nicholas.jones@philosophy.ox.ac.uk 24 January 2022

Against Representational Levels

Nicholas K. Jones

The final version of this paper is forthcoming in *Philosophical Perspectives*, Volume 36, Metaphysics, 2022, edited by John Hawthorne and Jason Turner.

Abstract

Some views articulate reality's hierarchical structure using relations from the fundamental to representations of reality. Other views instead use relations from the fundamental to constituents of non-representational reality. This paper argues against the first kind of view.

According to the *hierarchical view of reality*, some phenomena somehow give rise to others. For example, interactions amongst particles give rise to molecules; interactions amongst molecules give rise cells; interactions amongst cells give rise to humans; interactions amongst humans give rise to societies. In short, reality exhibits hierarchical structure. Two kinds of view about the nodes of this structure are available.

The first view holds that the nodes all involve representation, possibly except for nodes at the base of the hierarchy. On this view, the hierarchical structure concerns how various phenomena give rise to the truth-values of representations. Those various phenomena might themselves all involve representation, or alternatively there might be a non-representational base. Either way, my initial examples become misleading. For example, since neither molecules nor cells belong to the base of the hierarchy, interactions amongst molecules don't give rise to cells. Instead, the truth-values of representations of interactions amongst molecules give rise to the truth-values of representations of cells.

The second view holds that the nodes need not involve representation, including nodes above the base of the hierarchy. On this view, the hierarchical structure concerns how various phenomena give rise to others that need have nothing whatsoever to do with representation. My initial examples need not then be misleading, and molecules may indeed give rise to cells.

One central part of metaphysics seeks to locate phenomena within the hierarchy. This part of metaphysics asks what, if anything, gives rise to consciousness, morality, and mathematics for example. The two views just outlined correspond to two different conceptions of the proper subject matter of this central part of metaphysics: whether it

may properly concern the non-representational world or only representations thereof. I ague against the representational view below. Nodes of the hierarchy need have nothing to do with representation.

§1 introduces the issue more carefully, including some notable recent examples of the representational view I argue against. §2 presents my argument. §3 defends my argument's premises. §4 examines a fallback position for those attracted to the representational view.

1 Introduction

I begin with some terminology to help state my target question. The *level connectors* are the relations of *giving rise to* that supply reality's hierarchical structure, whatever those relations may be.[1] A debate about what it is for one thing to give rise to another—conversely: for the other to arise from the one—is a debate about what the level connectors are. So *a* (or the *a*s) bears a level connector to *b* (or the *b*s) iff *a* (or the *a*s) gives rise to *b* (or the *b*s) in the metaphysically interesting sense at issue here. (I omit the plural qualifications henceforth.)

A *representational phenomenon* is anything that involves representation. For example, a photograph or painting may be of a particular individual or location; a city map represents the city's streets as standing in certain spatial relations; a belief that it's sunny today represents that it's sunny today. A *non-representational phenomenon* is anything else, that doesn't involve representation. Two clarifications of my target notion of representation follow.

Firstly, I mean the ordinary commonplace notion of representation according to which some things represent and others don't. Photographs, pictures, maps, diagrams, and scale models all represent in this sense. Going beyond these commonplace paradigms, other candidate examples include scientific models, beliefs and other psychological attitudes, propositions, and expressions of interpreted languages like English. By contrast, sticks, stones, tables, chairs, planets, and galaxies typically do not represent, though they could in principle be used to do so. Although the ordinary commonplace notion of representation is not completely precise, the distinction it draws is sufficiently well understood for present purposes.

Secondly, the ordinary commonplace notion of representation comes in many varieties and forms. For example, something might represent the individual Tibbles, or represent the property *being a cat*, or represent that Tibbles is a cat, or represent Tibbles as a cat. Other notions that involve representation include properties like truth and falsity, as well as the property of being about a particular subject matter.

Here's the last piece of new terminology needed to state my target question: a level connector *R* is a *representational level connector* iff, necessarily, whenever something *a* gives rise to something *b* via *R* then *b* is a representational phenomenon; otherwise, *R* is *worldly*. That is, phenomena that arise via representational level connectors always

[1] I remain neutral whether level connectors are properly expressed by sentential connectives (Fine, 2012), or relational predicates (Schaffer, 2009). Both kinds of theory use a relational expression to connect levels, differing only over whether its argument positions require names or sentences. There is thus a sense in which both views admit a level connector relation, differing only over its location in the type hierarchy. For simplicity, I write as if the relational predicate view were true. §3.7 revisits the issue.

involve representation. By contrast, phenomena that arise via worldly level connectors may involve only constituents of non-representational reality and no representations of such constituents. Note that both representational and non-representational phenomena may arise via worldly level connectors, whereas only representational phenomena may arise via representational level connectors.

Candidate worldly level connectors include ground (Schaffer, 2009; Fine, 2012; Rosen, 2010), ontological dependence (Fine, 1995), metaphysical definability (Bacon, 2019), and supervenience understood as a relation on facts or properties (Kim, 1985).

Candidate representational level connectors include truthmaking (Armstrong, 1996; Heil, 2003, 2012; Fine, 2017),[2] requirements on reality (Williams, 2010), ground understood as a relation on propositions (Cameron, 2018),[3] and inter-theoretic reduction (Oppenheim and Putnam, 1958). In Theodore Sider's (2011, ch. 7) framework of metaphysical semantics, candidate representational level connectors also include semantic relations like (the converse of) the relation between sentences and their fundamental truth-conditions.[4]

We can now state my target question: which of the following two views is true?[5]

Representational Levels All level connectors are representational.

Worldly Levels Some level connector is worldly.

---

[2]Truthmaking is representational for Armstrong and Heil because it relates entities—states of affairs for Armstrong; substances, properties, and arrangements thereof for Heil—to truth-bearers, and the entity thereby gives rise to the truth-bearer's truth, which is a representational phenomenon. By contrast, the truthmaking relation employed in truthmaker semantics isn't clearly representational; different interpretations of the semantic formalism may go different ways. Nonetheless, Fine's (2017, p255) interpretation of this formalism makes truthmaking representational: "The idea of truthmaking is the idea of something on the side of the world [...] verifying, or making true, something on the side of language or thought."

[3]Theorists differ over whether propositions represent. According to (Richard, 2013) and (Stalnaker, 2012) propositions are truth-conditions, not entities that possess or represent truth-conditions; (Prior, 1971) and (Jones, 2019) recommend a similar view. If ground relates propositions thus understood, it's worldly: what's grounded is reality being a certain way. According to (King, 2007), (Soames, 2010), and (Merricks, 2015), by contrast, propositions represent reality and possess truth-conditions. If ground relates propositions thus understood, it seems representational: what's grounded is the proposition's truth. Cameron appears to treat propositions as representations of reality when he writes about what they say and their content in (Cameron, 2018, e.g. note 7 and pp338–339). He also introduces ground as "the phenomenon of *one truth being true* in virtue of some others" (p335, my emphasis), suggesting that what arises via ground is the instantiation of a representational property (i.e. truth) by a proposition. Cameron thus appears to treat ground as a representational level connector.

[4]I should note a slight terminological mismatch: the relations mentioned in the main text aren't strictly level connectors in my defined sense, although they are closely related to level connectors. Truthmaking provides a good example. Although truthmaking is a relation between entities and truth-bearers, the entity doesn't give rise to the truth-bearer itself. For example, in (Armstrong, 1996), truth-making relates states of affairs to truth-bearers, but the states of affairs don't give rise to the truth-bearers themselves. Rather, the state of affairs give rise to the truth of the truth-bearer. So truth-making itself isn't strictly a level connector in my defined sense. The level connector relates the entity not to the proposition made true, but to that proposition's truth. The real level connector is thus not ($\lambda xy.x$ makes true y) but ($\lambda xy.x$ makes true the z such that y is the truth of z). Similarly for the other representational level connectors mentioned in the main text. Having noted this terminological mismatch, I'll ignore it henceforth and for simplicity write as if those relations really were level connectors.

[5]List (2018) makes a similar distinction between levels of description and ontological levels.

I will argue against Representational Levels and for Worldly Levels. More carefully put, I will argue that Representational Levels has surprising and unattractive commitments. Although I regard those commitments as sufficient for rejection, others might take a different view. My primary goal is to reveal Representational Levels' true cost, not to argue that this price should not be paid. I invite defenders of Representational Levels to see my argument as a diagnostic tool to aid development of their view, by identifying exactly which premise(s) fails and why.

To be clear, I will not argue against representational level connectors themselves. I will argue against Representational Levels, the view that all level connectors are representational. My argument shows that representational level connectors alone cannot accommodate all cases of one phenomenon giving rise to another. One central lesson is that if there are representational level connectors, then each representational connection $a \Rightarrow b$ requires a counterpart worldly connection $c \Rightarrow d$, where d involves the (typically non-representational) entities represented by the constituents of b.

I argue against Representational Levels in §2 and defend my argument's premises in §3. I use the rest of this section to discuss the topic's interest.

The question of Representational Levels or Worldly Levels is a foundational question about the proper subject matter of a central part of metaphysics. The hierarchical view of reality appears to concern non-representational reality itself, not merely representational phenomena. For example, interactions amongst molecules give rise to cells, and cells are not representational phenomena. Similarly, *questions of fit* such as the following appear to concern reality itself, not merely representations: how do people and societies fit into the world of fundamental physics? Both the hierarchical view of reality and questions of fit are central to metaphysics and appear to often concern the non-representational world.

Assuming Worldly Levels, this appearance may be correct: questions of fit and the hierarchical view really do concern non-representational reality. By contrast, Representational Levels entails that questions of fit and the hierarchical view are either already about representational phenomena or ought to be reformulated in those terms; these central parts of metaphysics should be understood as concerning representations of reality, not only the reality represented. This is a striking, substantive, and revisionary view.

Representational Levels is not merely a hypothetical option; for some prominent metaphysicians appear to endorse it. Three examples follow.

Firstly, truthmaking is the only level connector in the metaphysical systems of David Armstrong (1996; 2004) and John Heil (2003; 2012). Because truthmaking is a representational level connector—truth is a representational property—these views are forms of Representational Levels.[6]

Secondly, according to Sider's (2011, §7.4) framework of metaphysical semantics, the

[6]This attribution of Representational Levels to Armstrong is complicated by his notions of an "ontological free lunch" and "no addition in being". On the one hand, he sometimes appears to allow supervenience as a worldly level connector alongside the representational level connector of truthmaking. On the other hand, he regards the supervenient as an "ontological free lunch" that constitutes "no addition in being" beyond the subvenient base. But is the supervenient real, for Armstrong? Does it exist? It's not clear. If Armstrong says so, then supervenience is indeed a worldly level connector and his view entails Worldly Levels. If Armstrong says not, then supervenience is not a worldly level connector and his view entails Representational Levels.

connections between levels are captured by a certain kind of semantic theory. Some more terminology is useful for characterising what sort of semantic theory. The *fundamental* comprises the base of the hierarchy induced by the level connectors, i.e. that which gives rise to all else and arises from nothing else.[7] The *derivative* comprises everything else, i.e. that which is not fundamental, and hence arises from something else. The kind of semantic theory that Sider uses to capture the connections between levels is for an object-language that contains vocabulary for derivative phenomena, e.g. 'Tibbles' and 'cat'. The theory assigns meanings to this object-language and expresses those meanings in metaphysically fundamental, joint-carving terms. This makes the level connectors representational because the phenomena given rise to all concern representations of the fundamental by linguistic items.

Thirdly, Ross Cameron (2018) employs a relation of ground understood as a relation on propositions, whereby one proposition's truth arises from that of some others.[8] Although he doesn't explicitly deny that there are other, non-representational level connectors, he is naturally interpreted in that way. He argues (on p335) that one of ground and truthmaking reduces to the other; since the argument generalises to ground and any other level connector, the view is naturally taken as one on which ground is both representational and the only level connector, contrary to Worldly Levels.

I've argued that Representational Levels is a striking, substantive, and prominently endorsed revisionary view about a central part of metaphysics. Let's now turn to my argument against it.

2 Argument

Here's the main idea behind my argument. Consider an ordinary true sentence like 'Tibbles is a cat'. Because the sentence is true, we can disquote to conclude that Tibbles is a cat. This conclusion is not about a representational phenomenon. It's about Tibbles and whether she's a cat. They're not fundamental, and must therefore arise via some level connector. Yet since they're also not representational phenomena, this level connector must be worldly, contrary to Representational Levels. The sentence's truth can arise via a representational level connector, unlike Tibbles herself and whether she's a cat. Any representational level connection in which the sentence's truth arises thus requires a counterpart worldly connection through which arise the non-fundamental entities represented, i.e. Tibbles and the fact that she's a cat. To put it vividly if also uncharitably, the attraction of Representational Levels rests on a conflation of use and mention. The rest of this section develops this argument more carefully. The next section defends my argument's premises.

My argument begins with the assumption that there are truths about the non-fundamental; I spell this out more precisely very shortly. Notable candidates include certain sentences (e.g. 'Tibbles is a cat'), propositions (e.g. the proposition that Tibbles is a cat), and beliefs (e.g. my belief that Tibbles is a cat). It will simplify exposition to focus on an arbitrary example of a truth about the non-fundamental. So let $\langle Ct \rangle$ be an

[7] The existence of fundamentalia is controversial; see (Bliss and Priest, 2018, Introduction and Part II) for more.

[8] See note 3 for more about Cameron on propositions, representation, and ground.

arbitrary entity satisfying all of the following:

Represents $\langle Ct \rangle$ represents that Tibbles is a cat.

Truth $\langle Ct \rangle$ is true.

Non-Fundamental The following are not fundamental (even if they exist): the fact that Tibbles is a cat, Tibbles, and *being a cat*.

These collectively say that $\langle Ct \rangle$ is a truth about the non-fundamental, with Represents determining which non-fundamentalia in particular. $\langle Ct \rangle$ might be the sentence 'Tibbles is a cat', the proposition that Tibbles is a cat (if propositions represent), my belief that Tibbles is a cat, or something else entirely, such as a map showing the locations of cats with one node labelled 'Tibbles'.

My next premise is a schema connecting representation and truth:

Link For any truth-evaluable entity x, (x represents that $\phi \rightarrow$ (x is true $\rightarrow \phi$)).[9]

Intuitively, Link says that what is truly represented as the case is the case.

What is this notion of representation-that in Link and Represents? Do sentences, propositions, beliefs, and all other truth-evaluable entities represent that p in the very same sense? My argument doesn't need them to do so. Think of representing that p as a generic or determinable notion, which different views about truth-evaluability flesh out in different ways. For example, sentences might mean that p, or express that p, or express the proposition that p, or say that p, or have the truth-condition that p; propositions might say that p, or have the truth-condition that p; beliefs might be a belief that p, or present things as such that p, or have the content that p. Perhaps these are all distinct notions, though I doubt it. Regardless, every view of truth-evaluability needs some such notion connecting truth-evaluable entities with how they represent things to be. Each such notion is a way of representing that p in the sense that concerns me.

A version of my argument can be formulated without employing representation-that. To do so, replace Represents and Link with these alternative principles, where representation-as replaces representation-that:

Represents* $\langle Ct \rangle$ represents Tibbles as a cat.

Link* For any truth-evaluable entity x, (x represents o as $F \rightarrow$ (x is true $\rightarrow$ o is F)).

The ordinary notion of representation includes representation-as: maps represent geographical features as spatially related in various ways, photographs and pictures represent their subjects as being various ways. We can also represent Tibbles as a cat in both language and thought. If you don't like representation-that or think it's not a form of ordinary representation, you should replace Represents and Link with their starred counterparts throughout the below.

The next premise is a schema concerning the existence of facts, objects, and properties:

[9] Potential counterexample: if you disbelieve that p, your disbelief plausibly represents that p yet is not true if p. I ignore this complication henceforth, focussing on truth-evaluable representations that aim to represent things as they are, such as declarative sentences and beliefs.

Existence If *o* is *F*, then there exist the fact that *o* is *F*, the object *o*, and the property *being F*.

The remaining two premises are:

Generation Anything that's not fundamental is given rise to by something.

Non-Representational The following are non-representational phenomena (even if they exist): the fact that Tibbles is a cat, Tibbles, and *being a cat*.

With my premises now in place, I'll say two things about them before presenting the argument.

Firstly, on Non-Fundamental, Existence, and Non-Representational. Each principle concerns facts, objects, and properties. Any one of those categories would suffice for my argument. So think of the argument as coming in three distinct versions, one for facts, one for objects, and one for properties. To simplify presentation, I will focus on the version about facts. The other versions are obtained by replacing 'the fact that Tibbles is a cat' with 'Tibbles' or '*being a cat*' throughout. If any version of the argument is sound, Representational Levels is false.

Secondly, on Non-Fundamental and Generation. Although each principle mentions fundamentality, my argument does not require that anything is fundamental. It requires only that certain things are not fundamental. My argument is therefore compatible with views on which nothing is fundamental and there are just increasingly fundamental levels each arising from the next.

Here's my argument against Representational Levels:

(1) By Link, Represents: $\langle Ct \rangle$ is true $\rightarrow$ Tibbles is a cat.

(2) By (1), Truth: Tibbles is a cat.

(3) By (2), Existence: There exists the fact that Tibbles is a cat.

(4) By (3), Non-Fundamental: The fact that Tibbles is a cat is not fundamental.

(5) By (4), Generation: Something gives rise to the fact that Tibbles is a cat.

(6) By (5), definition of 'level connector': Some level connector holds from something to the fact that Tibbles is a cat.

(7) By (3), Non-Representational: The fact that Tibbles is a cat is a non-representational phenomenon.

(8) By (6), (7): Some level connector holds from something to a non-representational phenomenon (specifically, to the fact that Tibbles is a cat).

(9) By (8), definition of 'representational': Some level connector is not representational.

(10) By (9): Representational Levels is false.

This argument shows that Representational Levels is false if Represents, Truth, Non-Fundamental, Link, Existence, Generation, and Non-Representational are all true.

3 Defence of the Premises

I now defend my argument's premises. As ever in metaphysics, this defence will not be decisive; one can respond by biting appropriate bullets. As noted in §1, however, my primary goal is to reveal Representational Levels' unattractive commitments, not to argue that those commitments are absolutely unacceptable.

Note that because $\langle Ct \rangle$ was an arbitrary truth about the derivative, an adequate response can't just reject a premise about $\langle Ct \rangle$; it must reject the corresponding claim about any other putative truth about the derivative.

3.1 Represents

Represents $\langle Ct \rangle$ represents that Tibbles is a cat.

Saying and thinking that Tibbles is a cat are ways of representing that Tibbles is a cat. So if Represents is untrue, we cannot say or think that Tibbles is a cat. Generalising beyond this particular example, the non-fundamental cannot be represented at all, in speech, thought, or otherwise. This is an unattractive view. We're used to the realist idea that the fundamental might be unknowable for limited beings such as ourselves. The idea that the non-fundamental cannot even be represented, whether by humans, omniscient deities, or propositions, is more radical. There are also more specific problems with the view.

If Represents isn't true, then seemingly trivial schemas like the following have false instances:[10]

> 'ϕ' represents that ϕ.
>
> The proposition that ϕ represents that ϕ.
>
> Every belief that ϕ represents that ϕ.

In the case at hand, none of 'Tibbles is a cat', the proposition that Tibbles is a cat, and my belief that Tibbles is a cat represents that Tibbles is a cat. This is especially puzzling for the proposition and belief. How could something deserve the label 'proposition that Tibbles is a cat' or 'belief that Tibbles is a cat' without representing that Tibbles is a cat? Although I know of no plausible answer to this question, there isn't space here to examine the theoretical roles of propositions and beliefs in detail. I leave it as a challenge for defenders of Representational Levels to provide views of propositions and beliefs compatible with the falsity of Represents.

Consider now the case where $\langle Ct \rangle$ is the sentence 'Tibbles is a cat'. What does this sentence represent, if not that Tibbles is a cat? There are two options. It might represent only something other than that Tibbles is a cat. Or it might represent nothing at all. I argue against these in turn.

Suppose first that 'Tibbles is a cat' represents only something other than that Tibbles is a cat. Then we have a new and surprising failure of the disquotational schema:

[10] Instances of the schemas in this section are obtained by uniformly replacing all occurrences of 'ϕ' (inside and outside quotation marks) with a meaningful declarative sentence.

'ϕ' represents that ϕ.

Failures of disquotation due to differences of context between utterance and report are familiar. For example, although you uttered, under your breath, 'I can't wait for the end of this paper', your words didn't represent that *I* can't wait for the end of this paper; they represented that *you* can't wait for the end of this paper. The present case is not of this kind. The present question is whether 'Tibbles is a cat' as used by me right here right now in this and the preceding sentences represents that Tibbles is a cat. It's hard to see how that could fail to be the case if 'Tibbles is a cat' (as used by me right here right now and in the preceding sentences) represents anything at all.

Now suppose instead that 'Tibbles is a cat' doesn't represent anything at all. Generalising, no sentence purporting to represent the non-fundamental represents anything at all. Almost every sentence uttered in the history of humanity is of this kind. Because meaning that *p* is a way of representing that *p*, it follows that the vast majority of utterances in the history of humanity have lacked meaning, contrary to our evident success in linguistic communication.

3.2 Truth

Truth $\langle Ct \rangle$ is true.

I see two main reasons not to reject Truth.

Firstly, given that $\langle Ct \rangle$ represents that Tibbles is a cat, if $\langle Ct \rangle$ is untrue, then Tibbles is not a cat.[11] Since Tibbles is in fact a cat, $\langle Ct \rangle$ is true. The point generalises: in order for all putative truths about the non-fundamental to be untrue, the non-representational world of cats, dogs, tables, and chairs must be radically different from how it actually is.

Secondly, because $\langle Ct \rangle$ was arbitrary, rejection of Truth provides an adequate response to my argument only if no representations of non-fundamentalia are true. It follows that almost all thought and talk throughout human history is untrue. This near-global error-theory is both unattractive and incompatible with extant versions of Representational Levels on which the level connectors associate ordinary non-fundamental truths with fundamentalia that give rise to their truth.[12]

3.3 Non-Fundamental

Non-Fundamental The following are not fundamental (assuming they exist): the fact that Tibbles is a cat, Tibbles, and *being a cat*.

Because $\langle Ct \rangle$ was arbitrary, rejection of Non-Fundamental provides an adequate response to my argument only if every fact, object, and property about which there are truths is fundamental. Since there are truths about every fact, object, and property, this trivialises fundamentality.

[11] This follows from Represents and a principle that reverses the embedded conditional in Link: *x* represents that $\phi \rightarrow$ ($\phi \rightarrow$ *x* is true).

[12] Versions of Representational Levels intended to accommodate ordinary, non-fundamental truth include some of those mentioned in §1: Armstrong's and Heil's *truth*maker theory, and Sider's metaphysical semantics.

This trivialisation of fundamentality will be unattractive to many. It entails that there are no differences in respect of fundamentality between facts that many regard as clearly differing in that way. For example, consider the fact that Tibbles is a cat, or the fact my coffee cup is red. These facts are multiply realizable: there are many specific ways in which Tibbles could be a cat, by having different microphysical configurations and histories; there are many specific ways in which my coffee cup could be red, by having different microphysical configurations and spectral reflectance profiles . Now consider those facts' actual realizations: the actual way in which Tibbles is a cat, by having a certain maximally specific material configuration and history; the actual way in which my coffee cup is red, by having a surface with a certain maximally specific microphysical configuration and spectral reflectance profile. It is natural to regard these maximally specific actual realization facts as more fundamental than, because they give rise to, the corresponding mutliply realizable facts. Yet if Non-Fundamental is false, this is not so; those facts are all equi-fundamental and none gives rise to any other.

For another example, note that the actual realization facts just mentioned are complex facts. We can see them as conjunctive facts, where each conjunct assigns some value of some fundamental quantity to some fundamental particle(s). It is natural to regard these individual conjuncts as more fundamental than, because they give rise to, the corresponding conjunctive facts. Yet if Non-Fundamental is false, this is not so; those facts are all equi-fundamental and none gives rise to any other.

I see one way of developing the view in order to mitigate these unattractive consequences. The idea is that although fundamentality trivialises on non-representational phenomena, it doesn't trivialise on representational phenomena. On this view, non-representational reality is all fundamental. Yet the level connectors nonetheless induce an interesting hierarchy amongst representational phenomena. Claims about non-representational phenomena giving rise to other non-representational phenomena, are all false, including the natural examples above. Yet there are nearby truths about representational phenomena giving rise to other representational phenomena. For example, the fact that Tibbles has such-and-such maximally specific history and microphysical configuration doesn't give rise to the fact that Tibbles is a cat; but the truth of the proposition that Tibbles has such-and-such maximally specific history and microphysical configuration does give rise to the truth of the proposition that Tibbles is a cat. These nearby truths about representational phenomena might be thought to mitigate the unattractive falsity of the corresponding claims about non-representational phenomena.

The main problem with this view is that it renders the metaphysical interest of the level connector-induced hierarchy amongst representational phenomena unclear. Abstracting slightly, the view says that (a) the truth of the proposition that p gives rise to the truth of the proposition that q, even though (b) the fact that p and the fact that q are both fundamental and neither gives rise to the other. The problem is to understand the significance of (a) in conjunction with (b). I see how the fact that p could give rise to the truth of the proposition that p, and likewise for q. But I can't understand how the truth of one proposition could arise from that of the other, given the existence of fundamental facts corresponding to both. The metaphysical interest of this supposed relation between representations of facts is mysterious, given that no level connectors

relate the fundamental facts represented.

Might one simply accept that every non-representational fact is fundamental, without mitigation? I do not think so. There is a conception of fundamentality on which it's plausible that every fact is fundamental. However, this paper concerns a different conception of fundamentality, on which it is not plausible that every fact is fundamental. Let me explain.

The first conception of fundamentality begins with the idea that some representations are more metaphysically perspicuous than others: they better capture the deep metaphysical structure of what they they represent. The fundamental representations are those that most perspicuously capture the deep metaphysical structure. Call this notion *structural fundamentality*.

The second conception of fundamentality begins, like this paper, with the idea that some phenomena give rise to others. The fundamental then comprises the base of the hierarchy induced by the level connectors: that which gives rise to all else and is not itself given rise to. Call this notion *base fundamentality*. Base fundamentality is what's primarily operative in the literature on ground, for example; it's also the conception of fundamentality that has concerned me throughout this paper.

Return to the view that every fact is fundamental. This view is plausible only for structural fundamentality, not for base fundamentality.

Here's why it's plausible that every fact is structurally fundamental. Note first that structural fundamentality is a property of representations: structurally fundamental representations capture the deep metaphysical structure of what they represent. This is a problem for the view that every fact is fundamental, because many facts do not involve representation. To overcome the problem, we need to extend the notion of structural fundamentality from the representational to the non-representational. The most natural way of doing so is as follows: to be structurally fundamental in the extended sense is to have (or possibly have) a representation that is structurally fundamental in the unextended sense. Every fact presumably does have (or could have) a structurally fundamental representation. So every fact is structurally fundamental in the extended sense.

Here's why it's not plausible that every fact is base fundamental. Consider a complex fact, for example a fact with structurally fundamental conjunctive representation.[13] Many complex entities are not naturally regarded as base fundamental, including typical conjunctive facts: the facts corresponding to their constituent conjuncts give rise to conjunctive facts. What gives rise to the conjunctive fact that $\phi \wedge \psi$ are both the fact that ϕ and the fact that ψ. Conjunctive facts are therefore not base fundamental, since they arise from their conjuncts. Similar arguments could be given for other kinds of complex facts, such as facts with structurally fundamental disjunctive or doubly negated representations.

Let us take stock. I've been considering a response to my argument that rejects Non-Fundamental. I argued that this response is incompatible with many natural views about what gives rise to what because it unattractively trivialises fundamentality. I then considered a way of mitigating these unattractive consequences, which I rejected because it renders the metaphysical interest and import of the level connec-

[13] If there are any true conjunctive (or otherwise complex) structurally fundamental representations, Existence implies that there are corresponding conjunctive (or otherwise complex) facts.

tors mysterious. Finally, I distinguished two conceptions of fundamentality, structural fundamentality and base fundamentality. Although Non-Fundamental plausibly fails for an extended notion of structural fundamentality, it holds for base fundamentality. And base fundamentality is what's at issue in this paper; for base fundamentality concerns level connectors, whereas levels connectors play no central role in the theory of structural fundamentality.

I said in §1 that my primary goal is to reveal Representational Levels' unattractive commitments, not to argue that those commitments are absolutely unacceptable. We're now at one point where defenders of Representational Levels could resist my argument. They could accept an unattractive trivialisation of fundamentality. Or they could explain the import and metaphysical role of representational level connections between representations of fundamental facts. Although neither option seems viable to me, others may perhaps take a different view.

3.4 Link

Link For any truth-evaluable entity x, (x represents that $\phi \rightarrow$ (x is true $\rightarrow \phi$)).

My argument is valid if '$\rightarrow$' expresses the material conditional. So counterexamples to instances of Link have the form:

$$x \text{ is true} \wedge x \text{ represents that } \phi \wedge \neg\phi.$$

In the case at hand, 'Tibbles is a cat', the proposition that Tibbles is a cat, and my belief that Tibbles is a cat all truly represent that Tibbles is a cat, even though Tibbles is not a cat. This makes truth and falsity mysterious. How could something truly represent that p when it's not the case that p? I know of no plausible answer to this question; especially so for the relatively straightforward instance of Link employed in my argument, which is far removed from the paradoxes.

3.5 Is Tibbles a cat?

I've now defended Link, Represents, Truth, and Non-Fundamental. There is also a general problem for any response to my argument that rejects only (all or some of) the first three of those principles. I used them only to derive that Tibbles is a cat at (2). So even if they're all false, the argument from (2)–(10) shows that Representational Levels is false nonetheless, provided only that Tibbles is a cat. Defenders of Representational Levels therefore cannot only reject Link, Represents, or Truth. They must also reject some other premise or deny that Tibbles is a cat, i.e. (2). Before turning to the other premises, let's briefly consider the idea that Tibbles is not a cat.

Recall that $\langle Ct \rangle$ is an arbitrary truth about the non-fundamental. A fully general response to my argument must therefore also reject the counterpart of (2) for each other truth about the non-fundamental. In particular, none of the following is the case: Tibbles exists, the Earth is a planet, Fido is a dog, you exist, you're a person, you're conscious. On this unattractive and revisionary view, the familiar world of cats, dogs, tables, chairs, people, and even our own conscious lives is not even remotely as we ordinarily take it to be.

3.6 Existence

Existence If *o* is *F*, then there exist the fact that *o* is *F*, the object *o*, and the property *being F*.

Instantiating the antecedent with a true derivative predication, as in my argument, yields the existence of derivative facts, objects, and properties. I offer two motivations for instances of Existence in this section. The next section also presents a version of my argument that does not make use of Existence. Existence serves only to make my argument simpler to express.

Firstly, cats are living, breathing, flesh and blood, spatiotemporally located inhabitants of the material world. It's not possible to be such a thing without existing. More generally, many instances of the following restricted version of Existence hold because an object cannot be *F* without existing:

> If *o* is *F*, then *o* exists.

Instances of this schema that can be justified in this way include those where *F* is 'purring', 'conscious', 'red', 'causally efficacious', 'living at 221b Baker Street', 'a winged horse', and many others. Although this doesn't justify Existence itself, it does just instances of the schema sufficient for an object-theoretic version of my argument.[14]

Secondly, instances of Existence about facts and properties follow from minimalist views on which facts and properties satisfy the following schemas:[15]

> For the fact that ϕ to exist is for it to be that ϕ.

> For *o* to instantiate *being F* is for it to be that *o* is *F*.

Instances of Existence that follow from these minimalist schemas suffice for fact-theoretic and property-theoretic versions of my argument.

Minimalist views about facts and properties are endorsed by Stephen Schiffer (2003, esp. ch.2) and Amie Thomasson (2015), sympathetically discussed by Agustín Rayo (2013, esp. ch.1), and play a prominent role in meta-ethical expressivism's problem of creeping minimalism (Dreier, 2004). While certainly controversial, minimalism can be motivated via high-level reflection on the relationship between language and reality (e.g. Rayo and Thomasson), via more localised considerations about the workings of fact-discourse and property-discourse (e.g. Schiffer), or via dissatisfaction with more metaphysically demanding views on which the existence of a fact that *o* is *F* or property of *being F* involves more than just *o* being *F*.

Suppose you accept a non-minimalist theory of facts and properties, on which *o* can be *F* without there being a fact that *o* is *F* or property of *being F*. It does not follow that you should reject Existence. If the minimalist schemas characterise *a* legitimate notion of fact, and *a* legitimate notion of property, Existence follows. These minimalist notions will likely even be compatible with your preferred non-minimalist notion: they may just be different notions, not rivals. Yet the minimalist notions suffice for my argument.

[14] Recall from §2 that my argument comes in three varieties, one about objects, one about properties.

[15] These minimalist schemas are stronger than my argument requires. The argument is valid if both sides of the schemas are materially equivalent.

3.7 Dispensing with Existence

A version of my argument can be formulated that does not make use of Existence. Existence just makes the argument simpler to express.

To reformulate the argument, we first replace predicates of facts with sentential operators. Those predicates include relational predicates for level connectors; for example, we might replace 'the fact that ϕ grounds the fact that ψ', we might write 'ψ because ϕ' where 'because' is interpreted ground-theoretically as in (Fine, 2012). I'll use 'because' as a placeholder for whatever our preferred binary level connector operator happens to be.

Another predicate of facts employed in my argument is 'is fundamental'. I'll replace this with the monadic operator 'it is fundamentally the case that', intended for characterising how fundamental reality is configured. For example, it is not fundamentally the case that Tibbles is a cat, but it is fundamentally the case (let's suppose) that particle *p* has value *v* of quantity *q*.

Our final predicate of facts is 'is a representational phenomenon'. I'll replace this predicate with the monadic operator 'it is at least partly representational whether'. Where I'd previously have said the fact that *p* involves representation, and is therefore a representational phenomenon, I will now say that it's at least partly representational whether *p*. It's at least partly representational whether *p* iff whether or not it's the case that *p* turns at least in part on how things are represented.

We can use these operators to formulate the following principles:

Non-Fundamental* It is not fundamentally the case that Tibbles is a cat (even if Tibbles is a cat).

Generation* If ($\phi\wedge$ it is not fundamentally the case that ϕ), then (ϕ because something, i.e. $\exists p(\phi$ because $p)$).

Non-Representational* It is not even partly representational whether Tibbles is a cat.

Representational Levels* If ϕ because something (i.e. $\exists p(\phi$ because $p)$), then it is at least partly representational whether ϕ.

Each starred principle is a reformulation of its unstarred counterpart, using sentential connectives instead of predicates to express level connectors, fundamentality, and representationality. This makes it hard to imagine a stable or well motivated view on which one of the starred principles held while its unstarred counterpart failed. In particular therefore, those attracted to Representational Levels should also be attracted to Representational Levels*. The arguments in §3.3, §3.8, and §3.9 for Non-Fundamental, Generation, and Non-Representational also carry over straightforwardly to their starred counterparts.

Both Generation* and Representational Levels* employ quantifiers binding variables in sentence position. How should this higher-order quantification be understood?

One historically popular view says that higher-order quantification is semantically equivalent to restricted first-order quantification over some special kind of individual, such as facts, propositions, states of affairs, or ways for reality to be. On this view,

the consequent of Generation* means something like: ϕ because there is a fact that *p*. Using Generation* as a premise will then make my reformulated argument depend on Existence (or some similar principle) to ensure an appropriate supply of this special kind of individual. We won't then have dispensed with Existence.

I know of no good arguments for the preceding view of higher-order quantification. An alternative primitivist interpretation of higher-order quantification is both viable and increasingly popular.[16] On this view, higher-order quantification is a *sui generis* kind of (non-substitutional) quantification, not quantification over values of first-order variables that happen to be facts or propositions or whatever. On this view, '$\exists p$ ...' is not semantically equivalent to a first-order generalisation such as 'there is a fact/proposition such that ...' or anything similar. Using Generation* as a premise will not then make my reformulated argument depend on Existence.

We can now argue against Representational Levels* thus:

(1*) By Link, Represents: $\langle Ct \rangle$ is true $\rightarrow$ Tibbles is a cat.

(2*) By (1*), Truth: Tibbles is a cat.

(3*) By (2*), Non-Fundamental*: Tibbles is a cat $\wedge$ it is not fundamentally the case that Tibbles is a cat.

(4*) By (3*), Generation*: Tibbles is a cat because something (i.e. $\exists p$(Tibbles is a cat because *p*)).

(5*) By (4*), Non-Representational*: Tibbles is a cat because something (i.e. $\exists p$(Tibbles is a cat because *p*)) $\wedge$ it is not even partly representational whether Tibbles is a cat.

(6*) By (5*), Representational Levels* is false.

Existence plays no role in this argument. It shows that if the starred counterparts of Non-Fundamental, Generation, and Non-Representational all hold as well as Link, Truth, and Represents, then Representational Levels* is false. Since Representational Levels* reformulates Representational Levels to replace predicates of facts with sentential operators, this is an argument against Representational Levels too.

3.8 Generation

Generation Anything that's not fundamental is given rise to by something.

Generation* If ($\phi \wedge$ it is not fundamentally the case that ϕ), then $\exists p(\phi$ because *p*)).

I will focus on the unstarred principle for simplicity. My discussion carries over straightforwardly to the starred principle.

Counterexamples to Generation have the form:

[16] On primitivist interpretations of higher-order quantification, see (Prior, 1971, ch3), (Boolos, 1975), (Rayo and Yablo, 2001), and (Williamson, 2003). For metaphysical applications of the view, see e.g. (Dunaway, 2013), (Williamson, 2013), (Dorr, 2016), (Jones, 2018), (Trueman, 2021), (Skiba, 2021), and (Fritz and Jones, 202X).

x is not fundamental $\wedge$ nothing gives rise to x.

Views that permit such counterexamples arguably change the subject. The core theoretical role for the fundamental is to be the minimal base that gives rise to all else. As Sider (2013, p105), puts it, "that the fundamental must in some sense be responsible for everything [...] seems definitive of fundamentality". Now, suppose that nothing gives rise to x. Then no base that excludes x is responsible for everything, since it's not responsible for x (whether by including x or giving rise to x). So if nothing gives rise to x, then x is fundamental. The core theoretical role for fundamentality thereby precludes counterexamples to Generation.

One could respond by adopting a different theoretical role for fundamentality, which doesn't verify Generation. The fundamental isn't then responsible for everything. So what is the fundamental responsible for?

I see only one natural answer to this question: fundamentality's core theoretical role is for the fundamental to give rise to the truth-values (and perhaps other representation-involving features) of all representations. On one version of this answer, the non-representational is all fundamental, and the derivative comprises only representational phenomena. However, this is not the view needed here. Indeed, Generation may hold on this view. By contrast, the view needed here admits non-representational non-fundamental phenomena that arise from nothing, such as the fact that Tibbles is a cat. Yet since those phenomena arise from nothing, it's mysterious why they don't count as fundamental. I don't know what more could be required for fundamentality.

3.9 Non-Representational

Non-Representational The following are non-representational phenomena (even if they exist): the fact that Tibbles is a cat, Tibbles, and *being a cat*.

Non-Representational* It is not even partly representational whether Tibbles is a cat.

I focus on the unstarred principle for simplicity. My discussion carries over straightforwardly to the starred principle.

Surely some phenomena are not representational. If the fact that Tibbles is a cat is not amongst them, let's just use a different example in my argument. An adequate response must therefore hold that only representational phenomena exist: everything whatsoever involves representation somehow or other. This is a form of idealism, with reality itself somehow constituted by representations of it. Let's set this radical view aside.

4 A Fallback Position

I've now defended each premise of my argument against Representational Levels. To close, I consider and reject a fallback position for those attracted to Representational Levels. The idea is that although Representational Levels itself is false, an alternative principle captures much of the spirit of the view.

This fallback position accepts that there are worldly level connectors. However, the resulting hierarchy amongst the non-representational is not fundamental. Rather, it arises from a hierarchy involving only representational level connectors:

Derivative Worldly Levels There are worldly level connectors, but facts about what they relate always arise from facts about what representational level connectors relate.

For example, suppose the fact f_1 that there are particles on the mat arranged cat-wise gives rise to the fact f_2 that there is a cat on the mat. That requires a worldly level connector holding from f_1 to f_2: $f_1 \Rightarrow_W f_2$. The existence of this worldly level connector $\Rightarrow_W$ is compatible with Derivative Worldly Levels but not with Worldly Levels. However, Derivative Worldly Levels requires that the fact $f_1 \Rightarrow_W f_2$ itself arises from some facts about representational level connections. For example, perhaps the truth of the proposition p_1 that there are particles on the mat arranged cat-wise gives rise to the truth of the proposition p_2 that there is a cat on the mat: $p_1 \Rightarrow_R p_2$. In that sense, the derivative worldly connection $f_1 \Rightarrow_W f_2$ is merely a projection onto the non-representational world of the fundamental representational connection $p_1 \Rightarrow_R p_2$ between representations of f_1 and f_2.

Derivative Wordly Levels is similar to a view I rejected in §3.3. On that view, the level connectors induced no hierarchy amongst non-representational phenomena, for consistency with Representational Levels. By contrast, the present view admits a hierarchy amongst non-representational phenomena, albeit derivative from a hierarchy amongst representational phenomena. Derivative Worldly Levels therefore avoids both of §3.3's objections to the earlier view. Firstly, Derivative Worldly Levels is compatible with standard views about how non-representational phenomena arise, provided the facts about how they do so are themselves derivative. Secondly, the metaphysical interest of level connections between representational phenomena is clear: they give rise to level connections between the (often non-representational) phenomena represented.

The main problem with Derivative Worldly Levels is that the fundamental level connections cannot concern only representational phenomena. To see why, recall the example above. Fact f_1 gives rise to fact f_2: $f_1 \Rightarrow_W f_2$. Call that complex fact about the worldly level connector $\Rightarrow_W$, w. Also part of the example, proposition p_1, which represents f_1, gives rise to proposition p_2, which represents f_2: $p_1 \Rightarrow_R p_2$. Call that complex fact about the representational level connector $\Rightarrow_R$, r.

According to Derivative Worldly Levels, something gives rise to fact w. Since w is not a representational phenomenon, the level connector involved cannot be representational. When introducing the example, I suggested that r might give rise to w: $r \Rightarrow w$. We can now see that this alone doesn't secure consistency with Derivative Worldly Levels. Since w is not a representational phenomenon, the level connector here must be worldly: $r \Rightarrow_W w$. Derivative Worldly Levels now requires that something X gives rise to the fact $(r \Rightarrow_W w)$ involving the worldly level connector $\Rightarrow_W$, i.e.: $X \Rightarrow (r \Rightarrow_W w)$. Can the level connector $\Rightarrow$ here from X to $(r \Rightarrow_W w)$ be representational, in line with Derivative Worldly Levels? It depends whether the fact $(r \Rightarrow_W w)$ counts as a representational phenomenon.

On the one hand, $(r \Rightarrow_W w)$ has a part r that clearly involves representation: r specifies how the representational properties (specifically truth) of proposition p_2 arise

from those of proposition p_1. On the other hand, $(r \Rightarrow_W w)$ has another part w to which representation is wholly irrelevant. The fact $(r \Rightarrow_W w)$ in $X \Rightarrow (r \Rightarrow_W w)$ thus involves more than just representational matters. So we have not vindicated the idea that only representational phenomena arise via the fundamental connections: the fact $(r \Rightarrow_W w)$ is partly representational but also partly not representational.

The argument generalises. Many facts about worldly level connectors do not involve representation. So if anything gives rise to those facts, level connections to non-representational phenomena are required. And if anything gives rise to those facts—i.e. to facts about what gives rise to facts about what gives rise to what—level connections to at least partly non-representational phenomena are required. And so on. This regress shows that the bedrock facts about what gives rise to what cannot concern only representational phenomena. The hierarchy amongst non-representational phenomena induced by worldly level connectors is not merely a projection from a more fundamental hierarchy that involves only representational phenomena.

Acknowledgements

For comments and discussion that significantly improved the paper, I'm grateful to Matti Eklund, Salvatore Florio, Dan Marshall, Eliot Michaelson, Matt Parrott, Martin Pickup, Alex Roberts, Gonzalo Rodriguez-Pereyra, Maja Spener, Andreas Stokke, Scott Sturgeon, Henry Taylor, and Al Wilson, as well as audiences at Lingnan and Uppsala.

References

Armstrong, D. (2004). *Truth and Truthmakers*. Cambridge University Press.

Armstrong, D. M. (1996). *A World of States of Affairs*. Cambridge University Press.

Bacon, A. (2019). Is reality fundamentally qualitative? *Philosophical Studies*, 176(1):259–295.

Bliss, R. and Priest, G., editors (2018). *Reality and its Structure: Essays in Fundamentality*. Oxford University Press.

Boolos, G. (1975). On second-order logic. *Journal of Philosophy*, 72(16):509–527.

Cameron, R. P. (2018). Truthmakers. In Glanzberg, M., editor, *The Oxford Handbook of Truth*, pages 333–354. Oxford University Press.

Dorr, C. (2016). To be F is to be G. *Philosophical Perspectives*, 30:39–134.

Dreier, J. (2004). Meta-ethics and the problem of creeping minimalism. *Philosophical Perspectives*, 18.

Dunaway, B. (2013). Modal quantification without worlds. In Bennett, K. and Zimmerman, D. W., editors, *Oxford Studies in Metaphysics Volume 8*, chapter 4. Oxford University Press.

Fine, K. (1995). Ontological dependence. *Proceedings of the Aristotelian Society*, 95:269–290.

Fine, K. (2012). Guide to ground. In Correia, F. and Schnieder, B., editors, *Metaphysical Grounding: Understanding the Structure of Reality*, chapter 1. Cambridge University Press.

Fine, K. (2017). Truthmaker semantics. In Hale, B., Wright, C., and Miller, A., editors, *A Companion to the Philosophy of Language*, pages 556–577. John Wiley & Sons, 2 edition.

Fritz, P. and Jones, N. K., editors (202X). *Higher-Order Metaphysics*. Oxford University Press. In preparation.

Heil, J. (2003). *From an Ontological Point of View*. Oxford University Press.

Heil, J. (2012). *The Universe as We Find It*. Oxford University Press.

Jones, N. K. (2018). Nominalist realism. *Noûs*, 52(4):808–835.

Jones, N. K. (2019). Propositions and cognitive relations. *Proceedings of the Aristotelian Society*. Forthcoming.

Kim, J. (1985). Concepts of supervenience. *Philosophy and Phenomenological Research*, 45(2):153–176. Reprinted in (Kim, 1993, 53–78).

Kim, J. (1993). *Supervenience and Mind: Selected Philosophical Essays*. Cambridge University Press.

King, J. C. (2007). *The Nature and Structure of Content*. Oxford University Press.

List, C. (2018). Levels: descriptive, explanatory, and ontological. *Noûs*. DOI: 10.1111/nous.12241.

Merricks, T. (2015). *Propositions*. Oxford University Press.

Oppenheim, P. and Putnam, H. (1958). Unity of science as a working hypothesis. *Minnesota Studies in the Philosophy of Science*, 2:3–36.

Prior, A. N. (1971). *Objects of Thought*. Oxford University Press.

Rayo, A. (2013). *The Construction of Logical Space*. Oxford University Press.

Rayo, A. and Yablo, S. (2001). Nominalism through de-nominalization. *Noûs*, 35(1):74–92.

Richard, M. (2013). What are propositions? *Canadian Journal of Philosophy*, 43(5–6):702–719.

Rosen, G. (2010). Metaphysical dependence: grounding and reduction. In Hale, B. and Hoffman, A., editors, *Modality: Metaphysics, Logic and Epistemology*, pages 109–136. Oxford University Press.

Schaffer, J. (2009). On what grounds what. In Chalmers, D. J., Manley, D., and Wasserman, R., editors, *Metametaphysics: New Essays on the Foundations of Ontology*, pages 347–383. Oxford University Press.

Schiffer, S. (2003). *The Things We Mean*. Oxford University Press.

Sider, T. (2011). *Writing the Book of the World*. Oxford University Press.

Sider, T. (2013). Symposium on *Writing the Book of the World*. *Analysis*, 73(4):751–770.

Skiba, L. (2021). Higher-order metaphysics. *Philosophy Compass*, 16(10).

Soames, S. (2010). *What is Meaning?* Princeton University Press.

Stalnaker, R. C. (2012). *Mere Possibilities: Metaphysical Foundations of Modal Semantics*. Princeton University Press.

Thomasson, A. L. (2015). *Easy Ontology*. Oxford University Press.

Trueman, R. (2021). *Properties and Propositions: The Metaphysics of Higher-Order Logic*. Cambridge University Press.

Williams, J. R. G. (2010). Fundamental and derivative truths. *Mind*, 119(473):103–141.

Williamson, T. (2003). Everything. *Philosophical Perspectives*, 17:415–65.

Williamson, T. (2013). *Modal Logic as Metaphysics*. Oxford University Press.