The Structure of an Unknowable System

Justin P. Holder[*]

October 2024

One may demarcate at least three major forms of philosophical antirealism: scientific, sceptical, and transcendental. Each regards certain systems as unknowable. And for each, it has been argued that even if the nature of a system is unknowable, its structure may still be known. Structuralism in this sense is as yet ununified: even though structure is supposed to solve a similar problem in each case, proposals tend to be formulated in ways that make them specific to their respective contexts. Here I present a unified framework for making claims about only the structure of a target system which are not undermined by any of the three forms of antirealism. The framework may therefore provide a baseline for our knowledge of the world which is safe from most philosophical antirealism. The defining characteristic of my approach is how it leverages the role that an unknowable system plays in underlying experiences. I develop a formula for sentences making claims about the structures of target systems and I account for the representation of systems by mathematical structures as well.

1. Introduction

I look out my window and see an avocado pear tree. No one can tell the full story of what is happening here. Even the staunchest scientific realist must admit that physics, chemistry, biology, and botany are all currently incomplete sciences, so the ultimately true and complete tale of how my experience came to be cannot yet be told. For others, there is no reason to go

[*] holder.justinp@gmail.com

this far: they do not think that science, or anything else, will ever give us an ultimately true and complete story. Consider the following characters:

> The **Scientific Antirealist** says that we lack convincing reasons to think that our best theories about the unobservable systems which underlie my experience are true.
>
> The **Sceptical Antirealist** says that I can never be sure whether or not some undetectable weird system, like an evil demon's magic or a computer simulation, underlies my experience.
>
> The **Transcendental Antirealist** says that human cognitive capabilities can never grasp the mind-independent nature of the reality which underlies my experience.

"System" as I use the term can refer to any part of reality whatsoever, including possibly a single object, a group of objects and the relations between them, or anything else. I will discuss the philosophical background of each antirealist character shortly but, before that, I want to comment on some commonalities. None of the three antirealists will dispute that I am having an experience with certain characteristics – characteristics that can be appropriately reported, according to the conventions of my language, as "I see an avocado pear tree through the window". Furthermore, none of the three antirealists will deny that there is a real system underlying my experience: they are not inclined to think that experiences spring forth spontaneously from nothing. So, all parties to this discussion should have these two commitments in common to serve as points of departure for discussion:

Experiential Givenness (EG): experiences occur and can be known to occur.

Experiential Embeddedness (EE): the occurrence or characteristics of any experience are dependent on a system or systems which are not identical with that experience.

EG may be controversial not because many people deny that experiences occur but because this phrasing seems too vague. There are various theories about the nature of experience, and some people are disjunctivists about perception, meaning that they hold (at minimum) that veridical and non-veridical experiences are essentially different sorts of things. So there can be disagreement about what is actually occurring in an instance where each party is apt to use the word 'experience' to label it. To accommodate this issue, we should interpret EG to be stating that there are occurrences which can be known to occur, and there is agreement that these occurrences are the referents of the word "experiences" despite disagreement over their nature.

It is difficult to imagine disagreement over EE. Even if solipsism were true, experiences would depend on the constitution of the one person in existence. The characteristics of experiences in Berkeleyan idealism depend on people's minds and the mind of God. So, even in extreme antirealism, there doesn't seem to be any tendency to hold that each experience is utterly independent of anything else. Perhaps EE would be challenged by a metaphysical monist like Śaṅkara, who held that there was ultimately only one thing in existence and that any appearance of multiplicity is an illusion (see Dalal 2021). I won't have space to discuss this idea so I must ask the reader to put it aside, along with any philosophical doctrine that seeks to eliminate experience from reality altogether.

What I will show in this paper is that these two innocuous commitments – EG and EE – are the only assumptions one needs to mount a structuralist response to all three antirealists we met above. A 'structuralist response' is meant to be something of a middle-way between traditional realism and traditional antirealism. It grants to the antirealist that the *nature* of the system in question is unknown or unknowable, but it grants to the realist that the *structure* of the system is known or knowable. Structuralism is already present in the literature on all three forms of antirealism I mentioned and I will review that literature where appropriate throughout the paper. It will be evident that even though broadly structuralist responses exist for each form of antirealism that I have identified, those responses are insulated from each other by their respective contexts. For example, scientific structuralists tend to assume that the observable world is unproblematically real in a way that the sceptical or transcendental antirealist would object to. And transcendental structuralism tends to be presented as something which follows from Kant's transcendental idealism, so it may be dismissed by philosophers unwilling to take that as a starting point. This friction between different forms of structuralism undermines how significant structuralism really is (or could be) to epistemology by keeping it parochial. A unified structuralism which is equally amenable to each context promises something of great value. It could allow us to say (for example) "*this* is what we know about the system underlying the author's experience of the pear tree" where our knowledge of "this" is not threatened by the arguments of *any* of our antirealists. If that is defensible, structuralism can provide a baseline for our knowledge of the world which is safe from most philosophical antirealism. But structuralism has not yet been articulated in a way that could fulfil this context-neutral purpose. My aim is to do so here, using the combination of EG and EE as structuralism's context-neutral starting-point.

My strategy begins with the existence of a model which presents a system as underlying some experience. From there, depending on which is appropriate, we may either appeal to a mathematical structure in the model as representing a target real system, or we can opt for a linguistic description of a real system based on the model. The linguistic route should conform to a predicate logic formula that I will develop here. The formula builds upon the well-known

Ramsey sentence strategy, but it does not produce modified Ramsey sentences. Using my strategy and based (let's say) on models of our best scientific theories, someone may claim something very significant and perhaps true about what underlies my pear-tree experience without being undermined by the possibility that those scientific theories are completely wrong about the nature of reality, or the possibility that I am a brain-in-a-vat, or the possibility that the nature of reality transcends human understanding.

In §2 I will give an overview of some structuralist literature that is crucial for context and reference. After that, there will be a three-tiered development and discussion of my formula, with each tier corresponding to one of the three forms of antirealism. In §3 I develop my approach as a form of scientific structural realism. In §4, I adapt the strategy to sceptical antirealism, showing that it can be used to answer both scientific and sceptical antirealism equally. In §5 I discuss transcendental antirealism, explaining what interpretive moves must be made to think that my approach can successfully articulate the structure of even transcendent systems.

2. Structural realism in perception and science

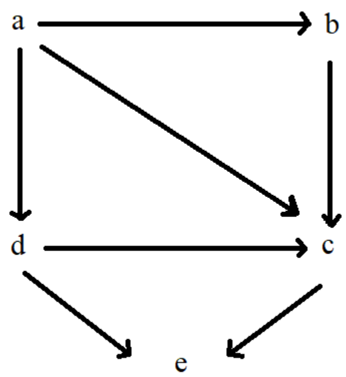

Figure 1

In the early 20th century, Bertrand Russell was convinced that he had defused a certain sort of transcendental antirealism – at least insofar as it applies to mathematical sciences – through appeal to "structure". He explained that structure is an abstraction from relations: it can be understood as a "map" which charts the pattern a relation or relations instantiate in some domain (Russell 1919/1993, 60). His example is a relation whose extension is the couples $\langle a, b\rangle, \langle a, c\rangle, \langle a, d\rangle, \langle b, c\rangle, \langle c, e\rangle, \langle d, c\rangle, \langle d, e\rangle$. The structure of that relation is the pattern pictured in Figure 1. The key feature of structure here is that it is abstract and multiply realizable: there are many relations whose extensions in some domains will form the same pattern (i.e., structure) pictured in Figure 1. This was the introduction of a technical understanding of 'structure' to epistemological problems in philosophy (c.f. Carnap 1928/2003, p.30). It will do as a working understanding of what 'structure' means for our purposes for now.

According to Russell, "a great deal of speculation in traditional philosophy" might have been avoided if philosophers had realised the importance of structure. He claims that if (as is "often said") subjective phenomena "are caused by things in themselves" then, contrary to the notion that the things in themselves would be utterly mysterious, "the objective counterparts [of phenomena] would form a world having the same structure as the phenomenal world, and allowing us to infer from phenomena the truth of all propositions that can be stated in abstract terms and are known to be true of phenomena." According to him, "the only difference must lie in just that essence of individuality which always eludes words and baffles description, but which, for that very reason, is irrelevant to science" (Russell 1919/1993, 61). Clearly, Russell has in mind Kantians and European idealists influenced by Kant.[1] We will take a closer look at this when we discuss transcendental antirealism in §5. What I want to draw attention to now is the fact that this structuralism is presented as a theory of *perception*. Russell is appealing to structure as an account of our knowledge specifically of the things which affect our senses and stimulate representations within our minds. This is thus a structuralist response to transcendental antirealism in that it grants that perception is incapable of informing us of the nature of the world around us but affirms that perception does inform us about that world's structure.

[1] Other philosophers closer to Russell's stance, like Schlick (1918/1974, §25-26), wrote of "things-in-themselves" but did not emphasize their mysteriousness, which seems to be Russell's core grievance.

Building on Russell's work, structuralism as a causal theory of perception was furthered by Grover Maxwell (1968; 1970a; 1970b). Maxwell was not only the first (I think) to coin the term "structural realism", but he was the first to use Ramsey sentences to formalize the idea. Ramsey sentences would go on to be a pillar of philosophical research on structural realism. The significance of Maxwell's deliberate and argued adherence to *realism* should not be ignored. Many philosophers will be familiar with Ramsey sentences due to Carnap, but questions of realism and antirealism were anathema to Carnap, for whom ontology was a matter of convenience (Carnap 1950). Carnap's use of Ramsey sentences can be appropriately called structuralism but had nothing to do with knowledge of metaphysically real, mind-independent systems. It therefore does not have the same goal as structural realism and does not suffer the same difficulties (see Friedman 2011). The use of Ramsey sentences for structural realism was a novel application of the device by Maxwell. Let's look at how it works.

Using Maxwell's (1970a, 186) own example will be instructive. Say we want to make the claim that if a radium atom decays then there will be a click "in an appropriately located and properly constructed Geiger counter". We might express this as follows:

$$\text{G:} \qquad \forall x[(\mathrm{A}x \wedge \mathrm{D}x) \rightarrow \exists y \mathrm{C}y]$$

Where $\mathrm{A}x$ means "x is an appropriately located radium atom", $\mathrm{D}x$ means "x decays", and $\mathrm{C}y$ means "y is a click on an appropriate Geiger counter". To 'Ramsify' this sentence, we first note which predicates are 'observational' and which are 'theoretical' (i.e., unobservable); here, C is in the former category and A and D are in the latter. We then replace only the theoretical predicates with existentially bound variables, which gives us G's Ramsey sentence, R(G).

$$\text{R(G):} \quad \exists \mathrm{X} \exists \mathrm{Y} \forall x[(\mathrm{X}x \wedge \mathrm{Y}x) \rightarrow \exists y \mathrm{C}y]$$

R(G) is supposed to say that if something instantiates both of two specific properties then there will be an appropriate Geiger counter click. This is a great deal less descriptive than G but, argues Maxwell, we may still interpret the Ramsified terms as representing real, mind-external relations and relata. Keep in mind that a property is logically a monadic relation, so when a Ramsey sentence expresses a logical pattern between predicates of any adicity, it expresses a pattern of relations, i.e., a structure in accordance with Russell's definition that we touched on at the beginning of this section.

The reader may notice that this does not look like an example from a causal theory of perception. It describes clicks in Geiger counters, not internal appearances of clicks in Geiger counters. That is correct: in the piece in which Maxwell gives this example, though he clearly expresses that he does not believe that anyone genuinely observes anything outside of their own mind, he puts that idea aside for the sake of discussion (1970a, 181). This results in the awkward situation where Maxwell's stated motivation for his structural realism does not match up well with his formalisation of it. One might think that this isn't a big deal: just reinterpret C as a mind-internal click within a suitably located and able observer. In this case, that seems doable. But it is not generally true that observational terms within a Ramsey sentence can be interchangeably interpreted as referring either to external observables or internal appearances without causing problems one way or the other. Take, for example, the simple claim (don't mind whether it is true) that stones are constituted only by atoms. We might express this as,

$$\text{S:} \qquad \forall x \forall y[(\mathrm{T}x \wedge \mathrm{N}yx) \rightarrow \mathrm{A}y]$$

where Tx means "x is a stone", Nyx means "y constitutes x" and Ay means "y is a set of atoms". It's obvious that if we reinterpret T to describe a subjective appearance then the topic of the sentence will change entirely, Ramsified or not. To bring this sentence in line with a Russell/Maxwell causal theory of perception, we would have to reconstruct it something like this:

$$\text{S*:} \qquad \forall x \forall y \forall z[(\mathrm{U}x \wedge \mathrm{B}xy) \rightarrow (\mathrm{N}zy \rightarrow \mathrm{A}z)]$$

Here, Ux means "x is the subjective appearance of a stone" and Bxy means "x arises through sense affectation by y". N and A are the same as before. Let's be generous and assume that complications associated with illusions, photographic images, etc., can be dealt with appropriately so that, in models of S*, y should always pick out a real stone. To Ramsify S* in the Russell/Maxwell way, we need to leave B interpreted since it is a fundamental element of their theory of perception:

$$\text{R(S*)}: \exists X \exists Y \forall x \forall y \forall z[(Ux \land Bxy) \rightarrow (Xzy \rightarrow Yz)]$$

This, according to Russell/Maxwell, would be the philosophically correct way to make the realist assertion that stones are constituted only of atoms.[2] Logically, there is nothing about this kind of Ramsey sentence that is intrinsically more problematic than a Ramsey sentence with non-subjective observational terms like R(G), but a big problem with the Russell/Maxwell approach will become evident in the next section when we discuss responses to what is known as Newman's objection.

The Russell/Maxwell causal theory of perception never attracted great attention in the philosophical literature. Structural realism instead took off when Worrall (1989) explained its potential to accommodate both the most influential argument for scientific antirealism and the most influential argument for scientific realism. The latter is the no miracles argument which claims that the incredible success of science would be a miracle if realism were false;[3] the former is the pessimistic induction which points out that past successful scientific theories have been entirely replaced so we should expect the current ones to be replaced as well (Laudan 1981). Structural realism, argued Worrall, can be the "best of both worlds" by showing that only the non-structural properties of past successful theories were fully replaced in newer

---

[2] To do this properly, the existential quantifiers in R(S*) should also bind X and Y across Ramsified expressions of well-confirmed physical theories in which predicates for "is constituted by" and "is an atom" have been substituted with X and Y respectively.

[3] This line of reasoning has a longer history but the "miracle" phrasing is due to Putnam (1975, 73).

theories. The pessimistic induction therefore has no force against commitment to the structural claims of current theories, and a realist interpretation of those structural claims alone should satisfy the no miracles argument. Worrall thus folded structural realism into a major pre-existing debate within the philosophy of science community, guaranteeing its relevance. He acknowledges in his paper that Maxwell coined "structural realism" before him, but he regards the Russell/Maxwell project as arising from different, "more 'philosophical'" concerns than his own, which he attributes genealogically to Poincaré; he calls for further research into the relationship between these projects (Worrall 1989, 117 note 32). There is no space to discuss Poincaré here except to say that I don't think his developed opinions should be considered structural realism at all (see Holder 2023). As regards the distinction between the Russell/Maxwell project and the Worrall project, the only material difference lies in the fact that Worrall was unconcerned with epistemological difficulties in perception and thus formulated his structural realism in a way consistent with naïve realism, as is the norm with scientific antirealism. That is borne out by the fact that Worrall uses Ramsey sentences for his structural realism in exactly the form that Maxwell suggested when his causal theory of perception was bracketed (Zahar and Worrall 2001).

Now, how do these ideas relate to my goal of formulating a structuralism which is not undermined by scientific, sceptical, or transcendental antirealism? Before considering the use of Ramsey sentences in any form, one must confront the problem known as 'Newman's objection' which threatens to undermine them.

2.5. Newman's objection

In 1928, the mathematician M.H.A. Newman published a criticism of Russell's causal theory of perception. This was before Ramsey sentences had been introduced to the topic, so Russell's claim was simply that perception informed us only of the world's structure. Newman argued that if all we are allowed to know is that such-and-such a structure exists in the world, then all our knowledge is trivially true given that the world has enough things in it. The idea is that for any set of objects, one can define relations instantiated by the objects which satisfy any structure consistent with the cardinality of the set. So, knowledge of only the world's structure is equivalent to knowledge of the minimum number of things in the world. This is known as Newman's objection. Russell conceded to Newman's objection without developing an alternative formulation of his structuralism (Demopoulos and Friedman 1985, 632).

With structural realism expressed through Ramsey sentences, it is not obvious that Newman's 1928 argument has the same force since Ramsey sentences do not contain *only* structural information. But Demopoulos and Friedman (1985) adapted Newman's work to the developments in the field, arguing that once a Ramsey sentence's observational consequences are true, "then the truth of the Ramsey-sentence *follows* as a theorem of set theory or second-order logic, provided our initial domain has the right cardinality" (635, emphasis in original). This was reinforced by Ketland (2004). The way to defeat Newman's objection is to somehow specify *which* relations supposedly form the target structure without thereby admitting enough knowledge of the nature of those relations to undermine the entire point of the structural realism project. Notably, Psillos (1999) has argued that this cannot be done and Ainsworth (2009) surveyed various attempts to do it, concluding that they all fail.

A possibility I want to address now is that, if we bite the bullet on quantifying over 'subjective appearances' and have a decent solution to Newman's objection, Ramsey sentences in the Russell/Maxwell vein like R(S*) could achieve the goal I laid out for this paper. The best regarded solution to Newman's objection in the existing literature is Melia and Saatsi's (2006). Their proposal is that Ramsey sentences be augmented with intensional modal operators which declare a law-like relation between the sentence's predicates, such as L_P: "it is physically necessary that…". This allows the sentence to assert the important causal/explanatory implications of the target system according to the original theory, and it is not trivial that relations exist which form the target structure *and* satisfy the modal implications of those in the original theory. Newman's objection therefore fails to take hold.

Here is Melia and Saatsi's solution applied to R(S*):

$$\exists X \exists Y \forall x \forall y \forall z[(Ux \wedge Bxy) \rightarrow \mathbf{L_P}(Xzy \rightarrow Yz)]$$

L_P is positioned this way so that it applies to the physical modality of stones themselves, not subjective appearances.[4] This is an impressive strategy for guarding Ramsification against Newman's objection, but it has some characteristics which make it less than ideal for my purposes. I want my structuralist proposal to function as a baseline for our knowledge of the world by being agreeable to virtually any antirealist or realist. But a standalone Ramsey sentence (however modified) will either quantify over 'subjective appearances' or not, and either option will prejudice the sentence against some theory of perception which either requires or rejects the existence of objects fitting such a description. Furthermore, Melia and Saatsi's strategy requires commitment to modal facts and invites metaphysical and epistemological debate over their existence or nature – it would be best for me to avoid this and I think that I can. See also Smithson's (2017) discussion of a difficulty with Melia and Saatsi's use of modality that my strategy will avoid.

Nevertheless, I think that Melia and Saatsi had the right idea in invoking the causal/explanatory implications of relations to deflect Newman's objection. And my own proposal cannot afford to assume that perception gives us more knowledge of reality than Russell and Maxwell thought it did. I have rejected the prospect of simply modifying a Ramsey

[4] If this still seems too uninformative, consider my footnote 2.

sentence to accommodate these two considerations. In what follows I will explain a different approach which I believe does the job elegantly.

3. The structure of an unobservable system

I will be explaining my structuralist proposal in stages. In this section, I will be focused only on addressing scientific antirealism, having put aside sceptical and transcendental antirealism. That is, we will assume, for now, that there are no epistemological problems with perception. We are currently concerned only with the structure of systems which are unobservable even by naïve realist standards.

The first concern I want to address is whether the structure of an observable system should be formulated linguistically at all. Most ontic structural realists – philosophers for whom structural realism is a metaphysical, not an epistemological position (see Ladyman 1998; Ladyman and Ross 2007; French 2014) – disavow linguistic devices like Ramsey sentences as appropriate vehicles for structural realism since, for them, ultimate physical reality can only be represented in mathematical models of our best theories of fundamental physics. Wallace (2022) offers a helpfully thorough presentation of this perspective, calling it 'math-first' (contrasted with 'language-first') realism.[5] According to the math-first view, the theories of physics tell us about the world not with sentences but through representing real systems within mathematical models. A scientific model can be any constructed entity which exemplifies theoretical axioms or equations. Beyond mathematical models, we may also speak of physical models such as material or digital artifacts in which a set of axioms or equations are true. But mathematical models are what lie at the heart of modern physics.

[5] The language-first/math-first distinction mirrors the more common syntactic/semantic distinction in approaching scientific theories. I prefer Wallace's terminology for the reasons he gives (2022, 348-9).

Math-first realism is necessarily structural in that mathematical structure is taken to be the knowledge we have of the relevant systems. This is a departure from predicate logic and the object-relation schema, and therefore from Russell's definition of structure as a pattern of relations. Any logico-mathematical pattern can be considered 'structure' from this point of view, which I am happy to embrace as an expansion of Russell's definition. Wallace also clarifies that the math-first view can be used for either epistemic structural realism according to which we know that the represented system has the target structure (its nature is mysterious) or ontic structural realism according to which the mathematical structure represents the very nature of the system. I am concerned in this paper only with epistemic structural realism.

I am sympathetic to the math-first approach. It is not a problem for or challenge to my project even though I will be developing a linguistic device for the representation of structure. I am trying to develop a maximally general structuralism, and while the math-first approach might be the best one for fundamental physics, it won't be very convenient if one wants a structuralist alternative for saying that there is a tree outside one's window. Even the most convicted math-first proponent should acknowledge that when it comes to high-level patterns in worldly structure – like trees – language is an appropriate tool for humans navigating those patterns in most situations (Wallace acknowledges this throughout his paper). That means linguistic devices must be a central part of a generalist structuralism meant to be accessible throughout epistemology. However, fundamental physics must also be included, so the math-first view must be accommodated. In my strategy, this is done by making models the starting-point of the approach while describing how to articulate a linguistic or mathematical structuralist claim from that common basis.

Before getting to it, I must define my technical term, 'observational fact'. An observational fact is something that is the case in the observable world. "Observable world" should be understood in a naïve realist way, and nothing rests on the precise demarcation between observable and unobservable here; suffice it to say that things like bicycles and grapes are observable and things like electrons and electromagnetic fields are unobservable. Note that 'observational fact' does not mean something *observed to be* a fact. For example, "the sun has risen every morning for the past five thousand years" and "it is impossible for light to pass through lead" might be considered facts about the observable world even though they are not verifiable through observation. An observational fact may consist simply in the existence of some phenomenon (e.g., a reading on an apparatus, a sound or movement), or it may consist in a regularity (e.g., the sun is never visible in Bridgetown at 1:00am local time; under *x* conditions, the reading on this apparatus is always a value greater than *y*). Models can be used to represent and help explain any sort of observational fact.

The following summarises my strategy. Note the stipulations (i – iv):

For some model, *M*,

i. O is an observational fact presented[6] in *M*.
ii. In *M*: certain specific relations are instantiated by specific relata such that a certain structure S is satisfied OR a mathematical structure T is presented.
iii. The fact described in (ii) underlies the fact that O in *M*.
iv. O represents an observational fact in reality, O'.

Our claim is as follows:

1. In reality: certain specific relations are instantiated by specific relata such that S is satisfied OR there is a certain system represented by T.
2. The fact described in (1) underlies the fact that O'.
3. For any logico-mathematical connection found between the facts mentioned in (2), approximately the same logico-mathematical connection can be found between the facts mentioned in (iii).

[6] I follow Katherine Brading's and Elaine Landry's helpful distinction between the system that a model *presents* and the system that it *represents* (Brading and Landry 2006, p.573; Brading 2011, p.48). The former is the system which is described by the model and which does not necessarily exist beyond it; the latter is the real system that can be identified as what the model has described.

Shortly, I will provide an example to demonstrate how this works in action. First, I will clarify a few things.

The disjunctions in (ii) and (1) are inclusive. Moreover, S could be a predicate precisification of T (it needn't be the only possible such precisification). Newman's objection is deflected for the linguistic (S) route here because we are picking out the instantiation of specific relations by specific relata – there is no assertion of there simply *being* relations among specific relata which form S (that would be trivial with the right cardinality). Nevertheless, this would go nowhere without providing more information about the specific relations and relata we are targeting. That information is provided in (2) where we claim that the fact that those relations are instantiated by those relata underlies the fact that O'. That the instantiation of relations forming S underlies O' is certainly not trivially the case – so not subject to Newman's objection. Newman's objection is not a serious worry for the math-first (T) route because being represented by a mathematical structure is not subject to that objection (though the 'representation' relation needs philosophical attention I cannot give it here). (3) exists to block a Newman-adjacent problem which I will describe after the example.

I must clarify what "underlies" means. This is where EE from §1 does its work. One of the two fundamental unargued assumptions I have asked of the reader, EE, asks us to accept that there is a system or systems which bear causal responsibility for the characteristics of an experience. It is my intention to leave the details of this causal responsibility as thin as possible. Any reasonable interpretation of "dependent" as it appears in EE counts as a viable interpretation of causal responsibility in this sense; and denying causal responsibility in this sense entails denying EE under any reasonable interpretation. Some X underlies some Y iff X is causally responsible for Y in the sense just explained.

I now move on to the example. Say there is some solid substance (we'll just call it 'the substance') which mostly dissolves when submerged in a certain liquid ('the liquid'). It is observed that when any quantity of the substance is exposed to the liquid, the remaining solids (which look like the substance) always have a quarter of the mass and volume of the original quantity.

There is a model of a successful theory which explains this fact. In the model, the substance is identified with a certain molecule, and the liquid is identified as a catalyst which reacts with the 'substance-molecule' to ultimately produce five products: four new molecules and the catalyst itself. Let's label these new molecules A-products, B-products, C-products, and D-products. They are produced in equal measure and are equal in mass. A-products in close proximity form strong bonds with each other very quickly, making it almost impossible to keep them in a homogenous mixture with a liquid. B-, C-, and D-products, however, have no such tendency and readily mix with the catalyst. According to the model, then, the solids which remain when the substance is submerged in the liquid constituted of A-product particles.

Here, our observational fact is "when some quantity of the substance is submerged in the liquid, only one quarter of that quantity remains solid". To explain it, a theorist who subscribes to the model we have just described might point to two unobservable relations as underlying this fact: A-products *solidify much more readily than* B-, C-, or D-products, and A-, B-, C-, and D-products are *produced in equal measure to* each other. These generate the structure in Figure 2, which I have represented in the diagram style used by Russell and

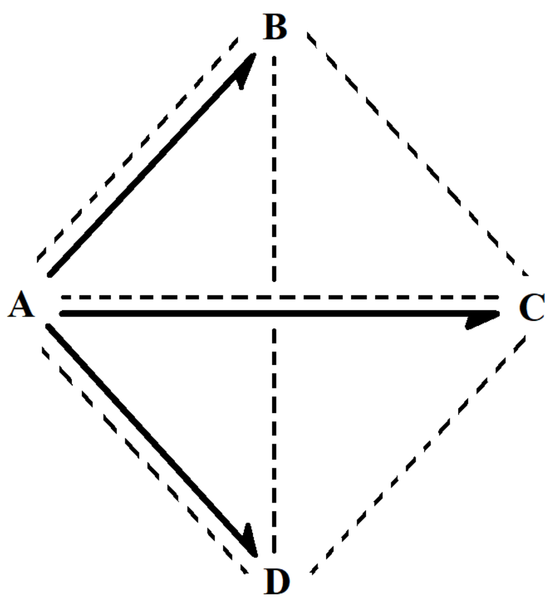

Figure 2

Newman.[7] In the structure, the solid lines represent the *solidify much more readily than* relation and the dotted lines represent the *produced in equal measure to* relation. This is an entirely non-empirical substructure of the model.

Now let us say that one of the theorists is a structural realist. They are unconvinced of the existence of any of the molecules we have referred to so far, and any of the relations which ostensibly hold between them. But they do commit to the belief that there are relations instantiating the abstract structure depicted in Figure 2, that this fact underlies the observational fact in question, and that any logico-mathematical connection between their counterparts in the model holds between this fact and the observational fact in reality. It is difficult to imagine states of affairs which satisfy all of these criteria, says the structural realist, and not remotely trivial to say that some such state of affairs is not just imaginary but actual. *That*, they say, is the knowledge the theory has given us via the model.

Now let us say that the structural realist's suspicions were correct. A-, B-, C-, and D-products do not exist. The liquid is not a catalyst but a reagent which only reacts with three-quarters of any given quantity of the substance. This is because the substance is composed not of one kind of molecule but of four isomers. Isomers are molecules which contain exactly the same number and type of atoms as one another, but have those atoms arranged in different structures.[8] We will call the isomers which constitute the substance A-isomers, B-isomers, C-isomers, and D-isomers. Each one contains one each of the atoms W, X, Y, and Z. Figure 3 shows the structure of each isomer.

[7] The structure formed by the relation represented by the solid line is the same structure that Newman uses as an example in his paper.
[8] This describes a 'structural' or 'constitutional' isomer, which is only one type. For the sake of simplicity, I am using the word 'isomer' as synonymous with 'structural isomer'.

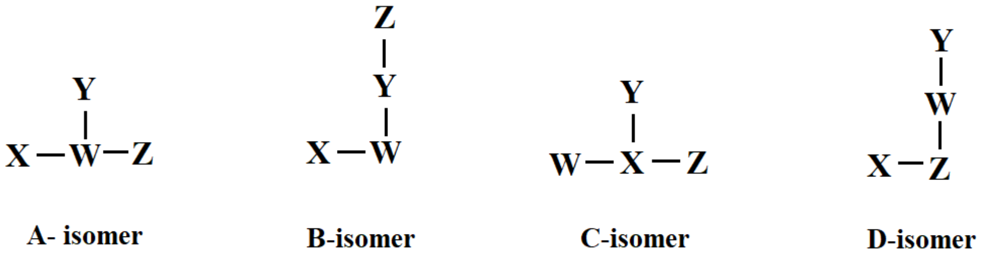

Figure 3

The process which creates the substance generates these isomers in equal quantities and in even distribution. They are *all* almost impossible to keep in a liquid state – they solidify rapidly when in close proximity to each other. What is happening is that the liquid is reacting with B-, C-, and D-isomers to produce different molecules altogether – ones which are happy to exist in a liquid mixture. A-isomers, however, happen to be much more stable than the others. They do not react with the liquid at all. When the substance is submerged in the liquid, all of the B-, C-, and D-isomers in the substance react with the liquid and the resultant products form a homogenous mixture with the liquid. Since the A-isomers do not react, they quickly bond with each other and form solids.

Here, the relation *solidifies much more rapidly than* does not exist between any of the isomers. However, we still have relations between them which underlie the observational fact and instantiate the structure in Figure 2. The relations are: *is much more stable than* (solid line) and *constitutes the substance in equal proportion with*[9] (dotted line). The letters in the diagram denote the correspondingly labelled isomers. It is *these* relations which really underlie the fact that "when some quantity of the substance is submerged in the liquid, only one quarter of that quantity remains solid" and they do indeed instantiate the structure in Figure 2. So far, the

[9] Technically, this should be *is an instance of the type of isomer of which instances constitute the substance in equal proportion with instances of the type of isomer instantiated by*, but this would make the discussion difficult to read. I hope the reader will appreciate my intent.

structural realist is vindicated. But what of (3), the logico-mathematical connection criterion? Now it can be looked at more closely.

We can see from Figure 3 that the atoms in an A-isomer also instantiate the structure in Figure 2. The bonding relation between atoms instantiates the substructure drawn by the solid lines and the relation *constitutes an A-isomer in equal proportion with* instantiates the substructure drawn by the dotted lines. Given that it is this particular structure which has resulted in the A-isomer's relative stability, it could be reasonably argued that the fact that these relations are instantiated also underlies the observational fact. These 'background' relations therefore also satisfy (1) and (2) of my proposal. That is worrisome: if it were limited to (1) and (2), the structural realist's claim in this example would come out true for *any* observational fact underlied by A-isomers. It would therefore fail to pick out the causal/explanatory facts that the structural realist is interested in. Call this 'the problem of background structure'.

To see how (3) deflects this problem in this case, one should start by looking for a logico-mathematical connection between the underlying relations and observational fact presented in the model. That only a **quarter** of the original quantity remains solid is a consequence (in the model) of the fact that there are **four** products instantiating the *solidify much more readily than* and *produced in equal measure to* relations. In reality, only a quarter of the original solid remaining is a consequence of the fact that there are four isomers which instantiate the relations *is much more stable than* and *constitutes the substance in equal proportion with*. The quarter/four logico-mathematical connection between the target structure and observational fact in the model is thus preserved in reality, even though the model's ontology is incorrect. Looking at the problematic 'background structure' we identified among the atoms in A-isomers, however, the fact that only a quarter of the substance remains solid is not a logical consequence of there being four atoms in an A-isomer. The logical connection from the model is not preserved for that structure, so (3) is not satisfied and the problem of background structure is blocked in this instance.

Since the *underlies* relation is so conceptually thin, a great many things could be said to underlie almost anything. And since my strategy intentionally gives great freedom to the structuralist in deciding what structure they want to commit to, the structure of their target system could be extremely simple. Those two points press the seriousness of the problem of background structure. Key to (3)'s effectiveness in blocking it is the claim that the preservation of logico-mathematical connections from model to world holds for *any* logico-mathematical connection that can be identified between the target system and the observational fact in the model. Thus the more the model tells us about the connection between target structure and observational fact, the less likely background structure will be a problem. This is natural: the richer the model, the less likely our structural realism fails to be sufficiently specific.

That concludes my example for now. In it, structure was represented pictorially and given a linguistic interpretation. Before ending the chapter, I will comment briefly on the math-first route. We can expect the models of mathematised sciences present precise, complex mathematical systems with carefully studied properties. Being mathematical, the presented system is inherently structural, so the claim that it represents a real system – while refraining from applying a predicate precisification interpreted in terms of (metaphysical) objects, properties, and relations – is structural realism implicitly. Asserting that the target system underlies an identified observational fact removes the threat of triviality. And exploring the logico-mathematical connections between two parts of a mathematical model is a much more easily understood task than exploring such connections between parts of a non-mathematical model or between sentences describing a model. So, my strategy reinforces the idea, frequently touted by ontic structural realists, that the math-first approach is more naturally suited to structural realism than the language-first approach.

For math-first realists, the biggest problem with my proposal is likely the requirement (i) to present the observational fact O within the model M itself, since models of fundamental physical systems may not have any explicit presentation of the observable world. My strategy requires, in van Fraassen's (1981) well-known terms, a single model which has both empirical and non-empirical substructures. O must be presented (van Fraassen would say 'embedded') in an empirical substructure of the model, and the system which has S will be presented in a non-empirical substructure. When working with a model lacking the requisite empirical substructure, theorists might believe that the model could be expanded to include such a substructure uncontroversially within established science; in such a case, my structural realist option is available even if it would take additional work to apply it rigorously. Why insist on having *M* present O in the first place? Because that is the only way to satisfactorily include (3) and deflect the problem of background structure.

The basics of my strategy are hopefully now understood. For the linguistic route, though, I must still provide a suitable formal account of how the target structure can be articulated and the appropriate structuralist claim expressed.

3.5. The formal account

The first thing we need to do with reference to some model is to identify a fact which we take to underly some observational fact. Let us do this now, using the example in the previous section.

> **Observational fact:** "When some quantity of the substance is submerged in the liquid, only one quarter of that quantity remains solid."

Underlying fact: "A-products form solids very quickly while B-, C-, and D-products do not, and each kind of product constitutes a quarter of the new molecules produced under catalysis."

Now we will formulate the underlying fact using a sentence, θ, in a second-order language.

A*x*: *x* is an A-product.

B*x*: *x* is a B-product.

C*x*: *x* is a C-product.

D*x*: *x* is a D-product.

S*x*: *x* forms solids very quickly.

Qx: *x* is an instance of one of the four types of molecule which each represent a quarter of the new molecules produced under catalysis when the substance is exposed to the liquid.

φ: $\forall w \forall x \forall y \forall z((Aw \wedge \mathrm{B}x \wedge \mathrm{C}y \wedge \mathrm{D}z) \rightarrow (Sw \wedge \neg Sx \wedge \neg Sy \wedge \neg Sz \wedge \mathrm{Q}w \wedge \mathrm{Q}x \wedge \mathrm{Q}y \wedge \mathrm{Q}z))$

I will borrow from Smithson (2017) his notation for "the fact that…". ⟦θ⟧ means "the fact that θ". Furthermore, we will refer to the observational fact listed above as O (there's no need to put this in a second-order language). With that, we will define:

$\mathbb{U}$⟦φ⟧⟦O⟧: the fact that φ underlies the fact that O

For traditional realists, $\mathbb{U}$⟦φ⟧⟦O⟧ is all that needs to be said. Structural realists, however, will be tempted to Ramsify φ to R(φ):

R(φ): $\exists \mathrm{X}_1 \exists \mathrm{X}_2 \exists \mathrm{X}_3 \exists \mathrm{X}_4 \exists \mathrm{X}_5 \exists \mathrm{X}_6 \forall w \forall x \forall y \forall z((\mathrm{X}_1 w \wedge \mathrm{X}_2 x \wedge \mathrm{X}_3 y \wedge \mathrm{X}_4 z) \rightarrow (\mathrm{X}_5 w \wedge \neg \mathrm{X}_5 x \wedge \neg \mathrm{X}_5 y \wedge \neg \mathrm{X}_5 z \wedge \mathrm{X}_6 w \wedge \mathrm{X}_6 x \wedge \mathrm{X}_6 y \wedge \mathrm{X}_6 z))$

This, of course, is trivial given that there are at least four things in the world. To fix this, we use the same basic idea as Melia and Saatsi (2006): invoke the causal/explanatory implications of the system we are interested in. But instead of using a modal operator to attempt this as Melia and Saatsi do, my strategy is to pick out the target system by tying it to the observational fact O using the *underlies* relation. We start by writing a new sentence, Ɍ(φ), which is just R(φ) without existential quantification:

$$Ɍ(\varphi){:}\quad \forall w \forall x \forall y \forall z((X_1 w \wedge X_2 x \wedge X_3 y \wedge X_4 z) \rightarrow (X_5 w \wedge \neg X_5 x \wedge \neg X_5 y \wedge \neg X_5 z \wedge X_6 w \wedge X_6 x \wedge X_6 y \wedge X_6 z))$$

We can then make use of Ɍ(θ) and the other tools we have developed to make the following claim, which is not trivial:

$$\exists X_1 \exists X_2 \exists X_3 \exists X_4 \exists X_5 \exists X_6 \mathbb{U}[\![Ɍ(\varphi)]\!][\![O]\!]$$

Based on the explanation of the example in the previous section, this turns out to be true even though $\mathbb{U}[\![\theta]\!][\![O]\!]$ is false. We can see that via the following possible interpretation of terms:

X_1x: *x* is an A-isomer.

X_2x: *x* is a B-isomer.

X_3x: *x* is a C-isomer.

X_4x: *x* is a D-isomer.

X_5x: *x* is very stable.

X_6x: *x* is an instance of one of the four types of isomer which each represent a quarter of the molecules which constitute the substance.

This strategy is enough to defeat Newman's objection, but there is still the problem of background structure. This can be solved in a similar way as before: additionally claim that any logico-mathematical connection found between the facts picked out by the terms Ɍ(φ) and O in $\exists X_1 \ldots \exists X_6 \mathbb{U}[\![Ɍ(\varphi)]\!][\![O]\!]$ can also be found between the facts picked out by the terms φ and O in $\mathbb{U}[\![\varphi]\!][\![\mathrm{O}]\!]$.[10] I do not see how this notion of logico-mathematical connection could be represented formally except as a primitive. That does not seem helpful or elegant, so I suggest leaving it as an informal addendum which only needs to be invoked to handle the problem of background structure.

To sum up, where:

θ is a sentence describing an unobservable system in a second-order language,

$\mathrm{R}(\theta)$ is the Ramsification of θ,

$\exists \mathrm{X}_1 \ldots \exists \mathrm{X}_\mathrm{n}$ are the existential quantifiers binding the predicate variables in $R(\theta)$,

$\not{R}(\theta)$ is $R(\theta)$ with those existential quantifiers removed,

O is an observational fact,

$[\![\theta]\!]$ means 'the fact that θ', and

$\mathbb{U}[\![\theta]\!][\![\psi]\!]$ means 'the fact that θ underlies the fact that ψ',

We can make a non-trivial claim about the structure of that unobservable system using a sentence of the form $\exists \mathrm{X}_1 \ldots \exists \mathrm{X}_\mathrm{n} \mathbb{U}[\![\not{R}(\theta)]\!][\![\mathrm{O}]\!]$.

§3 and §3.5 have been all about responding to scientific antirealism, so we have not challenged the ability of perception to inform us of non-structural features of the external world. We will do so in the next section. The math-first approach had us very gently question the ideology of conceiving of reality as ultimately comprised of objects and relations. We will take that up again in §5.

[10] $\not{R}(\varphi)$ is an open formula, but the fact that $[\![\not{R}(\varphi)]\!]$ can be parsed in the full expression shows that it picks out a fact in that context even though it is not a declarative sentence in isolation.

4. The structure of a deceptive system

Sceptical antirealism challenges our knowledge of the world by arguing that we cannot know that our perceptual experience is not caused by a weird, deceptive system such as the Cartesian evil demon manipulating our mind, or a computer feeding simulated sensory information to us, where it turns out we are disembodied brains in vats. Some of these scenarios (like the evil demon) may also stipulate that our grasp of basic logic and mathematics is a deception. Structuralism probably cannot get around this latter challenge, unless such cognitively confused agents could manage to construct a model with some structural correspondence to reality and recognise its value despite their corrupted intuitions. I will put that issue aside and focus only on how my structuralist proposal could function if one's *perceptual* access to the world is globally perverted by a deceptive system.

Structuralism as a response to sceptical scenarios has only, to my knowledge, been seriously explored by Chalmers (2018). But there are at least two aspects of Chalmers' approach which make it incompatible with my project. The first is that Chalmers, being focused on sceptical antirealism, ignores scientific antirealism. His discussion makes no distinction between scientific knowledge of the observable vs. unobservable world, which is the central axis of friction in the scientific realism literature. That may be fine for his paper since it is concerned only with sceptical antirealism. But my project requires me to take scientific antirealism into account, so our basic assumptions differ. The second issue is that Chalmers' account relies on the idea that our ordinary physical claims are inherently structural; this is not my approach.

The first move is to express perceptual judgments in a way which does not contradict naïve realism but also does not contradict the idea that perceptual experiences occur entirely within the mind. This is possible by accepting, first, that perceptual experiences *seem* like what the naïve realist believes them to be: presentations of the world as it is. We can accept that seeming without committing ourselves to the naïve realist view by saying that it *appears* that the world is such-and-such a way. Let's define an operator, $\mathbb{a}$, accordingly. $\mathbb{a}$(P): 'it appears that P'. There's no need to quantify over sense data or anything like that.

We're sticking with language-first structuralism for now, which is the more important approach here since sceptical antirealism primarily challenges our ordinary beliefs, not fundamental physics. We need a modified sort of Ramsey sentence. These should replace *all* predicates in the original sentence (not just the 'observational' ones) with variables and refrain from existentially binding them. Call them "Quietened Ramsey sentences" or "Q-Ramsey sentences" because, being open formulae, the do not actually say anything, and even if we did existentially bind their predicate variables, the sentences could say nothing of interest on their own. We will represent the Q-Ramsification of θ as Q(θ). Note that in cases where θ does not contain observational predicates (as in the previous discussion), Q(θ) and $\mathcal{R}$(θ) will be identical.

With these resources added to those presented at the end of the previous section, and where $\exists X_1 ... \exists X_n$ are quantifiers binding the predicate variables in $Q(\theta)$, we can make use of statements of the form $\exists X_1 ... \exists X_n \mathbb{U}[\![Q(\theta)]\!][\![\mathbb{a}(O)]\!]$. This is a non-trivial statement describing the structure of a system which relies on no knowledge of reality besides what appears in experience and that there is a system of relations underlying what appears in experience. As a demonstration, I will show that this can enable even a brain-in-a-vat to say something justified and true (and non-trivial) about reality outside of their simulated experiences.

Jalesa is the scientist who created the model which explains the interaction between the substance and the liquid in terms of A-products, etc. A bit of a sceptic, she does not want to commit herself to knowledge of the nature of the unobservable objects presented in her model. She chooses to believe the structural realist claim $\exists X_1 ... \exists X_6 \mathbb{U}[\![\mathcal{R}(\varphi)]\!][\![O]\!]$ instead of $\mathbb{U}[\![\varphi]\!][\![O]\!]$. However, in this version of the thought experiment, there are no A-isomers, etc., either. In fact, neither the substance nor the liquid exists as a physical object. In this version of the thought experiment, Jalesa is – in the real world – a brain suspended in a vat of life-sustaining fluid who is (and has always been) experiencing a simulation that is running on a computer and is transmitted to her neurons through electrical impulses indistinguishable from those that would normally come from sense faculties. Thus, one might reasonably think that $\exists X_1 \ldots \exists X_6 \mathbb{U}[\![\mathcal{R}(\varphi)]\!][\![O]\!]$ turns out to be false because the key terms in O fail to refer. If the reader is inclined to think that O remains true in these circumstances because 'the substance' and 'the liquid' will refer, here, to the simulated substance and liquid, not to imagined mind-external physical objects, then so much the better for me – my proposal will work as is with the interpretation of terms given below. But we can skip this potential point of contention by attending to the following strategy.

One day, Jalesa is convinced that she should be even more sceptical than she normally is. She decides that she will no longer trust that the objects of her perceptions exist except as mere appearances in her consciousness. She notes that, while the system presented in her model underlies O in the model, that same system also underlies $\mathbb{a}(O)$ in the model if she expands the model to include herself as an observer. In her model, then, one may pick either $\mathbb{U}[\![\varphi]\!][\![O]\!]$ or $\mathbb{U}[\![\varphi]\!][\![\mathbb{a}(O)]\!]$ as the point of departure for their structural realism. Having now opted for $\mathbb{U}[\![\varphi]\!][\![\mathbb{a}(O)]\!]$, Jalesa can take the steps we outlined above to weaken this claim into the structuralist $\exists X_1 \ldots \exists X_6 \mathbb{U}[\![\mathcal{R}(\varphi)]\!][\![\mathbb{a}(O)]\!]$. Satisfied, she decides that *this* is what she believes about the real system targeted by her theory.

On the present version of the thought experiment, the following is what is really going on. There is a team of mad scientists responsible for this unethical brain-in-a-vat research project. In a meeting, the mad principal investigator declared that the substance (which has no real counterpart outside of the simulation) must have four distinct tastes if Jalesa puts it in her simulated mouth. No one is sure why, especially since the substance is not a food, but such is the nature of this project. Unfortunately, the simulation engine only allows one to code one flavour per virtual object. The mad graduate student who had the task of coding the substance came up with the following workaround: every instance of the substance would actually be a continuous, even distribution of four objects, with the filenames sub1, sub2, sub3, and sub4, which should have identical properties other than their tastes. Incidentally, in the simulation, every object is assigned a 'dissolution resistance value' which, along with some other parameters, determines what liquids it will dissolve in when in a solid state. The substance is supposed to have a dissolution resistance value of 9, but the overworked mad graduate student accidentally hit the '9' key twice when coding sub1, resulting in a value of 99 for that file only. The result in the simulation was that, when the substance (which is really a uniform distribution of sub1, sub2, sub3, and sub4) comes into contact with the liquid, the objects sub2, sub3, and sub4 dissolve as they are supposed to while sub1 does not.

We can also say that, at some point, for whatever reason, Jalesa tastes the substance as well as the residue left after its exposure to the liquid. She finds that the substance originally has an overpoweringly complex taste, but the residue just tastes something like cherries. She finds that her theory goes a long way in explaining this fact: the residue is really made up of A-products, which makes it a distinct substance that might be expected to have other properties which differ from the original substance. Why not taste? Jalesa considers this a modest novel prediction of her theory, which strongly adds to its credibility.[11]

Jalesa's claim, $\exists X_1 \ldots \exists X_6 \mathbb{U}[\![\mathcal{R}(\varphi)]\!][\![\mathbb{a}(O)]\!]$, comes out true because of the following possible interpretation of terms:

X_1x: x has the filename sub1.
X_2x: x has the filename sub2.
X_3x: x has the filename sub3.

[11] Leplin (1997, 77) gives an account of novel prediction that can be used here.

X_4x: x has the filename sub4.

X_5x: x has a very high dissolution resistance value.

X_6x: x is an instance of one of the four virtual objects which constitute every instance of the virtual substance in even distribution.

The problem of background structure can be deflected with the same strategy as before.

This works whenever my scientific structural realist strategy works assuming that an observer can be modelled. Notice, too, that e.g., even though we cannot observe that light cannot pass through lead, it can *appear* to us that light cannot pass through lead – so this doesn't limit us to observational facts that are verifiable in experience. Furthermore, this can be used to formalise the Russell/Ramsey causal theory of perception. If all one wants is a formal way of saying that there is something (there's no telling what) which causes one's experience of a duck, one can say $\exists X \mathbb{U}[\![\exists x(Xx)]\!][\![\mathbb{a}(D)]\!]$ where D means "there is a duck". Importantly, a naïve realist would also consider that sentence true.

To sum up as we did before, where:

θ is a sentence describing a system in a second-order language,

$Q(\theta)$ is θ where variables have been substituted for *all* predicates,

$\exists X_1 \ldots \exists X_n$ are existential quantifiers binding the predicate variables in $Q(\theta)$,

O is an observational fact,

$\mathbb{a}$(O) means "it appears that O",

$[\![\theta]\!]$ means 'the fact that θ', and

$\mathbb{U}[\![\theta]\!][\![\psi]\!]$ means 'the fact that θ underlies the fact that ψ',

We can make a non-trivial claim about the structure of a system, even if our senses do not provide objective knowledge of the nature of reality, using a sentence of the form $\exists X_1 \ldots \exists X_n \mathbb{U}[\![Q(\theta)]\!][\![\mathbb{a}(O)]\!]$. On the math-first approach, we can just say that some mathematical structure in our model represents a system which underlies the fact that it appears that O. It should go without saying that there is no guarantee of these claims being true, just like any structural realist claim. But we shouldn't expect being in a sceptical scenario to prevent them from being true or stop us from figuring them out.

5. The structure of a transcendent system

The final topic is transcendental antirealism, the idea that the nature of reality is inaccessible in principle to human cognitive capacities. We saw in §2 that Russell's structuralism began as a response to the Kantian notion that things-in-themselves are mysterious and unknowable. For Russell's response to work, though, Kantians would need to think that there is a one-to-one correspondence between phenomenal relations and their relata and objective relations and their relata (i.e., things-in-themselves). This aligns with an especially realist interpretation of Kant such as Rae Langton's, in which phenomenal relations in space and time are nothing less than formal representations of the real relations that exist between things-in-themselves (Langton 2001, 214). Since all we know about real relations on that view is derived from their merely formal representation in our cognitive architecture, all we know about them is their structure. Other efforts to pursue structuralism through a Kantian lens (Bitbol 2010; Kauark-Leite and Neves 2016) are less realist, focusing on structures as formal preconditions for experience. The reason such literature cannot simply be appropriated for my purposes is the same as we have seen before: too parochial. One needs to accept a basically Kantian point of view as a starting-point for embracing any of the versions of 'transcendental structuralism' I have mentioned so far. What I want instead is to apply the device I developed in the previous section to transcendental antirealism, showing that the single device can be used regardless of one's starting assumptions.

There won't be any formal developments; the formula presented at the end of §4 is the final version. What will be discussed in this section is the philosophical interpretation of the quantifiers and terms of a sentence using that formula, or of the claim that a mathematical structure represents a real system. This is a big topic for an already-long paper so I will focus strictly on what I consider the two most serious concerns associated with using my strategy for transcendental antirealism: Can a transcendent system cause? Can predicate logic correspond with reality if there are no ultimately real objects or relations?

For the causal question, the issue is that transcendental systems are supposed to operate outside of our conceptual schema, which means that it might be inappropriate to attribute causal relations to them. Since underlies is a causal relation, this difficulty threatens to make my proposal incompatible with transcendent systems. This calls to mind the classic difficulty with Kant's transcendental idealism: if things-in-themselves are unknowable, how could we know that they affect us and cause appearances? I do not have space to discuss this literature here but see (Allison 2004, chapter 3; Langton 1998, chapter 1) and references therein.

I think the best option for me is to sidestep the issue by affirming that we are interested in the structure of transcendent systems *only insofar* as they underlie experiences. I indicated in §1 that EE is one of the initial buy-ins for my strategy. With that in mind, we should think that for any experience, either it is somehow dependent on a transcendent system or it is not. And if not a single experience is somehow dependent on a transcendent system, then I am comfortable with not accommodating such systems for similar reasons to a physicist dismissing a rival theory that differs from their own only by asserting the existence of undetectable particles which affect and are affected by nothing. We are interested in transcendent systems only insofar as experiences are dependent on them per EE.

The second question is harder. Even if causation is dealt with, we can't assume that any of our other concepts apply to a transcendent reality. It is frequently assumed that when a predicate logic sentence is meant to express a fact about the world, the values of its terms should be interpreted according to some metaphysical concept of relations, properties, and objects. But a transcendent system might not be composed of such things; it might not be composed of anything that we can imagine. Can we be structural realists about something that we cannot conceptualise at all? Math-first epistemic structural realism (see §3) accommodates this in principle: mathematics (we might think) can represent the structure of a system whose nature might not be representable by our cognitive endowment. The reason to believe that mathematics could successfully do that is the no miracles argument (see §2.5) which tries to convince us that our best fundamental physics must have gotten something right about the deep structure of the world, even if our metaphysics cannot. But what about the world beyond fundamental physics? We may be led to believe, based on the philosophy of Kant, Bradley (1893), Nāgārjuna (see Siderits and Katsura 2013), or others, that the failure of concepts to grasp objective reality is not limited to reality's "deep structure" but extends to reality as whole, including the parts that we usually navigate with language. If mathematics can represent a system which transcends metaphysical concepts (a big 'if'), could it be true that predicate logic can do so as well?

Perhaps that is a problematic comparison: the math-first realist thinks a mathematical model represents reality as a map does some territory (Wallace 2022, 350); the connection between a sentence and reality is usually thought to be different – involving truth, reference, etc. Nevertheless, sentences *also* connect to reality something like a how a map does. Consider this quote from Quine on Maxwell's causal theory of perception:

> One central plank in Professor Maxwell's platform is that our knowledge of the external world consists in a sharing of structure. This is to my mind an important

truth, or points toward one. Structure, in the sense of the word that is relevant to this important truth, is what we preserve when we code information.

> Send a man into another room and have him come back and report on its contents. He comes back and agitates the air for a while, and in consequence of this agitation we learn about objects in the other room which are very unlike any agitation of the air. Selected traits of objects in that room are coded in traits of this agitation of the air. The manner of the coding, called language, is complicated and far-fetched, but it works; and clearly it is purely structural, at least in the privative sense of depending on no qualitative resemblances between the objects and the agitation. (Quine's contribution to (Maxwell 1968, 161))

Whatever else might be going on with the connection between language and world, I think we should agree with Quine that the pragmatically successful use of language to communicate about the world involves the preservation of structure *at least*. Now let's imagine that there exist no objects or relations in any metaphysical sense that a human could comprehend. We would be unable to conceptualise the real-world values for the variables appearing in a sentence following my formula, and it would be unclear what quantification over the variables might signify outside of formal notation. I suggest that it is yet possible that the full sentence formally encodes a structure that is satisfied by reality in a way we can't comprehend. That is, the purely formal structure of the appropriate Q-Ramsey sentence 'maps' onto reality much in the way that math-first realists think certain mathematical structures do, and the quantifiers present in the full sentence can be interpreted as simply asserting that such a mapping is successful. This at least reduces the predicate logic problem to the mathematical one.

The remaining question is, would a transcendent system be the sort of thing that any formal structure could correspond to? I said before that, when it comes to fundamental physics, it is the no miracles argument that could potentially lead one to an affirmative answer to that question. For general claims about the world, we might use a similar abductive argument: if objective reality is transcendent, the view that formal linguistic and mathematical structures can structurally correspond with transcendent systems is the only view that doesn't make our ability to successfully navigate the world a miracle. One might counter that a transcendent system, being by definition something we cannot understand, may very well be implicated in 'miracles'. To this there is no possible rebuttal.

6. Conclusion

My goal in this paper has been to present a single approach to structuralism which can articulate knowledge we potentially have of objective reality despite the challenges associated with scientific, sceptical, and transcendental antirealism. Such a project comes into contact with a vast number of philosophical issues. It would be impossible to discuss each in-depth but I hope my choice of emphases was appropriate enough to convince the reader that the project is plausible.

The goal is not, of course, for the proposed strategy to be immune to objections. No philosophical view is. The point is rather that none of the identified antirealist traditions is incompatible in principle with my proposed strategy. From naïve realism to the causal theory of perception to brains-in-vats to things-in-themselves – there are no inherent and indefeasible incompatibilities between those notions and claims about objective reality made in the way I have suggested. That, as far as I know, is unique in epistemology.

Acknowledgements

I am grateful to the following people for useful feedback on this paper and earlier versions: Adam Caulton, James Read, Jan Westerhoff, Allison Aitken, David Liggins, and an anonymous reviewer.

References

Ainsworth, Peter M. 2009. "Newman's Objection." *The British Journal for the Philosophy of Science* 60 (1): 135–71. https://doi.org/10.1093/bjps/axn051.

Allison, Henry E. 2004. *Kant's Transcendental Idealism*. Revised and Enlarged edition. New Haven, Conn. ; London: New Haven, Conn; London : Yale University Press.

Brading, Katherine. 2011. "Structuralist Approaches to Physics: Objects, Models and Modality." In *Scientific Structuralism*, edited by Alisa Bokulich and Peter Bokulich, 43–65. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-90-481-9597-8_3.

Brading, Katherine, and Elaine Landry. 2006. "Scientific Structuralism: Presentation and Representation." *Philosophy of Science* 73 (5): 571–81. https://doi.org/10.1086/518327.

Bradley, Francis H. 1893. *Appearance and Reality: A Metaphysical Essay*. London: Swan Sonnenschein & Co.

Carnap, Rudolf. 1950. "Empiricism, Semantics and Ontology." *Revue Internationale de Philosophie* 4 (11): 20–40.

———. 1928/2003. *The Logical Structure of the World and Pseudoproblems in Philosophy*. Open Court edition. Chicago: Open Court.

Chalmers, David J. 2018. "Structuralism as a Response to Skepticism." *The Journal of Philosophy* 115 (12): 625–60. https://doi.org/10.5840/jphil20181151238.

Dalal, Neil. 2021. “Śaṅkara.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2021. Metaphysics Research Lab, Stanford University.

Demopoulos, William, and Michael Friedman. 1985. “Bertrand Russell’s The Analysis of Matter: Its Historical Context and Contemporary Interest.” *Philosophy of Science* 52 (4): 621–39. https://doi.org/10.1086/289281.

French, Steven. 2014. *The Structure of the World: Metaphysics and Representation*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199684847.001.0001.

Friedman, Michael. 2011. “Carnap on Theoretical Terms: Structuralism without Metaphysics.” *Synthese* 180 (2): 249–63.

Holder, Justin P. 2023. “Poincaré’s Radical Ontology.” *HOPOS: The Journal of the International Society for the History of Philosophy of Science* 13 (1): 151–79.

Ketland, Jeffrey. 2004. “Empirical Adequacy and Ramsification.” *The British Journal for the Philosophy of Science* 55 (2): 287–300.

Ladyman, James. 1998. “What Is Structural Realism?” *Studies in History and Philosophy of Science Part A* 29 (3): 409–24. https://doi.org/10.1016/S0039-3681(98)80129-5.

Ladyman, James, and Don Ross. 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press.

Langton, Rae. 1998. *Kantian Humility: Our Ignorance of Things in Themselves*. Oxford : New York: Clarendon Press ; Oxford University Press.

Laudan, Larry. 1981. “A Confutation of Convergent Realism.” *Philosophy of Science* 48 (1): 19–49.

Leplin, Jarrett. 1997. *A Novel Defense of Scientific Realism*. Oxford, New York: Oxford University Press.

Maxwell, Grover. 1968. “Scientific Methodology and the Causal Theory of Perception.” In *Problems in the Philosophy of Science*, edited by Imre Lakatos and Alan Musgrave, 3:148–77. Proceedings of the International Colloquium in the Philosophy of Science, London, 1965. Amsterdam: North-Holland Publishing Company.

———. 1970a. “Structural Realism and the Meaning of Theoretical Terms.” In *Analyses of Theories and Methods of Physics and Psychology*, edited by Michael Radner and Stephen Winokur, IV:181–92. Minnesota Studies in the Philosophy of Science. Minneapolis: University of Minnesota Press.

———. 1970b. “Theories, Perception, and Structural Realism.” In *The Nature and Function of Scientific Theories*, edited by Robert G. Colodny, 4:3–34. University of Pittsburgh Series in the Philosophy of Science. University of Pittsburgh Press.

Melia, Joseph, and Juha Saatsi. 2006. “Ramseyfication and Theoretical Content.” *The British Journal for the Philosophy of Science* 57 (3): 561–85.

Newman, M. H. A. 1928. “Mr. Russell’s ‘Causal Theory of Perception.’” *Mind* 37 (146): 137–48.

Russell, Bertrand. 1919/1993. *Introduction to Mathematical Philosophy*. 2nd ed. New York: Dover Publications.

Schlick, Moritz. 1918/1974. *General Theory of Knowledge*. Translated by Albert E. Blumberg. Library of Exact Philosophy. Wien: Springer-Verlag.

Siderits, Mark, and Shōryū Katsura. 2013. *Nāgārjuna’s Middle Way: The Mūlamadhyamakākarikā*. Classics of Indian Buddhism. Boston: Wisdom Publications.

Smithson, Robert. 2017. "Newman's Objection and the No Miracles Argument." *Erkenntnis* 82 (5): 993–1014. https://doi.org/10.1007/s10670-016-9855-1.

Wallace, David. 2022. "Stating Structural Realism: Mathematics-First Approaches to Physics and Metaphysics." *Philosophical Perspectives* 36 (1): 345–78. https://doi.org/10.1111/phpe.12172.

Worrall, John. 1989. "Structural Realism: The Best of Both Worlds?" *Dialectica* 43 (1/2): 99–124.

Zahar, Elie, and John Worrall. 2001. "Appendix 4." In *Poincaré's Philosophy: From Conventionalism to Phenomenology*. Chicago: Open Court.