Against Alien Structure

Abstract. Is the metaphysical structure of the world familiar or alien? That is, can the structure of the world be presented using familiar linguistic resources, or would that require radically new, alien resources? Eklund (2024) has recently argued for the epistemic possibility of alien structure. My aim in this paper is to argue against that possibility.

Keywords: alien structure; alien language; perspicuity; higher-order metaphysics.

Is the structure of the world *familiar* or *alien*? That is, can the structure of the world be presented using familiar linguistic resources, or would that require radically new, alien resources?

Take the sentence 'Socrates is wise'. This sentence presents a certain kind of metaphysical structure:[1] there is an object (i.e. Socrates) and it is a certain way (i.e. wise). 'Plato loves Socrates' presents a different metaphysical structure: there are two objects (i.e. Plato and Socrates) and they are related to each other in a certain way (i.e. the former loves the latter). These sentences are perfectly familiar, and so the structures that they present count as familiar too. But mightn't there be other, alien structures that cannot be presented in familiar terms? I obviously cannot give any intuitive examples of alien structures, but it is easy to see why you might suspect that they exist. Familiar languages evolved to suit the parochial needs of their speakers, operating in the limited environments that they found themselves in. It is at least doubtful that such languages would be capable of presenting absolutely every structure to be found in the world.

Eklund (2024) has recently developed this line of thought into a book-length defence of alien structure. Or, at very least, he defends the *epistemic possibility* of alien structure. It is not really his aim to show that there actually are alien structures, but he does want to establish that this is a live epistemic possibility.[2] In that sense, then, Eklund's aim is somewhat modest. My aim, by contrast, is quite immodest: I want to argue that it is not at all a live possibility that any structure is alien.

[1] In what follows, I will always mean *metaphysical structure* by 'structure', unless I clearly indicate otherwise.

[2] Eklund is also inclined to think that alien structures are *metaphysically* possible. However, as he emphasises, 'a main aim [of his book] is just to address the coherence of alien structure' (Eklund 2024: 12), and that is the aim that I will focus on here.

1 Familiar versus alien languages

In this section, I will make some preliminary comments about the distinction between familiar and alien languages. Here is how Eklund draws this distinction:

> alien languages as I will speak of them are semantically different from familiar languages in a broadly structural way: they are different with respect to the semantic categories of their expressions or with respect to modes of composition. (Contrast: languages differing from familiar ones only in what they have expressions for would not be structurally different.) An alien language does not merely have, for example, words for things we don't have words for. (Eklund 2024: 5)

As this quotation makes plain, Eklund places a lot of weight on the idea of semantic categories, or, in my preferred terminology, *semantic types*. Intuitively, two expressions belong to the same type iff they play the same kind of semantic role. It is not easy to make this intuitive thought precise, but we can at least illustrate it with some good examples. All names belong to the same type, because they all play the same kind of semantic role: they all purport to refer to objects. Of course, different names might purport to refer to different objects, and so in that sense they might have different semantic roles. But they still have the same *kind* of semantic role. Sentences, on the other hand, play an entirely different kind of semantic role from names: names purport to refer to objects, and sentences say that thing are thus-and-so.[3] So names and sentences belong to different types. Following one standard convention, I will use *e* for the type of names, and *t* for the type of sentences.

A language is alien iff it includes *alien types* of term or *alien modes of composition*, i.e. types of term or modes of composition that do not appear in any familiar languages. But when does a language count as familiar? Eklund is understandably reluctant to offer anything like necessary and sufficient conditions for familiarity. However, in this paper, I will operate with the following sufficient condition:

Basic Familiarity: Any language you understand is familiar to you.

Eklund adopts *Basic Familiarity* as a working assumption, but his official position is that it is not obviously true. In particular, he takes seriously the thought that English might not be familiar to English-speakers (Eklund 2024: 6). As he notes, it is a real possibility that, at some level of linguistic analysis, English might turn out not to function at all as English-speakers tend to assume. If so, Eklund thinks it might be reasonable to describe English as unfamiliar.

I want to take *Basic Familiarity* more seriously than Eklund. There is undoubtedly some sense in which English should count as 'unfamiliar' in the scenario Eklund describes. But we are here co-opting the word '(un)familiar' as a term of art, and it is up to us to restrict its use in ways we think will be theoretically fruitful. And I do think

[3] I think that names and sentences provide the clearest example of a difference in type that we could hope for, but even this example can be challenged. Frege (1893: §2), for example, infamously treated sentences as a variety of name.

that adopting *Basic Familiarity* is fruitful. The idea of an 'alien metaphysical structure' is philosophically enticing in large part because it suggests that the structure is inaccessible to us, at least for now; perhaps we might later enrich our expressive powers in a way that will allow us to present that structure, but, as things are, that is simply beyond us. However, even if, deep down, English behaves in unexpected ways, English-speakers can still *use* it to present whatever structures it lets them present. So I will adopt *Basic Familiarity* as a stipulation to help home in on the sense of 'familiar' I think is most interesting.

I will also rely on another assumption about familiarity:

Functional Familiarity: If α and β are familiar types, then $\alpha \to \beta$ is also a familiar type.

When α and β are both types, $\alpha \to \beta$ is a *functional type*. In particular, $\alpha \to \beta$ is the type of terms that can be applied to a term of type α to produce a term of type β. Exactly how one term is applied to another varies from language to language, but in formal languages it is usually just a matter of concatenation: $(X^{\alpha \to \beta} Y^{\alpha})^{\beta}$.[4]

Here is a quick argument for *Functional Familiarity*. Suppose that α and β are both familiar. In that case, terms of type $\alpha \to \beta$ merely map us from one familiar type to another; we can think of these terms as expressing functions from the values of type-α terms to the values of type-β terms. It is hard to see how a language could ever count as alien merely by including terms like that. Terms of type $\alpha \to \beta$ do not provide us with a radically new way of representing the world. They just shuffle us around types of representation we were already familiar with.

You might not yet be convinced by that quick argument. Imagine a community who speak a first-order language. The people in this community will be familiar with a variety of types of term (sentences, names, predicates, quantifiers, and connectives), but they might be ignorant of the general idea of functional types. Of course, we can represent their predicates, quantifiers and connectives with functional types. For example, monadic predicates are type $e \to t$. However, although the people in this community are familiar with terms that combine with *names* to produce *sentences*, it might not have occurred to them that you can form type $\alpha \to \beta$ for *arbitrary* types α and β. For example, they might never have dreamed of forming terms of type $((e \to t) \to t) \to t$, i.e. the type of second-order quantifiers.

This objection is absolutely right, so far as it goes. The general idea of functional types was hard won, and so we cannot assume that *Functional Familiarity* holds for all communities. However, *we* are familiar with functional types — and if you happened

[4] Eklund (2024: 129) briefly considers, but does not endorse, the suggestion that some alien types are so alien that we cannot form functional types from them. However, his case in favour of this suggestion strikes me as weak. He imagines an alien language with the basic types f and t, rather than e and t. He notes that, if we are allowed to form functional types, we will then get the type $f \to t$. He then worries that, since this functional type is structurally identical to the familiar type $e \to t$, the difference between f and e is a difference in name only. However, the mere fact that we can form a type which maps f to t and a type that maps e to t does not even *hint* that type-f terms play the same kind of semantic role as type-e terms.

not to be, you are now — and so *Functional Familiarity* surely holds *for us*: if α and β are familiar *to us*, then $\alpha \longrightarrow \beta$ is also familiar *to us*.[5] In an effort to keep things readable, I will tend not to explicitly relativise familiarity to us; but, when context does not clearly indicate otherwise, that relativisation should be taken as tacit.

That will do for the preliminaries. In the next section, we will shift our focus from alien *languages* to alien *metaphysical structures*, and I will begin making my argument that there are no alien structures. Before we get going, though, I want to be clear now that I do *not* reject the (epistemic) possibility of alien languages. I agree with Eklund (2024: 74) that it would 'be a miracle if all possible languages were found among what we find familiar'. However, I do not think that goes any way toward showing that any *metaphysical structures* might be alien.

2 The naïve approach

A *metaphysical structure* is alien iff it can be presented only by an alien language. But what it is it for a language to *present* a structure? Eklund actually offers two ways of approaching the notion of presentation. Both approaches are introduced right at the start of the book (Eklund 2024: 14–7), and they are run in parallel all the way to the end. According to the first approach, presentation is just representation; according to the second, presentation is *perspicuous* representation. For now, we will focus on the former, which we can think of as the *naïve approach*; we will turn to the latter *sophisticated approach* in §§4–6.

Eklund illustrates the naïve approach with this example:

> one might suggest that the object/property distinction is presented by familiar languages in the alethic sense that there are sentences of familiar languages that cannot be true unless there are both objects and properties. (Eklund 2024: 14)

This example is not *quite* what we need, because we are after an explanation of what it is to present a *structure*, and the example is about presenting a *distinction*. However, I think that what Eklund has in mind is clear enough. 'Socrates is wise' cannot be true unless some object (i.e. Socrates) is some way (i.e. wise).[6] So that sentence cannot be true unless the *some-object-is-some-way* structure is exemplified in the world. That is the sense in which 'Socrates is wise' represents, and thus on the naïve approach *presents*, that structure.

[5] I am here assuming that we understand functional type theories, even when they are not readily translatable into a natural language like English. This assumption is common among the new wave of higher-order metaphysicians (see e.g. Fritz and Jones 2024), but I should warn that it is also rejected by some dissenters (e.g. Hofweber 2022). However, it is noteworthy that Eklund (2024: 14) is happy to grant that we are familiar with 'standard' higher-order logic, by which he presumably means functional type theory.

[6] We are obviously meant to hold the meaning of 'Socrates is wise' fixed when evaluating this modal claim.

So, according to the naïve approach, a metaphysical structure is *familiar* if the truth of some familiar sentence requires that it be exemplified; otherwise, the structure is *alien*.[7] But this approach immediately runs into trouble.[8] Imagine that you are confronted by an alien sentence. The terms that build up that sentence might be alien, but the sentence as a *whole* is a familiar type of term. (This follows from *Basic Familiarity*, along with the fact that you understand languages that include sentences, i.e. terms of type t.) In principle, then, you could expand your language by adding a new atomic sentence letter, e.g. P, on the stipulation that P is to represent the very same state of affairs as the alien sentence;[9] this expanded language would still count as familiar, since it would still deploy only familiar types of term and modes of composition. And, given the naïve approach, P would present whatever structures were presented by the alien sentence. So those structures are familiar after all. Moreover, we could pull the same trick with any alien sentence, and so no structure is alien at all.[10]

Following Eklund (2024: 237), we can call this the *piggyback argument* against alien structure.[11] You might be tempted to dismiss the piggyback argument out of hand. It is obviously unsatisfying to suggest that we could present a metaphysical *structure* with an *unstructured* sentence letter! However, it is important to remember that, for now, we are working with the *naïve* idea that presentation is representation, not the *sophisticated* idea that presentation is perspicuous representation. It might not be perspicuous, but when a sentence letter is stipulated to represent the same state as

[7] Strictly speaking, this definition classifies *ineffable structures*, i.e. structures that could not be presented by any language, as 'alien' (Eklund 2024: 17 & 234–41). However, I am sceptical about the very idea of ineffable structures: what could it be about a structure that would make it resistant to being presented by *any* language, even an alien one? At very least, ineffable structures would be their own special category, with their own special problems, and so I will set them aside for their own special discussion.

[8] Eklund (2024: 14–5 & 199-200) raises a different worry about his initial gloss on the naïve approach: imagine a language that included some *false* sentences that presented a structure in the alethic sense explained above, but no *true* ones; Eklund is inclined to deny that such a language would really present that structure in the sense that matters. Now, personally, I disagree with Eklund here. I would rather say that the language did present the structure, even though it failed to provide any true examples of it. However, it is easy enough to accomodate Eklund's reaction: we just need to say that a language presents a structure only when the truth of some *true* sentence requires that the structure be exemplified.

[9] Throughout the paper, I will describe sentences as *representing states of affairs*(i.e. *ways for things to be*). Nothing hangs on this terminology. If you prefer, you can say that sentences *express propositions*. (In my own view, there is literally no difference between these two ways of putting things; see Trueman 2021: ch. 14, 2022.)

[10] To be clear, I am concerned with what is *in principle* possible here. There may be practical obstacles to your piggybacking on alien sentences. Perhaps you could not piggyback if you were never confronted with any alien sentences, or if you could not recognise an alien sentence when you saw one. (Compare Eklund 2024: 239.) But, for any alien sentence, there could *in principle* be a familiar language that included an atomic sentence letter which represented the same state as that alien sentence, even if it would not be a language we are actually in a position to speak.

[11] Eklund does not present the piggyback argument as an objection to the naïve approach to alien structure, but as an objection to the idea that some states of affairs are ineffable.

an alien sentence, the naïve approach will imply that it presents whatever structures the alien sentence presents.[12]

As far as I can tell, there is only one good way of trying to block the piggyback argument, without abandoning the naïve approach. The argument relies on the assumption that alien 'sentences' are sentences *in the familiar sense* (i.e. terms of type t); that is what guarantees that we can piggyback on alien sentences with familiar unstructured sentence letters. The argument would stop in its tracks if we rejected that assumption. Of course, alien 'sentences' must still be capable of 'representing' structures in some sense, and so there must still be a sense in which they can be 'true' or 'false'. But, according to this reply, these are not the same senses in which familiar sentences represent structures, are true, or are false.

Now, it is undeniably hard to provide any concrete details about how these senses might differ, but it is hard to give *any* concrete details about the inner workings of *any* alien language. Moreover, Eklund (2024: ch. 10) offers various strategies for easing ourselves into the idea that there might be alien languages that 'represent' the world without sentences. So, for the sake of argument, I am happy to grant that this is a viable reply to the piggyback argument.[13]

However, this reply works only if the semantic values of alien 'sentences' are not admissible values for familiar sentences; otherwise, the piggyback argument could resume as before. More generally, they cannot be admissible values for *any* familiar type of term. To see why, suppose that some familiar term, a^{α}, can take the same value as some alien 'sentence'. If a is not a sentence (i.e. $\alpha \neq t$), then we cannot use a to make a statement about how things are — we do that with sentences in familiar languages — and so a cannot represent anything in the alethic sense. Nonetheless, we could introduce a new predicate, $T^{\alpha \to t}$, by stipulating that it is to express whatever sense of 'truth' is applicable to the value of the alien sentence we started with. (We might not be able to say much about what that sense amounts to, but that does not stop us from picking it out by this description. Indeed, we cannot understand the reply to the piggyback argument unless we understand that description!) Since α and t are both familiar types, *Functional Familiarity* implies that $\alpha \to t$ is familiar too. So Ta is a familiar sentence, which predicates 'truth' of the value of a. Importantly, this is not to say that Ta has the same value as the original alien 'sentence'; after all, Ta is a familiar sentence, not an alien 'sentence'. However, Ta surely still presents the same structures as the alien 'sentence', in the alethic sense: the truth condition for Ta is precisely the 'truth' condition for the alien 'sentence'.

[12] Eklund (2024: 238) makes this very point in his discussion of ineffable states.

[13] Richardson (2026: §3) also presents a version of the piggyback argument against the naïve approach. However, he does not consider the reply that I have just laid out. More generally, Richardson does not seem to take the idea that alien languages contain alien types of term and/or modes of composition as seriously as I think we should. Nonetheless, I should emphasise that I absolutely agree with the spirit of Richardson's arguments, even if I disagree with some of the details. (Although I do think that Richardson slightly misrepresents his conclusion: he sometimes claims to have shown that there can be no alien *languages* (e.g. pp. 277 & 286), but his conclusion actually concerns alien *structures*.) For discussion of Richardson's argument against the sophisticated approach, see fn. 36.

As I mentioned above, I think that distinguishing between familiar and alien 'sentences' is the *only* way to halt the piggyback argument. So, I take the naïve approach to entail the following principle:

Alien Values: There is an alien structure only if there is something that cannot be the value of any familiar type of term.

However, as I will argue in the next section, *Alien Values also* implies that there are no alien structures. So, in the end, distinguishing between kinds of 'sentence' is only a way of delaying the conclusion of the piggyback argument, not of avoiding it altogether.

Before I actually present that argument, though, I should say a little more in defence of my claim that the naïve approach entails *Alien Values*, since that principle looks dangerously like something that Eklund explicitly rejects:

> When discussing the alien, I should also distinguish between ontology and ideology. Often when I speak of structure, both alien and familiar, I will speak of ontological categories, and the categories will be discussed as categories of entities. This talk may encourage thinking of alien structure simply as a matter of the existence of alien kinds of entities. But how representations represent the world is not only about what entities these represent the world as containing, but more generally how they represent the world as being [...]
>
> Taking the distinction between the what and the how into account, we should be open to the idea that a language might be alien not just in terms of what entities it presents the world as containing but also in terms of how it presents the world as being. Let us say that *ontological* alienness is alienness when it comes to what kinds of entities there are, and *ideological* alienness is alienness when it comes to the kind of way the world is presented as being that goes beyond differences in what kinds of entities there are. In principle, at least, each kind of alienness can obtain without the other. (Eklund 2024: 11-2)

On the face of it, *Alien Values* might appear to be the claim that there cannot be alien structure unless there is alien ontology. However, that appearance would be misleading. I intend the idea of the *value* of a term to be taken in a very broad sense. Roughly, the value of a term is whatever that term is used to talk about.[14] But I do not want to assume that every value of every term is an object (i.e. the sort of thing that can be referred to by a type-e name), or is in any way object-like. As I mean it, the *value* of a constant is whatever that constant refers to/expresses/represents/etc., in whatever sense is suitable for constants of that type; the admissible values of an (unrestricted) variable of type-α are all the things that can be values of constants of type-α. So, for example, even if you are a nominalist who thinks that the role of a type-$e \rightarrow t$ predicate is not to *refer* to an object but to *express a condition* on objects, you still

[14] This use of 'value' is standard, but it is not the *only* standard use. Sometimes, the 'semantic value' of a term is just an entity which somehow encodes the semantic role of that term, regardless of whether that entity is what the term is used to talk about. In *this* sense, the semantic value of a predicate, for example, might well be a set, even if we do not talk about sets when we use predicates. There is obviously nothing wrong with this way of using 'semantic value', but it just isn't the way I am using it here.

should not deny that predicates 'have values'; you should just say that the value of a given predicate is the particular condition it expresses, and second-order quantifiers quantify over conditions-on-objects in general. Of course, you might worry that this encompassing way of thinking about the values of terms is incoherent, but that is not something that worries me (Trueman 2013, 2021: ch. 7).[15] More importantly, it is not something that worries Eklund (2024: 128–9), either.

So, if ontological commitments are commitments to objects, or even to things that are sufficiently object-like, then *Alien Values* does not imply that there can be no alien structure without alien ontology. Of course, you might prefer to use 'ontology' just as broadly as we are using 'value', so that ontological commitments are commitments to any type of value whatsoever. However, it does not strike me as at all objectionable to say that alien structure requires alien 'ontology' in this broad sense. At very least, it does not seem to be what Eklund was objecting to.

3 Against the naïve approach

Let's grant, then, that the naïve approach to alien structures implies *Alien Values*. Here is a (semi-)formalisation of the consequent of *Alien Values* (where you can choose whatever type you prefer for α):[16]

(1) $\exists x^{\alpha}$(x is not a value of any familiar type of variable).

We can now offer another simple argument against the epistemic possibility of alien structures. Since we understand (1), *Basic Familiarity* implies that type-α terms are familiar. And in that case, every value of x is obviously a value of a variable of familiar type (namely, x). So, (1) is false, and, given *Alien Values*, it follows that there are no alien structures.[17]

This argument against alien structure assumes that we understand (1), and in principle you could try to resist it by denying that assumption. But that would just make

[15] This way of thinking about the values of terms is sometimes called *neutralism*. For discussions of neutralism, see: Prior 1971: ch. 3; Rayo and Yablo 2001; MacBride 2006: 445; Wright 2007. Neutralism is an important assumption behind much of the recent work in higher-order metaphysics. For a small sample of that work, see: Williamson 2003, 2013; Krämer 2014; Trueman 2015, 2021, forthcoming(b); Dorr 2016, 2025; Jones 2018; Skiba 2021; Fritz 2023; Bacon 2024; Fritz and Jones 2024.

[16] There are some subtle questions about how to *fully* formalise (1). The questions are difficult because we need to establish whether different types of term can 'have a value' in the same sense as each other. (My own view is that they cannot, see Trueman 2021: chs 4 & 5.) Fortunately, though, we do not need to answer those questions here. However you think they should be answered, the following line of thought will carry over.

[17] You might wonder whether Eklund's (2024: ch.3) discussion of different approaches to ontological categories might help here. (Thanks to an anonymous reviewer for raising this possibility.) It is important to be clear that it cannot. Eklund's question at this stage of the book is: how should we divide entities we can quantify over into categories like *object*, *property*, etc. However exactly we answer that taxonomic question — which Eklund (2024: 54) himself describes as 'bookkeeping' — the point remains that *no* entity of *any* category could possibly witness generalisation (1).

things worse. (1) is a semi-formalisation of the consequent of *Alien Values*, which we derived by intelligible means from the naïve approach to alien structure. If we did not understand (1), then that could only be because we did not understand the naïve approach in the first place. So, granting that we do in fact understand it, we can conclude that the naïve approach implies that there are no alien structures.

Eklund himself is aware of this problem. Here is how he expresses it:[18]

> It is arguably implausible that there are expressions of some other language that refer to some entities that cannot be referred to at all in any familiar language: note the awkwardness in the above paragraph of using 'it' for something that allegedly cannot be referred to in familiar language. (Eklund 2024: 21)

There is, though, something peculiar with how Eklund puts things. He takes it to be merely *implausible* to say that there are things that cannot be the values of familiar terms, rather than disastrously self-defeating.

My best guess is that Eklund is being cautious because this style of argument is well known from the literature on *absolute generality*, and that literature also provides a well known response. Here is how the argument goes in the generality debate. *Generality restrictivists* claim that no (first-order) quantifier is *absolutely* general; every quantifier must be interpreted as quantifying over some restricted domain (i.e. a domain that does not include absolutely every object). That applies to the restrictivist's own quantifiers, and so they are committed to the claim 'There is something that I am not quantifying over', which we can semi-formalise as:

(2) $\exists x^e$(I am not quantifying over x).

But they cannot truly assert this sentence: when they assert it, every value of x is trivially something that they are quantifying over.[19]

To my mind, the most promising response to this problem is to go modal. Restrictivism should not be articulated in non-modal terms, as the thesis that every domain *is* restricted. Rather, restrictivism is inherently modal: *necessarily*, for every domain, there *could be* an object not included in that domain. That way, restrictivists would not be committed to (2), but to (2′):

(2′) $\Diamond\exists x^e$(I am not actually quantifying over x).

Importantly, the $\Diamond$ in (2′) does not express a standard *circumstantial* modality: (2′) is not a claim about how things would have been in alternative possible circumstances. Rather, the modality concerns the *interpretation* of $\exists x^e$. We *actually* interpret $\exists x^e$ as quantifying over some particular domain, but we *could* interpret it as quantifying

[18] We should note, though, that in this quotation, Eklund is explicitly concerned only with the suggestion that there are things that no familiar term can *refer to*, not with the more general suggestion that there are things that no familiar term can *take as its value*. I will skate over this in the main text, since Eklund's point obviously generalises.

[19] This objection was forcefully developed by Williamson (2003: §5).

over some other domain. (2′) is true iff we could interpret $\exists x^e$ as quantifying over something that it does not quantify over on its actual interpretation.[20]

I do not want to take a stand on whether this reply is ultimately enough to save generality restrictivism.[21] What matters here is that it certainly cannot be repurposed as a response to the argument I just presented against alien structure. To see why, imagine that we tried to modalise (1) in a way analogous to (2′):

(1′) $\Diamond\exists x^\alpha$(x is not a value of any actually familiar variable).

How exactly are we meant to understand the $\Diamond$ in (1′)? Presumably, $\Diamond\phi$ should be true iff there is some way of expanding the range of types available in our language so that ϕ comes out true. That way, even if x^α is unfamiliar, (1′) could still be true. However, the whole idea of a modal operator like that is confused. If $\Diamond$ is a genuine modal operator, then ϕ must be *used*, not merely *mentioned*, in $\Diamond\phi$. The variable x^α is therefore used in (1′), and so, given that we understand (1′), *Basic Familiarity* implies that α is a familiar type. Thus, (1′) must be false, just like (1).

You might worry that this objection leans too heavily on the idea that the $\Diamond$ in (1′) is a *modal* operator. Maybe it's better seen as a *metalinguistic* operator, so that ϕ is mentioned, not used, in $\Diamond\phi$. Presumably, $\Diamond\phi$ would then say that the sentence ϕ is true in some language that expands on the range of types available in our language.[22] We would then be able to understand (1′) even if type-α were unfamiliar. However, the point remains that, if x^α were an unfamiliar type of variable, then we still could not understand the sentence mentioned after the $\Diamond$, i.e. '$\exists x^\alpha$(x is not a value of any actually familiar variable)'; after all, x^α is still being used in *that* sentence. But in that case, we could not read (1′) as a way of saying anything about the values of any alien terms; it would merely say that a certain string of symbols is true in some language that we did not speak. (It would be a bit like telling a monolingual English speaker that 'La neige est blanche' is true in French; that English speaker would not be in a position to hear you as telling them about the colour of snow.)[23]

Of course, it may well be true that we could learn to speak a language that includes types that are currently alien to us. If we did, then we could use the new types of variables to say that there are things that are not values of any of our old variables.

[20] For some different implementations of this modal reply, see: Fine 2006; Linnebo 2010, 2018: ch. 3; Studd 2019: ch. 6; Berry 2022.

[21] For the record, I suspect that it is not, for the reasons given by Button (forthcoming).

[22] If the $\Diamond$ in (1′) is now being understood as a metalinguistic operator, then it would be natural to understand the 'actually' in the same way. However, it would then be unclear how to make sense of the attempt to quantify into its scope. This strikes me as a potentially serious objection to the metalinguistic proposal, but the one I am about to present in the main text seems more central.

[23] What if someone who spoke an alien language as well as English assured us that (1′) says something about the values of alien terms? (Thanks to Chris Jay for asking this question.) The problem is our understanding what this person could possibly mean. In particular, we will not be able to understand them unless their phrase 'the values of alien terms' is familiar to us. So, if we understand them, then we know that they are not quantifying over anything that cannot be quantified over in familiar terms.

But, by (the necessitation of) *Basic Familiarity*,[24] those things would no longer be alien to us: we could present metaphysical structures involving them just by using the new variables at our disposal.

Here is one last attempt to get around the problem I have raised for the naïve approach. What if, instead of (1) or (1′), we tried this?[25]

(1*) $\exists\alpha\exists x^{\alpha}$($x$ is not a value of any familiar type of variable).

The answer depends on how we understand the quantifier $\exists\alpha$. We could read it *substitutionally*, so that (1*) is true iff some substitution instance of (1) is true. However, if the substitution instances were all drawn from our own language, then (1*) would still be false, for exactly the same reason as (1). And if the substitution instances were drawn from some enriched language that includes more types than our own, then (1*) would be in exactly the same boat as (1′). (Whether it would be problematic in the same way as the modal reading or the metalinguistic reading of (1′) would depend on your preferred approach to substitutional quantification: if you took (1*) to be, in effect, the infinite disjunction of the relevant substitution instances, then (1*) would be analogous to the modal reading; if you took (1*) to be the explicit statement that at least one of those substitution instances is true, then it would be analogous to the metalinguistic reading.)

Alternatively, we might try to read $\exists\alpha$ as a genuine, non-substitutional, quantifier that ranges over *types*. Now, my own view is that this kind of quantification is incoherent, when types are understood in the way that I described in §1. However, I do not want to lean too heavily on that conviction here, and it should be noted that quantification over types is permitted in some type theories (for examples, see Crosilla 2024; Wilhelm forthcoming). More importantly for present purposes, if you understand a language that permits quantification over types, then variables belonging to any of those types should count as familiar to you. After all, that language allows you to quantify over the values of those types of variable, as demonstrated by (1*). So (1*) would *still* be false.

I am not aware of any other replies that show any potential as ways of saving the naïve approach. As a result, I do not think that we need to share in Eklund's caution: if we understand the naïve approach at all, then it surely implies that there are no alien structures.

[24] It seemed needlessly distracting to necessitate *Basic Familiarity* when I first presented it in §1; however, I take it that the motivation I offered for adopting *Basic Familiarity* also extends to its necessitation.

[25] Thanks to Daniel Berntson for suggesting that I discuss this option.

4 The sophisticated approach

We turn now to the *sophisticated approach* to alien structure. Consider the following pairs of sentences:[26]

- The number of dinosaurs is zero / There are no dinosaurs.
- Susan is a sibling / Susan shares a parent with someone else.
- Socrates' death took place / Socrates died.
- Susan instantiates the property of running / Susan runs.
- The average adult has 2.4 children / The number of children divided by the number of adults equals 2.4.

One reasonable view is that both sentences in each pair represent the same state of affairs, but only the second sentence represents it *perspicuously*: in some sense, the second sentence offers greater insight into the state than the first. According to the sophisticated approach, there is alien structure iff some state cannot be perspicuously represented by any familiar sentence: no familiar sentence offers any genuine insight into that state.

Now, the examples given above are undoubtedly suggestive, but if we really want to develop the sophisticated approach, then we will need to offer a more systematic theory of perspicuity. Eklund (2024: §1.4) takes a pluralistic approach here. He considers a number of accounts of perspicuity, but does not settle on any one as his official choice. However, he is at least attracted to the following account, which treats perspicuity as a matter of form or structure:

> A sentence perspicuously represents reality when the logical form of the sentence corresponds in the right sort of way to the aspect of reality that it represents […]
>
> Talk of capturing in the right sort of way is not precise. What does it mean? What sorts of things settle whether it is "Socrates is wise" or "Socrates instantiates wisdom" that corresponds in the right sort of way to what both these sentences represent? If one takes there to be such entities as facts, one may think that it is the number of constituents of the fact that matter. If the relevant fact only has constituents corresponding to "Socrates" and "is wise" (and where the "is" of predication is understood not as standing for a relation but as having a different function), the former sentence perspicuously represents the fact. If it has constituents corresponding to all of "Socrates", "instantiates", and "wisdom", it is the latter sentence that perspicuously represents the fact. (Eklund 2024: 16)

Something like this account of perspicuity has been offered by many philosophers.[27] One standard way of filling out the details is to think of sentences as representing

[26] Eklund (2024: 15) credits Rayo (2013: 3) for the first four examples; the final pair appears often in discussions of perspicuity.

[27] Eklund cites Hawthorne and Cortens (1995: 154–6) and J. S. Russell (2018: 566–8); however, while they sympathetically discuss this account of perspicuity, none of these authors quite accepts it outright. For some more sympathetic discussions of this account, see: B. Russell 1918; Fine 2001: 25; Rubenstein 2024: 1112–3, forthcoming: 6; Jones forthcoming: §4.2.

structured states of affairs: every state has a unique structure, and a sentence is perspicuous iff the syntactically simple constituents match up with the constituents in the corresponding state. However, the idea that states have unique structures faces well known formal difficulties,[28] and so it is important to emphasise that this account of perspicuity can be developed without it. As Jones (forthcoming: §4.2) has recently explained, it is enough to assume that every state of affairs has a *metaphysically privileged* decomposition: even if states can be decomposed in all sorts of *ad hoc* and gruesome ways, it might still be that a select few (or even *just one*) of the decompositions are somehow chosen by the world, perhaps because they best capture the 'underlying nature' of the state (Jones forthcoming: 20).[29]

With Jones's point in mind, we can offer the following:[30]

Structural Perspicuity: A sentence is perspicuous iff:

(i) it is transparent (i.e. the state represented by the sentence is a function of the values of its syntactically simple constituents); and
(ii) there is a privileged decomposition of the state represented by the sentence such that
  (a) the value of each syntactically simple constituent of the sentence is a constituent of the state on that privileged decomposition, and
  (b) each constituent of the state on that privileged decomposition is the value of some syntactically simple constituent of the sentence.

Clause (i) is there to ensure that the sentence genuinely presents the values of its simple constituents as being arranged in a metaphysical structure:[31] if a sentence is non-transparent, then what it says is not determined purely by composing the values of its simple constituents; it is determined, at least in part, by how we *represent* those values.[32] Clause (ii) then captures the idea that the values of the simple constituents

[28] It is obviously inconsistent with the standard higher-order principle β-Conversion; less obviously, the Russell-Myhill paradox also reveals that it is inconsistent with a weaker principle, Extensional β-Equivalence. See Trueman forthcoming(b) for a philosophical justification of β-Conversion, and Fritz 2023: §3.2 for an exceptionally clear formulation of the Russell-Myhill paradox.

[29] Jones assumes that every state has a unique maximally privileged decomposition, but I do not need to make that assumption here. (I am sceptical about the idea that states have *any* privileged decompositions, but I will set that scepticism aside for now.)

[30] *Structural Perspicuity* defines a *non-scalar* notion of perspicuity: a sentence is either perspicuous or not; one sentence cannot be *more perspicuous* than another. It might be helpful, then, to think of it as defining *maximal perspicuity*. We should note, though, that structurally different sentences could still represent the same state of affairs maximally perspicuously, since that state could have multiple privileged decompositions.

[31] Importantly, assuming that a sentence is transparent is not the same as assuming that the sentence itself deploys functional application as a mode of composition, in the sense of applying the value of one term to the value of another. As Eklund (2024: 63–4) emphasises, we have no reason to think that all (complex) alien sentences are composed by functional application. (Thanks to an anonymous reviewer for urging me to be clear on this point.)

[32] You might worry that, if a sentence cannot present a structure unless it was formed by a transparent mode of composition, then generalisations like $\forall x Rax$ cannot present structures, since there

of a perspicuous sentence are exactly the constituents of the corresponding state of affairs, on some privileged decomposition.

Before moving on, though, I should also note that this account of perspicuity could easily be generalised. Following Eklund, I have taken *sentences* to be the sorts of things that can be structurally perspicuous. However, *Structural Perspicuity* could be generalised to cover any type of term whatsoever. This is worth noting because, as I mentioned in §2, Eklund (2024: ch. 10) is open to the possibility of alien languages that somehow present metaphysical structures, but not via sentences in the familiar sense. Still, to keep things from getting even more complex than they already have to be, I will stick with sentences for now, and leave the generalisation to any interested reader.[33]

5 Against the sophisticated approach

My aim in this section is to argue that, if we adopt *Structural Perspicuity*, then the sophisticated approach to alien structure is no better than the naïve approach: it also entails *Alien Values*, and thus also implies that there are no alien structures.

To show this, let $\dagger(a_1^\alpha, b_1^\beta)$ be a perspicuous sentence in an alien language; a_1 and b_1 are the only simple constituents in the sentence, and $\dagger$ is some mode of composition available in the alien language.[34] Also suppose that the value of a_1^α can be the value of some familiar type of term, $a_2^{\alpha'}$, and the value of b_1^β can be the value of some familiar type of term, $b_2^{\beta'}$.[35]

is a notorious result that the value of $\forall x Rax$ is not a function of the values of its simple constituents (in the sense of 'value' employed here). However, that result depends on a particular definition of the *constituents* of $\forall x Rax$. For an alternative definition that avoids this result, see Trueman forthcoming(a).

Similarly, you might worry that this transparency requirement will prevent grounding-claims (e.g. 'P grounds $P \vee Q$') from presenting structures. However, while it is standard to assume that grounding is hyper-intensional, it is much less common to assume that it is non-transparent. (We can combine transparency with hyper-intensionality so long as we are willing to individuate states of affairs more finely than by necessary equivalence. See Rosen 2010: 124–6.) If grounding-claims turned out to be non-transparent, I do not think it would be at all unreasonable to deny that they present metaphysical structures in the sense at play here.

[33] Here is one detail worth highlighting. Suppose that $\star(a, b)$ is a perspicuous term of some familiar type $\alpha \neq t$, and that it has the same value as some alien 'sentence'. Since $\star(a, b)$ is not a sentence, we will need to introduce a new term, $T^{\alpha \to t}$, to predicate the appropriate notion of 'truth' to its value, just as we did in §2. Crucially, though, there is no reason to assume that the sentence $T(\star(a, b))$ will also perspicuously represent the value of the original alien 'sentence'. Nonetheless, familiar language still has the means to perspicuously pick out the value of the alien 'sentence' with $\star(a, b)$, and then predicate 'truth' of that value with T. This is surely enough to classify the structure presented by the alien 'sentence' as *familiar* on the sophisticated approach.

[34] $\dagger$ may be basic or derived, in the sense that $\dagger(a_1, b_1)$ is the result of applying other basic modes of composition in some particular order.

[35] Elsewhere, I have argued that, if the value of a_1^α can also be the value of $a_2^{\alpha'}$, then it must be that $\alpha = \alpha'$ (Trueman 2021: esp. ch. 5). However, I will not help myself to this assumption here.

We can use a_2 and b_2 to construct a familiar sentence that also perspicuously represents the very same state of affairs as $\dagger(a_1, b_1)$.[36] The first step is to construct a familiar sentence that represents the same state as $\dagger(a_1, b_1)$, without loss of any structure. The key to pulling off this trick is to note that, by *Functional Familiarity*, $\alpha' \rightarrow (\beta' \rightarrow t)$ is a familiar type.[37] We could, then, introduce a familiar type of term, $R^{\alpha' \rightarrow (\beta' \rightarrow t)}$, on the stipulation that $Rx_2^{\alpha'} y_2^{\beta'}$ is to have the same value as $\dagger(x_1^{\alpha}, y_1^{\beta})$, whenever x_2 has the same value as x_1, and y_2 has the same value as y_1.[38] (R is well-defined because, by hypothesis, $\dagger(a_1, b_1)$ is perspicuous and thus transparent.) Ra_2b_2 will then be a familiar sentence that represent the very same state as $\dagger(a_1, b_1)$.

It is obvious that Ra_2b_2 must satisfy clause (iib), since $\dagger(a_1, b_1)$ is assumed to be perspicuous, and each simple constituent in $\dagger(a_1, b_1)$ corresponds to a simple constituent with the same value in Ra_2b_2. That is the sense in which we have not *lost* any of the alien sentence's structure. Unfortunately, however, we have no guarantee that it will satisfy clause (iia): we have introduced a new simple constituent, R, not present in the perspicuous sentence $\dagger(a_1, b_1)$. In that sense, we have introduced *excess* structure into our familiar sentence.

But this obstacle is easily overcome. Once you are familiar with a particular functional term, it is easy to introduce a corresponding mode of composition. For example, suppose you are familiar with a term, $U^{e \rightarrow (e \rightarrow t)}$, that expresses the *is underneath* relation. If you wished, it would then be easy to introduce a new mode of composition, $\binom{y^e}{x^e}$, on the stipulation that $\binom{y^e}{x^e}$ is always to represent the same state as $Ux^e y^e$. This mode of composition may not have been present in your original language, but it would be wildly implausible to suggest that adding it would result in an *alien* language: all it lets us do is express a familiar relation between familiar objects with one less term! (Eklund (2024: 81–2) makes similar remarks about Sellars's (1962) imaginary language *Jumblese*.)

In just the same way, we could introduce a new mode of composition, $\dagger'$, on the stipulation that $\dagger'(x^{\alpha'}, y^{\beta'})$ is always to represent the same state as $Rx^{\alpha'} y^{\beta'}$. It is impossible to see why Ra_2b_2 should count as familiar but $\dagger'(a_2, b_2)$ shouldn't. So $\dagger'(a_2, b_2)$ is a familiar sentence that satisfies every clause of *Structural Perspicuity*: the simple constituents of $\dagger'(a_2, b_2)$ have exactly the same values as the simple constituents of the perspicuous alien sentence $\dagger(a_1, b_1)$, and it is obvious by its construction that $\dagger'(a_2, b_2)$ must be transparent if $\dagger(a_1, b_1)$ is.[39]

[36] Richardson (2026: §5) suggests that we could introduce a familiar term simply by *stipulating* that it is to perspicuously represent the same thing as some alien term. However, as Richardson (2026: 286) acknowledges, this is not obviously the sort of thing that you can stipulate. By contrast, I will show how to construct a familiar sentence that *demonstrably* satisfies the clauses of *Structural Perspicuity*.

[37] I am here adopting the standard Currying-method to transform polyadic functions into monadic functions. So $\alpha' \rightarrow (\beta' \rightarrow t)$ stands-in for dyadic functions from α' and β' to t.

[38] It does not matter what value Rx_2y_2 has when these conditions are not met. One natural option would just be to leave it undefined; another would be to let it be some canonical contradiction.

[39] Although this argument shows that $\dagger'(a_2, b_2)$ satisfies every clause of *Structural Perspicuity*, it does not follow that any other sentence of the same form also satisfies these clauses. This might suggest a possible response to my argument: perhaps we should say that a sentence is perspicuous iff *every* sen-

This argument is painfully abstract, and so it would be helpful to give an example. Unfortunately, however, I obviously cannot give any examples of alien sentences (let alone *perspicuous* ones). Nonetheless, we can consider an example involving only familiar types of terms that still illustrates the process described above. Let's suppose that $F^{e\to t}a^{e}$ is perspicuous. So there is a privileged decomposition of the state of affairs represented by Fa that splits it up into a property, an object, and nothing else. On this privileged decomposition, the property attaches directly to the object, without any additional glue to hold them together.

Now imagine that the predicate F has the same value as the name f^{e}; you might think of F as the formal counterpart of the verb phrase 'is wise', and f as the counterpart of the noun phrase 'wisdom'. (In my opinion, this is all a terrible mistake, but let's not worry about that now.) My claim is that, if Fa is perspicuous, then we can perspicuously represent the very same state using f instead of F. Now, we cannot just substitute f for F. Fa was formed by applying the predicate F to the name a, and the rules of application will not allow us to apply one name to another. But we can always introduce a new predicate, $I^{e\to(e\to t)}$ for *instantiation*, and stipulate that $Ix^{e}y^{e}$ is to have the same value as $Y^{e\to t}x^{e}$ whenever y has the same value as Y. Now Iaf will represent the very same state as Fa. Of course, this sentence may not yet be perspicuous: we have introduced a new simple constituent, I, but, on the privileged decomposition we started with, F *directly* attaches to a, without any intermediary. However, once we have the term I available to us, there is nothing to stop us from introducing a new mode of composition with the following stipulation: $\overline{x^{e}y^{e}}$ is always to have the same value as $Ix^{e}y^{e}$. This new mode of composition determines which state $\overline{af}$ represents from the values of a and f, just as ordinary predication determines which state Fa represents from those same values. So, by *Structural Perspicuity*, $\overline{af}$ perspicuously represents the same state as Fa.

Setting this example aside, the main point is that the argument laid out in this section can be generalised to show that *Structural Perspicuity* implies the following: if the values of alien terms can always be taken as the values of familiar terms, then any state perspicously represented by an alien sentence can also be perspicously represented by a familiar sentence.[40] So, if we combine *Structural Perspicuity* with the sophisticated

tence formed via the same mode of composition satisfies the clauses of *Structural Perspicuity*? (Thanks to Nick Jones for proposing this reply.) However, there are two objections to this proposal. First, it is not obviously well motivated. If our aim is to figure out whether the syntactic structure of a given sentence matches up in the right way with a privileged decomposition of the state it represents, why should we look at the structural relationships between *different* sentences and the states *they* represent? Second, adopting this proposal would have unhappy consequences elsewhere. Suppose that the sentence Rab satisfies the clauses of *Structural Perspicuity*. It follows that Fa cannot also satisfy those clauses when F expresses the property *bears R to b* (i.e. $\lambda x.Rxb$). The proposal under consideration then implies that *no* sentence of the same form as Fa can be perspicuous. But, surely, the fact that some dyadic predication is perspicuous should not imply that no monadic predication is!

[40] In order to generalise the argument, we must assume that no alien sentence is longer than every familiar sentence. In particular, if all familiar sentences must be finitely long, then so must all alien sentences. On the other hand, if we do permit infinitely long familiar sentences, then we need to generalise *Functional Familiarity* to cover functional terms that take infinitely many inputs.

approach to alien structure, we are led back again to *Alien Values*. But we have already seen that *Alien Values* implies that there are no alien structures. (Or at least it does if we grant that we understand *Alien Values* in the first place, and things would only get worse for the sophisticated approach if we denied *that*.)

6 Other notions of perspicuity?

The argument that I just presented relied on a particular account of perspicuity, namely *Structural Perspicuity*. One way for Eklund to respond, then, would be to reject that account. And, in fact, Eklund does consider two other accounts of perspicuity.

It is not uncommon for philosophers to complain that the *surface form* of some sentence obscures its *logical form*. (Consider, for example, sentences like 'Nobody loves everybody'.) One way for a sentence to be 'perspicuous' is for its surface form not to be misleading in this way; in Eklund's (2024: 17) terminology, a sentence that is 'perspicuous' in this sense *displays* its logical form. I do not doubt that this notion of perspicuity has an important role to play in philosophy. However, it is much less clear that it has any role to play in discussions of alien metaphysical structures. How a sentence represents reality is determined by its logical form. So two sentences could represent reality in *exactly* the same way, even if only one of those sentences displayed its logical form; the difference between these sentences would only concern the ease with which we could *discern* the way that each represented reality. Moreover, this seems to be Eklund's (2024: 17) own view too: he largely sets this variety of perspicuity aside, on the grounds that, 'if all sentences have familiar logical forms, there will be no alien logical forms for sentences to display'.

So now let's consider Eklund's *other* other account of perspicuity, which he sketches as follows:

> [*Structural Perspicuity*], even assuming it is helpful as far as it goes, makes perspicuous representation strictly a matter of logical form. At the risk of complicating things, I want to allow for the possibility that representations with the same form, in any natural sense of 'form', can differ in how perspicuously they represent reality. For example [...] if there are different word–world semantic relations (different reference-like relations), then maybe two true sentences '$F(a)$' and '$G(b)$' can be such that 'F' and 'G' have the same worldly correlate and so do 'a' and 'b', but because the expressions stand in different reference-like relations to their correlates, one sentence more perspicuously represents the world than the other does. I am not saying that I commit myself to this being so; all I want to do is to leave the possibility open. (Eklund 2024: 16–7)

One thing that makes this passage difficult to square with the rest of Eklund's book is that, elsewhere, he is at pains to emphasise that alien languages are *structurally* different from familiar languages. So, if some alien sentence perspicuously represents a state of affairs that no familiar sentence can perspicuously represent, then shouldn't that be because of the alien sentence's unfamiliar logical form? However, this might

just go to show that, for Eklund, the 'logical form' of a sentence is not straightforwardly determined by the types (or 'semantic categories') of its constituents. For example, Eklund (2024: ch. 4) discusses the idea that there might be alien counterparts to names that do not refer, but do still do something *very similar*, namely *refer**. He seems happy to say that these alien terms would not count as names (type e), since they would play a different kind of semantic role, but they would still get to count as names* (type e^*), since their role would be so similar to that of real names. Perhaps Eklund is then willing to say that $F^{e\to t}(a^e)$ would have the same 'logical form' as $G^{e^*\to t}(b^{e^*})$, despite the differences in type.

My first real concern with this conception of perspicuity is just that it is not yet well developed. As Eklund (2024: 107–9) is the first to admit, it is not at all clear whether there are any terms that play a semantic role *very similar to*, but still *different from*, names.[41] But, more importantly, this conception of perspicuity does not seem like a good fit for the sophisticated approach to alien structure.

We started §4 with this intuitive thought: a perspicuous sentence should offer some insight into the state of affairs that it represents. That intuitive gloss might not be appropriate for *all* conceptions of perspicuity, but it seems correct for the conception involved in the sophisticated approach. But *how* might a perspicuous sentence provide this insight? The obvious idea is that the sentence behaves as a *model* of the state.[42] We should be able to read-off important properties of a state from the properties of a sentence that perspicuously represents it. (Of course, we might not know which sentences are perspicuous, but we can at least explore what would follow *if* one sentence or another were perspicuous. In other words, we can consider different hypotheses about which sentences are the best models for which states.)

However, the whole idea that a perspicuous sentence is a model of a state seems intimately connected with a structural approach to perspicuity. For a sentence to be a model of a state (in the relevant sense), the simple syntactic constituents of the sentence must go proxy for constituents of the state, and so each of the former must stand in some reference-like relation to one of the latter. But it does not seem to matter *which* reference-like relation it is (assuming that any of the constituents of the state could be 'referred to' in more than one sense in the first place). After all, what we are trying to model is the state represented by the sentence, not anything to do with the *representing* relation itself.

I do not want to pretend that any of this establishes, once and for all, that *Structural Perspicuity* is the only way of spelling out the notion of perspicuity involved in the sophisticated approach to alien structure. However, I do think it is enough to push the

[41] I think that names and predicates both 'refer', but do so in completely different senses from each other (Trueman 2021: chs 2–6). However, I think that these senses are *completely different*; more than that, they are *incommensurable*. As a result, I would never suggest that a name and a predicate might each refer in their own way to the *same thing*, and nor would I say that $F^{e\to t}(a^e)$ has the same logical form as $G^{(e\to t)\to t}(b^{e\to t})$.

[42] In other words, something like the Tractarian *picture theory* should hold for perspicuous sentences, even if it does not hold for sentences in general.

onus back onto Eklund here. If the epistemic possibility of alien structure ultimately rests on a novel conception of perspicuity, then much more needs to be said about the viability of that conception. In the meantime, it seems we have good reason to remain sceptical about that possibility.

7 Conclusion

Let's suppose that I am right: insofar as we understand the hypothesis that there are alien structures, we know that it cannot be true. It is important to emphasise that this is only a conclusion about alien *structures*, not alien *languages*. Nothing I have said has gone any way toward suggesting that there cannot be alien languages, or that we cannot truthfully *say* that there are.

The reason for this difference is that the terms in alien languages can all just be objects (i.e. the sorts of things that we can quantify over with first-order quantifiers), just like the terms in familiar languages.[43] So there need be no difficulty in quantifying over alien terms in a familiar language. Nor need there be any difficulty in applying familiar relations to alien terms. In particular, then, we can say that there are terms that *play different kinds of semantic roles from* any familiar terms. In other words, we can say that there are alien types of term, and thus that there are alien languages.

Nonetheless, the point remains that we cannot say that alien terms take different values from familiar terms. There is a sense, then, in which there is a principled limit on the *resolution* of our talk about alien languages. We can make the low-resolution claim that alien terms play different semantic roles from familiar terms, but we cannot make high-resolution claims about what that difference consists in.

This limit leaves us in a slightly odd position. Imagine that we actually meet a community of aliens, speaking an alien language. Even though we cannot understand their language, it might still be obvious to us that they really are speaking a *language*. We can imagine, for example, that these aliens are far more technologically advanced than us, and that their remarkable innovations always follow complex and lengthy exchanges in their language. It would perhaps be natural for us to conclude that the aliens' technological advantage is a direct product of the fact that they use an alien language rather than a familiar one. But, for all that, we would still know that we would end up saying something false if we tried to say that alien structures are presented by the alien language. (This limit matters; for example, it means that there cannot be a project of accommodating alien structures in our metaphysical theorising of the sort envisioned by Eklund (2024: ch. 2).)

I think it is helpful to compare this conclusion with Putnam's (1981: ch. 1) (in)famous refutation of brain-in-a-vat scepticism. According to Putnam, a brain in a vat would be so expressively impoverished by its circumstances that it could not even express the thought that it is a brain in a vat. (Or, at least, it could not if it had never

[43] Some Fregeans have suggested that predicates should not be thought of as objects (e.g. Geach 1953: 410, 1976: 59–60). However, see Oliver (2010) for a response.

interacted with anyone or anything outside of the vat.) Since you can express that thought, you can infer that you are not a brain in a vat. Of course, this is not to deny that you might undergo a remarkable experience tomorrow which, on long enough reflection, you might describe truly by saying, 'I was once a brain trapped in a vat, but then I was freed and brought out into the real world'. However that string of words would have a meaning for you in those circumstances that they cannot have for you now.

Similarly, if you learned an alien language tomorrow, you might truthfully declare, 'This language allows me to present structures that I couldn't present before'. But that string of words would have a meaning for you then — expressed using the alien types of terms — that it cannot have for you now.

One of the reasons that I think this comparison is helpful is that it hints at something important about the status of type theory in philosophy. More and more metaphysicians are theorising in type-theoretic terms, and this is sometimes presented as an exciting new development in the field. But when metaphysicians theorise *about* type theory, they tend to approach it in the same, broadly abductive way that they approach other topics in metaphysics. Different type theories are presented as different hypotheses about how reality is structured, and we are encouraged to choose our preferred type theory on the basis of a cost-benefit analysis.[44] However, questions about which types are available to us are so fundamental to our expressive powers that we might suspect that they do not belong to the ordinary domain of abductive metaphysics. Like brain-in-a-vat scepticism, it may be that they belong in the realm of *transcendental* considerations. But, for now, I will just leave that as suggestive speculation, to be pursued some other time.[45]

References

Bacon, Andrew (2024). *A Philosophical Introduction to Higher-Order Logics*. New York, NY: Routledge.

Berry, Sharon (2022). *A Logical Foundation for Potentialist Set Theory*. Cambridge: Cambridge University Press.

Button, Tim (forthcoming). 'Why are all the sets all the sets?' *Noûs*.

Crosilla, Laura (2024). 'Constructive type theory, an appetizer'. In: Fritz and Jones 2024, pp.220–242.

Dorr, Cian (2016). 'To be F is to be G'. *Philosophical Perspectives* 30, pp.39–134.

— (2025). 'Higher-order quantification and the elimination of abstract objects'. *Disputatio*.

Eklund, Matti (2024). *Alien Structure: Language and Reality*. Oxford: Oxford University Press.

Fine, Kit (2001). 'The question of realism'. *Philosophers' Imprint* 1, pp.1–30.

[44] For examples of this outlook, see: Florio and Jones 2021; Bacon 2024.

[45] Thanks to Daniel Berntson, Tim Button, Matti Eklund, Owen Hulatt, Chris Jay, Nick Jones, Lukas Skiba, and audiences at Bergen, Uppsala, and York. Thanks also to the anonymous reviewers for their helpful comments.

Fine, Kit (2006). 'Relatively unrestricted quantification'. In: Rayo and Uzquiano 2006, pp.20–44.
Florio, Salvatore and Nicholas K. Jones (2021). 'Unrestricted quantification and the structure of type theory'. *Philosophy and Phenomenological Research* 102, pp.44–64.
Frege, Gottlob (1893). *The Basic Laws of Arithmetic*. Vol. I. Jena: Pohle.
Fritz, Peter (2023). *The Foundations of Modality*. Oxford: Oxford University Press.
Fritz, Peter and Nicholas K. Jones, eds. (2024). *Higher-Order Metaphysics*. Oxford: Oxford University Press.
Geach, Peter Thomas (1953). 'Quine on classes and properties'. *The Philosophical Reivew* 62, pp.409–412.
— (1976). 'Saying and showing in Frege and Wittgenstein'. *Acta Philosophica Fennica* 28, pp.54–70.
Hawthorne, John and Andrew Cortens (1995). 'Towards ontological nihilism'. *Philosophical Studies* 79, pp.143–165.
Hofweber, Thomas (2022). 'The case against higher-order metaphysics'. *Metaphysics* 5, pp.29–50.
Jones, Nicholas K (2018). 'Nominalist realism'. *Noûs* 52, pp.808–35.
— (forthcoming). 'Opacity in the book of the world?' *Philosophical Studies*.
Krämer, Stefan (2014). *On What There Is For Things To Be*. Vol. 1. Studies in Theoretical Philosophy. Frankfurt: Klostermann.
Linnebo, Øystein (2010). 'Pluralities and sets'. *The Journal of Philosophy* 107, pp.144–164.
— (2018). *Thin Objects*. Oxford: Oxford University Press.
MacBride, Fraser (2006). 'Predicate reference'. In: *The Oxford Handbook of Philosophy of Language*. Ed. by Ernest Lepore and Barry Smith. Oxford: Oxford University Press, pp.422–475.
Oliver, Alex (2010). 'What is a predicate?' In: *The Cambridge Companion to Frege*. Ed. by Michael Potter and Tom Ricketts. Cambridge: Cambridge University Press, pp.118–148.
Prior, Arthur (1971). *Objects of Thought*. Oxford: Oxford University Press.
Putnam, Hilary (1981). *Reason, Truth and History*. Cambridge: Cambridge University Press.
Rayo, Agustín (2013). *The Construction of Logical Space*. Oxford: Oxford University Press.
Rayo, Agustín and Gabriel Uzquiano, eds. (2006). *Absolute Generality*. Oxford: Oxford University Press.
Rayo, Agustín and Stephen Yablo (2001). 'Nominalism through de-nominalization'. *Noûs* 35, pp.74–92.
Richardson, Kevin (2026). 'How to speak an alien language'. *Journal of the American Philosophical Association* 12, pp.276–293.
Rosen, Gideon (2010). 'Metaphysical dependence: Grounding and reduction'. In: *Modality: Metaphysics, Logic, and Epistemology*. Ed. by Bob Hale and Aviv Hoffmann. Oxford: Oxford University Press, pp.109–135.
Rubenstein, Ezra (2024). 'Two approaches to metaphysical explanation'. *Noûs* 58, pp.1107–1136.
— (forthcoming). 'Generalism without generation'. *The Journal of Philosophy*, pp.1–27.
Russell, Bertrand (1918). 'The philosophy of logical atomism'. In: B. Russell 2007, pp.174–281.
— (2007). *Logic and Knowledge*. Nottingham: Spokesman.

Russell, Jeffrey Sanford (2018). 'Quality and quantifiers'. *Australasian Journal of Philosophy* 96, pp.562–577.
Sellars, Wilfrid (1962). 'Naming and saying'. *Philosophy of Science* 29, pp.7–26.
Skiba, Lukas (2021). 'Higher-order metaphysics'. *Philosophy Compass* 16, pp.1–11.
Studd, James P. (2019). *Everything, More or Less: A Defence of Generality Relativism*. Oxford: Oxford University Press.
Trueman, Robert (2013). 'Neutralism within the semantic tradition'. *Thought* 1, pp.246–251.
— (2015). 'The concept *horse* with no name'. *Philosophical Studies* 172, pp.1889–1906.
— (2021). *Properties and Propositions: The Metaphysics of Higher-Order Logic*. Cambridge: Cambridge University Press.
— (2022). 'Truth without dependence'. *Aristotelian Society Supplementary Volumes* 96, pp.89–121.
— (forthcoming[a]). 'Determinate compositionality'. *Inquiry*.
— (forthcoming[b]). 'Fregean predication'. *Mind*.
Wilhelm, Isaac (forthcoming). 'Talk about types'. *Mind*.
Williamson, Timothy (2003). 'Everything'. *Philosophical Perspectives* 17, pp.415–65.
— (2013). *Modal Logic as Metaphysics*. Oxford: Oxford University Press.
Wright, Crispin (2007). 'On quantifying into predicate position: Steps towards a new(tralist) perspective'. In: *Mathematical Knowledge*. Ed. by Mary Leng et al. Oxford: Oxford University Press, pp.150–74.