



What are empirical consequences? On dispensability and composite objects

Alex LeBrun¹

Received: 10 March 2021 / Accepted: 17 August 2021
© The Author(s) 2021

Abstract

Philosophers sometimes give arguments that presuppose the following principle: two theories can fail to be empirically equivalent on the sole basis that they present different “thick” metaphysical pictures of the world. Recently, a version of this principle has been invoked to respond to the argument that composite objects are dispensable to our best scientific theories. This response claims that our empirical evidence distinguishes between ordinary and composite-free theories, and it empirically favors the ordinary ones (Hofweber 2016, 2018). In this paper, I ask whether this response to the dispensability argument is tenable. I claim that it is not. This is because it presupposes an indefensible thesis about when two empirical consequences are distinct or the same. My argument provides some insight into what our empirical consequences are, and I conclude that empirical evidence is radically metaphysically neutral. This gives us some insight into the significant content of our scientific theories—the content that a scientific realist is committed to—and I show how this insight relates to questions about theoretical equivalence more broadly.

Keywords Composition · Dispensability · Empirical equivalence · Theoretical equivalence

1 Introduction

In Lorentz’s ether theory, ether was postulated as a substance that acted as the medium for the transmission of light through space. Philosophers of science tell the following story for why physicists no longer accept the existence of ether (see, e.g., Norton, 2008).

✉ Alex LeBrun
lebrun@ucsb.edu

¹ University of California, Santa Barbara, USA

Lorentz's ether theory had many predictions and observations, including length contraction, which is the phenomenon that a moving object's measured length will be shorter than its proper length. Einstein proposed an alternative theory, special relativity, that has the exact same predictions and observations (Bradley, 2020, p. 9), including length contraction, without needing to posit the existence of ether. Special relativity thus showed physicists that ether is dispensable.¹

Einstein showed that ether is dispensable to theories of light by providing an attractive alternative theory which (i) does not appeal to ether and (ii) is empirically equivalent to Lorentz's. Physicists took the dispensability of ether as good reason to abandon ontological commitment to it. Today, philosophers of science and metaphysicians operate on the same understanding of dispensability. Some claim that numbers are *indispensable* to scientific theories by arguing that one cannot provide nominalistic alternatives that meet these two conditions (cf. Colyvan, 2001; Field, 2016). Philosophers generally take the dispensability of an entity to have ontological consequences; if an entity is dispensable, we should abandon ontological commitment to it.

One argument within the metaphysics of ordinary objects is best understood as a dispensability argument. This Composite Object Dispensability Argument (CODA) concludes that ordinary composite objects like metal bars are dispensable to our best scientific theories (cf. Rosen & Dorr, 2002; Sider, 2013; Brenner, 2018, p. 660). The idea is that the only things needed to explain phenomena like conduction are the microphysical particles that "make up" the metal bar. The CODA offers a strategy for constructing a variant of any scientific theory, and these variants supposedly meet the criteria for dispensing with composites. They are meant to be attractive theories that (i) do not appeal to composites and (ii) are empirically equivalent to the ordinary theories that do appeal to composite objects. Proponents of this dispensability argument take themselves to have shown that composites are dispensable to any scientific theory which they appear, and this is meant to be evidence that there are no composite objects like metal bars.

Here I am concerned with empirical equivalence and its relation to the CODA. Two theories are empirically equivalent in virtue of sharing the same empirical evidence or content, which is understood as having the same empirical consequences.² According to the CODA, the *composite-free* theories have the same empirical consequences as the *ordinary* theories that appeal to composites; our empirical evidence is neutral between them. For example, an ordinary theory might have the empirical consequence that there is a metal bar in the lab. A composite-free alternative would have the empirical

¹ It is common for philosophers of science to tell a similar story for why physicists no longer believe in absolute space. See Friedman (1983, p. 112). Additionally, absolute space and ether are taken to play roughly the same role in Lorentz's ether theory—as providing a privileged inertial frame.

² Here we are presupposing some rough distinction between the empirical and non-empirical. Such a presupposition raises questions and concerns about the theory-ladenness of observation (see Fodor, 1984). There are difficult questions about any particular distinction between the empirical or non-empirical, and whether consequences like *There is an electron in the bubble chamber* are empirical. Here, we work with an intuitive distinction between theory and observation, and we rest easy knowing there are difficult boundary cases. Unlike the logical positivists, we are not drawing the boundary between meaningfulness and nonsense, and so the question of whether some particular consequence is empirical or not is less pressing than it was for them. See Lewis (1988, p. 4) for a similar motivation.

consequence, roughly, that some microphysical particles “arranged metal bar-wise” are in some particular place. These two empirical consequences are taken to be the same.

Some, though, disagree with the CODA’s claim of empirical equivalence. For example, Hofweber claims that composite-free theories are *trivially* empirically inequivalent to their ordinary counterparts. He says,

There is lots of evidence that supports the [composite] object theory over the things arranged object-wise theory. The object theory predicts that there is a bar of metal in the lab, the object-wise theory doesn’t predict it. That there is such a bar can be confirmed with the observation that there is such a bar of metal in the lab. (Hofweber, 2016, p. 199)

Hofweber is claiming that the empirical evidence for our scientific theories—the observations and predictions—is “clearly in favour of [the existence of composite] objects” (Hofweber, 2018, pp. 321–322). The claim is not merely that we have more or better *scientific* reasons to prefer ordinary theories, but rather that we have better *empirical* reasons to prefer ordinary theories. He is clear that “[empirical] scientific evidence does in fact distinguish” between composite-free and ordinary theories, and it favors the ordinary ones (*ibid.*). Hofweber here presupposes a thesis about the individuation conditions of empirical consequences. In particular, he is committed to the thesis that differences in the “thick” mereological content between two empirical consequences suffices for a difference between those empirical consequences. As a result, if two theories’ empirical consequences differ in their “thick” mereological content, then they are empirically inequivalent. Otherwise it could not be that our empirical evidence supports the ordinary theory over the composite-free one. In this way, Hofweber is committed to the empirical significance of “thick” mereological content; he is presupposing that “thick” content *matters* empirically. Accordingly, the empirical consequences of a composite-free theory are trivially inequivalent to the empirical consequences of our ordinary theories. Call this the trivial response to the CODA.

At this point, the dialectic is brought to a halt. The CODA claims that composite-free theories are clearly empirically equivalent to their ordinary counterparts. The trivial response claims that composite-free theories are trivially empirically *inequivalent* to their ordinary counterparts. Without some clear understanding of the individuation conditions for empirical consequences, we cannot adjudicate this disagreement.

My topic is the individuation conditions of empirical consequences. My proximate aim is to settle whether the trivial response to the CODA is tenable. As we will see, my ultimate target is anyone who claims that “thick” content makes an *empirical* difference. I will argue that there is no good conception of empirical consequences that will permit the trivial response to the CODA. My argument proceeds by considering successful cases of dispensing, like when we rid our physics of ether. My thesis supports the position I call empirical quietism, which entails that we cannot settle any distinctively metaphysical disputes by appealing to empirical evidence. Although quietism may seem obvious to some, there are two additional philosophical payoffs that the following discussion yields.

First, because empirical equivalence plays a prominent role in many ontological arguments, it is imperative that we understand the conditions under which two theories are empirically equivalent. For any pair of empirically equivalent theories, there are at least pressing three ontological questions. First is the underdetermination question, which asks whether we should be committed to either theory's ontology; it seems that the existence of an empirically equivalent alternative should threaten to undermine our confidence in the theory we accept (cf. van Fraassen, 1980; Laudan, 1990; Stanford, 2009; Worrall, 2011). Second is the theoretical equivalence question, which asks whether the two theories are fully equivalent; it may be that two theories seem to have different ontological commitments, but that this is a merely apparent difference (cf. North, 2009; Curiel, 2014; Barrett, 2015, 2019, 2020; Weatherall, 2019a). Third is the dispensability question, which asks whether we have reasons to prefer one theory's ontology to another; it may be that we have reasons similar to parsimony to accept one of two empirically equivalent theories. For philosophers of science and metaphysicians, determining exactly when two theories have the same empirical consequences is important for the role that empirical equivalence plays in these arguments, and my conclusion will entail that some conceptions of empirical equivalence—ones antagonistic to quietism—should be rejected.

Second, my discussion will allow us to draw broader lessons on theoretical equivalence. The literature on theoretical equivalence is concerned with the conditions under which theories are fully equivalent, in the sense of saying the same thing about the world. Some people in that literature endorse a position adjacent to Hofweber's. Whereas Hofweber argues that an empirical consequence's "thick" content is relevant for individuation, these folks argue that a *non-empirical* consequence's "thick" content is likewise relevant for individuation. For example, North (2009) argues from ostensible differences in the structure of two formulations of classical mechanics—a difference in "thick" content—to the inequivalence of those formulations. If my arguments against those who ascribe empirical significance to "thick" content are correct, we should tread lightly. If we will have learned anything, it is that two theories having the same or different consequences is a complicated matter, not to be decided by only considering metaphysically rich content.

2 Preliminaries

The trivial response claims that composite-free theories are trivially empirically inequivalent to ordinary scientific theories, and it rejects the CODA on those grounds. We begin by investigating exactly what the trivial objection is committed to. The CODA, as an argument in its own right, is rarely discussed in the literature. It is usually implied by the claim that appeal to ordinary objects is "pragmatic" (cf. Healey, 2013, p. 53; Brenner, 2018, p. 660³). I hope the following explanation of the CODA shows that there is philosophical value in pursuing it explicitly.

³ Healey (2013) does not explicitly endorse the conclusion of the CODA—that we ought to reject the existence of composite objects—but he does accept that what scientists regard as composed is partially determined by the context in which the scientist is operating.

2.1 The CODA

Dispensability arguments are given by those who draw ontological commitments from an entity's dispensability. Consider the following dispensability principle: If some entity is dispensable to our best scientific theories, then we ought not be committed to its existence. Scientific realists may be inclined to accept the dispensability principle if they believe that science is the best guide to answering ontological questions. This dispensability principle seems to be what drove physicists to abandon commitment to ether.

To show that an entity is dispensable to some scientific theory, we must provide an alternative theory that dispenses with that entity. A dispensing theory is one that fits the following account, adapted from Colyvan (2001, p. 77):

Dispensability. An entity (or structure) X is dispensable to a theory T if and only if T has an attractive variant T^- for which:

- (i) T^- does not appeal to X s, and
- (ii) T^- has the same empirical consequences as T .

If there's a theory that appeals to some entity (or structure⁴) we suspect to be dispensable, to show its dispensability we provide an attractive theory that does not appeal to the entity (or structure) and has the same empirical consequences as the original. (i) requires that the variant does not appeal to the dispensable entity; if it does, we haven't shown that the entity is unnecessary. (ii) requires that the variant has the same empirical consequences; if it doesn't, then that suggests that the entity does play an explanatory role.⁵ These conditions are each necessary and jointly sufficient for showing an entity to be dispensable.

Empirical consequences, intuitively, are the observations and predictions of a theory. Theories make observations and predictions about the world; they tell us what it is like and what it will be like. Sameness of empirical consequences, or empirical equivalence, occurs when two theories make the same observations and proffer the same predictions about the world. For the time, we will operate on this intuitive notion of empirical consequences and empirical equivalence.

The CODA argues that all composite objects are dispensable to our best scientific theories. One might be sympathetic to the CODA because they consider appeal to composite objects to be merely pragmatic. Composites like iron bars may just be heuristics, allowing us to better understand complex scientific explanations but not serving a genuinely explanatory role in those explanations (Brenner, 2018, p. 660).⁶ Instead, the complex physical phenomena are fully explained by partless microphysical

⁴ Cf. North (2009, p. 64) and Barrett (2020, pp. 2–3).

⁵ What if T^- explained more than T did? Does the definition entail that T' does not dispense with the entity in question? Per the definition provided, it seems that the entity is not dispensable. Some may find this problematic, since it seems like we ought to prefer T^- to T . But note that we have reasons beyond the dispensability of the entity to prefer T^- —it explains more! So we have ordinary, empirical reasons to prefer T^- , rather than *a priori* reasons of dispensability.

⁶ Cf. Osborne (2016).

entities and processes.⁷ If this is persuasive, then it seems that ordinary scientific theories need not appeal to composite objects, since all we need to explain everything is microphysical. In accordance with the standard account of dispensability, the CODA provides variants for each scientific theory that appeals to composites. Instead of meticulously constructing these variants, the CODA offers a strategy to construct, for every ordinary scientific theory that appeals to composite objects, a composite-free variant that has the same empirical consequences. That strategy, which may be familiar, works like this.⁸

Consider the theory of rust R : rust, a reddish-brown substance that is the result of corrosion, is an iron oxide that forms on iron in the presence of oxygen together with water or air moisture. The iron is a reducing agent, giving up electrons, while the oxygen is an oxidizing agent, gaining electrons, resulting in iron oxide—rust. As an ordinary scientific theory, R appeals to composite objects; *rust* is an *iron oxide* that forms on *iron* in the presence of *oxygen* and *water*.

The CODA presents a strategy for constructing a variant of R that does not appeal to composites and has its same empirical consequences. Let's simplify and consider only one empirical consequence and try to rid R of just one composite object. Take the following observation delivered by the theory when some particular iron bar rusted after being exposed to moisture-rich air.

(b) This iron bar rusted.

(b) is an empirical consequence that appeals to a composite: this iron bar.

Proponents of the CODA think all genuine explanatory work is done by the partless microphysical particles that make up the “iron bar.” If they are correct, then to reveal the actual explanatory structure of our theories, we ought to replace all appeal to *iron bars* in R with appeal to only the partless, microphysical mereological *simples* and the complex ways in which they are arranged to be iron bar-wise.⁹ (We must also replace appeal to properties that are realized by composite objects with collective properties that are realized by arrangements. Though it is much more complicated than this, I will simply refer to the collective property variant of any ordinary property by appending it with the prime symbol. This way *rusting* becomes *rusting'*, where *rusting'* is realized by simples in arrangements.) This is how to construct a variant of R where appeal to iron bars is replaced with appeal to simples (and appeal to rusting is replaced with appeal to rusting'). Call this new theory R^- , which has the following empirical consequence:

(b^-) These simples arranged iron bar-wise rusted'.

The CODA claims that R^- is empirically equivalent to R and that (b) is the *same empirical consequence* as (b^-). The idea is this: When we observe what we'd ordinarily describe as an iron bar rusting, the two theories can equally sufficiently explain it. R will

⁷ A similar motivation is found in Sider (2007), where he argues that a composite object does not afford a thing with any causal powers beyond those had by the parts of that thing.

⁸ Here I follow Dorr (2002) and Rosen and Dorr (2002).

⁹ The ‘arranged X -wise’ locution is from Inwagen (1990). Here I assume both that the composite objects that science appeals to are not extended simples (Cf. McDaniel, 2007) and that the world is not *gunky* (Cf. Sider, 1993).

explain that the iron bar went through the process of rusting, and R^- will explain that the simples arranged iron bar-wise collectively went through the process of rusting'. Sure, the iron bar-free variant will be more difficult to comprehend, since it appeals to philosophical entities like simples and is cognitively cumbersome, but this is not a mark against its empirical adequacy.

Though the CODA claims that (b) and (b^-) are the same empirical consequence, we note that they have different “thick” metaphysical content. This notion of thick metaphysical content is to be understood as a consequence’s associated underlying metaphysical picture. (b) is associated with an underlying metaphysical picture where there are composites, and (b^-) is associated with one where there are no composites. According to the CODA, the mere fact that (b) and (b^-) have different thick metaphysical content is not sufficient for them being distinct empirical consequences. As we’ll see, Hofweber demurs.

R^- purportedly meets the conditions for showing iron bars to be dispensable. First, R^- does not appeal to iron bars. Second, if the CODA’s reasoning is correct, R^- is empirically equivalent to R . Given these facts about R^- together with the conditions for showing an entity to be dispensable, we have shown that iron bars are dispensable to theories of rust.¹⁰ Accordingly, the CODA concludes that we ought to reject the existence of iron bars. This is the general strategy for dispensing with composite objects that appear in different theories, which the CODA takes to show that all composite objects in all scientific theories are dispensable.

2.2 The empirical significance of thick content

Hofweber straightforwardly asserts that our empirical evidence decides in favor of the existence of composite objects. The idea seems to be that our theories’ empirical consequences come “pre-loaded” with a particular mereological picture and that this mereological picture, which presents the world as containing composite objects, is representationally significant. In brief, Hofweber is committed to the thesis that thick mereological content is empirically significant.

There are philosophers who are committed to similar theses. They presuppose or otherwise argue that some other thick metaphysical content is empirically significant. It would behoove us to see a few of these other arguments in metaphysics. Here’s a test for whether a philosopher is committed to the thesis that some thick content is empirically significant: if someone claims that a theory’s evidence is incompatible with one but not all sides of some distinctively metaphysical debate, then they believe that the metaphysical content of that debate, when it appears in theories, is empirically significant. If one was an evidential quietist, who does not think that thick content is

¹⁰ Some suggest that providing a dispensing theory also requires that the variant one provides is more attractive than the original theory. Colyvan (2001) and Field (2016) make these claims. If one accepts this, one might be tempted to reject the CODA on the grounds that the composite-free variant is not sufficiently attractive to show that composite objects are dispensable. There are rumblings of this response in the objects literature already. For example, Parsons (2013, p. 332) denies composite-free scientific theories because “composite objects play a crucial role in the best explanations of my experience.” I take this to mean that showing an entity to be dispensable requires offering an entity-free theory that *best* (or better) explains the phenomena, and not offering an entity-free theory that simply explains the phenomena.

empirically significant, they would deny that a theory's evidence could be compatible with one but not all sides of a distinctively metaphysical debate. Note that adherence to the empirical significance of some thick content (e.g., mereological) does not entail that one adheres to the empirical significance of any other thick content (e.g., identity). Here are three instances of philosophers who are committed to the empirical significance of thick metaphysical content.

First, in the literature on personal identity, Blatti (2012) argues from evolutionary theory to the thesis that human persons are identical to organisms. His idea is that any non-organism metaphysical position will be inconsistent with the empirical consequence of evolutionary biology that *my ancestor is an organism*. Blatti here straightforwardly presupposes that a theory's thick content regarding *identity* is empirically significant; our empirical evidence for the theory of evolution apparently settles the debate over personal identity. Second, Williamson (2007, p. 223) argues that scientific theories that are composite-free will not be supported by the same evidence that our current theories are; this is because the evidence for our current theories is committed to the existence of composite objects.¹¹ E.g., the evidence for *R* consists of claims like the *hygrometer* measured such-and-such humidity levels, which appeals to a composite object. Williamson is committed to the empirical significance of mereological content in a manner weaker than Hofweber's. He is presupposing that there is a *prima facie* evidential problem for composite-free scientific theories, whereas Hofweber argues that there is an open-and-shut evidential problem for composite-free theories. Finally, Lowe (2003, 2005) gives an argument against composite-free theories, where one interpretation of this argument is that such theories which do not appeal to properties like *mass* and *momentum* trivially cannot explain what ordinary physics theories explain because they do not appeal to the exact properties of *mass* and *momentum* as such. For Lowe, the structure of mass and momentum as such are empirically significant.

In each of these examples, it is natural to think that philosophers are committed to the empirical significance of some thick metaphysical content. They are giving arguments which proceed from considerations of empirical evidence to some conclusion about purely metaphysical matters.¹²

2.3 The trivial response

Let's return to the trivial response. Hofweber claims that our empirical evidence favors the existence of composite objects. Here we will examine how this claim entails a rejection of the CODA and unravel its commitments.

We focus on the claim of the CODA that ordinary and composite-free theories are empirically equivalent. The trivial response reasons as follows. Among the empirical consequences of *R* is that *this iron bar* rusted, whereas *R⁻* has no such empirical consequence. *R⁻* has the empirical consequence (*b⁻*):

(*b⁻*) These simples arranged iron bar-wise rusted'.

¹¹ Cf. Bagwell (2020).

¹² See Bailey and Brenner (2020) for additional, similar examples.

But this is not the *same* empirical consequence, according to the trivial response. Even if we can construct a composite-free variant in the way provided above, and bypass any other objections, Hofweber thinks the predictions and observations of the two theories are trivially different. This is because he claims that our empirical evidence favors a particular mereological picture. If our empirical evidence supports the existence of composite objects, then trivially any empirical evidence that does not support the existence of composite objects is not the *same* empirical evidence. For if our observations confirm (b) but not (b^-) , then (b) and (b^-) are distinct empirical consequences. The trivial response thus rejects the claims of the CODA that R and R^- are empirically equivalent. It is trivially impossible to provide an empirically equivalent, but composite-free, alternative to any ordinary scientific theory, says the trivial response.

Hofweber's trivial response requires a particular thesis about the individuation conditions of empirical consequences (as does any other variety of the empirical significance thesis like those outlined in Sect. 2.2). The trivial response succeeds only if we can empirically distinguish between R and R^- , and presupposes a position on how finely empirical consequences are individuated. In particular, the only difference Hofweber points to as a distinguishing feature between (b) and (b^-) is the thick metaphysical content—the underlying mereological picture associated with each empirical consequence.

Let us formally define this thesis on empirical consequence individuation. For any theories T_A and T_B , where T_A has the empirical consequences $(a_1), (a_2), \dots$ and T_B has the empirical consequences $(b_1), (b_2), \dots$, the trivial response is committed to the following:

Fine Grained If the underlying mereological picture associated with $(a_1), (a_2), \dots$ is different from the underlying mereological picture associated with $(b_1), (b_2), \dots$, then T_A and T_B are not empirically equivalent.

Fine Grained entails that empirical consequences may be individuated by the particular mereological pictures associated with those empirical consequences; in this sense, it is a *fine-grained* understanding of the individuation conditions of empirical consequences. (b) paints a picture where there is a composite object, an iron bar, that behaved in a certain manner; it rusted. (b^-) is not associated such a picture; instead, is only associated with there being simples in arrangements that rusted'. Because these two empirical consequences are associated with different mereological pictures—i.e., they have different thick mereological content—it follows from Fine Grained that they are not the same empirical consequence. Thus the trivialist can give a simple argument for the empirical inequivalence of R and R^- .

If Fine Grained is correct, then the CODA is trivially unsound. The observations and predictions of ordinary scientific theories are about, and thereby appeal to, composites like planets, iron bars, and organisms. These empirical consequences present a particular mereological picture—one where there *are* planets, iron bars, and organisms. And were we to construct theories with empirical consequences associated with different mereological pictures—no planets, iron bars, nor organisms as such—then Fine Grained entails that these empirical consequences are necessarily distinct from those of ordinary scientific theories. Because the CODA attempts to achieve precisely this,

the claim that composite-free theories are empirically equivalent to ordinary scientific theories is trivially false if Fine Grained is true.

3 Rejecting fine grained

The trivial response claims that composite-free theories are trivially empirically inequivalent to their ordinary counterparts. This is because it presupposes Fine Grained, and Fine Grained entails such empirical inequivalence. To me, Fine Grained is neither obviously true nor obviously false. To adjudicate the dispute about empirical equivalence, we must look into theories of empirical consequences that would vindicate Fine Grained. If we find that there is no good, obvious way to vindicate Fine Grained, then we can tentatively reject the trivial response to the CODA. I argue that we find this and more: any candidate theory of empirical consequences will not vindicate Fine Grained. Accordingly, we ought to reject the trivialist response to the CODA.

In this section, I will propose and reject two theories of the individuation conditions of empirical consequences that entail Fine Grained. Because we are determining standards of empirical equivalence, we cannot rely upon *intuitions* of empirical equivalence. Otherwise we are at the stalemate indicated at the outset. Instead, I rely on the possibility of successful *dispensing*, and I argue that from these cases of dispensing we can infer facts about empirical equivalence. This is because a dispensing theory, per the standard account of dispensability, must be empirically equivalent with the original theory. The basic idea of my argument is that Fine Grained precludes the possibility of dispensing with entities that might be dispensable.

3.1 Semantic individuation

The most straightforward way to vindicate Fine Grained is a theory of empirical consequences where they are individuated according to their semantic content. Though many would find this independently implausible, it is instructive to see why it fails. Here is Semantic Individuation:

Semantic Individuation Empirical consequences (c_1) and (c_2) are the same empirical consequence if and only if the expressions of (c_1) and (c_2) have the same semantic content.

With Semantic Individuation, we have a test for whether two empirical consequences are the same—and, accordingly, a test for whether two theories have the same empirical consequences. Namely, whether their expressions are synonymous.

We can show how Semantic Individuation will entail Fine Grained. Recall R and R^- and the expressions of their empirical consequences in (b) and (b^-). Under any usual standard of synonymy, (b) and (b^-) are not synonymous. Given Semantic Individuation, (b) and (b^-) are not the same empirical consequences. Moreover, there are no empirical consequences of R^- that are synonymous with (b), since R^- is explicitly formulated in terms that are not synonymous with composite-terms. Accordingly, R and R^- are empirically inequivalent. More generally, semantic content itself is finely

discriminating, and so individuation according to semantic content will be finely discriminating too. This is how Semantic Individuation entails Fine Grained.

But we have reason to think that Semantic Individuation is independently problematic: it entails the trivial impossibility of dispensing with entities that are intuitively dispensable.

Consider an alternate history of astronomy. As we currently think of them, constellations are apparent groupings of stars seen only according to the Earth's relative position to them. For the astronomer, there aren't constellations; there are only stars. But suppose that modern astronomers, after they fully understood that stars are Sun-like entities that are often light-years away from each other, nonetheless still thought that stars sometimes formed *a constellation*. And suppose they still believed in the independent existence of constellations for no good reason—just a superstitious holdover from antiquity. Semantic Individuation entails that we could not show that constellations are dispensable. This is because Semantic Individuation bars the possibility that constellation-free theories could have the same empirical consequences.

For consider what a dispensability argument would look like. First, we would take the current alternate theory of astronomy C that has the following empirical consequence, referring to a particular constellation (together with some background conditions):

(c) This constellation is visible in August.

Suppose someone thought that constellations are dispensable to our theories. This constellation dispenser might suggest that we adopt the constellation-free theory C^- that has a different empirical consequence.

(c^-) These stars are visible in August

C^- seems to dispense with constellations. We can explain everything we want in astronomy without appealing to constellations if we accept C^- instead of C . But if Semantic Individuation is true, then because (c) and (c^-) are not synonymous—and because there is no constellation-free empirical consequence that is synonymous with (c)—the constellation-free theory is empirically inequivalent to the alternate astronomy theory.

This is a problem for Semantic Individuation. Constellations *are* dispensable to C . And they're dispensable in the usual way, where we conduct an *a priori* investigation into what parts of a theory are necessary to explain what the theory explains, and we realize that constellations are just a vestigial aspect of C . So, constellations are dispensable to C , and we can provide a constellation-free theory C^- , but Semantic Individuation entails that the C^- is trivially empirically inequivalent to C . Thus, because empirical equivalence is a necessary condition for a dispensing theory, and because we have a dispensing theory, there is a counterexample to Semantic Individuation and we ought to reject it.

It seems as though Semantic Individuation is *too* fine-grained. Note here that we are not consulting our intuitions about whether (c) and (c^-) are the same empirical consequence. The argument is more general than that. The argument is that Semantic Individuation entails that *no* entity is dispensable if that entity appears in our empirical consequences. For to dispense with that entity, we must provide an alternative that has

synonymous empirical consequences. And if the empirical consequences are synonymous then we haven't dispensed with the entity. *C* is just an illustrative instance of the restrictions that Semantic Individuation places on which entities are possibly dispensable. Semantic Individuation implausibly entails that no entity which appears in our empirical consequences is dispensable. Accordingly, we should reject it for placing overly restrictive boundaries on the kinds of entities that can be dispensed with.¹³

3.2 Representational individuation

One might naturally think that empirical consequences are individuated according to their representational contents. Scientists approach the external world through their perceptual capabilities, and it is reasonable to think that the empirical consequences of scientific theories are (at least informed by) the deliverances of our perceptual systems; call these representational contents. Representational contents are data about the world that are present to an organism. When the representational contents of two empirical consequences are different, that seems to be the mark of individuation we are seeking.

Here we explore such a theory of the individuation conditions of empirical consequences, which I will call Thick Representational Individuation (TRI). The basic idea is that, with the addition of some theses about humans' representational contents, there are ways to distinguish between the empirical consequences of ordinary scientific theories and their object-free variants because of the differences in representational content. In Sect. 3.4, I argue that TRI fails because of its adherence to one of the additional theses about the representational contents of humans.

One small caveat about my approach. A representational theory of individuation claims that empirical consequences are individuated according to the representational contents of those empirical consequences. This articulation leaves open the question of how representational contents are related to empirical consequences. It is unspecified whether empirical consequences are *constituted* by representational contents, or *partially constituted* by representational contents, or simply *informed* by representational contents. It is consistent with the present theory that there is more to empirical consequences than representational contents. The only commitment is that differences in representational contents is sufficient for differences in empirical consequences. I will refer to *an empirical consequence's representational contents* as a stand-in for whatever particular relationship one wishes to commit to.

Many philosophers have argued that the representational contents of perception, when that perception is veridical, bears a non-representational or external relationship to the thing being represented.¹⁴ Call this thesis External:

External All veridical representational contents bear external, non-representational relationships to what is being represented.

¹³ Thanks to an anonymous referee for helping me clarify the import of this argument.

¹⁴ This conception of the contents of perception includes any externalist theory. See, e.g., Dretske (1997) and Stalnaker (2003).

The idea here is that in order to account for a variety of phenomena regarding perception, it must be that veridical representations bear non-purely-representational relations to the things being represented. (Hereafter, I will drop the word ‘veridically’, and unless otherwise specified, representational contents are veridical.) There must always be some external or worldly relation between representation and represented. Exactly what relation is disputed. Some think that the representational content of a perception of some tree is *constituted*, in part, by the tree (e.g., Fish, 2009; Johnston, 2004); others think that this same content is *causally explained* by the tree (e.g., Burge, 2010), perhaps together with the evolutionary conditions of the perceiving creature. External is similar to externalist theories of semantic content, where cases of successful reference require a non-semantic relation between the referring term and the thing referenced. (Consider here Kripke’s famous causal theory of reference (Kripke, 1981).)

The first step of a Representational Individuation theory is thus an adherence to the thesis of External. The idea is that the empirical consequences of a scientific theory bear a non-representational relation to the world—that, for example, the actual iron bar in (*b*) explains, in some way, the representational content of (*b*). Here, then, is a first pass at a Representational Individuation theory:

Representational Individuation Empirical consequences (c_1) and (c_2) are the same empirical consequence if and only if the representational contents of (c_1) and (c_2) are the same, where representational contents are External.

When empirical consequences differ, this is sometimes explainable by differences in the way the world is. That this particular bar of metal rusted in the presence of moisture-rich air partially explains the empirical consequence (*b*). Consider the empirical consequence ($\neg b$).

($\neg b$) This iron bar did not rust.

On this understanding of empirical consequences, ($\neg b$) is distinguishable from (*b*) because the representational contents of ($\neg b$) and of (*b*) are distinct. We can, for example, *perceive* whether or not the iron bar rusted. Moreover, we can chalk this difference in representational contents up to a difference in the way the world actually is—in one case, the bar of metal rusts, and in the other, it fails to rust.

The claim that representational contents are External is an elegant and powerful understanding of empirical consequences, and one which, I will suggest, might vindicate Hofweber’s trivial response.

3.3 Thick representational individuation (TRI)

Representational Individuation as presented is not enough to entail Fine Grained. In particular, it is unclear whether the representational contents of perception are *robust* enough to distinguish between an iron bar and simples arranged iron bar-wise. Here, we examine a thesis that, when paired with Representational Individuation, will entail Fine Grained. This thesis is Thick:

Thick Representational contents present an underlying mereological picture of what is being represented.¹⁵

According to Thick, our representational contents are mereologically detailed. It claims that representational contents present a mereological picture of the thing represented. The idea here is that we can distinguish between representational contents where the only difference between two contents is a difference in their mereological pictures. This is because, per Thick, our representational contents actually present these two as having distinct metaphysical pictures. If Thick is true, we will only veridically perceive something *as* a composite object if it is a composite object. And when presented with a genuine composite object, so long as there are no perceptual errors, we will perceive it *as* a composite object. Our perceptual capacities have bequeathed us with the ability to not be fooled into representing simples arranged object-wise as an object, nor representing a composite object as simples arranged object-wise (so long as there are no errors). We can *perceptually* distinguish between the two. Byrne (2019), for example, holds this view.

Some may find Thick evolutionarily implausible—it seems like there is no evolutionary reason for an organism’s perceptual capacities to be able to distinguish between simples in arrangements and composites. Others simply disagree with Thick, and claim that our representational content “would be the same whether or not the atoms arranged ...[object]-wise composed something” (Merricks, 2001, p. 9).¹⁶ But for the sake of argument, we will for now grant Thick alongside External.

These two features of empirical consequences and representational contents yield the following individuation conditions for empirical consequences:

Thick Representational Individuation Empirical consequences (c_1) and (c_2) are the same empirical consequence if and only if the representational content of (c_1) and (c_2) are the same, where representational contents are Thick and External.

TRI says that two empirical consequences are distinct when they have distinct representational contents, where representational contents are External and Thick. In short, the idea here is that representational contents are quite representationally dense—not only are we presented with some coarse-grained information about the world, but also with an underlying mereological picture. Thick content, in other words, is part of the total representational content of an empirical consequence.

TRI entails Fine Grained straightforwardly. If TRI is true, then empirical consequences are individuated according to their representational contents. According to Thick, representational contents present an underlying mereological picture of what is being represented. If this is so, then empirical consequences are finely mereologically individuating. Our empirical consequences can distinguish between those cases where the only difference is the mereological facts. Here we finally have a theory of individuation conditions that vindicates Fine Grained.

¹⁵ Thick should strictly be read as follows: All (or all relevant) empirically-relevant representational contents present an underlying mereological picture of what is being represented. We will leave Thick quantifier-less in the main body because this precise articulation adds complications about what Hofweber is committed to.

¹⁶ See also Korman (2014, p. 4) and Thomasson (2014, p. 16, p. 157).

Here is how a trivialist would use TRI to argue against the CODA. The representational content of (*b*) is that some iron bar behaved in a certain way. This representational content presents a thick metaphysical picture where the iron bar is represented *as* an iron bar (and not simply simples arranged iron bar-wise). The representational content of (*b*⁻), on the other hand, is that some simples arranged iron bar-wise behaved in a certain way. Thus, if Thick is true, we can *perceptually* distinguish between the empirical consequences of *R* and those of *R*⁻. Per TRI, because we can perceptually distinguish between (*b*) and (*b*⁻), they are not the same empirical consequence. Consequently, *R* and *R*⁻ are trivially empirically inequivalent. It is a consequence of the theory of TRI that the CODA is unsound, independent of any of the other problems with the dispensability argument. Note that Thick is the crucial premise here; in order to vindicate the trivial response, it must be that our representational contents present a mereological picture.

3.4 Against thick representational individuation

Again, we will reject TRI because it precludes the possibility of dispensing with entities that might be dispensable. We begin the counterexample with a fact about humans. We perceive faces as faces—we do not *judge* an array of features to be a face, but instead perceive it that way.¹⁷ In this way, humans actually have Thick representational contents when faces are involved; we recognize not only *arrangements* of facial features, but a face. This is, at least to some extent, a biological capacity, but there is some debate over the extent to which it is learned.¹⁸

It is a rather mundane mereological picture that is presented when we represent something as a face rather than as arrangements of facial features, but it is a mereological picture nevertheless. Moreover, here we are granting to the trivial response that there *are* cases where we perceptually distinguish things according to their underlying mereological picture.

We continue with another alternate history of astronomy. Suppose that in the night sky in August, there were some stars that looked exactly like a human woman's face—call this Phoebe's Face. When scientists look into the night sky, they seem to perceive Phoebe's face in the arrangement of stars. Their representational contents are actually as of Phoebe's Face, not merely some stars arranged face-wise. The perceptual evidence is so convincing, and the image so detailed, that everyone in this alternative history genuinely believes that Phoebe's Face exists. The theories of astronomy (together with background conditions) might well contain the following empirical consequence.

(*p*) Phoebe's Face is visible in August.

¹⁷ There is empirical evidence that suggests that this is a perceptual capacity—even if it is slightly informed by culture. Individuals can perceive faces in complex pictures or scenes in a short enough time frame that there is no time for substantial cognitive influence—that is, they didn't think before *seeing* the faces. See VanRullen and Koch (2003).

¹⁸ Siegel (2010) argues that it is largely learned—that we can perceive doubt on a person's face if we know that person well enough and know when they doubt something. Block (2014) is unconvinced.

(p) is an empirical consequence of this alternate theory of astronomy P . And if TRI is correct, then (p) is committed to Phoebe's Face being a *face*, much like how (b) commits to that iron bar being a *composed* iron bar.

Suppose, though, that a scientist was conducting *a priori* investigations into whether there are any parts of our theories that are dispensable, and they nominate Phoebe's Face. After all, they think, we know that humans have a proclivity for seeing things as faces, and sometimes it is hyperactive. Maybe, then, our perceptions of Phoebe's Face are illusory in a sense. It seems that Phoebe's Face is dispensable to P ; we can explain everything by appealing to arrangements of stars that seem to look like Phoebe's Face.

Accordingly, the would-be Phoebe's Face dispenser would construct variants of all astronomy theories without appealing to Phoebe's Face. Among the empirical consequences of a variant P^- would be the following:

(p^-) Those particular stars arranged face-wise are visible in August.

In (p^-), the particular stars referenced are those that "make up" Phoebe's Face. P^- shows that Phoebe's Face is dispensable to P . We do not need to appeal to a face in order to explain all of what P explains, since P^- can explain everything just as well.

If Phoebe's Face is dispensable to P in the way just described, then according to the standard account of dispensability, P^- is empirically equivalent to P . And that would signify that (p) and (p^-) are the same empirical consequence. However, per TRI, (p) and (p^-) are *not* the same empirical consequence.

This is because Thick discriminates between different mereological pictures. And as we know, humans perceive faces *as* faces. Because of this, the representational contents of a *face* and of *Xs arranged face-wise* are distinct. We do not represent something as a face-like arrangement, but as a face. Since the representational contents of (p) present something as a face, and the representational contents of (p^-) present some things as looking like a face, these representational contents are distinct. Compare this to the case of an iron bar and simples arranged iron bar-wise. Thick entails that these two are distinct: the representational contents when perceiving an iron bar, if Thick is true, present the thing perceived *as* an iron bar. The representational contents when perceiving just simples arranged iron bar-wise, on Thick, do not present the things perceived as an iron bar. It is the same here as with Phoebe's Face. So, the same considerations that allow TRI to vindicate Fine Grained show that (p) and (p^-) are distinct empirical consequences.

However, this is a problem. Phoebe's Face *is* dispensable, and we can show this by offering P^- . And yet TRI, because it entails that the empirical consequences of a Phoebe's Face-free astronomy theory are trivially distinct from the ordinary theory, tells us that Phoebe's Face is not dispensable. So, TRI ought to be rejected. Once again, the problem is not merely that P seems empirically equivalent to P^- . Rather, the problem is that TRI entails the impossibility of dispensing with mereologically-rich entities that appear in our empirical consequences. If there is a mereologically-rich entity that is doing no genuine explanatory work in our theory, but we can perceptually distinguish that theory from one that is identical except it lacks the mereologically-rich entity, then TRI entails that the entity is not dispensable. But by hypothesis the entity is not doing any genuine explanatory work. Any entity that appears in our theories that

is not genuinely explanatory should be dispensable, but TRI precludes the possibility of dispensing with some such entities. That is why we reject it.

3.5 Rejecting trivial responses

The trivial response to the CODA, as it stands, cannot be defended. For, to do so, one must provide a theory of the individuation conditions that will entail the trivial difference between the empirical consequences of object-free and ordinary theories. And no good theory of individuation conditions vindicates Fine Grained. A theory where semantic differences individuate empirical consequences will not do, nor will a theory where thick representational content differences individuate empirical consequences. These theories fail because there are simple cases of dispensing that are not possible by fiat.

The thrust of my argument against Hofweber's response is that there might be cases where some composed entity that appears in empirical consequences is dispensable, and any theory of individuation that vindicates Fine Grained trivially rules these cases out. This is because Fine Grained entails that any two theories with empirical consequences that have different thick mereological content cannot be empirically equivalent. Fine Grained, then, precludes the *possibility* of dispensing with a composed entity that appears anywhere in a theory's empirical consequences, and this seems absurd. Surely there could be a dispensable composed entity in our empirical consequences. So we must reject Fine Grained.

We can generalize my objection to other philosophers who adhere to the position that thick metaphysical content is empirically significant. The mereological version of this position presupposes that empirical consequences can be individuated along mereological lines. But if it is in principle *possible* to dispense with some mereologically-rich entity in our empirical consequences, e.g., Phoebe's Face, then it is false that empirical consequences can be individuated along mereological lines. *Mutatis mutandis* for other versions of this position: A different version presupposes that empirical consequences can be adjudicated along some particular metaphysical line. But if it is in principle possible to dispense with some metaphysically-rich entity in our empirical consequences, then it is false that empirical consequences can be individuated along such metaphysical lines. I have not yet given counterexamples for non-mereological versions of Fine Grained, but it is intuitive that for any variety of this position, there will be an analogue to Phoebe's Face. All we need is a possible scenario where we have mistakenly inferred that some thick metaphysical content in an empirical consequence is significant and where our theories are just as good when we rid our theories of that metaphysical content.

We should reject Hofweber's trivial response to the CODA because there is no good theory of empirical consequences' individuation conditions that could serve to show that ordinary and composite-free theories are trivially empirically inequivalent. We can also give similar arguments for any other variety of the position that thick metaphysical content is empirically significant. Accordingly, I have made a case for evidential quietism: our empirical evidence is radically silent on distinctively metaphysical disputes.

4 Upshots for full equivalence

Let us consider in more detail the position that thick content is scientifically significant. I argued that thick content is not *empirically* significant. Here I make a tentative case that thick content is sometimes not scientifically significant at all.

In the literature on theoretical equivalence, which examines the conditions under which two theories are fully equivalent, there are philosophers who endorse a position similar to Hofweber's. These philosophers seem to presuppose that thick metaphysical content, as it appears in *non-empirical* consequences, is scientifically significant. There are cases where the only relevant difference between two theories is a difference of thick content, and philosophers claim that such a difference suffices for those two theories being inequivalent. Here I will present one recent instance and suggest that we should be careful about concluding that some thick metaphysical difference is scientifically significant.

Usually, philosophers in the full equivalence literature use equivalence as a means to understanding the significant content of a scientific theory (Barrett, 2019, pp. 1186–1192; Weatherall, 2019b, Sect. 5). The idea seems to be that when we have a good understanding of when two theories are equivalent, we can better gauge which parts of those theories should be taken literally from a scientific realist perspective. Here, I approach from the reverse end. I am arguing that we can learn some things about *equivalence* by examining which parts of our theories are representationally significant. We already have some reason to think that thick content is not empirically significant, and I aim to show that in some cases, thick content is not at all significant. Moreover, recall above that I examined whether thick content was empirically significant by using an alternative notion, dispensability, as a proxy. For full equivalence, there is no analogous proxy to judge whether some content is significant. Accordingly, the conclusions reached here are more tentative than the preceding, but I hope to show that there is headway to be made by approaching full equivalence from this angle.

Consider Hamiltonian and Lagrangian formulations of classical mechanics. It is commonly held among physicists and philosophers of physics that these two formulations are equivalent: the two theories say the same thing about the world and are mere notational variants. Recently, North (2009, 2021a) has pushed back against this received view. She argues that there are differences between the structures of these two theories that are significant in the sense that they show that the theories are in fact *inequivalent*. This argument has generated much discussion, and has led many philosophers to discuss the conditions under which two theories are equivalent in general.

Let me simplify the debate. Briefly, in Hamiltonian mechanics, the state of a classical physical system is specified by the particles' positions and momentum, whereas in Lagrangian mechanics, the state of a classical physical system is specified by the particles' positions and *velocity*. Lagrangian state-spaces have metric structure, whereas Hamiltonian have merely symplectic structure. This difference in structure, North argues, calls into question the equivalence between the two formulations. North claims that because these two theories do not have a structure-preserving mapping between them—the structure of one is literally not present in the other—they are not equivalent.

North claims that this structural difference is significant; it is enough to show that these theories say different things about the world from the scientific perspective. North's argument that this apparent structural difference is sufficient for inequivalence requires an answer to the following question about equivalence in general: in virtue of what can difference in structure be sufficient for inequivalence? Or, as she says, "The question is whether they are equivalent, full stop. The answer depends on whether what differences there are *matter* in any way" (North, 2021a). We need some measure for when an apparent difference in structure is significant. Surely there are cases where structural differences are sufficient for inequivalence; for example, one difference between Newtonian and Galilean spacetime is the structure of absolute rest, and this difference seems significant. On the other hand, the same theory formulated in two different natural languages, say French and English, leads to two theories that have *some* structural difference, though we think that this is a case of mere notational variance. As a result, we need some answers to what kinds (or degrees) of structural differences are significant. Regarding the question of whether what differences there are matter, North might answer that even thick structural differences matter.

The difference between Hamiltonian and Lagrangian mechanics seems to be a merely thick structural difference. One formulation presents an underlying structure that is distinct from the other's; that one theory says a system is specified by momentum and the other velocity is strictly speaking a different underlying structural picture. The structural difference between these two theories is merely in their thick structural content. This is because, in most normal cases, a metric structure can be "recovered" within Hamiltonian mechanics. (More precisely, there is a mapping from any hyperregular model of Lagrangian mechanics to a model of Hamiltonian, and *vice versa*. Cf. Barrett (2019).) If so, then while North is correct that there is not a structure-preserving mapping between the two theories, the difference between the structures of the two formulations is merely in the thick structural content. Hamiltonian mechanics is not committed to metric structure *as such*, but the role that metric structure plays within Lagrangian mechanics can be recovered within Hamiltonian.

The present interpretation of North's argument is that she is pointing to a *thick* structural difference between the two formulations and presupposing that even thick structural differences are enough for inequivalence. (North might want to push back against this interpretation, but a natural reading of her argument presents such an interpretation.) I will suggest that this answer to the above question, that even thick structural differences matter, leads to problems similarly to how Hofweber's presupposition that thick content is empirically significant leads to problems. Consequently, North must offer a precise account of when differences in structure entail inequivalence; otherwise, her objection to the standard view cannot stand.

Let us consider a case of theories which have a difference in merely thick structural content but intuitively this difference is not significant. If there is such a case, then merely thick structural differences are not sufficient for inequivalence. Consider the theory of linear orders.¹⁹ We can formulate the theory using the concept of a nonstrict order *less than or equal to*, signified the binary predicate \leq , or we can formulate it using the concept of a strict order *less than*, signified by the binary predicate $<$. These

¹⁹ Cf. Winnie (1986) and Barrett (2020, p. 1187).

formulations have different axioms; for instance, the first has the axiom that everything bears \leq to itself, whereas the latter has the axiom that nothing bears $<$ to itself. There is a difference in thick structure between these two formulations. The first avails itself of the property *is less than or equal to*, whereas the latter avails itself to the property *is less than*. The second, strictly speaking, does not appeal to *less than or equal to* as such. This is a thick structural difference. Moreover, there is no structure-preserving mapping between them since they trivially and explicitly have different structures.

However, there's an obvious sense in which these are equivalent formulations of linear orders. They both ascribe, in some sense, the same structure to sets. We take it that the difference between formulating linear orders with either \leq or $<$ is not a significant difference; it does not make a difference to the content of the theory. This is a *mere* thick structural difference that we ought not interpret as scientifically significant. We noted in Sect. 3 that there are many cases where thick differences are not scientifically significant; in particular, we cannot appeal to thick differences to conclude that two theories are empirically inequivalent. Likewise, here we cannot point to the thick difference between these two formulations of linear orders as significant; just because there is a difference in the thick structural content, we cannot infer that the two formulations are inequivalent. I take it that one can consult their intuitions to tell that the difference between *less than* and *less than or equal to* does not matter in the scientific sense. Moreover, we can easily recover the structure of \leq on the theory that only has $<$ and *vice versa*.

Here we have a case where, intuitively, thick structural differences are scientifically *insignificant*. Some thick structural differences do not themselves entail that two theories are theoretically inequivalent. If so, then North's claim that the structures of Hamiltonian and Lagrangian mechanics are sufficient for their inequivalence cannot hinge only on there being a thick structural difference between the two.

This argument is not meant to be conclusive.²⁰ If one truly wishes to distinguish between Hamiltonian and Lagrangian mechanics on the basis of thick structural differences, then they may simply accept that these two formulations of linear orders are inequivalent and that even thick structural differences matter. My argument is meant to show simply that North's presupposition about the significance of structural differences, though seemingly innocuous, leads to unpalatable consequences. It seems unintuitive to count all thick structural content as significant, since it leads us to conclude that many theories which we take (or should take) to be equivalent are trivially

²⁰ The present paper had been finished by the time North published her most recent book, North (2021b). Much of what has been said in this section would have changed in light of her Chapter 7 of that book. In particular, North has provided a more general explanation of what is happening between Hamiltonian and Lagrangian mechanics. In that chapter, North considers a theory's "picture of the world", which corresponds directly with what I call a theory's thick content. She identifies multiple pairs of theories that she takes to be *informationally* equivalent but *metaphysically* inequivalent; these theories disagree on "what there is, what it is like, and how and why it behaves in certain ways to give rise to what we observe" (North, 2021b, p. 196). She clarifies that the difference between Hamiltonian and Lagrangian mechanics is a difference between the metaphysical pictures they present (North, 2021b, p. 224). North suggests that metaphysical inequivalence is sufficient for theoretical inequivalence. In this we agree, though—as has become clear throughout this paper—it is not always clear when two theories *are* metaphysically inequivalent. For we need some measure of when two theories make different metaphysical claims about the world. It seems to me that the example about linear orders is *not* a case of metaphysically inequivalent theories, though they explicitly differ in structure. Still, then, North must answer whether what differences there are matter.

inequivalent. Here is the lesson we should learn: North claims that there is a significant difference between Hamiltonian and Lagrangian mechanics which entails that the two formulations are inequivalent. One way to articulate the difference between Hamiltonian and Lagrangian mechanics is a difference in thick structural content, in the underlying structural picture associated with each formulation. The latter appeals to metric structure and the former merely symplectic structure. Yet *thick* structural differences between two theories seem not to be sufficient for inequivalence. We have already seen that two theories which differ only in thick content might still be empirically equivalent. Likewise, we should think that two theories which differ only in thick content might still be fully equivalent. Moreover, we have good reason to think that there *are* theories which differ in their thick structural content but are fully equivalent. In order to conclude that Hamiltonian and Lagrangian mechanics are inequivalent, North must provide a difference-maker between the two that is more than a mere thick structural difference.

5 Conclusion

What I have shown is this: First, that many philosophers adhere to some thesis that thick metaphysical content is empirically significant, and that any variety of this thesis requires an indefensible presupposition about the individuation conditions of our empirical consequences. Second, that the trivial response to the CODA is not tenable. It is not permissible to infer empirical inequivalence from differences in the underlying metaphysical pictures of two theories. Third, that these arguments have some purchase in the theoretical equivalence literature. We should be suspicious of philosophers who argue for inequivalence solely on the basis of differences among thick metaphysical content. There is much more work to be done on determining when the consequences of theories are the same or different, and I have shown that there are implications across philosophy of science and metaphysics for these determinations.

Acknowledgements I am especially grateful to Thomas Barrett for his perceptive comments on multiple drafts. Thanks also to Jeff Bagwell, Kevin Falvey, Dan Korman, Teresa Robertson, and two anonymous referees for their help on previous versions.

Funding No funding was received to assist with the preparation of this manuscript.

Declaration

Conflict of interest The author has no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bagwell, J. N. (2020). Eliminativism and evolutionary debunking. *Forthcoming in Ergo*.
- Bailey, A. M., & Brenner, A. (2020). Why composition matters. *Canadian Journal of Philosophy*, 50(8), 934–949.
- Barrett, T. W. (2015). On the structure of classical mechanics. *The British Journal for the Philosophy of Science*, 66(4), 801–828.
- Barrett, T. W. (2019). Equivalent and inequivalent formulations of classical mechanics. *The British Journal for the Philosophy of Science*, 70(4), 1167–1199.
- Barrett, T. W. (2020). How to count structure. *Forthcoming in Noûs*.
- Blatti, S. (2012). A new argument for animalism. *Analysis*, 72(4), 685–690.
- Block, N. (2014). Seeing–As in the light of vision science. *Philosophy and Phenomenological Research*, 89(3), 560–572.
- Bradley, C. (2020). The non-equivalence of Einstein and Lorentz. *Forthcoming in the British Journal of Philosophy of Science*.
- Brenner, A. (2018). Science and the special composition question. *Synthese*, 195(2), 657–678.
- Burge, T. (2010). *Origins of objectivity*. Oxford University Press.
- Byrne, A. (2019). Perception and ordinary objects. In J. Cumpa & B. Brewer (Eds.), *The nature of ordinary objects*. Cambridge University Press.
- Colyvan, M. (2001). *The indispensability of mathematics*. Oxford University Press.
- Curiel, E. (2014). Classical mechanics is Lagrangian; it is not Hamiltonian. *The British Journal for the Philosophy of Science*, 65(2), 269–321.
- Dorr, C. (2002). *The simplicity of everything*. Ph.D. thesis.
- Dretske, F. (1997). Naturalizing the mind. *Noûs*, 31(4), 528–537.
- Field, H. (2016). *Science without numbers* (2nd ed.). Oxford University Press.
- Fish, W. (2009). *Perception, hallucination, and illusion*. Oxford University Press.
- Fodor, J. A. (1984). Observation reconsidered. *Philosophy of Science*, 51(March), 23–43.
- Friedman, M. (1983). *Foundations of space-time theories: Relativistic physics and philosophy of science*. Princeton University Press.
- Healey, R. (2013). Physical composition. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 44(1), 48–62.
- Hofweber, T. (2016). *Ontology and the ambitions of metaphysics*. Oxford University Press.
- Hofweber, T. (2018). Replies to Eklund and Uzquiano. *Analysis*, 78(2), 315–334.
- Inwagen, P. V. (1990). *Material beings*. Cornell University Press.
- Johnston, M. (2004). The obscure object of hallucination. *Philosophical Studies*, 120(1–3), 113–83.
- Korman, D. Z. (2014). Debunking perceptual beliefs about ordinary objects. *Philosophers' Imprint*, 14, 1–25.
- Kripke, S. (1981). *Naming and necessity*. Blackwell.
- Laudan, L. (1990). Demystifying underdetermination. In C. W. Savage (Ed.), *Scientific theories* (pp. 267–297). University of Minnesota Press.
- Lewis, D. (1988). Statements partly about observation. *Philosophical Papers*, 17(1), 1–31.
- Lowe, E. J. (2003). In defense of moderate-sized specimens of dry goods. *Philosophy and Phenomenological Research*, 67(3), 704–710.
- Lowe, E. J. (2005). How are ordinary objects possible? *Monist*, 88(4), 510–533.
- McDaniel, K. (2007). Extended simples. *Philosophical Studies*, 133(1), 131–141.
- Merricks, T. (2001). *Objects and persons* (Vol. 3). Oxford University Press.
- North, J. (2009). The 'structure' of physics: A case study. *The Journal of Philosophy*, 106, 57–88.
- North, J. (2021a). Formulations of classical mechanics. In E. Knox & A. Wilson (Eds.), *Forthcoming in A companion to the philosophy of physics*. Routledge.
- North, J. (2021b). *Physics, structure, and reality*. Oxford University Press.
- Norton, J. D. (2008). Must evidence underdetermine theory? In M. Carrier, D. Howard, & J. Kourany (Eds.), *The challenge of the social and the pressure of practice* (pp. 17–44). Pittsburgh, PA: University of Pittsburgh Press.
- Osborne, R. C. (2016). Debunking rationalist defenses of common-sense ontology: An empirical approach. *Review of Philosophy and Psychology*, 7(1), 197–221.
- Parsons, J. (2013). Conceptual conservatism and contingent composition. *Inquiry: An Interdisciplinary Journal of Philosophy*, 56(4), 327–339.

- Rosen, G., & Dorr, C. (2002). Composition as a fiction. In R. Gale (Ed.), *The Blackwell companion to metaphysics* (pp. 151–174). Blackwell.
- Sider, T. (1993). Van Inwagen and the possibility of Gunk. *Analysis*, 53(4), 285–289.
- Sider, T. (2007). Parthood. *Philosophical Review*, 116(1), 51–91.
- Sider, T. (2013). Against parthood. *Oxford Studies in Metaphysics*, 8, 237–293.
- Siegel, S. (2010). *The contents of visual experience*. Oxford University Press.
- Stalnaker, R. (2003). What might nonconceptual content be? In Y. H. Gunther (Ed.), *Essays on nonconceptual content* (pp. 95–106). MIT Press.
- Stanford, K. (2009). In E. Zalta (Ed.), *Underdetermination of scientific theory*, *Stanford encyclopedia of philosophy*. Stanford University.
- Thomasson, A. (2014). *Ontology made easy*. Oxford University Press.
- van Fraassen, B. C. (1980). *The scientific image*. Oxford University Press.
- VanRullen, R., & Koch, C. (2003). Is perception discrete or continuous? *Trends in Cognitive Sciences*, 7(5), 207–213.
- Weatherall, J. O. (2019a). Part 2: Theoretical equivalence in physics. *Philosophy Compass*, 14(5), e12591. <https://doi.org/10.1111/phc3.12591>
- Weatherall, J. O. (2019b). Why not categorical equivalence? In J. Madarász & G. Székely (Eds.), *Forthcoming in Hajnal Andreka and Istvan Nemeti on unity of science*.
- Williamson, T. (2007). *The philosophy of philosophy*. Wiley-Blackwell.
- Winnie, J. (1986). Invariants and objectivity: A theory with applications to relativity and geometry. In R. G. Colodny (Ed.), *From quarks to quasars* (pp. 71–180). Pittsburgh University Press.
- Worrall, J. (2011). Underdetermination, realism and empirical equivalence. *Synthese*, 180(2), 157–172.

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