *Relativity: The Special and General Theory*. Pi Press.
- Ellis, Brian (2001): *Scientific Essentialism*. Cambridge University Press.
- Ellis, Brian and Peter Bowman (1967): Conventionality in distant simultaneity. *Philosophy of Science* 34, 116–36.
- Ellis, Brian and Caroline Lierse (1994): Dispositional essentialism. *Australasian Journal of Philosophy* 72: 1, 27–45.

- Esfeld, Michael (2009): The modal nature of structures in ontic structural realism. *International Studies in the Philosophy of Science* 23, 179–94.
- (2010): Humean metaphysics versus a metaphysics of powers. In G. Ernst and A. Hüttemann (eds.), *Time, chance and reduction*, Cambridge University Press, 119–35.
- Essen, L. and A. C. Gordon-Smith (1948): The velocity of propagation of electromagnetic waves derived from resonant frequencies of a cylindrical resonator. *Proceedings of the Royal Society A* 194, 348–361 (1948).
- Ewing, A.C. (1934): *Idealism: A Critical Survey*. Methuen & Co. Ltd.
- Feenberg, Eugene (1974): Conventionality in distant simultaneity. *Foundations of Physics* 4, 121–26.
- Feinberg, G. (1967): Possibility of faster-than-light particles. *Physical Review* 159: 1089–1105.
- Fine, Kit (2001): The question of realism. *Philosopher's Imprint* 1, 1–30.
- Fodor, Jerry A. (1998): *Concepts: Where Cognitive Science Went Wrong*. Clarendon Press.
- Foster, John (1982): *The Case for Idealism*. Routledge & Kegan Paul.
- (2003): *The Divine Lawmaker*. Clarendon Press.
- Friedman, Michael (1977): Simultaneity in Newtonian mechanics and special relativity. *Minnesota Studies in the Philosophy of Science* 8, 403–32.
- Froome, K. D. and L. Essen (1969): *The Velocity of Light*, Academic Press.
- Gosse, Philip Henry (1857): *Omphalos*. John van Voorst.
- Grice, H. P. (1989): *Studies in the Way of Words*. Harvard University Press
- Grünbaum, A. (1956): Relativity of simultaneity within a single Galilean frame: a rejoinder. *American Journal of Physics* 24, 588–90.
- Grünbaum, Adolf (1955): Logical and philosophical foundations of the special theory of relativity. *American Journal of Physics* 23, 450–64.
- (1963): *Philosophical Problems of Space and Time*. Alfred A. Knopf.
- (2010): David Malament and the conventionality of simultaneity: a reply. *Foundations of Physics* 40, 1285–97.
- Grünbaum, Adolf, Wesley C. Salmon, Bas C. van Fraassen, and Allen I. Janis (1969): A panel discussion of simultaneity by slow clock transport in the special and general theories of relativity. *Philosophy of Science* 36, 1–105.
- Gunn, D. and I. Vetharanim (1995): Relativistic quantum mechanics and the conventionality of simultaneity. *Philosophy of Science* 62, 599–608.
- Hartz, Glenn A. (1992): Leibniz's phenomenalisms. *The Philosophical Review* 101: 3, 511–49.
- Hawthorne, John (2001): Causal structuralism. *Philosophical Perspectives* 15, 361–78.

- Jackson, Frank (1987): *Conditionals*. Basil Blackwell.
- (1998). *From Metaphysics to Ethics*. Oxford University Press.
- Jackson, Frank and Robert Pargetter (1977): Relative simultaneity in the special relativity. *Philosophy of Science* 44, 464–74.
- Jammer, Max (2006): *Concepts of Simultaneity*. The John Hopkins University Press.
- Janis, Allen (1983): Simultaneity and conventionality. In R. Cohen and L. Laudan (eds.), *Physics, Philosophy and Psychoanalysis*, D. Reidel, 101–10.
- (2010): The conventionality of simultaneity. *Stanford Encyclopedia of Philosophy*.
- Karlov, Leo (1970): Does Roemer's method yield a unidirectional speed of light? *Australian Journal of Physics* 23, 243–58.
- Lange, Marc (2009): *Laws and Lawmakers*. Oxford University Press
- Langton, Rae (1998): *Kantian Humility*. Clarendon Press.
- Leibniz, G.W. (1989): *Philosophical Essays* (ed. R. Ariew and D. Garber). Hackett.
- Lewis, David (1973): *Counterfactuals*. Blackwell.
- (1986): Causation. In his *Philosophical Papers II*, Oxford University Press, 159–72.
- (1994): Humean supervenience debugged. *Mind* 103, 473–490.
- (2000): Causation as influence. *Journal of Philosophy* 97:4, 182–97.
- Liebowitz, B. (1956): "Relativity of simultaneity in a single Galilean frame. *American Journal of Physics* 24, 587–88.
- Livanios, Vassilios (2008): Bird and the dispositional essentialist account of spatiotemporal relations. *Journal for General Philosophy of Science* 39: 2, 383–94.
- Loewer, Barry (2004): Humean Supervenience. In J. Carroll (ed), *Readings on Laws of Nature*, University of Pittsburgh Press, 176–206.
- Lowe, E. J. (1998): *The Possibility of Metaphysics*. Clarendon Press.
- (2010): On the individuation of powers. In A. Marmodoro (ed.), *The Metaphysics of Powers*, Routledge, 8–26.
- Luce, A. A. (1954): *Sense without Matter*. Nelson.
- Macdonald, A. (2005): Comment on 'The role of dynamics in the synchronization problem,' by Hans C. Ohanian. *American Journal of Physics* 73: 454–55.
- Mackie, J. L. (1973): Dispositions and powers. In his *Truth, Probability, and Paradox*, Clarendon Press, 120–53.
- Malament, David (1977): Causal theories of time and the conventionality of simultaneity. *Nous* 11, 293–308.
- Marinov, S. (1974): The velocity of light is direction-dependent. *Czechoslovakian Journal of Physics B* 24, 965–70.

- Marmodoro, Anna (2009): Do powers need powers to make them powerful? *History of Philosophy Quarterly* 26: 4, 337–52.
- (forthcoming): Aristotelian powers at work. In J. Jacobs (ed.), *Putting Powers to Work*, Oxford University Press.
- Martin, C.B. (1994): Dispositions and conditionals. *The Journal of Philosophy* 94, 187–211.
- Maudlin, Tim (2007): *The Metaphysics Within Physics*. Oxford University Press.
- (2012): *Philosophy of Physics: Space and Time*. Princeton University Press.
- McLaughlin, Brian P. and Karen Bennett (2011): Supervenience. *Stanford Encyclopedia of Philosophy*.
- Merricks, Trenton (2007): *Truth and Ontology*. Clarendon Press.
- Minguzzi, E. (2002): On the conventionality of simultaneity. *Foundations of Physics Letters* 15:2, 153–69.
- Mittelstaedt, Peter (1977): Conventionalism in special relativity. *Foundations of Physics* 7, 573–83.
- Møller, C. (1962): New experimental tests of the special principle of relativity. *Proceedings of the Royal Society of London A* 270, 306–11.
- Molnar, George (2003): *Powers*. Oxford University Press.
- Mumford, Stephen (2004): *Laws in Nature*. Routledge.
- Northrop, F. S. C.: Whitehead's philosophy of science. in P. A. Schilpp (ed.), *The Philosophy of Alfred North Whitehead*, Northwestern University Press, 165–208.
- Norton, John (1986): The quest for the one way velocity of light. *The British Journal for the Philosophy of Science* 37: 1, 118–20.
- Ohanian, H. (2004): The role of dynamics in the synchronization problem. *American Journal of Physics* 72, 141–8.
- Øhrstrøm, Peter (1980): Conventionality in distant simultaneity. *Foundations of Physics* 10, 333–43.
- Petkov, Vesselin (1989): Simultaneity, conventionality and existence. *British Journal for the Philosophy of Science* 40, 69–76.
- (2009): *Relativity and the Nature of Spacetime*. Springer.
- Pirani, F. (1970): Noncausal behavior of classical tachyons. *Physical Review D* 1, 3224–5.
- Podlaha, M. P. (1980): On the impossibility to measure one-way velocity of light, *Lettere al Nuovo Cimento* 28, 216–21.
- Price, H. H. (1950): *Perception* (2nd ed.). Methuen.
- Putnam, Hilary (1967): Time and physical geometry. *Journal of Philosophy* 64: 240–47.
- Putnam, Hilary (1978): *Meaning and the Moral Sciences*. Routledge & Kegan Paul.

- Quinton, A. M. (1972): *Absolute Idealism*. Oxford University Press.
- Reichenbach, Hans: *The Philosophy of Space and Time*. Dover Publications.
- Robinson, Howard (1982): *Matter and Sense*. Cambridge University Press.
- , (2009): Idealism. In B. P. McLaughlin, A. Beckermann, S. Walter (eds.), *The Oxford Handbook of Philosophy of Mind*, Clarendon Press, 189–204.
- Ruderfer M. (1960): First-order terrestrial ether drift experiment using the Mössbauer radiation. *Physical Review Letters* 5, 191–2.
- Ruebenbauer, K. (1980): Isotropy of the velocity of light. *International Journal of Theoretical Physics* 19, 217–9.
- Salmon, Wesley C. (1977): The philosophical significance of the one-way speed of light. *Nous* 11, 258–92.
- Sarkar, S. and J. Stachel (1999): Did Malament prove the non-conventionality of simultaneity in the special theory of relativity? *Philosophy of Science* 66, 208–20.
- Schaffer, Jonathan (2009): On what grounds what. In D. Chalmers, D. Manley, R. Wasserman (eds.), *Metametaphysics*, Clarendon Press, 347–83.
- Scott-Iversen, Paul A. (1944): Introductory notes on a reformulation of the special theory of relativity. *Philosophical Magazine* 35, 105–20.
- Sellars, Wilfrid (1963): Phenomenalism. In his *Science, Perception and Reality*, Routledge & Kegan Paul, 60–105.
- Shea, J. H. (1998): Ole Roemer, the speed of light, the apparent period of Io, the Doppler effect, and the dynamics of Earth and Jupiter. *American Journal of Physics* 66, 561–569.
- Shoemaker, Sydney (1980): Causality and properties. In P. van Inwagen (ed.), *Time and Cause*, Reidel, 109–35.
- Sjödin, T. (1979): Synchronization in special relativity and related theories. *Nuovo Cimento* 51B, 229–46.
- Sprigge, Timothy (1983): *The Vindication of Absolute Idealism*. University of Edinburgh Press.
- Sprigge, Timothy L. S. (1983): *The Vindication of Absolute Idealism*. University of Edinburgh Press.
- (1993): *James & Bradley*. Open Court.
- (2006): *The God of Metaphysics*. Oxford University Press.
- Stern, Robert (2009): Hegel's idealism. In his *Hegelian Metaphysics* (Oxford University Press), 45–76
- Stolakis, George (1986): Against conventionalism in physics. *The British Journal for the Philosophy of Science* 37, 229–32.
- Tola, Fernando and Carmen Dragonetti (2004): *Being as Consciousness: Yogācāra Philosophy of Buddhism*. Motilal Banarsidass.
- Tooley, Michael (1977): The nature of laws. *Journal of Philosophy* 7, 667–98.

- Torretti, Roberto (1979): Jackson and Pargetter criterion of distant simultaneity. *Philosophy of Science* 46, 302–5.
- Townsend, Burke (1980): Jackson and Pargetter on distant simultaneity. *Philosophy of Science* 47, 646–55.
- Ungar, Abraham (1986): The Lorentz transformation group of the special theory of relativity without Einstein's isotropy convention. *Philosophy of Science* 53, 395–402.
- van Cleve, James (1999): *Problems from Kant*. Oxford University Press.
- van Fraassen, Bas C. (1987): Armstrong on laws and probabilities. *Australasian Journal of Philosophy* 65:3, 243–60.
- van Inwagen, Peter (2006): *The Problem of Evil*. Oxford University Press.
- Weingard, R. (1972): Relativity and the reality of past and future events. *British Journal for the Philosophy of Science* 23, 119–21.
- (1985): Remark on 'The special theory of relativity and the one-way speed of light.' *American Journal of Physics* 53, 492.
- Winkler, Kenneth P. (1989): *Berkeley: An Interpretation*. Clarendon Press.
- Winnie, John A. (1970): Special relativity without one-way velocity assumptions. *Philosophy of Science* 37, 81–99, 223–238.
- Yates, David (2013): The essence of dispositional essentialism. *Philosophy and Phenomenological Research* 87: 1, 93–128.
- Zangari, M. (1994): A new twist in the conventionality of simultaneity debate. *Philosophy of Science* 61, 267–75.

Long contents

Introduction	1
1 The concept of idealism	4
1.1 Problems about defining idealism	5
1.2 Introducing a definition	8
1.2.1 The basic criterion	10
1.2.2 Historical comparisons	10
1.2.3 Clash with physical realism	12
1.2.4 The story so far	13
1.2.5 Too thin?	14
1.2.6 Trivially false?	14
1.3 Complications	16
1.3.1 Self-observation	17
1.3.2 Worldbound objects	17
1.3.3 Essentiality of origin	18
1.3.4 Psychophysical laws	19
1.3.5 Thoroughly psychofunctional matter	20
1.3.6 Alien observers	21
1.3.7 Summary of 1.3.1–1.3.6	22
1.3.8 The fine-tuned definition of ideality	23
1.3.9 P.S. Idealism about abstracta	24
1.4 The problem of grounding	25
1.5 Summary of Chapter 1	28

2	The ideality of matter and space:	
	Three contemporary arguments	29
2.1	The Power Regress	30
2.1.1	Infinite C-trees	32
2.1.2	Finite nomic sequences	36
2.1.3	Infinite epistemic descent	38
2.1.4	Objections	41
2.2	From the Power Regress to idealism	43
2.2.1	Are fundamental physical properties dispositions?	45
2.2.2	The problem of spatiotemporal position	47
2.2.3	Unknown categorical bases	48
2.2.4	Summary of 2.1 and 2.2	50
2.3	Scrambled and gappy worlds	51
2.3.1	The modal argument against spatial realism	51
2.3.2	The abductive argument against spatial realism	56
2.3.3	Summary of 2.3	60
2.4	Summary of Chapter 2	61
3	Real spacetime as excess structure	64
3.1	Albert the ant	64
3.2	The conventionality of simultaneity	70
3.2.1	Measuring the one-way velocity of light	70
3.2.2	Consequences of deviance	76
3.3	Physical realist solutions	79
3.3.1	Experimental tests	79
3.3.2	Appeals to simplicity	82
3.3.3	Mathematical arguments	84
3.3.4	Eternalism	85
3.3.5	Gauge freedom	86
3.3.6	Summary of 3.3	90

3.4 The idealist solution	92
3.4.1 Sketch of the solution	92
3.4.2 Elaborating the solution	92
3.4.3 Two objections	97
3.4.4 The superiority of the idealist solution	100
4 Truthmakers and laws in idealism	102
4.1 Truthmakers	103
4.1.1 Truthmakers for prehistory	104
4.1.2 Truthmakers for the present	108
4.1.3 Summary of 4.1	110
4.2 Laws	110
4.2.1 Humean laws	113
4.2.2 Laws as second-order universals	117
4.2.3 Laws from powers	119
4.2.4 Summary of 4.2	122
4.3 Summary of Chapter 4	122
Conclusion	123
Appendix:	
Clock transport synchrony	126
References	130

This dissertation contains no material accepted for any other degrees in any other institutions. Neither does it contain material written and/or published by another person, except where noted in the form of bibliographical references.

Kodaj Dániel

7 July 2014