

Following logical realism where it leads

Michaela Markham McSweeney¹

Published online: 22 November 2017
© Springer Science+Business Media B.V., part of Springer Nature 2017

Abstract Logical realism is the view that there is logical structure in the world. I argue that, if logical realism is true, then we are deeply ignorant of that logical structure: either we can't know which of our logical concepts accurately capture it, or none of our logical concepts accurately capture it at all. I don't suggest abandoning logical realism, but instead discuss how realists should adjust their methodology in the face of this ignorance.

Keywords Logical realism · Structure · Unknowability · Ineffability · Fundamentality · Relations

This paper is about *logical realism*, the view that there is mind-and-language-independent logical structure in the world.¹ I think that logical realism is true, but I won't argue for that here. I explore a different question: if logical realism is true, which (if any) of our logical concepts 'carve nature at its joints'? That is, which of our logical concepts most accurately reveal and respect worldly logical structure?

¹ I am intentionally leaving the boundaries of logical realism blurry here, though they will become clearer in what is to come. It is hence hard to say exactly which philosophers count as realists. But Almog (1989), who says that there are logical, structural, permutation-invariant, 'pre-facts' in the world, and Sider (2011), who argues that logical terms are among the 'structural' terms—they figure into a perfectly fundamental description of the world—certainly hold the view. 'Logical realism' gets used in very messy ways in the philosophy of logic literature (e.g. the positions that Resnik (2000) characterizes as realist aren't quite what I have in mind here, though some of them may count), and for that reason I set much of that literature aside here, in order to keep the argument I make here relatively clear.

✉ Michaela Markham McSweeney
michaelamcs@gmail.com

¹ Philosophy Department, Boston University, 745 Commonwealth Ave, Boston, MA 02215, USA

Many metaphysical realists think that two theories can be true, and in some sense equivalent, despite one being *metaphysically better* than another at describing the world. Indeed, it is plausible that realism commits us to this. Consider the theory that describes the world in terms of ‘grue’ and ‘bleen’ and the theory that uses ‘green’ and ‘blue’. If one of these theories is true, so is the other. But the latter clearly seems to better respect the structure of reality than the former does—that is, the latter seems to carve nature at its joints better than the former does. If realism is true in a particular realm, then it matters, metaphysically, what choices we make about language when stating our theories of that realm.²

If the world has a kind of structure, it must make sense to ask which of our concepts, terms, and descriptions respect that structure. So, if the world has logical structure—if logical realism is true—it must make sense to ask which of our *logical* concepts and terms respect that structure. (We needn’t think we should *quantify* over this structure to be realists about it.)

Logical realists are faced with the questions: are ‘&’ and ‘~’ perfectly joint-carving? ‘v’ and ‘~’? ‘↓’ (the ‘neither/nor’ connective)? What about quantifiers?

I argue that none of the obvious answers to these questions are correct, and that realists must accept one of two views. The first, *Privileged*, grants that some of our familiar logical constants are perfectly joint-carving, but says that we cannot know which. The latter, *Unfamiliar*, says that none of our logical constants or concepts respects worldly logical structure: we can’t talk about that structure in a joint-carving way with terms like ‘&’ and ‘~’. While I will offer some reasons for thinking that *Unfamiliar* is preferable to *Privileged*, the main claim I will motivate is simply that one or the other of the two views is true.

Both *Privileged* and *Unfamiliar* avoid a serious objection to logical realism. The logical realist must, it seems, distinguish between two otherwise equivalent theories, T, which employs ‘ \forall ’, ‘&’, and ‘~’, and T’, which employs ‘ \exists ’, ‘v’, and ‘~’. But something has gone wrong if we are in a position in which we are forced to ask and answer which of these theories is joint-carving, indeed if we are forced to think there is any worldly difference between them at all. The theories seem to be paradigmatic *mere notational variants*: the differences between them don’t reflect any differences between their metaphysical commitments, and so the question of which is more joint-carving seems like a bad one. If logical realism implies that the question is a good one, then, perhaps, so much the worse for logical realism.³

² While I won’t defend the particular conception of realism at stake here, it is worth noting that I am thinking of realism in a very general sense: the initial idea is just that we should take the questions of which logic corresponds to reality, and *how* that logic corresponds to reality, to be a substantive, non-trivial, serious one, with an objective answer. In Sect. 1, I will say more about the two notions of logical realism I want to consider in the paper. The kind of realism I have in mind in this introduction is much more general. (I admit that, as Jenkins (2010) argues, there are multiple possible notions of realism at stake here, but I think the vague and general notion of metaontological realism can do the job here, and in Sect. 1, I will get clearer on the two views I wish to consider in more detail.)

³ This challenge to logical realism is not original to me, and was posted by Sider (2011, ch. 10) as an objection to his own view. Sider argues for *egalitarianism*—the view that all of our logical constants are equally and also perfectly “structural” (joint-carving). I don’t engage directly with Sider’s argument in this paper, but I should point out that, if my argument works, it entails that Siderian egalitarianism can’t

Privileged and *Unfamiliar* avoid this objection by providing explanations for why the question is misguided, even for the logical realist. According to *Privileged*, the sense in which the question is misguided is that it is unanswerable. The proponent of *Privileged* thinks that one or the other collection is indeed joint-carving, but that we could never know which. For the proponent of *Unfamiliar*, the question is misguided because none of ‘ \forall ’, ‘ \exists ’, ‘ \sim ’, ‘ \vee ’, and ‘ $\&$ ’ respect worldly structure. We can’t accurately capture the logical structure of reality using any of our familiar logical constants, and so none of them are joint-carving.⁴

While *Privileged* and *Unfamiliar* both have the virtue of explaining what is wrong with the “bad question”, both views are surprising, and perhaps initially disconcerting. They both (though in different ways) entail that we are deeply ignorant about what the world is really like. If one wants to take this as a reductio of logical realism, one is free to; though I think logical realism is true, the argument here is conditional.

The structure of the paper is as follows: in section one, I say a bit more about logical realism, distinguish two forms it might take, and show that the logical realist faces a choice between four incompatible positions. In section two, I motivate (using the more familiar case of pairs of converse asymmetric relations) two principles. The first, *non-arbitrariness*, says that we shouldn’t believe a theory, T, if for every reason we have for preferring T to a distinct theory T’, there is an exactly parallel reason for preferring T’ to T. The second, *Weak Non-Redundancy*, says that there are no unexplained necessary connections between fundamental facts. In section three, I show that both forms of logical realism, together with these two principles, entail that either *Privileged* or *Unfamiliar* must be true. Finally, in section four, I more tentatively argue that *Unfamiliar* is preferable to *Privileged*.

For simplicity’s sake, I focus, in what follows, on classical logic. While similar issues will arise for most logics, some will do better than others.

1 Logical realism

Before I argue for the disjunction of *Privileged* and *Unfamiliar*, I need to distinguish two kinds of logical realism.

Footnote 3 continued

be true. Specifically, egalitarianism is ruled out by one of the principles, *Weak Non-Redundancy*, that I defend in section one. Warren (2016) argues that Sider’s treatment of theories like T and T’ as being the kinds of things that might be structurally distinct is problematic by his own lights (since he takes theory choice in metaphysics to be continuous with theory choice in science, and T and T’ are equivalent on any reasonable understanding of what is going on in science). Donaldson (2014, sec. 4) argues that Sider’s approach of embracing all of the (standard first-order) logical constants as structural is at odds with mathematical and logical practice, in the course of arguing that we have no reason to think that first-order quantifiers (as opposed to Quinean predicate functors) are structural. My argument significantly differs from both Warren’s and Donaldson’s; I assume that metaphysicians might have more resources for theory choice at their disposal than scientists do; I am far friendlier to Sider’s general project; and the conclusion I come to is not considered as a serious option by either Warren or Donaldson.

⁴ We might think that *Unfamiliar* better satisfies our intuition that the question is misguided; I think this is right, and will say something about it in the final section of the paper.

First, one might think that expressions like ‘&’, ‘v’, ‘~’, etc. *refer* to individual entities. The most likely candidates are truth functions, or the “worldly” equivalent of truth functions (perhaps: whatever the *logical* constituents of states of affairs are). I’ll call this kind of view *ontological realism*.

One might instead think that while ‘&’, ‘v’, ‘~’, etc. are *syncategorematic*—that they don’t refer at all—that they nevertheless play an important role in carving up reality. They are bits of *ideology*. On this view, the constants don’t just carve reality up in a way that’s convenient for our own purposes, or that depends on the way our minds or language are structured. Rather, certain logical constants perfectly represent the way reality is—they are *joint-carving* bits of ideology.⁵ One might think, for example, that there is conjunctive fundamental structure, that is, that conjunction carves nature at its joints, but that there is no Sheffer stroke (the “not/and” connective) structure—the Sheffer stroke does not. I will call this kind of view *ideological realism*.

Ontological and ideological realism are analogues of more familiar positions. Consider the predicates ‘green’, ‘blue’, ‘bleen’, and ‘grue’. Ontological realism is like a view which says that at least some of these predicates refer to genuine properties. There is then a question of whether some of these properties are more or less fundamental than others (or alternatively, which predicates genuinely refer). The ontological realist faces a similar question: whether certain logical entities are more fundamental than others.

Ideological realism is like the view that says that ‘green’, ‘blue’, ‘bleen’ and ‘grue’ are predicates which don’t refer to properties, but yet still do better and worse jobs at carving nature at its joints. Just as the property nominalist can think that ‘green’ and ‘blue’ are more *natural* or *fundamental* or *structural* predicates than ‘grue’ is, so the ideological realist might think that certain logical constants carve the world up better than others.⁶ So Ideological realists are nominalists, but they are nominalists who take ideology metaphysically seriously.⁷

For both kinds of logical realists, *being a correct fundamental theory* is an extremely fine-grained matter. Whether the theory that is formulated using \exists , \sim , and v (call it T) is the correct fundamental theory comes apart from whether the theory that we would naturally call equivalent to T, but which is formulated using \forall , \sim , and $\&$ (call it T’) is the correct fundamental theory. The theories have different fundamental metaphysical commitments (either ideologically or ontologically). Fun(T) isn’t committed to there being any fundamental conjunction at all. And Fun(T’) is. So T and T’ are fundamentally distinct.^{8,9}

⁵ For discussion of such a view see Sider (2011) and Turner (2016, introduction).

⁶ For more on this sort of view see Lewis (1986) and Sider (2011).

⁷ See, e.g., Lewis (1983), Sider (2011) and Turner (2016).

⁸ Don’t read much into what I mean by ‘theory’ here. I might as well have said “description in a language”, or something like that.

⁹ One might take this fact—the fact that ontological realism is committed to T and T’ saying fundamentally distinct things about the world—to itself be a *reductio* of ontological realism. If one is so inclined, then one should read this paper as a way of resisting the *reductio*—for, I show, the right version of logical realism is *not* committed to this.

If $\text{Fun}(T)$ is a distinct claim from $\text{Fun}(T')$, then it makes a *metaphysical* difference, and not just a pragmatic one, which one we choose to accept. So we are faced with a familiar problem: there doesn't seem to be any possible reason we could have for believing $\text{Fun}(T)$ rather than $\text{Fun}(T')$. And vice versa. What do we do? Accept $\text{Fun}(T)$? Accept $\text{Fun}(T')$? Neither? Both? There are four things we might believe.

Both-fun Both $\text{Fun}(T)$ and $\text{Fun}(T')$.

Neither-fun Neither $\text{Fun}(T)$ nor $\text{Fun}(T')$.

One-fun One of $\text{Fun}(T)$ and $\text{Fun}(T')$, and I can know which.

One-fun(ignorance) Either $\text{Fun}(T)$ or $\text{Fun}(T')$, but I can't know which.¹⁰

In Sect. 3, I will argue that if realism is true, then only *One-fun(ignorance)* or *Neither-fun* could be right. And these are essentially just the views I introduced earlier as *Privileged* and *Unfamiliar*. *Privileged* is the view that some subcollection of the logical constants is perfectly joint-carving, but that we can't know which. And that is just what *One-fun(ignorance)*, suitably generalized, amounts to. *Unfamiliar* is the view that the fundamental logical structure of the world looks nothing like any of our logical constants, and so none of our constants joint-carvingly represent reality. *Neither-fun*, suitably generalized, just gets us *Unfamiliar*. (We need logical realism, as well, to get *Unfamiliar* from *Neither-fun*, since *One-fun* is compatible with there being no fundamental logical properties, objects, or concepts at all).

First, in Sect. 2, I motivate the two principles I need, by exploring a different, but structurally similar, case: pairs of converse asymmetric relations.

2 Asymmetric relations, fundamentality, and non-redundancy

The goal of this section is to motivate two principles, *Non-Arbitrariness* and *Weak Non-Redundancy*. I do so by considering *converse asymmetric relations*; I then return, in Sect. 3, to logic and use these principles to argue for the disjunction of *Privileged* and *Unfamiliar*.

Suppose we are considering whether either of two converse asymmetric relations (such as *is beneath* and *is on top of*), R and Q , is fundamental. There are various possibilities. I will use almost the same names for them as in the logic case, prefixing them with "(CR)" to avoid confusion.

(CR)*Both-fun* Both R and Q are fundamental.

(CR)*Neither-fun* Neither R nor Q is fundamental.

(CR)*One-fun* One of R and Q is fundamental, and I can know which.

(CR)*One-fun(ignorance)* One of R and Q is fundamental, and I can't know which.

¹⁰ Following Sider, I am using 'joint-carving' and 'fundamental' largely interchangeably here (See Sider 2011, introduction). In Sect. 3.3, I will briefly consider a possible version of logical realism on which we distinguish between the two.

Before I begin the argument, I want to fend off a potential confusion that will otherwise re-arise in Sects. 2 and 3. Readers might worry I've left a plausible view out: that 'R' and 'Q'—the predicates that purportedly name those asymmetric converse relations—in fact are just two different ways to refer to a single relation. I'm sympathetic to this view. But I haven't ignored it. Rather, I've *stipulated* that 'R' refers to an asymmetric relation like *is above* and 'Q' refers to an asymmetric relation like *is beneath*. Given this stipulation, the view that 'R' and 'Q' refer to the *same* relation must be either a version of *Neither-fun* (e.g. if the underlying relation is not identical to *is above* or to *is beneath*) or *One-fun* or *One-fun(ignorance)* (e.g. if the underlying relation is identical to one of *is above* or *is beneath*).

In Sect. 2.1, I motivate a principle, *Non-Arbitrariness*, which rules out *(CR)One-fun*. In Sect. 2.2, I motivate a principle, *Weak Non-Redundancy*, which helps rule out *(CR)Both-fun*.¹¹ The arguments about relations assume that relations are real entities. I claim, but do not argue here, that the argument is easily adaptable to target nominalism about relations. When I turn to logic, however, I will show that the argument that targets logic extends to both ontological and ideological realists.

2.1 Non-arbitrariness

In this section, I argue against *(CR)One-fun*. I motivate an epistemic principle that says that we are not justified in believing something when that belief is *arbitrary*; more specifically, when we have equally good reason to believe an incompatible alternative. If we are in such a situation when it comes to asymmetric converse relations, then *(CR)One-fun* can't be right. But we need to clarify what it means for there to be equally good reason to believe an incompatible alternative.

Consider some standard cases of converse asymmetric relations: *above* and *beneath*, *loves* and *is loved by*, *to the west of* and *to the east of*. In typical cases, we really don't seem to have any reasons to think one member of the pair is fundamental for which there don't exist similar reasons in favor of taking the other member as fundamental. We lack reasons for privileging one relation or another.¹²

Suppose there are two theories, T^A and T^B , which are equivalent in every respect except that the former uses *is above* and the latter uses *is beneath*. Let $\text{Fun}(T)$ mean that T is fundamental. So long as we accept that there is a coherent, genuinely metaphysical distinction between the fundamental and the non-fundamental, we must accept that $\text{Fun}(T^A)$ and $\text{Fun}(T^B)$ may not both be true even if T^A and T^B are. $\text{Fun}(T^A)$ and $\text{Fun}(T^B)$ say different things about the world; if $\text{Fun}(T^A)$ is true, and $\text{Fun}(T^B)$ is not, then *aboveness* is a fundamental relation, and *beneathness* is not (or, if we want to nominalize, 'is above' carves nature perfectly at its joints, whereas 'is beneath' does not). But if we think we get *beneathness* "for free" from $\text{Fun}(T^A)$,

¹¹ While my discussion of relations is inspired by Dorr (2004) and Fine (2000), the argument I give is somewhat different from either of their arguments (though much closer to Dorr's, who appeals to something similar to *Weak Non-Redundancy* to motivate his argument), so I won't spend much time discussing them. Williamson (1985) also discusses this issue. And it can, of course, be seen as rooted in Russell.

¹² For further motivation see Dorr (2004).

then we should think that *both* T^A and T^B are true even when $\text{Fun}(T^B)$ is false. So being fundamentally true is a more fine-grained matter than being true. (None of this is to deny the possibility that *both* $\text{Fun}(T^A)$ and $\text{Fun}(T^B)$ are true; this would be the case if we thought both *aboveness* and *beneathness* were fundamental relations; and this would be to endorse *CR-BothFun*. I argue against this possibility in Sect. 2.2.)

Do we have any reasons for favoring $\text{Fun}(T^A)$ over $\text{Fun}(T^B)$? No. For each reason we might have for thinking that $\text{Fun}(T^A)$, we have an exactly parallel reason for thinking that $\text{Fun}(T^B)$, and vice versa.

To clarify, I am only talking about reasons that might have some bearing on the *metaphysical* status of these relations, and in particular, on their fundamentality. It is important to underscore that the problem that arises for asymmetric relations (as well, as I'll argue in section two, as for logical constants) is not a lack of empirical evidence favoring one relation or the other, or one constant or another. Nor is it that we can't evaluate the balance of reasons, for example because one theory scores better with respect to one theoretical virtue (e.g. parsimony) and another scores better with respect to a different theoretical virtue (e.g. explanatory power), and we don't know how to compare the two.¹³

Rather, in the cases this paper targets, all of our competing theories score equally well with respect to each theoretical virtue, and any argument we could make in favor of one would generate an exactly parallel argument for the others. When we are in such a situation, I claim, we can't possibly be justified in believing one theory over the others.

It follows that we should not believe (only) $\text{Fun}(T^A)$, and we should not believe (only) $\text{Fun}(T^B)$. The principle at play here is:

Non-Arbitrariness: If we have no reasons to favor $\text{Fun}(T)$ over $\text{Fun}(T')$ (and vice versa), and $\text{Fun}(T)$ and $\text{Fun}(T')$ are incompatible, then we should not believe $\text{Fun}(T)$ and we should not believe $\text{Fun}(T')$.

If *Non-Arbitrariness* applies in the case of relations, it will rule out *(CR)One-fun*.

If *(CR)One-fun* is false, this is an epistemic claim, and is distinct from the claim that it is not *true* that either R is fundamental or Q is. It is important to disentangle these two claims, even if in the end one thinks that we can move freely from the rejection of *(CR)One-fun* to the purely metaphysical claim that it can't be that one and only one of R or Q is fundamental. This is because, when we have evidence that one of the disjuncts of a disjunction is true, but no evidence as to which, it is

¹³ This issue comes up frequently in metaphysics. For example, Bennett (2009) argues that this comes up in disputes over constitution, Dorr and Rosen (2002) suggest something somewhat similar about composition. Though Korman (2010, p. 125) points out that at least in the composition case, it's not true that we have no reason to think that, e.g., there are snowballs but no snowdiscalls (a snowdiscall is "something made of snow that has any shape between being round and being disc-shaped and which has the following strange persistence conditions: it can survive taking on all and only shapes in that range."), whereas it does seem to be true that (as I will discuss momentarily) we have no reason to think that the number two is identical to $\{\{\emptyset\}\}$ and not $\{\emptyset, \{\emptyset\}\}$. In the first case we've got significant intuitive support for the claim that there are snowballs but no snowdiscalls, whereas there doesn't seem to be any prima facie intuitive support for one view or the other about which set-theoretic reduction of the numbers is true. The narrow cases I'm focused on here are like the set-theoretic case and not like the composition case in this (and many other) respects. (Dasgupta (2015) also discusses this issue more generally).

plausibly permissible to believe the disjunction but not either disjunct.¹⁴ It is an important question whether, and in what cases, we should endorse a stronger, metaphysical non-arbitrariness principle (one which would also rule out *(CR)One-fun(ignorance)*). I will discuss this further in Sect. 4. Note, though, that I only need the epistemic claim to generate the disjunctive conclusion of the paper.

2.2 Weak non-redundancy

I have argued that we cannot be justified in believing that *R* is fundamental (but *Q* is not), and that we cannot be justified in believing that *Q* is fundamental (but *R* is not). So *(CR)One-fun* cannot be true. I won't extensively argue that *(CR)Both-fun* is false. But I do want to examine and briefly motivate the principle that helps to rule it out.

It is often assumed that there is some kind of *non-redundancy* or *minimality* constraint on the fundamental. This includes and perhaps begins with Lewis, who says of the perfectly natural properties that “there are only just enough of them to characterise things completely and without redundancy” (1986, p. 60). Lewis' claim might rule out *(CR)Both-fun*. If one is already convinced that there is no redundancy at the fundamental level, one should already be convinced that it can't be that both *R* and *Q* are fundamental, because this would entail that there was overlap in the fundamental furniture of the world—that we had *more than we needed* to recover the rest of the world.¹⁵

But why have people been so quick to assume this principle? Perhaps it is motivated by Occamism. Perhaps it is nearly analytic on a certain conception of fundamentality. Claims about the fundamental such as that it is *only what God would have to make to get everything else for free* suggest a conception of fundamentality on which it simply couldn't be that *R* and *Q* are both fundamental. For surely God would need make at most one of *R* and *Q*!

Even if we endorse one of these motivations for non-redundancy, it's hard to say exactly what it means for there to be redundancy at the fundamental level. It is sometimes taken to mean that the fundamental must be a *minimal supervenience base*. Intuitively: that the fundamental properties, e.g., are those contained in “the” smallest set that everything supervenes on. However, as Sider (1996) points out, there is no single smallest set that everything supervenes on, which makes it difficult to appeal to minimal supervenience—which set do we choose? Arbitrariness rears its head again.

Strong minimality constraints rules out the possibility of both members of any pair of converse asymmetric relations being fundamental. But it's not obvious that we should accept this kind of constraint, and so we can't rule out *Both-fun* so quickly. Sider (1996, 2011), Eddon (2013), Cowling (2013), and Wang (2016) argue

¹⁴ Of course, whether we can accept disjunctions without accepting either disjunct is tendentious (see e.g. Dummett 1991), but this is just the kind of case where we might think such a thing is possible, so it is important to clearly separate the two questions.

¹⁵ The *underlying principle* that Lewis makes use of, of course, comes directly from Hume, but I say that this particular conception of it begins with Lewis because it is important here that our focus is fundamentality (or Lewisian ‘naturalness’), neither of which Hume would have been happy with.

against minimality constraints.¹⁶ And one of Sider's complaints about the constraint is that it rules out what he takes to be plausible cases of fundamental redundancy: logical constants and converse asymmetric relations!

Schaffer (2010) gives an argument for a condition on the fundamental *material objects* that he calls 'No Overlap'. He appeals to a Humean recombination principle: the fundamental entities should be freely recombinable. The fundamental entities should be *independent* of one another in that any way of combining them amounts to a genuine metaphysical possibility. But as Schaffer points out, there are modal constraints on overlapping entities: "Consider two overlapping homogeneously red circles, each of which individually could have been green." (2010, p. 40) We can see the problem: circle A can't retain its parts while being green if circle B retains its parts while remaining red (this would entail that the overlapping bits of A and B are both homogeneously red and homogeneously green). So we can't both endorse free recombination and allow for overlapping fundamental ontology.

Can a similar argument be made with respect to relations, and in particular, pairs of converse asymmetric relations like *R* and *Q*? Recombination is typically characterized as a principle about regions of space or spacetime (e.g. Lewis 1986, pp. 70–90) and how the entities that live in spacetime can be recombined. But we need something much more general, along two different dimensions: first, we want our principle to apply to different *kinds* of entities, and second, we don't want it to discriminate between nominalist views and platonic ones.

To accommodate both the nominalist and the platonist, the principle I will propose applies to facts, staying neutral on what kind of entity a fact is (state of affairs, true sentence, proposition, etc.). I am, however, committed to fundamental facts only containing fundamental entities (be they linguistic items, concreta, abstracta, etc.):

Weak Non-Redundancy: There are no unexplained necessary connections between fundamental facts.

Why 'non-redundancy' when the principle seems to be about unexplained necessary connections, and not redundancy? According to Schaffer, it is overlapping (or redundant) fundamental *individuals* that creates unexplained necessary connections. But it is not only fundamental individuals that cause problems: we get the same problems with states of affairs, properties and relations.

Weak Non-Redundancy is weaker than typical minimality constraints: it allows for necessary connections between fundamental facts—it just requires that those connections themselves have explanations.

I won't extensively argue for *Weak Non-Redundancy*. Indeed, I suspect that it is the sort of principle which must be, in some sense, assumed rather than argued for—

¹⁶ Cowling (2013) really argues for a particular conception of ideological parsimony, but it is one which he suggests can help us avoid arbitrariness worries by allowing us to embrace a multiplicity of ideology when such ideology is interdefinable, and hence by allowing for ideological redundancy in precisely the kinds of cases we're concerned with here. I don't take up Cowling's proposal here since my argument for non-redundancy doesn't run via an appeal to parsimony, but rather via a demand for metaphysical explanation.

but it's important to note that it is significantly weaker than standard recombination principles.¹⁷ *Weak Non-Redundancy* is consistent, for example, with there being metaphysical laws that constrain possible combinations of fundamentalia, and it is consistent with there being essence facts that explain why certain fundamental facts travel together through modal space. It is only inconsistent with it being the case *both* that some fundamental facts do travel together through modal space, *and* that there is no underlying explanation for why the necessarily co-obtaining facts co-obtain.

Instead of directly arguing for *Weak Non-Redundancy*, I want to examine its consequences for relations, and hopefully, along the way, convince readers that we ought to adopt the principle.

Suppose that we have an ontology of states of affairs, and suppose that *R* and *Q* are a pair of asymmetric converse relations, and are both fundamental. Then, for every fundamental state of affairs (e.g. [Rab]) that involves the instantiation of one of our relations, there will be a corresponding fundamental state of affairs that involves the instantiation of the other, [Qba]. And they will necessarily co-obtain. It is implausible that there is *no* explanation for this necessary connection. On such a view there is simply no reason at all why [Rab] and [Qba] travel together, and nothing to explain why *R* and *Q* are so intimately related. They are just two relations that happen to always travel together through modal space. This is the view—the only view!—that *Weak Non-Redundancy* rules out. It is much weaker than typical recombination principles, which don't allow [Rab] and [Qba] to be intimately related at all.

In order to rule out (*CR*)*Both-fun*, then, we would need more than just *Weak Non-Redundancy*. We would also need to rule out potential explanations for *why* [Rab] and [Qba] travel together through modal space.

One candidate explanation is to claim that [Rab] and [Qba] are identical—that they are the same state of affairs, described two different ways. It is of course generally open to us to say that '[Rab]' and '[Qba]' are just two names for the same state of affairs (indeed, this is what Fine (2000) claims). But recall my clarification at the beginning of Sect. 2: we are *assuming* that *R* and *Q* are distinct relations. And this assumption is incompatible with there being only one state of affairs here. Surely *if* being a renate and being a cordate are distinct properties, the state of affairs [Izzy is a renate] is distinct from the state of affairs [Izzy is a cordate] precisely *because* the properties are distinct.¹⁸ To take this as a *reductio* of our

¹⁷ See Wang (2016) for a convincing argument for a nearby claim (re: Humean Recombination principles), and Wilson (2010) for a much more in-depth discussion of the motivations (or lack thereof) for "Hume's Dictum" than I can engage in here.

¹⁸ Ramsey uses this kind of observation to argue against complex universals: he claims that there is only one proposition, *aRb*, that can be seen in three different ways: *a* and *b* are related by the two-place relation *R*; *a* has the one-place property of being *R-related to b*; *b* has the one-place property of being *R-related to a*. But, given that there is only one proposition there, it cannot be that there are these distinct properties and relations (1925, pp. 405–406). One way to resist this when it comes to states of affairs is to deny that they have relations as constituents, for example, by thinking that states of affairs are 'chunks' of reality that make sentences and propositions true. But on such a view, it's hard to see why we would need to posit fundamental relations in the first place, if states of affairs were fundamental.

initial assumption is to admit that there is a single underlying symmetric relation and a single state of affairs is to reject that either both R and Q are fundamental, or that either is (remember that we are *assuming* that they are asymmetric relations), so it rules out both $(CR)Both\text{-}fun$ and $(CR)One\text{-}fun$. So this explanation doesn't work.

There are other possible explanations. Perhaps it is a *metaphysical law* that R and Q always travel together through modal space. Perhaps it is part of the *essence* of R that wherever it is instantiated, so is Q , and vice versa. I won't argue for against these possibilities here, but I will argue against similar claims when I turn to logic.

I have argued that when it comes to pairs of converse asymmetric relations, *Weak Non-Redundancy* motivates rejecting $(CR)Both\text{-}fun$: we cannot think that both relations are fundamental. And I have motivated, but have not extensively argued for, the claim that *Non-Arbitrariness* motivates rejecting $(CR)One\text{-}fun$. I will now return to my main focus, logical realism, and show how these principles motivate structurally similar conclusions there.

3 Logical realism

In this section of the paper, I will argue that logical realists must accept either *Unfamiliar* or *Neither*. In Sect. 3.1, I show that both ontological and ideological realists should deny *One-fun*. In Sect. 3.2, I show that ontological realists should deny *Both-fun*. In Sect. 3.3, I show that ideological realists should deny *Both-fun*. If realism is true, the remaining options, *One-fun(ignorance)* and *Neither-fun*, are just *Unfamiliar* and *Neither*.

3.1 Non-arbitrariness and one-fun

I will begin by showing that the ontological logical realist should deny *One-fun*. Which of '&' or 'v' refers to a fundamental entity? We can't generate any reasons to favor one such that we don't have a parallel reason to favor the other, and so *Non-Arbitrariness* applies.

What would such reasons look like? I can think of three considerations that we might hope would give us non-parallel reasons, and none are successful at doing so. The first is parsimony. Perhaps we should take theories with the fewest number of logical primitives to be more likely to be fundamental. But then we will be stuck trying to compare a theory formulated with the down dagger (the single connective that expresses "neither/nor") with one formulated with the Sheffer stroke (the single connective that expresses "not/and"). Each is equally parsimonious with respect to logical concepts, and we have no way to arbitrate between the two.

The second concerns what concepts are easiest for us to work with. Perhaps, it is easier for us to reason using the connectives \sim and $\&$ than it is for us to reason using the connectives \sim and v . But it's very hard to see how (at least for the kind of realists under discussion) such pragmatic considerations about ease of use could generate *metaphysical* reasons for preferring one theory to another. Any such argument would be quickly debunked by considering (actual or possible) reasoners for whom \sim and v were easier to reason with than \sim and $\&$.

Finally, we might make an appeal to conjunction's *seeming* like the most natural connective as a reason to claim that, say, conjunction and negation are the fundamental constants. There are two responses to this.

First, even if we grant that, for some reason, conjunction and negation seem to be the most natural constants, it is hard to see why we then wouldn't have a problem deciding between conjunction and negation on the one hand and the Sheffer stroke on the other. If anything, it seems that parsimony considerations would push us towards treating the Sheffer stroke as fundamental. But then the intuitive motivation for treating conjunction and negation as fundamental is lost.

Second, while a case might be made for conjunction being more natural than the other constants, this is not obvious of negation. The insistence that it is conjunction and negation that are most fundamental bottoms out, I think, in an intuition that conjunction is the most natural, and then a realization that conjunction alone has very little expressive power, but that adding negation in gives us just as much expressive power as any other collection of (propositional) connectives.¹⁹

There is no metaphysically relevant reason to prefer one of T to T' for which there is not a parallel reason for preferring the other. So *Non-Arbitrariness* applies. We should reject that either Fun(T) or Fun(T'). So if ontological realism is true, *One-fun* is false.

Everything I have said so far applies to ideological realism. We needn't make any adjustments to *Non-Arbitrariness* in order to apply it to nominalistic views. So if ideological realism is true, *One-fun* is false.

It remains to be shown that *Weak Non-Redundancy* generalizes to logical constants, which means that we cannot accept *Both-fun*: that *both* Fun(T) and Fun(T'). I show that this is the case for ontological realism in Sect. 3.2, and for ideological realism in Sect. 3.3.

3.2 Weak non-redundancy and ontological realism

The logical realist thinks that our use of the constants reflects something worldly: either they directly refer to entities in the world, or they best capture the way the world is structured. One way this could be true is if we had an ontology of not just atomic but also logically complex states of affairs. That view might be unpalatable, but that doesn't matter here. I'll use it to argue against *Both-fun*, and it will be easy to see how to discharge the assumption that there are such states of affairs. (I will also assume that states of affairs are *truthmakers* for at least some true sentences.)

Let's suppose that *Both-fun* is true. Then T has a sentence of the form ' $\sim A \vee \sim B$ ' in it which is logically equivalent to a sentence of the form ' $\sim (A \& B)$ ' in T'. Take such a pair of sentences. What are the fundamental truthmakers for these sentences? There are four options. I will rule them all out.

Option 1 *One of $[\sim A \vee \sim B]$ and $[\sim (A \& B)]$ is the fundamental truthmaker for both sentences.* Suppose that the state of affairs $[\sim A \vee \sim B]$ (where, importantly,

¹⁹ Note the similarity between the issues with negation that arise here and those that arise for truthmakers, e.g. in Armstrong (2004).

what is inside the square brackets does not merely *name* a state of affairs but also *displays its internal structure*—the state of affairs itself is, in some sense, conjunctive) is the most fundamental truthmaker for both sentences. Could *Both-fun*, the claim that T and T' were both fundamental, be true? No. ' $\sim A \vee \sim B$ ' and ' $\sim(A \& B)$ ' are clearly not equally joint-carving representations of a state of affairs [$\sim A \vee \sim B$], which has disjunctive and not conjunctive internal structure. This generalizes to [$\sim(A \& B)$]. So option 1 is out.

Option 2 *Both states of affairs are fundamental truthmakers for both sentences.* This falls prey to a redundancy worry—we now have two redundant fundamental states of affairs. And redundancy is especially pernicious here, because it is explosive: on pain of arbitrariness, we will be forced to admit that highly gerrymandered but “logically equivalent” states of affairs, such as [$\sim \sim \sim (\sim A \vee \sim B) \& \sim(A \& B)$], make both sentences true, and also that [$A \rightarrow \sim B$] does, and so on. (Remember that the ontological realist must metaphysically distinguish between all of these states of affairs, since they all have different constituents and she is a realist about their constituents.) Hence, we should reject the claim that both our candidate states of affairs make both sentences true. So option 2 is out.

Option 3 *Neither state of affairs is a fundamental truthmaker for either sentence.* This immediately pushes us into *Neither-fun*. If neither [$\sim A \vee \sim B$] or [$\sim(A \& B)$] is the fundamental truthmaker, then neither T nor T' is fundamental. The most fundamental theory of the world would involve a *single* sentence that matches the structure of whatever the most fundamental truthmaker is for ' $\sim A \vee \sim B$ ' and ' $\sim(A \& B)$ '. Our more fundamental state of affairs ought to have a more joint-carving description than either of these sentences. Hence, if there is a single fundamental truthmaker for both ' $\sim A \vee \sim B$ ' and ' $\sim(A \& B)$ ', then *neither* T nor T' is the fundamental theory, for that theory would have a sentence the structure of which exactly matches the structure of its truthmaker. So, if option 3 is correct, neither T nor T' is fundamental. This contradicts the initial assumption of *Both-fun*. So option 3 is out.

Option 4 [$\sim A \vee \sim B$] is the fundamental truthmaker for ' $\sim A \vee \sim B$ ' and [$\sim(A \& B)$] is the fundamental truthmaker for ' $\sim(A \& B)$ '. (This is the natural view.) This faces us with an immediate problem: why do the states of affairs [$\sim(A \& B)$] and [$\sim A \vee \sim B$] necessarily co-obtain? If each of these states of affairs is fundamental, then *Weak Non-Redundancy* applies, and we need some sort of explanation for this necessity.

One might be tempted to appeal to the logical equivalence of the two states of affairs as such an explanation. But if logical equivalence explains anything, it explains only why the truth values of ' $\sim(A \& B)$ ' and ' $\sim A \vee \sim B$ ' necessarily travel together—not why two fundamental states of affairs necessarily travel together. Logical equivalence can serve as an explanation for why two sentences or propositions necessarily have the same truth value—but the *reason* it can serve as such an explanation is that it suggests that two sentences are, in some sense, fundamentally equivalent—that they describe the very same underlying state of affairs! Logical equivalence as a relationship between representational entities, not

worldly ones. But option 4 *assumes* that there is a worldly distinction between the two states of affairs.

So logical equivalence can't do the explanatory work here. What we need is some kind of *metaphysical law* that connected our two fundamental states of affairs. If one wants to call such a law "logical equivalence", one may, but this is misleading: it would have to be a metaphysical law that relates otherwise free-wheeling states of affairs, not something conventional or metaphysically harmless that concerns our *representations* of states of affairs.

One way to resist my argument is to posit such laws. Note, however, that these laws would have to do some metaphysical heavy lifting. They wouldn't merely tell us that whenever we have a single fundamental state of affairs we get all of its logically equivalent states of affairs "for free": that we are *automatically* committed to its logically equivalent states of affairs as derivative entities. Instead, they would have to explain necessary connections between *fundamental* states of affairs; for the only way for *Both-fun* to be true is if *both* states of affairs are fundamental (since they are *both* the most fundamental truthmakers for a sentence in a fundamental theory). This seems to me to require us to posit something unlike what I think of as a metaphysical law, which relates the non-fundamental to the fundamental in much the same way scientific laws seem to unfold across time or space-time.²⁰

Alternatively, we might want to claim that it is somehow part of, or follows from, the essences of conjunction, disjunction, and negation that these states of affairs necessarily travel together. In order for essences to do the explanatory work necessary here, they can't be *modal*—that is, it can't be that essence facts are analyzed in terms of modal facts. If we want to appeal to essence facts as *metaphysically explaining* modal facts—which is exactly what we are doing here—then we need a *non-modal* account of essence, like Fine's (1994a, b).

Conjunction's essence would have to *contain* something about disjunction, and vice versa, in order for us to get an explanation for the two states of affairs necessarily traveling together. (If we only had one, but not the other, of the essences containing information about the other, then presumably one would be more fundamental than the other—and we are assuming that they are equifundamental here and looking for an explanation of that fact.) So conjunction and disjunction would have to have *reciprocal* essences.

And there is something fishy about positing symmetrical reciprocal essences in this way, given that we are *using* the essence facts to do explanatory work. Indeed, to the extent that we can accept that there could be reciprocal essences, it seems that they fail to do the requisite explanatory work and instead support *Neither-fun*.²¹

²⁰ Wilsch (2015) articulates and argues for this thought: "Laws of metaphysics are akin to laws of nature in the sense that they guide the development of the world along a dimension. Whereas the natural laws work along the temporal dimension, the metaphysical laws work along the axis of fundamentality: from the truths of fundamental physics via the truths of chemistry, biology, and so on, all the way up. According to the specific conception I develop in this paper, the metaphysical laws characterize 'construction-relations,' which include composition, set-formation, and property-determination, among many others (p. 3294)".

²¹ Correia (2012) is worried about, and proposes a solution to, this issue of reciprocal essences having to do explanatory work for Fine. While he is not targeting the logical case in particular, along the way he

Intuitively: if all of the logical constants have facts about their relationships with the other logical constants *built into their essences*, then this would be because there was something *more fundamental* than all of them that grounded *all* of these necessitation relations between them. In other words, such a picture should push us either towards a kind of holism (on which all of logic taken together is fundamental, but no individual constant is), or towards a kind of structuralism about logic (where constants are mere nodes in a structure, and the structure is more fundamental than those nodes). Neither picture is one on which conjunction and disjunction are fundamental. Both, however, are consistent with *Neither-fun*.²²

If we accept a fairly natural connection between fundamentality and essence-dependence, we should reject that it could be the case both that this was true, and that conjunction and disjunction were individually more fundamental than conjunction and disjunction taken together. For accepting both would mean that we would have an essence-dependence chain that bottoms out in a non-fundamental entity (the plurality of conjunction and disjunction). And we might think that essence-dependence chains either must bottom out at the fundamental level, or, if we deny that there is such a level, continue forever. So, if the necessary connection is explained by something about the plurality of the natures of conjunction and disjunction, but not partly in the nature of conjunction, then conjunction is not more fundamental than the plurality is.²³

There is more to say here, but I'm happy to leave open that there *might be* some account of the fundamental logical-metaphysical laws, essences of logical constants, or some alternative, that could explain these states of affairs necessarily co-obtaining. *Weak Non-Redundancy* applies if we take option 4, and so the central point here is that the burden is on the proponent of *Both-fun* to produce a plausible explanation of these necessary connections between fundamental facts. Without such an explanation, option 4 is out.

Footnote 21 continued

does motivate the idea that collections (typically sub-collections of the whole collection) of logical constants together ground some kinds of necessities. I am unclear, though, about exactly what the metaphysical status of the logical constants is supposed to be on his view.

²² I don't mean to suggest that the positive view I argue for here is committed to either holism or structuralism about logic. (See Koslow (1992) for an argument for the structuralist picture.) But one form of structuralism, at least, is one way to cash out *Neither-fun*: the idea is that what is fundamental is the structural relations that hold between the constants, rather than the constants themselves; those structural relations explain the apparent dependencies between the constants themselves. I don't find the structuralist picture entirely satisfying, for reasons that I partly gesture at in Sect. 4, but otherwise do not present here.

²³ An alternative one might want to adopt is presented in Barnes (forthcoming). Barnes argues that fundamental entities can *symmetrically* depend on one another; roughly, the view is something like a metaphysical equivalent of coherentism. At first glance, this view might seem to violate *weak non-redundancy*; whether this is true depends on whether Barnes thinks that this symmetric dependence is itself explanatory of the necessary connections between fundamental entities that it entails. (The view is clearly inconsistent with stronger versions of the Humean thought, as she makes clear.) My suspicion is that she would want to object to my adoption of *weak non-redundancy* rather than use symmetric dependence to do the work it would need to do here.

To sum up: the only way the ontological realist can maintain that *Both-fun* is true, given our assumptions about states of affairs and truthmaking, is by claiming that both $[\sim(A \ \& \ B)]$ and $[\sim A \vee \sim B]$ are fundamental states of affairs. But then *Weak Non-Redundancy* applies, and we are faced with the same issues we were in the case of relations. Hence, if ontological realism is true, *Both-fun* is false. Since *One-fun* is also false, the ontological realist's only other options are *Privileged* and *Unfamiliar*.

I won't do so here, but one could reconstruct this story without talking of states of affairs at all. What is somewhat less straightforward is reconstructing it for the ideological realist. I do this in Sect. 3.3.

3.3 Weak non-redundancy and ideological realism

Ideological logical realism, recall, is the view that while '&', 'v', '~', etc. are *syncategorematic*—they don't refer—they nevertheless play an important role in carving up reality: they are the kind of terms that do better and worse at respecting metaphysical structure, where we aren't quantifying over that structure. In order to think about ideological realism, we need to think about candidate fundamental truths, and avoid thinking of logical constants as referring to *entities*. Moving to such a view creates one complication, but in the end we get the same conclusion as in Sect. 3.2: *Weak Non-Redundancy* applies, so *Both-fun* cannot be true.

Supposing that we accept $\text{Fun}(T)$ and that ' $\sim A \vee \sim B$ ' is a sentence in T , then, if we also accept $\text{Fun}(T')$, we are committed to both ' $\sim A \vee \sim B$ ' being a fundamental truth, and ' $\sim(A \ \& \ B)$ ' being a fundamental truth. This is because T' , by definition, contains those sentences expressible in its language that are logically equivalent to those T contains.

For the ideological realist, there are multiple disambiguations of *Both-fun*. First, it could be the case that these two truths are equally joint-carving because the world is structured in two distinct ways. Second, it could be the case that the world is structured in a single way, that these two truths are equally best ways to represent that structure, and that there is no more joint-carving way to represent it.

In the former case, it is easy to see how *Weak Non-Redundancy* applies: even if we don't want to commit to those bits of worldly structure being ontological, as soon as we ideologically commit to the idea that there are multiple fundamental ways the world is structured, we can ask why those bits of structure always "travel together": why it is that ' $\sim A \vee \sim B$ ' is true whenever ' $\sim(A \ \& \ B)$ ' is, and vice versa, given that they can't both perfectly represent the same structure. And, as for ontological realism, we'd be hard-pressed to come up with an explanation for this.

In the latter case, *Both-fun* just collapses into the same position as *Neither-fun*: there is a single underlying "way the world is", or bit of worldly structure. It is just that we are additionally committed, in this case, to the idea that ' $\sim A \vee \sim B$ ' and ' $\sim(A \ \& \ B)$ ' are *tied* as being the best ways to represent this single way the world is. Note that there will inevitably be a large number of other sentences that will also tie with these two, e.g. ' $\sim A \rightarrow B$ '.

Why does the claim that there is a single way the underlying world is, such that ' $\sim A \vee \sim B$ ' and ' $\sim(A \ \& \ B)$ ' are *tied* as being the best ways to represent it, collapse into *Neither-fun*? Because if these two descriptions are tied as most joint-carving,

then the underlying way the world is can't resemble *either* the structure of ' $\sim A \vee \sim B$ ' or the structure of ' $\sim(A \& B)$ ' more. The logical forms of these sentences are radically different. Since they have two radically distinct sentential structures but are representing the exact same bit of worldly structure, it cannot be that *either* of them mirrors that structure in a joint-carving way. For if ' $\sim A \vee \sim B$ ' was perfectly joint-carving, it would follow that ' $\sim(A \& B)$ ' was not at all; but then the two sentences would not be tied as best representations of the world, as one would be joint-carving and the other would not. So it must be that the structure of *neither* of these two sentences resembles the underlying way the world is. And that is just *Neither-fun*.

I just gave an argument that appeals to "structure" in explanation, quantifies over it, counts it, etc. This may seem problematic, given that we are assuming nominalism about that structure. But it's a presupposition of the fundamentality-friendly nominalist that we are forced to think about things this way, while keeping in mind that we don't really mean to be reifying "structures". So, insofar as my reader dislikes the argumentative strategy, her distaste ought to lie with fundamentality-friendly nominalism generally, and not with my argument here.²⁴

To see this, notice that there are two opposing views we might have about *representation*, and in particular, what it is to be a perfectly joint-carving *representation* of the world. One view says that a perfectly joint-carving representation of reality is joint-carving at least partly in virtue of the world actually being a certain way (where we needn't be reifying "ways", and where we are being very loose about what 'in virtue of' means): having a certain nature or structure. A second view says that nothing worldly contributes to which description of the world is perfectly joint-carving. And we can't hold the second view if we are taking our words metaphysically seriously, regardless of whether we think they refer.

Consider what it would be to endorse the second view if we were property nominalists: 'green' and 'blue' more joint-carvingly describe the world than 'grue' and 'bleen' do, but this has nothing to do with how (language-independent) reality *is*. This directly conflicts with the initial motivation for distinguishing between *metaphysically* more and less fundamental language: we thought we were doing so to reflect something about *the world*, not to reflect purely linguistic facts.

Moreover, the second view makes the fundamentality-friendly nominalist's theses metaphysically uninteresting: she thought she was telling us something about metaphysics, about the world, but it turns out that she is just reporting facts about our language, that in no sense have their source in the language-independent world. So the only way to make the fundamentality-friendly nominalist's thesis metaphysically substantive, and in line with her motivation for adopting it, is to

²⁴ See, for earlier ancestors, e.g., Quine (1948), Armstrong (1978). More directly on this issue, see Lewis (1983), who suggests that naturalness can apply to the nominalist's predicates, Sider (2011), and Turner (2016, section 1.1.3) for an especially clear discussion of this issue.

adopt the first view: a perfectly joint-carving representation of reality is joint-carving *in virtue of* the world being a certain way.²⁵

Let me consider one possible way that the logical realist might resist what I've argued for so far. She might insist that there is space between *joint-carvingness*—which one might take to track some sort of objective similarity—and *fundamentality*. She might endorse, for example, Rosen's (2010) account of logical constants, on which disjunctive facts are always grounded in their disjuncts, and conjunctive facts are always grounded in their conjuncts. There is one sense, on such a view, in which there are no fundamental logically complex facts. Might one combine such a view with the view that the logical constants are perfectly joint-carving? Yes.

I think this sort of view is best made sense of by ideological realism; there is nothing logical in the fundamental furniture of the world, but we must be realists in some other sense about the logical constants; they mark objective similarity, perhaps. If this is so, the ideological realist is still faced with the question of *which* description (' $(\sim A \ \& \ \sim B)$ ' vs. ' $\sim(A \vee B)$ ') better captures this objective similarity, or if they do equally well, whether there is a better way to capture the similarity. I suspect (though I have not argued for this here, so more work would need to be done to show this) that it does not matter whether she thinks these sentences represent fundamental or non-fundamental facts. What matters is only that logical equivalence is not sufficient for *being equivalent with respect to joint-carvingness*. It seems hard to deny this and maintain logical realism, and the argument I have given here will go through unless one can do so. I want to grant, though, that there is possible room for resistance here, and also room for exploring other versions of realism that don't fall prey to the argument I've given here.²⁶

I've shown that the first way for the ideological logical realist to endorse *Both-fun* turns out to fall prey to *Weak Non-Redundancy*, and the second collapses into *Neither-fun*, and hence, given that we are assuming realism, into *Unfamiliar*.

²⁵ This style of argument is perhaps uncharitable to the neofregean, and in particular is incompatible with Rayo's (2013) views. However, I have the same questions about his view as I do here. Rayo claims that there are certain "just is" statements, such as *for there to be a table there just is for it to tableize*, that are true, and that this view is compatible with the kind of "metaphysicalist" (in his terms) view I am discussing here—one which says that different languages are more and less apt for carving nature at its joints. Elsewhere, I argue that these two claims are jointly incompatible, for very similar (but expanded) reasons to those I give here. In particular, there's an extended argument to be given for the claim that the metaphysicalist ought to buy into the claim that *being more* or *being less* joint-carving, while a feature of languages or linguistic or representational entities, must have its source, in some sense, in the language-independent world, and that this is makes the metaphysicalist's views incompatible with (at least some of) Rayo's views.

²⁶ Thanks to an anonymous referee for raising this worry.

4 Choosing between *Privileged* and *Unfamiliar*

I have now argued that both ontological and ideological logical realists must accept either *Privileged* or *Unfamiliar*. I will conclude by asking which view is better. I suggest that the balance of reasons lies in favor of *Unfamiliar*, but that *Unfamiliar* needs to be elaborated more before we could decide between the views.

First, I'll discuss the issue of whether we should move from the epistemic version of *non-arbitrariness* that I appealed to in sections one and two, and a metaphysical version of the principle, which would rule out views like *One-fun(ignorance)* and *Privileged* as well as views like *One-fun*.

I will start by examining a well-known appeal to non-arbitrariness: Parfit's (1984) discussion of fission. Parfit asks us to consider a case where my brain is divided between two bodies, 'Lefty' and 'Righty'. The question, in this case, is what happens to me. There are four options: I survive as both Lefty and Righty, I survive as Lefty, I survive as Righty, or I do not survive. We can set aside Parfit's conclusions here. His argument against the second or third option is that it can't be that I determinately survive as either Lefty or Righty (but not both). And this is not an epistemic claim: there is nothing *metaphysical* (nothing in the psychological and physical facts) that decides between me surviving as Lefty or surviving as Righty, and that's why it can't be that I determinately survive as either.

On one interpretation of Parfit, he is not appealing to some general metaphysical non-arbitrariness principle. Rather, he is basing the claim that it can't be that I determinately survive as Lefty on the claim that there is no feature of reality that could make it the case that I determinately survive as Lefty (and not Righty). There being a brute fact that I determinately survive as Lefty would count as being such a feature of reality. So it must be ruled out. In other words: the argument is not that *because* it would be arbitrary to believe that I determinately survive as Lefty (and not Righty), it must be the case that I do not determinately survive as Lefty (and not Righty). Rather, there are additional premises: first, it is not the case that it is a brute matter that I survive as Lefty, and second, that there are no other facts (e.g. psychological or physical facts) that could make it the case that I survive as Lefty.

If this is the right interpretation of Parfit, then he seems to be bridging the gap from an epistemic to a metaphysical non-arbitrariness claim: he is suggesting that if our reasons run out for deciding between Righty and Lefty, then the only remaining option is its being a brute fact that Righty or Lefty is me; but there are no such brute facts (and so, I take it, he is suggesting that if our reasons run out, metaphysics runs out too).²⁷

We might think something similar about relations if we were convinced that it simply couldn't be a brute fact that *is on top of* is fundamental and *is beneath* is derivative.²⁸ And similarly, we might think something similar about logical

²⁷ It's a messy issue whether this is the correct interpretation of Parfit. One might instead think: the analytic principle is restricted to non-brute, grounded facts. But if that's right, then Parfit's argument is going to fail if personal identity facts are the kinds of facts that could be brute.

²⁸ Fine (2000) also makes this sort of transition. Fine's argument runs through the claim that there is only one *state of affairs* that the sentences 'a is on top of b' and 'b is beneath a' picks out. If that's right, then

constants if we were convinced that it simply couldn't be a brute fact that (e.g.) '&' is fundamental and 'v' is derivative. This seems plausible, but not obviously true.

Regardless of whether there is a good general argument for doing so, philosophers often move quickly from epistemic to metaphysical non-arbitrariness principles in specific cases. Consider Benacerraf's argument against set-theoretic reductions of the natural numbers, which relies exactly on this sort of transition: there are multiple sets that are equally good candidates for being the number seven, there aren't any reasons to prefer one set-theoretic reduction over another, and so it follows that numbers are not sets. This argument seems very convincing to many of us, and that people have adopted it to similarly argue against set-theoretic reductions of e.g. pairs (Armstrong 1986, 1989), propositions (Jubien 2001), and counterparts (Merricks 2003), suggests that many of us at least implicitly endorse the move from the epistemic to the metaphysical claim. I don't mean to be appealing to authority here. Rather, I just want to point that lots of us already seem to accept arguments that move from epistemic to metaphysical non-arbitrariness, so we should either be swayed to do so here as well, or we should re-examine our acceptance of such arguments more generally.

The strongest reasons that we have to bridge the gap between epistemology and metaphysics in the case at hand are specific to logic, rather than general. The question of equivalence provides us with one such reason. *Privileged* has the consequence that the world discriminates between what is intuitively the same theory formulated first using '&', '~', and 'E', and second using 'v', '~', and 'A'. For *Privileged* entails that some small subcollection of our logical constants is fundamental. If '&' is in that subcollection and 'v' is not, then there is an important metaphysical distinction between these two theories—we just can't know which is better. Many of us will want to reject that there is a genuine difference between these two theories. But we might not want to reject logical realism altogether.

But *Unfamiliar* allows us to endorse logical realism without having to embrace the metaphysical distinctness of theories that seem to be mere notational variants. For it posits some unfamiliar structure that serves to anchor or ground our familiar logical constants. The proponent of *Unfamiliar* can hold that the '&' theory and the 'v' theory are indeed mere notational variants, since she can hold that they both make use of logical constants that are both derivative of a common, more fundamental, logical structure.

I've given both (weak) general and (stronger) specific reasons for embracing *Unfamiliar* and rejecting *Privileged*. Whatever our position is about gap-bridging,

Footnote 28 continued

there has to be a fact of the matter about what the relation is that is a constituent in that state of affairs. He hence rules out the redundancy option. But it's unclear exactly how he's ruling out the arbitrariness option—that it just happens that 'a is on top of b' is the joint-carving description of this state of affairs, because *is on top of* is, in fact, the relation that is a constituent in it. As far as I can tell, we don't get an explicit *argument* from the claim that we ought not believe that *is on top of* is the constituent relation or that *is beneath* is the constituent relation to the claim that neither actually is. But perhaps the idea is that it simply couldn't be a brute fact that *is on top of* was the constituent in this state of affairs, and that there is no possible *reason* that it rather than *is beneath* is.

any general principle that moves us from epistemic to metaphysical non-arbitrariness must be defeasible.

Consider the dispute between Newton and Leibniz about substantivalism and relationalism about space. There is reason to think that Newton was justified in maintaining substantivalism about space despite Leibniz's argument that substantivalism violates the principle of the identity of indiscernibles. We might characterize the issue as one about knowledge rather than the PII: one thing that Leibniz showed is that there's something fundamentally unknowable about the world if substantivalism is true. Newton was hence forced into the same sort of situation we are here—of choosing between a position very much like *Privileged* (substantivalism) and one very much like *Unfamiliar* (since he rejected relationalism as a candidate replacement for substantivalism).

But Newton's bucket argument seemed to show that the only available way of fleshing out the unfamiliar-like option—relationalism about space—failed to account for the empirical phenomena. And it is plausible that unknowability and a proliferation of possibilities is superior to not being able to account for the empirical phenomena. So perhaps, if there was no satisfactory relationalist theory, and there was no other alternative, then Newton's substantivalist position was a reasonable one.²⁹

Hence, Newton might have been justified in endorsing a position that looked like *Privileged*, because he believed that there simply was no acceptable candidate more fundamental theory one could move to.³⁰ (This is not definitive. We might instead claim that he was *not* so justified, because being in a *Privileged*-like situation at the least demands of us that we search for some more fundamental grounds for our theories.)

Of course, we don't have empirical evidence which contradicts *Unfamiliar*. Still, I think there is a general lesson here: avoiding unknowability claims (like *Privileged*) is, I think, a theoretical virtue to be traded off against others. It might only be worth moving to *Unfamiliar*-like options if we have some idea of how to flesh them out in a coherent fashion. So, return again to Benacerraf: the claim that numbers are not sets might be justified by the demonstration that it would be arbitrary for the number seven to be one candidate set vs. another; but only, perhaps, in conjunction with us being able to articulate *some* alternative view about what numbers are, if not sets. And if this is right, it's not so clear that we are justified in moving from *Privileged* to *Unfamiliar* without being able to say more about what worldly logical structure is like.

Put differently: if it turns out that all of the ways that we might flesh out *Unfamiliar* are theoretically inferior to *Privileged*, then we might want to reject

²⁹ This may not be an entirely accurate reconstruction of the debate. But that doesn't matter—if you object, just treat the case as presented as a toy one. The point stands. For discussion of related issues, see Dasgupta (2015), Maudlin (1993, 2012) and Sklar (1974).

³⁰ Whether he was right to think that there was no such theory is unclear, especially given advances in the philosophical literature on this dispute. Sklar (1974) points out that the relationalist *can* give an acceptable explanation of the bucket phenomenon. But that theory is both *indeterministic* and *non-local*. And at this point, we are simply trading theoretical costs against one another.

Unfamiliar! While *Unfamiliar* is in many ways more appealing than *Privileged*, more work must be done before we are justified in moving from the latter to the former.

In Sect. 1, I argued that we are forced to believe either that we can't know which of a pair of converse asymmetric relations is fundamental, or that neither is. What might replace a pair of relations as a single fundamental entity or bit of worldly structure that would eliminate this problem? We might look for single symmetrical relations that could somehow do the job (perhaps with some additional theoretical apparatus). This may be hard, but at least there are suggestions for how to do it.³¹ With logic, on the other hand, we are left at something of a loss when looking for a more fundamental theory to ground the logical constants

We might try to appeal to a primitive entailment relation as the more fundamental entity that the constants are derivative of. But there are problems with this strategy. First, it leaves us with a sort of mysticism both about what this entailment relation is, and about how it generates the logical constants that we know and love. Accounts of entailment either reference the constants qua truth-functions, or appeal to inference rules for the constants. Both take a 'constants-first' approach, and for a good reason: entailment is at least conceptually explicated in terms of other notions that seem to be *conceptually* prior. This does not, of course, mean that entailment is not *metaphysically* prior to these notions, but it does mean that more work would have to be done to demonstrate how entailment could be fundamental.

Second, entailment itself is an asymmetric relation of the sort that the first part of this paper raised problems for. The problem for logical constants and the problem for asymmetric relations are not just analogous—they are in this sense intimately entwined. If we are going to move to taking entailment as the primitive, fundamental entity or structure, much more work needs to be done to flesh out what that would look like, and how this issue would be avoided.

There are other strategies we might take, but I won't elaborate further on them here. I have argued that logical realism commits us to one of two surprising views: that the logical structure of the world is unknowable, or that it is deeply unfamiliar. These views have serious consequences for our philosophical and scientific theorizing, and we ought to consider them carefully.

Acknowledgements Special thanks to Shamik Dasgupta for many iterations of comments and discussion. Thanks also to Mark Balaguer, Ross Cameron, Rebecca Chan, Sam Cowling, Neil Dewar, Maegan Fairchild, Juliet Floyd, Hans Halvorson, Dan Korman, Kevin Lande, the late, wonderful, deeply missed Josh Parsons, Gideon Rosen, Jeffrey Sanford Russell, Jonathan Schaffer, Theodore Sider, Tuomas Tahko, Kelly Trogdon, Jason Turner, Jennifer Wang, Max Weiss, Nathan Wildman, an anonymous referee, and audiences at Oxford, Princeton and the Canadian Philosophical Association.

References

- Almog, J. (1989). Logic and the world. *Journal of Philosophical Logic*, 18(2), 197–220.
 Armstrong, D. (1978). *Universals and scientific realism, v. 1: Nominalism and realism*. Cambridge: Cambridge University Press.

³¹ E.g. in Dorr (2004), Fine (2000) and Williamson (1985) (perhaps).

- Armstrong, D. (1986). In defence of structural universals. *Australasian Journal of Philosophy*, 64, 85–88.
- Armstrong, D. (1989). *Universals: An opinionated introduction*. Boulder: Westview Press.
- Armstrong, D. (2004). *Truth and truthmakers*. Cambridge: Cambridge University Press.
- Barnes, E. (Forthcoming). Symmetric dependence. In R. Bliss, & G. Priest (Eds.), *Reality and its structure*. Oxford: Oxford University Press.
- Bennett, K. (2009). Composition, colocation, and metaontology. In D. J. Chalmers, D. Manley, & R. Wasserman (Eds.), *Metametaphysics*. Oxford: Oxford University Press.
- Correia, F. (2012). On the reduction of necessity to essence. *Philosophy and Phenomenological Research*, 84, 639–653.
- Cowling, S. (2013). Ideological parsimony. *Synthese*, 190(17), 3889–3908.
- Dasgupta, S. (2015). Substantivalism vs. relationalism about space in classical physics. *Philosophy Compass*, 10, 601–624.
- Dasgupta, S., & Turner, J. (2017). Postscript. In Barnes, E. (Ed.), *Current controversies in metaphysics*. London: Routledge.
- Donaldson, T. (2014). Reading the book of the world. *Philosophical Studies*, 172(4), 1051–1077.
- Dorr, C. (2004). Non-symmetric relations. *Oxford Studies in Metaphysics*, 1, 155–192.
- Dorr, C., & Rosen, G. (2002). Composition as a fiction. In R. Gale (Ed.), *The Blackwell companion to metaphysics* (pp. 151–174). Oxford: Blackwell.
- Dummett, M. (1991). *The logical basis of metaphysics*. Cambridge: Harvard University Press.
- Eddon, M. (2013). Fundamental properties of fundamental properties. *Oxford Studies in Metaphysics*, 8, 78–104.
- Fine, K. (1994a). Essence and modality. *Philosophical Perspectives*, 8, 1–16.
- Fine, K. (1994b). Senses of essence. In W. Sinnott-Armstrong, D. Raffman, & N. Asher (Eds.), *Modality, morality, and belief: essays in honor of Ruth Barcan Marcus*. Cambridge: Cambridge University Press.
- Fine, K. (2000). Neutral relations. *Philosophical Review*, 109(1), 1–33.
- Jenkins, C. (2010). What is ontological realism. *Philosophy Compass*, 5(10), 880–890.
- Jubien, M. (2001). Propositions and the objects of thought. *Philosophical Studies*, 103, 47–62.
- Korman, D. Z. (2010). Strange kinds, familiar kinds, and the charge of arbitrariness. *Oxford Studies in Metaphysics*, 5, 119–144.
- Koslow, A. (1992). *A structuralist theory of logic*. Cambridge: Cambridge University Press.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4), 343–377.
- Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.
- Maudlin, T. (1993). Buckets of water and waves of space: Why spacetime is probably a substance. *Philosophy of Science*, 60(2), 183–203.
- Maudlin, T. (2012). *Philosophy of physics: Space and time*. Princeton: Princeton University Press.
- Merricks, T. (2003). The end of counterpart theory. *Journal of Philosophy*, 100, 521–549.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Oxford University Press.
- Quine, W. V. O. (1948). On what there is. *Review of Metaphysics*, 2, 21–38.
- Ramsey, F. P. (1925). Universals. *Mind*, 34(136), 401–417.
- Rayo, A. (2013). *The construction of logical space*. Oxford: Oxford University Press.
- Resnik, M. D. (2000). Against logical realism. *History and Philosophy of Logic*, 20(3–4), 181–194.
- Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale & A. Hoffman (Eds.), *Modality: Metaphysics, logic, and epistemology*. Oxford: Oxford University Press.
- Schaffer, J. (2010). Monism: The priority of the whole. *Philosophical Review*, 119(1), 31–76.
- Sider, T. (1996). Intrinsic properties. *Philosophical Studies*, 83(1), 1–27.
- Sider, T. (2011). *Writing the book of the world*. Oxford: Oxford University Press.
- Sklar, L. (1974). *Space, time, and spacetime*. Berkeley: University of California Press.
- Turner, J. (2016). *The facts in logical space*. Oxford: Oxford University Press.
- Wang, J. (2016). Fundamentality and modal freedom. *Philosophical Perspectives*, 30(1), 397–418.
- Warren, J. (2016). Sider on the epistemology of structure. *Philosophical Studies*, 173(9), 2417–2435.
- Williamson, T. (1985). Converse relations. *The Philosophical Review*, 94, 249–262.
- Wilsch, T. (2015). The nomological account of ground. *Philosophical Studies*, 172(12), 3293–3312.
- Wilson, J. (2010). What is Hume's dictum, and why believe it? *Philosophy and Phenomenological Research*, 80(2), 595–637.