



Mind-Independence, Realism, and Reality

ABSTRACT: *Some metaphysically interesting notions fall outside of the realm of the mind-independent, and as such will erroneously be considered unworthy of our attention by any view that thinks only of realist metaphysics as substantive (Taylor, 2023). In this article I propose two ways of conceiving of substantive metaphysics that includes some mind-dependent phenomena. The first is to understand substantivity in terms of carving at the joints, but to take where the ‘joints’ are to depend in part on our own interests, preferences, and conceptual schemes (the joints are ours, not nature’s). The second is to think of ‘reality’ as a system of explanatory dependence that includes some suitably embedded mind-dependent phenomena, and thus to deny (if we continue to characterize realism in terms of mind-independence) that we should be realists about all of ‘reality’.*

1. Substantive Metaphysics and the RML Problem

Contemporary metaphysics has seen a surge of interest in *metaphysical structure*. The focus might be on fundamentality (as e.g. in Fine 2001; Schaffer 2010; Wilson in progress), building relations (as e.g. in Bennett 2017; Schaffer 2009) or eliteness (as e.g. in Lewis 1986; Sider 2011), but in each case the key idea is that reality has a structure that is independent of human thought and talk. Moreover, it has become common to think that the primary task for the metaphysician is investigating and attempting to discover this structure (see e.g. Schaffer 2009; Sider 2011). Because structure is construed realistically, it seems the limits of legitimate metaphysical enquiry have been set to exclude the mind-dependent (at least so long as realism is understood in terms of mind-independence).

Elanor Taylor (2023) calls a version of this problem the Realism-Mind-Independence-Legitimacy (RML) problem. Realist metaphysicians define substantivity, and hence genuine metaphysical inquiry, in terms of responsiveness to mind-independent metaphysical structure (Taylor 2023: 2). This allows Taylor to state the problem as follows (2023: 6):

Any view on which realism, defined in terms of mind-independence, marks the boundaries of legitimate inquiry, faces counterexamples from legitimate enquiry responsive to mind-dependent phenomena.

There are various plausible examples of investigations that seem both legitimate, and to concern phenomena that are not mind-independent. Social metaphysics (which is Taylor’s primary concern) seems to provide many such cases if we think, as is standard, that mind-dependence is a central aspect of social construction, and that social construction is vital to our understanding of the social world (see Taylor



2023: 7). But much of traditional metaphysics also concerns phenomena that might plausibly be considered mind-dependent and nevertheless deemed part of legitimate enquiry. In fact, *any* debate between realists and antirealists where mind-dependence is part of the antirealist's story about the relevant posit, can be considered an example of something that runs afoul of the RML problem, including debates about possible worlds, fictional characters, colors, moral properties, and mathematical objects (to name but a few). Taylor herself notes that there are legitimate explanations that don't give information about mind-independent determination relations and thus seem like instances of non-realist explanation (2023: 6). In fact, it seems enough that there are legitimate explanations that don't *only* give such information. Any conception of explanation according to which some mind-dependent factor is a necessary condition on something's counting as an explanation runs into the RML problem. Taylor notes further that there are instances of kinds central to scientific enquiry whose boundaries are not determined in a fully mind-independent way (2023: 6). If enquiry into such things is itself legitimate, we shouldn't be understanding legitimacy in terms of mind-independence.

Realist metaphysicians like Sider (2017) and Schaffer (2017) have responded to the objection (put to them by Barnes (2014) and Mikkola (2015)) that their approach to metaphysics threatens to exclude social metaphysics on the grounds that it deals with the non-fundamental, and have insisted that there is room in their approach for social metaphysics. Nevertheless, contemporary metaphysics has tended to conceive of 'reality' and thus of the subject matter of metaphysics as something discoverable, as being about what the world is like 'out there' beyond our ideas, and untainted by our ways of thinking and talking about it. This is so even when the relevant subject matter is plausibly counterfactually dependent on us, as in the case of social metaphysics (see e.g. Sider 2017: 2477).

Below, I suggest two different ways we might conceive of legitimate metaphysical inquiry, and thus respond to the RML problem. The first is to adopt a version of extant appeals to metaphysical structure, but to allow that structuring notions themselves might not be entirely objective and mind-independent, and hence that something's being fundamental or joint carving (for example) might not depend only on what the mind-independent world is like.¹ The second is directly to give a new account of 'reality' that includes the mind-dependent, based on the role things play in our best explanations.

2. Mind-dependent Metaphysical Structure

It is often claimed that the job of the metaphysician is to 'carve nature at the joints', where lurking behind the metaphor is a commitment to thinking that there is some objectively best way to divide up reality, and words and concepts that cleave most closely to these divisions are to be preferred. Joint-carving is a metaphor for the realist metaphysician: reality comes pre-divided, and the more effectively we respect

¹ For reasons of space I focus here on joint-carving in terms of Sider's notion of structure, but something similar could be achieved with a focus on a different notion of metaphysical structure such as grounding. I leave this for future work.

those divisions with the tools available to us (our words and concepts) the easier things will be.²

Once we have in play the idea that some things carve at the joints, our task is to discover which things those are. So, in a caricature of a philosophical debate between a mathematical Platonist and some kind of antirealist about mathematical entities, the Platonist claims whilst the antirealist denies that properties like *being prime* mark a genuine division in nature. We can think of the debate between the Platonist and her rival as substantive because there is something clearly at issue here; whether or not reality admits of primeness. Compare this with a different kind of debate, such as that between somebody who claims and somebody else who denies that Marmite is delicious. Presumably, deliciousness is not the kind of property that marks a genuine division in nature between the delicious and the non-delicious, and so insofar as we think of substantivity as connected to the question of what kind of joint-carving notions are in the vicinity, the debate about Marmite is nonsubstantive. Holding onto the idea that serious metaphysics deals only with the substantive, we get the result that the debate about primeness is a worthwhile pursuit (in Taylor's words, is 'legitimate enquiry') whilst the debate about the deliciousness of Marmite is not.

The traditional alternative to realism and its commitment to joint-carving is metaphysical deflationism, where deflationists (e.g. Hirsch 2002; Thomasson 2009; 2015) think that metaphysical disputes can be settled as a matter of convention, or depending on the choice we make of conceptual or linguistic framework. According to such views, whether or not we divide up the world into things that are prime and things that aren't depends on, e.g., which interpretation of the existential quantifier we subscribe to, or what kinds of platitudes ordinary speakers are willing to accept, and which kinds of inferences they are prepared to make on the basis of them. Crucially, apparently substantive debates in metaphysics come down to questions about how our concepts are applied, and are not (entirely) settled by what the mind-independent world is like. It follows that ontology is not the deep and substantive pursuit we might have thought it to be.

We might then consider the possibility of a middle way according to which we can make use of the metaphor of joint-carving and the idea that it conveys (that there are better and worse ways to carve things up) but deny that the relevant joints are entirely mind-independent. Dasgupta (2018: 280) makes the point that all parties agree that some kinds of theorizing are better (it is better, for example, for us to theorize in terms of green rather than grue).³ What is at issue is whether or not it is *objectively* better. On a view where the joints are not entirely mind-independent, some ways of carving things up are genuinely better than others. Part of what makes them better, though, is that they are better *for us*, given some or all of factors including our particular

² We might wonder in what sense things will be easier. Perhaps if we cleave closely to reality's joints we can expect that our explanations will be successful, we will be able to formulate laws in simple terms, be better able to formulate testable hypotheses, our words and concepts will be projectable, and so on. Sider (2011) defends such ideas in chapter three: 'Connections'.

³ Things are grue when they are observed to be green before some future time *t*, but blue when observed at or after *t*. The predicate is originally introduced by Goodman (see e.g. his 1983: 74).

interests, conceptual schemes, patterns of concern, and explanatory aims. Such a view might make the epistemology of joint carving less problematic; and provide us with a way to recognize joints beyond the natural, including in the special sciences and in the social realm. While Sider (2017: 2469) thinks his broad framework is also consistent with recognizing higher-level joints, his preferred way of filling out the details is one where any such joints are reducible to joints at lower levels. The appeal to less-than-fully-realist joint carving also has advantages over traditional deflationism in that it explains the deeply held intuition that some carvings are genuinely better than others, and guarantees the substantivity of debates in multiple different areas of metaphysics.

2.1. Sider's Structure

Sider's *Writing the Book of the World* (2011) articulates and defends a kind of 'hardcore' metaphysical realism; a view according to which (2011: 18):

the point of human enquiry...is to *conform* itself to the world, rather than to *make* the world. The world is "out there", and our job is to wrap our minds around it.

Sider's is perhaps the most detailed recent defence of metaphysical realism as a worldview, and the notion of being 'perfectly joint-carving' plays a central role. I'll focus on Sider's conception of joint-carving in what follows.

For Sider, being perfectly joint-carving is synonymous with being perfectly *structural*, where structure for Sider is an operator that attaches to any portion of a language, such that predicates, connectives, quantifiers and so on might all be considered candidates for carving nature perfectly at the joints. A language which uses only perfectly joint-carving terms will be a fundamental language, and this (for Sider) would be the best language for doing metaphysics.

Sider thinks of substantivity in terms of his notion of structure: metaphysical debates are substantive (roughly) when answers to the relevant questions differ with respect to how structural they are—when one candidate answer carves at the joints more closely than the alternatives. A sentence *S* is nonsubstantive for Sider if and only if *S* would have had different truth values under different candidate meanings, where those candidate meanings are equally joint-carving and equally faithful to *S*'s conceptual role (see Sider 2011, chapter 4). Sider thinks that joint-carvingness is a feature of the world "out there" and is independent of human thought and talk.

The way in which Sider takes structure to be objective and mind-independent leads to concerns about the epistemology of structure, which have been explored in the recent literature. Warren (2016) argues that Sider's realism about structure faces the Benacerraf challenge to realism about the causally inaccessible 'in its starkest form' (2016: 2427). The problem is that our attitude-forming mechanisms are causal mechanisms, but facts about structure are (according to Sider) causally inaccessible, mind-independent, and objective. If we assume that our beliefs about structure are reliable, this must just be by some kind of remarkable happy accident. If the structure-facts were different, there is no reason to think that our beliefs about

structure would differ accordingly (because the structure-facts are independent of our belief-forming mechanisms). And if our beliefs about structure were different than they in fact are (e.g. because the history of science played out slightly differently), Sider must maintain that the structure-facts remain the same (Warren 2016: 2428). Warren (2016: 2421–2) points out that Heisenberg’s matrix mechanics and Schrödinger’s wave mechanics are generally accepted to be mathematically (and thus empirically) equivalent, but have markedly different ideology. Different groups of scientists prefer different theories based on their different preferences and mathematical competencies, but Sider could not accept that this would make a difference to the structure facts.

It’s worth noting here that if this argument works, it seems to threaten *any* causally inaccessible, mind-independent metaphysical posit. But any worries about over-generalization should be tempered by the reminder that my broader proposal here is that we think of metaphysical structure in the wide sense (and not just of Siderian structure) as not entirely objective and mind-independent. Metaphysics is about metaphysical structure and fundamentality, but I will suggest that we might think of *those very notions* as dependent on us in some way, rather than as being features of the mind-independent world.

Sider himself does offer an epistemology for structure. He (2011: 12) subscribes to the now fairly dominant (see e.g. Paul 2012; Morganti and Tahko 2017) and vaguely Quinean conception of metaphysics as continuous with science, further claiming that we should believe not just the ontology, but also the ideology of our best theories. A good theory is more likely to have joint-carving ideology; part of the explanation of the success of our best theories is that they are expressed in joint-carving terms. First order theories of the world might employ joint-carving ideology, but they don’t include the ideology of joint-carving. Joint-carving is a second-order ideological commitment not present in these first order theories, as Sider himself acknowledges (2011: 138). Since Sider’s is a project in metametaphysics, perhaps the appropriate theories to consider are metametaphysical theories. Sider’s theory includes platitudes like ‘reality is structured’, and so carries an ideological commitment to structure.

It is beyond the scope of this article to adequately assess the Quinean methodology as it applies in metametaphysics, but it is worth noting the following. Let’s grant that Sider’s theory enjoys a high degree of success.⁴ There are two ways we might explain this success. The first is to hold with Sider that the theory is successful because realism about structure is (close to being) true. But we might also explain the success of our theory in a more antirealist-friendly way: given the nature of our conceptual schemes (for example), it is simplifying, unifying, and explanatorily powerful for us to embrace the ideology of structure, and this could be so whether or not realism about structure is true. The availability of an alternative explanation for the success of theories such as Sider’s (which seem to run afoul of safety and sensitivity conditions for knowledge) threatens to undermine the realist’s epistemology.

⁴ Success here cannot be empirical confirmation or predictive success (since metametaphysical theories are empirically equivalent), but could be understood in terms of something like explanatory power (where Sider (2011) claims a great deal of explanatory power for his notion of structure).

To be clear: the problem is not that the realist might be wrong (the possibility for error is one of the hallmarks of realism, and so it would be inappropriate to use this possibility against the realist). The problem is that it doesn't seem that we are in possession, even in principle, of a reliable method for detecting errors in our theorizing about structure on Sider's view. We might think that we can be reasonably confident of at least most of our judgments about structure because Sider's metametaphysics seems to be a good theory, but it's not really clear what makes for a good metametaphysical theory beyond being consistent with first order theories in science and metaphysics, and exhibiting relevant theoretical virtues like simplicity, explanatory power, and unification.

Sider takes his theory to score well on all of these counts. But *even if* a metametaphysical theory employing realist structure does better than its rivals, it's not clear that the best explanation for this is that there really is realist metaphysical structure. Virtues such as simplicity, explanatory power, and unification can be viewed not as guides to truth, but instead as pragmatic constraints that arise as a result of our particular cognitive makeup.

In the case of simplicity, for example, we can note that simple theories are easier to understand, and so simplicity is a theoretical virtue. But we have no reason to expect the mind-independent world to conform to such preferences. In a similar vein, we can note that unification (very roughly) involves explaining a large number of facts with a small number of argument patterns, and that this exploits a close connection between explanation and understanding (see, e.g., Kovacs 2020), where understanding may be taken to be a subjective notion. Explanatory power is a matter of explaining well, where the subjectivist highlights the importance of a good explanation meeting subjective constraints. These might include pragmatic constraints (the complexity of the explanation offered, how well it serves the interests of parties to the explanation, and how the explanation is formulated); psychological constraints (e.g., the relevant concepts possessed by agents who are party to the explanation; their psychological state) epistemic constraints (taking into account the background knowledge of any agents involved, the cognitive sophistication of those agents, and whether the explanation is appropriately connected to understanding). The subjectivist can insist that these constraints hold even for *metaphysical* explanations (see, e.g., Thompson 2023).

Given that such antirealist accounts of the relevant virtues are genuine rivals to the realist way of thinking here, it doesn't seem the realist can rely on the theoretical success of metametaphysical theories employing structure to secure an epistemology for structure. Sider could be completely wrong about structure, and yet his theory could still be an attractive metametaphysical theory.

This doesn't amount to an argument for the antirealist's way of explaining things, but it introduces a kind of dialectical stalemate. It's just not clear what either side can say to convince the other of their explanation of the success of a successful theory. Accordingly, it doesn't seem we have a reliable epistemology for realist structure. In the light of that it seems reasonable to consider alternatives. To refuse this on the grounds that it would render much of metaphysics 'illegitimate' would be to beg the question in the context of looking for a way to establish the boundaries of legitimate metaphysical inquiry.

2.2. Against Subjectively-infused Structure

Sider (2011: 18) says that the view that structure is tied to us in some way ‘is just incredible’. There’s little by way of argument here (as Sider himself acknowledges), but his main motivation seems to be that if we adopt what he calls ‘subjectivism’ about structure, we must accept that facts about similarity, intrinsicality, laws of nature and so on—those facts that are themselves (Sider thinks) to be understood in terms of structure—are also just ‘a reflection of our language (or whatever)’. Sider later claims that the main problem is that joint-carving languages and beliefs are better than their non-joint carving alternatives, and that if structure is subjective so is this betterness, and that this would be ‘a disaster’ (2011: 65).

I think Sider is being insufficiently careful here. One can think that structure is in part determined by something mind-dependent without succumbing to radical subjectivism. For example, in Lewis’ work on reference magnetism (see, e.g., 1984; 1986) by which Sider is heavily influenced, Lewis takes reference to be a matter of balancing charity to use against eligibility. Reference is constrained by which of our terms carve nature at the joints (whatever carves at the joints is more eligible to serve as a referent) but sometimes there are good reasons to ascribe less eligible referents to our expressions (see e.g. 1984: 227). That is, it is sometimes justified to take people to be referring to something relatively ineligible (for example, philosophers sometimes use the expression ‘grue’ to refer to grue, even though ‘green’ would be the more eligible referent in the vicinity). We can imagine a variant on Lewis’ position that builds charity to use *into* the eligibility of the referent rather than taking the two constraints to be independent, so that the eligibility of referents is determined both by objective facts about how natural they are, and by facts about how useful they are in a given context. Eligibility would thus not be fixed as it is in Lewis’ picture, but would vary by context. (I don’t claim such a view is especially attractive, but neither does it seem entirely implausible.) This would be a kind of subjectivism about eligibility, but not a radical subjectivism, and it wouldn’t make facts that draw on those referents (e.g. about laws of nature, similarity, etc.) radically subjective either. In fact, Sider himself claims that one of the ways in which some notions might be more structural than others is that they have a certain value to us (2011: 132). Absolute structure must, he says, be completely objective, but partial subjectivism about comparative structure is acceptable.

I’m not sure it can be reasonably maintained that which sorts of things matter to us are irrelevant in terms of structure at the fundamental level, but play a role in comparative assessments of structure. It would be like insisting that whether some food is healthy or not is an objective matter, but that when comparing different foods, how tasty something is might help us to characterize one dish as healthier than another. This strikes me as a mistake: taste is either a component of healthiness (including both absolute healthiness and comparative assessments of healthiness) or it isn’t. Similarly for fundamentality: if perfect joint-carvingness is entirely unaffected by what we care about, what we care about should have no influence on which of two imperfectly joint carving notions is to be considered the more joint carving.

Setting this aside, Sider’s reason for finding any degree of subjectivity amongst the absolutely fundamental objectionable is that substantivity, objectivity, explanatoriness

and so on depend on the objectivity of structure. Sider thinks that debates are substantive when they concern structural notions (2011, §4.2), that objectivity is a matter of carving more closely to the joints (2011, §4.6), and that true theories are more explanatory when they are stated in joint-carving terms (e.g. 2011: 23). If structure isn't entirely "out there" in the world, then neither are these other notions. 'No discourses would be objectively objective' (2011: 65)—disaster! One response is to deny that these notions should be understood in terms of structure in the first place, but it seems to me that we can make perfect sense of notions of objectivity, substantivity, and so on defined in terms of structure that is *not* entirely mind-independent.

It's important to be clear on what is at stake here. In his (2017) clarification of his position, Sider is at pains to point out that his account of structure does not preclude facts about the dependence of phenomena on human activity from featuring in the *content* of an objective description of reality (2017: 2477). That is to say: facts about things like social kinds might both be mind-dependent in the counterfactual sense (they wouldn't exist or would be different if humans didn't exist or were different), and nevertheless might be objective, feature in explanations, and be the subject of substantive debate.⁵ This brings out a subtle and often overlooked point about the relationship between objectivity and mind-independence: these two things do not always go together. Sider thinks that we can play a role in making the world, and that this is what happens in the case of social metaphysics. But once the social metaphysical facts are settled they are part of the world, and so there is an objective fact of the matter about which of a set of candidate meanings for relevant expressions carves most closely to the joints. So long as this is the case, Sider thinks, social metaphysics is substantive.

Sider's proposal allows that we (via our social practices, for example) might play a role in settling what we might think of as the first-order metaphysical facts, which might include facts like 'gender is distinct from sex' (Sider 2017: 2470). What Sider finds objectionable is that we might play a role in settling what we might think of as the second-order metaphysical facts, which might include facts about which of a set of candidate meanings carves more closely to the joints. Of course, there is a sense in which if our social practices play a role in settling objective first-order facts about the world, we *do* settle which of a set of candidate meanings carves more closely to the joints (by settling where the joints are). But this settling is done, as it were, in advance. *At the point of selection* (from candidate meanings), the theorist plays no role. By contrast, in the view I am developing facts about a community of theorizers may help to settle which of a set of candidate meanings carves most closely to the joints, because joint-carving is not merely a matter of latching on to objective worldly facts.

We can define notions of substantivity, objectivity and so on in terms of structure just as Sider suggests. It's only when we ask a further question: 'what makes it the case that some *x* is structural?' that we get an answer revealing that structure itself is not entirely mind-independent. What makes it the case that certain things are structural has partly to do with what is of central importance to us, because this is how structure is defined. On this kind of a view, whether some question is objective, substantive,

⁵ This article is a response to Barnes (2014) and Mikkola (2015) who argue, to the contrary, that Sider's focus on the fundamental threatens to exclude social metaphysics.

and so on depends on whether it is given in terms that carve at the joints (as Sider states). It's just that where the joints are is not just about what the world is like (whether or not we had any role in its coming to be that way). According to this position, some discourses are indeed objective (because sentences characteristic of the discourse have some privileged, joint-carving interpretation).

To claim that something about our judgments, preferences, conceptual schemes or explanatory projects constrains or partly constitutes notions like joint-carving, fundamentality, or structure is clearly not consistent with the kind of strong metaphysical realism preferred by people like Sider. However, if notions like structure are infected with some degree of subjectivity (as Sider might put it) but still play the roles that such realists take them to play in metaphysics, then this is the limit of realist inquiry; there is no further sense in which we might cleave more closely to nature's joints.⁶ Subjectivity-infused joints are the only joints there are, and we can carve more or less closely to them. The friend of subjectivity-infused structure can argue that structure without subjectivity is not something to aim for—that's just not what structure is like. This kind of thinking will presumably be intolerable to those who, like Sider, think it is built into the very idea of metaphysical structure that it be mind-independent, but this is a philosophical conviction that Sider himself admits he has no argument for (2011: 12), and so it is reasonable to consider alternatives.

Recently, Dasgupta (2018) has made a case that even if we grant metaphysical realists the absolute mind-independence of naturalness and/or structure (and cognate notions), this is not enough to motivate the kind of position that Sider defends. Sider (and less explicitly, other metaphysical realists) make not only a metaphysical claim to the effect that some properties, entities, quantifiers and so on are somehow elite or metaphysically privileged, but also a value-theoretic claim that theorizing in terms of those elite notions is objectively better than any alternatives. Dasgupta argues that this value-theoretic claim is seemingly impossible to defend, because there is no legitimate (mind-independent) explanation for why structural beliefs are objectively better than other beliefs that are not elite, but might better serve the interests of some communities.

The realist claims that it is *constitutive* of structure that (e.g.,) structural beliefs are better than others. Dasgupta gives two reasons for thinking that we shouldn't accept this. First, it is too tight an explanatory circle to say that we can explain why theorizing in terms of green is better than theorizing in terms of grue on the grounds that green is more structural than grue, *and* to hold that to be structural *just is* to be better for our theorizing. If that theorizing in terms of green is better is *constitutive* of green being more structural, we can't then appeal to the fact that being green is more structural than being grue in order to *explain why* theorizing in terms of green is better. The second reason not simply to assert that structure is theory guiding by nature 'labels the problem without solving it' (Dasgupta 2018: 300). By contrast, the antirealist goes through all the hard work of explaining why, for creatures like us, theorizing in structural terms is (subjectively) better than the alternative.

If Dasgupta is right, then pure realism about structure and its cognates (i.e., the claim that there is something objectively metaphysically privileged about some

⁶ We might think of this as a kind of 'internal realism' such as that defended by Putnam e.g. in his (1987).

properties, quantifiers, etc.) is toothless. Things only get interesting if we further accept that theorizing in terms of structure is *better*, but the only sense of betterness that we should accept is that offered by the antirealist, who thinks that betterness depends on our interests (broadly conceived). Thinking about structure not as ‘objectively objective’, but rather as picking out the carvings that are best *for us* allows the subjectivist to give a full account of what makes reasoning in terms of structure (subjectively) better, and avoids the epistemic concerns facing the realist about structure.

2.3. Subjectively-infused Structure

We can think of subjectively-infused structure as a kind of ‘vegetarian’ alternative to joint-carving (see Taylor (1993: 88), which comes with the promise that we can retain the central idea that some notions are privileged, but reject the claim that this privilege is a primitive feature of reality. Taylor (1993) develops an account of natural properties as those expressed by predicates that play ‘the more central and classificatory roles within a theory’ (1993: 88). Taylor’s idea is that a primitive predicate *F* is to be considered more central to a theory (and thus more elite) than another primitive predicate *G* when, by the theory’s lights, *G* can only be instantiated when *F* is. Non-primitive predicates derive their eliteness from that of the primitive predicates they can be expressed in terms of. Since it is common sense that ‘horse’ implies ‘animal’, for example, ‘animal’ is more central and thus more elite than ‘horse’ relative to ‘common sense’ (Taylor 1993: 90).

Like Taylor, the friend of subjectively-infused structure holds that the eliteness of some bearer of structure or naturalness is not a primitive fact about that entity or notion. It is instead determined by the role that entity or notion plays in our best theory, where what makes it the case that some theory is better than another is determined both by the degree to which it is empirically confirmed, and the extent to which it satisfies our demands for simplicity, explanatory power, unification and so on. Whilst the former offers a worldly constraint on eliteness, the latter offers a subjective constraint (when understood along the lines I suggested in §2.1 above).

Like Taylor, I think that in order to see which notions are elite relative to our best theory, we need only see which play the most central roles. We can think of these notions as the ‘joints’ in our theories. But instead of defining centrality in terms of implication, I think we should define it in terms of explanatory power, where notions that have a lot of explanatory power are those that balance simplicity with explaining a large number of things well (where ‘well’ is closely connected to understanding). This is very much the picture involved in unificationist accounts of (metaphysical) explanation (see Baron and Norton 2021; Kovacs 2020), but unlike many unificationists I think of explanation as subjectively constrained, as described earlier. Accordingly, we end up with a thoroughgoing subjectivism about structure, but it is not a radical subjectivism. Much of the data that feeds into the development of our best theory comes from the mind-independent world. Proponents of different theories will have different views on what are the proper bearers of eliteness, but this should be no surprise (and is not unique to thinking of structure as subjectively infused).

When we think of structure, or more generally of eliteness as subjectively infused, we will find that the kinds of things that count as elite won't only include those things metaphysicians have tended to want to be realist about, or only whatever is mind-independent. For example, if the notion of gender has a lot of explanatory power (it can provide a good explanation of many of our observations in a way we find appealing and easy to understand, for example) and it therefore plays a relatively central role in our best theory of the world (which itself is more theoretically virtuous than rival theories) then we should take the notion of gender to be fairly elite. Whether we think of gender as mind-independent or not isn't settled by what we've said so far (it might be that our best theory makes gender a social construct, or alternatively that it makes it a biological fact, or something else), but what certainly *isn't* mind-independent is whether gender counts as elite, because this comes down to questions about the role something plays in the theory that *we* (with our particular preferences, interests, and cognitive schemes) take to be the best one.

It is worth mentioning one further contrast between this view and extant positions in the literature. Elanor Taylor (2016; 2020) develops a notion of 'naturalness in context' ('c-naturalness'), which is a highly context-sensitive notion of eliteness relative to an activity. This account maximizes the balance of salience to the activity (which might be playing chess, or working as a nurse) with versatility (how widely applicable the property is) with respect to the activity. Since salience can be understood as something akin to 'importance', and versatility as connected to explanatory power, this account of naturalness in context is in some ways quite similar to the account I have outlined above.

Taylor takes c-naturalness to be grounded in (realist) metaphysical naturalness, where we can think of metaphysical naturalness as maximally general. The view I have developed here says that even at maximum generality, joint-carving or eliteness is relative to our best theory. The friend of subjectively-infused structure might well make use of c-naturalness for more highly specified contexts, and might understand naturalness or structure simpliciter to be a kind of limiting case of c-naturalness holding with maximal generality. What they deny is that there are any joints in reality extending beyond our best theory, or that there is any reason to think we could have any epistemic access to such joints if they did exist. So, while the two accounts have several things in common, they differ in that where Taylor's account recognizes two fundamentally different kinds of eliteness—one that is fully realist and another that is not—ours recognizes only one notion that is never fully realist.

2.4. The RML Problem Again

We introduced subjectively-infused structure in part as a way to answer the RML problem. We can now develop that answer in two ways. Both begin by denying that it is realism (understood in terms of mind-dependence) that marks the bounds of legitimate enquiry. The first says with Sider that substantivity (which we can take to be the hallmark of legitimacy) is understood in terms of structure. Substantive questions are those for which one candidate meaning carves at the joints better than alternative meanings (see Sider 2011 chapter 4, especially §4.2. Take a question like

‘are trans women women?’, and suppose for the sake of argument that there are two candidate meanings for the expression ‘women’, one according to which ‘women’ picks out a biological category such as having XX chromosomes, and one according to which ‘women’ picks out a complex social category determined by systematic patterns of oppression. (These are vastly simplified for brevity.) The second candidate meaning seems to carve more closely to the joints (in the sense of the subjectively-infused structure under discussion) than does the first. Gender on this way of understanding it plays an important role in our theorizing about the social world, which is an important part of the world. It has far more explanatory power than does the first candidate meaning, and interpreting questions about gender as questions about this complex social category fits with the observation that gender matters to us. Because one of these candidate meaning carves more closely to the joints than the alternative(s), questions about gender are substantive and are part of giving a serious account of the world; they are part of legitimate metaphysical enquiry.

One might respond that we have oversimplified matters, and that there are in fact multiple equally joint carving interpretations of the expression based on slightly different social constructivist accounts of gender. So long as all of these generate the same answer to our question as each other, and all carve more closely at the joints than do rival interpretations that would generate different answers, the original question still counts as substantive (see Sider 2011: 47). Settling which of a number of very similar theories is in fact the best one is tricky, but so long as gender so interpreted plays a central role in each of them, we are justified in taking that notion to be fairly structural. Richardson (2023) allows that various candidate meanings might be equally joint carving, and thus that we should take terms like ‘woman’ to be referentially indeterminate, but argues that this indeterminacy is compatible with relative (but not perfect) substantivity. Richardson’s proposal generates similar results to my own in this case, but he denies that the usefulness of a kind in a context is any indication of metaphysical priority (2023: 198).

The second way to answer the RML problem doesn’t subscribe to Sider’s account of substantivity, instead drawing a direct link between substantivity or giving a serious account of the world, and playing a central role in our best theory. The idea is that things that play central roles in our best theory are important (to us), explanatorily powerful, and consistent with empirical data. As such, they are the sorts of things that we need to consider when we give a serious account of what the world is like. Because it is *our* best theory of what the world is like (and also because we are part of the world) there is no guarantee that everything that we need to consider when giving a serious account of the world is mind-independent. In fact, it is very likely that at least some of the relevant things we must consider will be mind-dependent.

A possible advantage of the first kind of answer to the RML problem over the second is that it suggests a somewhat more precise characterization of significance in terms of Sider’s semi-technical notion of substantivity. On the other hand, one might not wish to subscribe to Sider’s account, or might argue that the vagueness in the second answer is not unwelcome. We shouldn’t expect sharp boundaries between what we consider a serious account of what things are like and what we consider not

to be, because notions like significance come in degrees. In fact, one might claim that the seemingly sharp boundary between the substantive and the non-substantive in the work of people like Sider is artificial, given that the notion of structure with which substantivity is to be characterized itself comes in degrees. In the subjectively-infused account of structure we lack the kind of significance placed by Lewis and Sider on the *perfectly* natural and the *perfectly* structural, but this can be seen as a merit of the view. As Sider himself notes, though he conceives of structure as absolute, many of the uses he puts structure to in fact concern the comparative notion of structure (2011: 129). The friend of subjectively infused structure lacks Sider's reason for taking structure to be absolute (which is that facts about structure are fundamental, and facts about comparative structure are nonfundamental) because facts about subjectively infused structure can be further explained in terms of facts about theoretical virtues, interests, and so on (and so are nonfundamental in this sense).

Taking structure to be subjectively-infused leaves open whether or not realism is to be understood in terms of mind-independence. Supposing that we continue to understand realism in terms of mind-independence, subjectively-infused structure is an antirealist view of metaphysical structure. The RML problem assumes at its outset that realism marks the bounds of legitimate enquiry, and so if we ought to be antirealists about metaphysical structure, enquiry into metaphysical structure is not legitimate. Something like this sort of worry is what drives Sider so strongly against subjectively-infused structure, but if, like Taylor (2023) we opt to deny the realism-legitimacy link, we can recoup the advantages of doing metaphysics in terms of structure whilst also having an answer to the RML problem.

Subjectively-infused structure bears some similarities to Barnes' (2017) interpretation of Haslanger's (e.g. 2012) view about social categories like race and gender, but there are important differences. Barnes takes Haslanger to be a realist about such categories, but also to hold that these categories arise out of the activities of human beings. Barnes argues that gender and race on Haslanger's view must count as non-substantive given Sider's understanding of substantivity, but that this is a mistake because (according to Haslanger) social categories are real, objective, casually efficacious, and so on. Barnes effectively argues that Sider's conception of substantivity is too austere because it counts as nonsubstantive some debates about things we ought to be realists about (contra Sider's 2017 account of his earlier work). In other words, Barnes can be seen as accepting that realism marks the bounds of legitimate enquiry, but denying that realism ought to be understood in terms of mind-independence. (Haslanger counts as a realist about social categories who nevertheless thinks that those categories are in some sense mind-dependent, and that debates about race and gender are substantive).

My view is more permissive than Barnes' interpretation of Haslanger about the kinds of things that can count as substantive, because the connection between realism and substantivity is severed (although not too permissive, because it is constrained by the empirical facts). It allows us to hold on to Sider's view of substantivity by modifying what it takes for something to carve at the joints. In Barnes' interpretation of Haslanger, the picture we end up with is one where 'the joints *in reality* might go beyond the joints *in nature*' (Barnes 2017: 2424). Barnes is thinking of reality here as whatever we are realist about. My view differs in that for me, *in*

reality characterizes the domain of serious metaphysics or legitimate enquiry. It is not required that we think of all of ‘reality’ as mind-independent. If that means that we should be antirealists about some parts of reality, so be it, but the linguistic difficulties involved might ultimately be reason to move away from a characterization of realism in terms of mind-independence.

3. Reality

An alternative approach to legitimate, mind-dependent metaphysics directly tackles the claim that ‘reality’ includes more than the things we might want to be realists about. In this way we can continue, should we want to, to define realism in terms of mind-independence. We can also think of serious metaphysics as concerned with all of what (for want of a better term) I’ll call ‘reality’, where reality includes e.g. aspects of the social world which we take to be socially constructed. Roughly, we can think of ‘reality’ as a system of explanatory dependence, where everything that either makes a sufficient explanatory difference or is itself sufficiently explained by other parts of the system is part of reality, and hence part of the domain of serious metaphysics. Below I begin to sketch what such a view might look like.

The key idea is that whatever is sufficiently embedded in our explanatory system—whatever forms part of our best theory of the world or is itself explained by that theory—would count as part of reality. An ‘explanatory system’ thus in some ways resembles a Kuhnian paradigm; it is characterized by what is accepted by a community of thinkers working with common views, concepts, and methods. A good question to ask then, especially in the context of thinking about social metaphysics where one might have (at least in part) an emancipatory aim, is whether and how we can make changes to an explanatory system by going against the accepted orthodoxy.⁷

A key thing to highlight is that something (like a social kind, for example) might be part of our explanatory system even if it is involved in explanations that are in fact false, if no better explanation is available. However, as the views of a community change, the kinds of explanations that count as good explanations will also change. So, for example, while in the past a White American might have offered a (racist) explanation for the disproportionately high number of African-Americans in US jails in terms of a propensity towards criminal behaviors, she might now take the same fact to be explained in terms of structural racism. Certain members of the community will have fought to make the second explanation heard, and will have done so by connecting it to evidence, which will itself come to seem more plausible as other explanations appealing to structural racism become better ingrained. The racist explanation was always false, but because the concept of structural racism was previously not salient among members of the explanatory community, it was not easy to dismiss that explanation in favor of a better alternative.

Sometimes though, a progressive explanation is offered that comes too soon to be accepted by the community, or is met with a lot of resistance from people who have

⁷ I’d like to thank an anonymous referee for drawing my attention to this question.

not accepted the common views or concepts of the more progressive group. In the former case, proponents of the progressive view could continue working to demonstrate the success of their explanation and its coherence with other beliefs in the community, and in time the explanation could become entrenched, rendering it part of reality. In the latter, one might reasonably suppose that we could end up with two paradigms; two explanatory systems that are incommensurable with one another.

This is a permissive conception both of reality and of the domain of metaphysics, but it is not overly permissive. Our account of reality will include some mind-dependent things like social kinds and categories on the basis that they form part of our best theory of the world; we might think of them as socially fundamental, and of the socially fundamental as part of our best theory. But we needn't go so far. Even thinking of social kinds and categories as determined or explained by further factors (perhaps a combination of biological and psychological factors, for example) will be sufficient for including them in the domain of serious metaphysics, or for thinking of them as part of legitimate enquiry.

Contrast this with, for example, astrological kinds such as star signs. A person's star sign is partly explained by their date of birth, and some astrological conjectures are partially explained by the relative positions of certain stars. However, the majority of purported astrological facts are insufficiently supported by aspects of our best theory of the world, and do not themselves play an important explanatory role. (This is because facts that have purported astrological explanations can be better explained in other terms, i.e. they can be given a complete explanation in terms of more central parts of our best theory, such as psychological, social, and physical facts.) As such, astrological notions are insufficiently embedded in our best explanatory system, and should be rejected as not part of reality. Astrology is not part of serious metaphysics, but we should maintain that it is legitimate to enquire into the distinction between science and pseudoscience, and so we might (in the way I describe at the beginning of the article) want to hold off from claiming that serious metaphysics marks the boundary of legitimate enquiry.

Phenomena count as part of reality insofar as they are assimilated in our best system of explanation. Phenomena featured in candidate explanations that are not so assimilated (e.g. phlogiston; star signs; witches) are not part of reality. The account is permissive in that many kinds of social, economic, and psychological phenomena will count as part of reality, but it is not unduly permissive. There will be plenty of scope for debates about whether things like fictional characters will count as part of reality, and the debate will concern how well they are assimilated into our best explanations. I expect that moral facts and properties, for example, will count as part of reality because of the central role they play in explaining features of our behavior. This is not to say that we must be moral realists; remember that we can think of the question of realism as expressed in terms of mind-independence, and many things that are part of reality in our inflated sense will not be mind-independent.

Here is one way we might make this idea more precise. In his seminal paper in which he introduces the notion of grounding into the contemporary debate, Fine (2001) addresses what he takes to be a serious problem for realist and antirealist alike: in order to respect the way in which ordinary speakers think and talk we must

distinguish between things that are *really* the case and things that are merely apparently the case. Only the former things are part of what he calls a *metaphysical* conception of reality (Fine 2001: 3). Let's call this M-reality. Fine thinks that grounding talk exchanges problematic questions about what exists in M-reality for more tractable questions about what grounds what (where for Fine, *ground* is very clearly an explanatory notion). The idea is that we can answer the question of what grounds a given proposition without comment on the metaphysical status of the components of that proposition (see Fine 2001: 11), and then that we can work out once we have the chain of grounds for the proposition at hand whether the proponent of that proposition is a realist or an antirealist. A realist about some domain of discourse will hold that all true propositions characteristic of that domain are either M-real (they are basic, i.e., ungrounded) propositions about M-reality) or grounded in M-real propositions (Fine 2001: 28).

One way to think of my proposal is that we use something like Fine's appeal to explanatory connections between propositions to give us a characterization of what is part of reality, where reality encompasses both M-reality and all of the propositions grounded in it. Importantly, this will include propositions that Fine would classify as antirealist, because such a proposition essentially includes components that are not components of any M-real proposition that grounds it (see Fine 2001: 18). But, on my proposal, we should separate the question of realism from that of grounding: we can look to the grounds of a given true proposition to see what makes it part of reality and what kind of explanatory role it plays. We can maintain that the propositions about which we are antirealist are those that are mind-dependent; looking at the grounds for a given proposition might help here too—we can see whether the proposition or some component of it essentially depends at some point in the right kind of way on mental activity. Excluded from reality are those propositions that are false, or that are not appropriately grounded in the set of true, basic propositions that are part of M-reality.

Fine's proposal as I have described it concerns propositions and their components, but this is easily translated back into talk of 'phenomena' if we think of the components of propositions as phenomena like moral properties, atoms, witches, and so on. As I have described it, this way of presenting the view relies on Fine's notion of M-reality, but this isn't essential. We could instead appeal to the notion of grounding as a robust form of explanatory connection between true propositions such that components of relevant true propositions that are part of this system are part of reality.⁸ The proposal also makes use of Fine's notion of grounding and of metaphysical explanation which are generally taken to be robustly realist. It is however consistent with an antirealist account of grounding and/or metaphysical explanation (see, e.g., Thompson 2019; 2022).

Again, this proposal is very permissive about what gets to count as part of reality, but this kind of permissive conception is defended in the literature. For example, Amie Thomasson (e.g., 2015) argues that we can establish the existence of all sorts of things on the basis of 'easy arguments', where we start with an uncontroversial truth

⁸ This way of putting it excludes us being committed to the reality of components of negative existential propositions, for example.

and reason through some trivial steps to an ontological conclusion. Though this is a deflationary approach to metaontology, Thomasson claims that we should be ‘simple realists’ about the entities established on the basis of these arguments: there is nothing deflated about the entities themselves. They are part of reality in the only sense there is of being part of reality (see Thomasson 2015: 153). Again, my proposal here divorces realism about a thing (which involves mind-independence, for the sake of argument) from being *part of reality*, which is a matter of being explanatorily connected in the right kind of way to other things, and in particular to those things which we take to be explanatorily central. Some of those things that feature very prominently in successful explanations might themselves be mind-dependent. This is key to responding to the RML problem.

4. Concluding Remarks

I have suggested two ways in which we might respond to the RML problem, both of which differ from the response Taylor (2023), who describes the problem, puts forward. The first is something of a middle way between traditional realism and deflationism, according to which we can accept that there are better and worse ways to carve things up but deny that the relevant ‘joints’ are entirely mind-independent. Some ways of thinking and talking are genuinely better than others, but part of what makes them better is that they are better *for us*, given some or all of factors including our particular interests, conceptual schemes, patterns of concern, and explanatory aims. This view allows that debates (e.g., in social metaphysics) can be substantive when they are cast in joint-carving terms, and it renders the problematic epistemology of the realist notion of joint-carving more tractable. The second is to think permissively about the notion of ‘reality’ such that reality is a system of explanatory dependence that includes some suitably embedded mind-dependent phenomena, and thus to sever the connection between realism and reality (at least, insofar as we continue to think of realism in terms of mind dependence).

NAOMI THOMPSON 

PHILOSOPHY, UNIVERSITY OF BRISTOL, BRISTOL, UNITED KINGDOM OF GREAT
 BRITAIN AND NORTHERN IRELAND,
naomi.thompson@bristol.ac.uk

References

- Barnes, E. (2014). Going Beyond the Fundamental: Feminism in Contemporary Metaphysics. *Proceedings of the Aristotelian Society*, 114, 335–51.
- Barnes, E. (2017). Realism and Social Structure. *Philosophical Studies*, 174(10), 2417–2433.
- Baron, S. and Norton, J. (2021). Metaphysical Explanation: The Kitcher Picture. *Erkenntnis*, 86, 187–207.
- Bennett, K. (2017). *Making Things Up*. Oxford: Oxford University Press.
- Dasgupta, S. (2018). Realism and the Absence of Value. *Philosophical Review*, 127(3), 279–322.
- Fine, K. (2001). The Question of Realism. *Philosopher's Imprint*, 1–30.

- Goodman, N. (1983). *Fact, Fiction and Forecast, fourth edition*. Cambridge, MA: Harvard University Press.
- Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press.
- Hirsch, E. (2002). Quantifier Variance and Realism. *Philosophical Issues*(12), 51–73.
- Kovacs, D. (2020). Metaphysically Explanatory Unification. *Philosophical Studies*, 177(6), 1659–1683.
- Lewis, D. (1984). Putnam's Paradox. *Australasian Journal of Philosophy*, 61(3), 221–236.
- Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.
- Mikkola, M. (2015). Doing Ontology and Doing Justice: What Feminist Philosophy Can Teach Us About Meta-Metaphysics. *Inquiry: An Interdisciplinary Journal of Philosophy*, 58(7–8), 780–805.
- Morganti, M., & Tahko, T. (2017). *Moderately Naturalistic Metaphysics*. Synthese, 2557–2580.
- Paul, L. A. (2012). Metaphysics as Modeling: The Handmaiden's Tale. *Philosophical Studies*, 160, 1–29.
- Putnam, H. (1987). *The Many Faces of Realism*. Open Court.
- Richardson, K. (2023). The metaphysics of gender is (relatively) substantial. *Philosophy and Phenomenological Research*(107), 192–207.
- Schaffer, J. (2009). On What Grounds What. In D. Chalmers, D. Manley, & R. Wasserman, *Metametaphysics: New Essays on the Foundations of Ontology* (pp. 347–383). Oxford: Oxford University Press.
- Schaffer, J. (2017). Social construction as grounding; or: fundamentality for feminists, a reply to Barnes and Mikkola. *Philosophical Studies*, 174(10), 2449–2465.
- Schaffer, J. (2010). Monism: the priority of the whole. *Philosophical Review*, 119(1), 31–76.
- Sider, T. (2011). *Writing the Book of the World*. Oxford: Oxford University Press.
- Sider, T. (2017). Substantivity in feminist metaphysics. *Philosophical Studies*, 174(10), 2467–78.
- Taylor, B. (1993). On Natural Properties in Metaphysics. *Mind*, 102(405), 81–100.
- Taylor, E. (2016). Naturalness in Context. *Inquiry*, 59(4), 319–342.
- Taylor, E. (2020). Social categories in context. *Journal of the APA*, 171–187.
- Taylor, E. (2023). Substantive Social Metaphysics. *Philosopher's Imprint*, 23(18), 1–18.
- Thomasson, A. (2009). The Easy Approach to Ontology. *Axiomathes*, 19(1), 1–15.
- Thomasson, A. (2015). *Ontology Made Easy*. Oxford: Oxford University Press.
- Thompson, N. (2019). Questions and Answers: Metaphysical Explanation and the Structure of Reality. *Journal of the American Philosophical Association*, 5(1), 98–116.
- Thompson, N. (2022). Setting the Story Straight: Fictionalism about Grounding. *Philosophical Studies*, 179, 343–361.
- Thompson, N. (2023). How to be an antirealist about metaphysical explanation. *Ratio*, 36(4), 260–273.
- Warren, J. (2016). Sider on the Epistemology of Structure. *Philosophical Studies*, 173(9), 2417–2435.
- Wilson, J. (in progress). *Fundamentality and Metaphysical Dependence*. Oxford: Oxford University Press.