A Lewisian Argument Against Platonism, or Why Theses About Abstract Objects Are Unintelligible

Abstract: In this paper, I argue that all expressions for abstract objects are meaningless. My argument closely follows David Lewis' argument against the intelligibility of certain theories of possible worlds, but modifies it in order to yield a general conclusion about language pertaining to abstract objects. If my Lewisian argument is sound, not only can we not know that abstract objects exist, we cannot even refer to or think about them. However, while the Lewisian argument strongly motivates nominalism, it also undermines certain nominalist theories.

1. Introduction

In this paper, I argue that there is no way we can have genuine concepts of abstract objects. Even if abstract objects exist, we cannot think about or refer to them. The common nouns for distinctive kinds of abstract objects, such as 'property', 'set', or 'proposition', and the expressions for the distinctive ways in which those kinds of abstract objects relate to things, such as 'instantiate', 'member of', or 'represent', are all meaningless. My argument is inspired by David Lewis' argument against the intelligibility of "magical ersatz" theories of possible worlds. It resembles Lewis' argument, but focuses on language rather than relations and yields a general conclusion: all common nouns and expressions like the ones mentioned above are unintelligible. Because of the similarity between my argument and Lewis', I call my argument the Lewisian argument. I will begin by reviewing Lewis' argument and the current state of the discussion surrounding it. I will then construct the Lewisian argument and discuss some possible responses. Afterward, I will note that several rejoinders to the standard epistemic arguments against platonism do not undermine it. I will then argue that if the argument is sound, it strongly motivates nominalism while

simultaneously undercutting some nominalist attempts to explain the tremendous usefulness of language about abstract objects.

2. Lewis Against Magic

In *On the Plurality of Worlds* (pp. 174-191), Lewis introduces a theory of possible worlds which he calls "magical ersatzism" and subjects it to a sustained critique. The theory goes as follows (pp. 174-176). There are some abstract entities, *elements*, that bear a relation, *selection*, to the concrete world. Neither 'element' nor 'selection' are defined: they are instead primitive notions. (Though Lewis does not suggest this possibility, one might define an element as something which is possibly selected [Nolan 2020, p. 4728], albeit this still leaves selection as a primitive.) The magical ersatzer also takes modality to be primitive. Each element u has the following feature: for some way the concrete world can be, necessarily, u is selected iff the concrete world is that way. An element can be said to *represent* a way that the concrete world can be just in case necessarily, if it is selected, then the concrete world is that way. For example, if a given element is such that necessarily, it is selected iff there is an invisible smiling cat, that element represents that there is a cat. An element x is *maximal* just in case for every other element y, either necessarily, if x is selected, y is selected as well, or necessarily, if x is selected, y is not selected. Possible worlds are simply maximal elements. The elements are diverse and plentiful enough that for every way the concrete world can be, some element has it for its selection condition. In turn, maximal elements either represent that so and so or represent that not so and so for any condition of the concrete world. Elements lack any internal structure, so their selection conditions cannot be explained in terms of anything like the familiar phenomena of linguistic structure or pictoral resemblance (in contrast to some theories that Lewis considers on pp. 142-174). So goes the magical ersatz theory.

Lewis proposes a dilemma for the magical ersatzer. Either selection is an internal relation or an external relation. An internal relation is one that obtains between its relata in virtue of their intrinsic natures, while an external relation is one that doesn't.[1] Suppose selection is internal. Then it is impossible for us to have any grasp of it. Lewis supports this point as follows (pp. 177-179). The intrinsic properties of elements are inaccessible to us. Because we are not able to have "acquaintance" with abstract objects, which would require some sort of causal interaction, we have no way of grasping what these intrinsic properties are. These intrinsic properties must be varied and plentiful enough to account for all the different selection conditions of the elements. Yet they are utterly mysterious. Moreover, selection is a relation that only obtains between the concrete world and elements. It isn't a relation that we can witness being manifested entirely within the concrete world, so we have no acquaintance with an example of it from which to draw our understanding. It would seem, then, that we would need to have some sort of grasp of the intrinsic nature of elements in order to grasp selection. Yet we don't have any such grasp. Thus it is a complete mystery as to how we can have genuine concepts of elements or selection in our thought, or genuine expressions for them in our language. The mystery here is intractable enough to strongly suggest that it's impossible for us to be able to think or talk about these things.

On the other hand, suppose that selection is external. One issue, again, is how it is that we can understand the primitives given that we cannot be acquainted with elements and that selection never obtains entirely within the concrete world (p. 179). The main difficulty, though, is why it should be that *necessarily*, a given element is selected iff the concrete world is so and so (pp. 179-181). It would seem to be a criterion of being an external relation that there is no necessary

---

[1] Technically Lewis defines an external relation as one that is not internal and depends on the intrinsic nature of the fusion of the relata (p. 177), but as Lewis himself points out (p. 182, number 3), nothing is really lost just by thinking of an external relation as a relation that isn't internal.

connection between some entity being a certain way and its bearing that relation to something that shares no part with it, yet here there is said to be a necessary connection between the concrete world being a certain a way and bearing an external relation (selection) to some disjoint object (the element). Lewis supports this point with an example (p. 180). If we know anything at all about metaphysical possibility, we know that the arrangement of some particles in a region of spacetime far away from Earth does not supervene on what is happening on Earth. Even if there were some nomological necessity here, there would be no metaphysical necessity. What goes for spatiotemporal relations goes for external relations in general: the fact that elements are abstract makes no difference.

So, if selection is an internal relation, we understand neither selection nor elements. However, it cannot be an external relation because it involves a necessary connection between the concrete world and elements, and external relations do not support necessary connections of this sort. Since selection must be internal or external, there is no such relation as selection, or at least no relation like it that we can possibly entertain in our thought or language. Likewise, there are no such things as elements, or at least we cannot think about or refer to them. In the end, magical ersatzism is not an intelligible theory.

Lewis' argument has had some influence. It has been observed before that the considerations that Lewis brings up apply to abstract objects beyond ersatz possible worlds. For example, Peter van Inwagen (1986, pp. 202-210) could not identify a flaw in Lewis' argument but pointed out that it generalizes to sets as well. Membership is either internal or external. If it is internal, we cannot grasp it for the same reasons as selection, but it cannot be external because of the necessary connection between sets and their members. Since sets play a central role in Lewis' philosophy, van Inwagen made a *tu quoque* argument against Lewis: if Lewis' argument is sound,

Lewis must give up on the possibility of meaningfully talking about sets.[2] Lewis himself ultimately arrived at a solution to this problem. By adopting the strategy of the Appendix to *Parts of Classes* (1991), Lewis eliminates any use of primitive abstract expressions in his exposition of set theory. This strategy, however, requires the existence of inaccessibly-many mereological atoms. Even worse, Lewis' overall theory of abstract objects is only available to those who endorse his modal ontology. Set theory alone is not enough to deliver properties, relations, and propositions: to follow Lewis, one needs concrete possibilia as well. This is a pill too bitter for most to swallow.[3]

More recently, Daniel Nolan (2020) has extended the problem to properties, relations, and instantiation, while recognizing that the problem has ramifications for theories of almost every kind of abstract object. A. R. J. Fischer (2018) has also noted the possibility that "Lewis' critique is so powerful that it wreaks havoc across the enterprise of metaphysics" (p. 2337), but chooses not to dwell on that fact and focuses on possible worlds in his discussion of the critique. However, the problem that Lewis raises has generally been seen as a puzzle to be solved rather than to constitute the basis of a sound argument for nominalism. Nolan, for example, attempts a solution. I take the latter approach. I state the problem in general form for all kinds of abstract objects and

---

[2] This is a modest simplification. Van Inwagen goes on to consider in more depth the possibility that membership is "range-internal" or "range-external", but these complications are unnecessary to go into here.

[3] A curious fact is that Lewis was bothered by van Inwagen's *tu quoque* argument despite at times regarding impure sets as being located where their members are and capable of causally interacting with things. Lewis even says that pure sets are the most dispensable and metaphysically suspicious of them all (1986, pp. 83-84). The reason for Lewis' concern is how van Inwagen identifies the heart of the problem. For van Inwagen, the key is that we can only distinguish between sets by means of the membership relation. Even if they are spatiotemporally located and causally involved in the world, we don't seem to be able to access the properties of impure sets that are relevant to the membership relation (Lewis 1991, pp. 29-35). In fact, we only come to know of a given set's location and causal entanglements by first knowing what its members are. "Acquaintance" is a necessary condition to be able to grasp the intrinsic natures of things, but it isn't a sufficient one. In the end, Lewis is ambivalent even about the locatedness and causal involvement of impure sets (1991, pp. 31-33). Since most abstract objects are not causally active, acquaintance is what I will focus on in my Lewisian argument. Still, the point is well-taken: the story of how an internal relation is understood doesn't automatically stop once causal relations are introduced. A view like Penelope Maddy's (1990) is not subject to the difficulty van Inwagen raises for identifying sets since Maddy explicitly views sets as being perceptible and having a variety of sensible qualities in addition to being spatiotemporally located and causally involved. Different sets, therefore, are apt to be distinguished by all sorts of properties on Maddy's view.

argue that all extant solutions are unsatisfactory. Rather than speaking of internal and external relations when introducing my Lewisian argument, I will speak in terms of linguistic expressions: it would raise thorny philosophical issues to rely on talk of relations in an argument that such talk is unintelligible. In the end, I propose, there is no genuine alternative to nominalism.

3. The Lewisian Argument

Let *s* be a common noun for a specific kind of abstract object, and let *e* be an expression for the distinctive way that abstracta that fall under *s* relate to the world. For example, if *s* is 'property', *e* is 'instantiates'; if *s* is 'set', *e* is 'member of'; and so on. Then I offer the following argument:

1. If no one understands *e*, then no one understands *s*.

2. If *e* is an internal expression, then no one understands *e*.

3. *e* is not external.

4. *e* is either internal or external.

5. Therefore, no one understands *s* and no one understands *e*.

The argument is logically valid. Premise (1) is supported by how we would explain what the common nouns mean: "A property is an instantiable thing, or at least of the same kind as instantiable things (some cannot be instantiated, such as the property of being tall and short, but they are the same kind of thing as properties that can be)"; "A set is something that has members, or in the case of the empty set is the same kind of thing as that which has members"; and so on.

Premise (4) is true by definition. Say that *e* is an internal expression if it applies to some things in virtue of how they are intrinsically (all of them, that is), and external otherwise. Thus 'cylindrical' and 'taller than' are plausibly internal, but 'purchased' and 'far away from' are not: whether something is cylindrical is intrinsic to it, and whether one thing is taller than another depends solely on how those two things are intrinsically (how tall each is), but whether someone

bought a thing is not intrinsic to it, and whether one thing is far away from another does not depend on how those two things are intrinsically (their relative distance). Since for an expression to be an external expression is for it to not be internal, every expression is internal or external, including *e*.

The controversial premises of the argument are (2) and (3). I will first defend premise (2), then defend premise (3). Suppose that *e* is internal. Then, similarly to the selection relation when it is taken to be internal, it cannot be explained how it is that we grasp *e*. For any *n*-ary expression that applies to some things in virtue of how they are intrinsically, to grasp it one needs to perceive or causally interact with some things to which *e* applies (jointly or individually, as appropriate to the expression), or *e* needs to be analyzable in some way in terms of already familiar notions.[4] For example, consider 'color'. One grasps it by seeing colored things. Or consider 'negatively charged'. One grasps it by physically interacting with electrons and discerning what causal profiles electrons seem to have. Finally, consider 'black all over and white all over'. Though one never perceives or interacts with such a thing – nothing can be all black and all white – one can grasp it by analyzing it in terms of 'black all over' and 'white all over', which themselves can be grasped by perceiving black and white things. To hold that there is an exception to this principle is to believe in magic: it is to think we can grasp a phenomenon (i) with an essential class of constituents such that we never had any input from the world concerning its members, and (ii) which cannot be analyzed in terms of other phenomena on which we have received input.

The trouble is that *e* is an expression for a distinctive way that certain abstracta relate to the world. Therefore, *e* applies to some things only if at least one of them is abstract, but no one perceives abstracta or causally interacts with them.[5] Moreover, it doesn't seem possible to analyze

---

[4] Probably, perception requires causal interaction, but in case that is not always true I add it as an additional option for gaining understanding.

[5] This assumes a standard platonic view of abstracta, wherein abstracta are not spatiotemporally located and are causally isolated from concreta. Those who hold, for example, that properties are spatiotemporally-located,

e in terms of already familiar notions. So no one understands *e* if *e* is internal.[6] That establishes premise (2).

An initial reaction one might have to this proposal is that it is a false trilemma to suppose that *e* can only be grasped through perception, causal interaction, or analysis. Two alternatives to these three options in particular are worth considering here, since they are very natural suggestions, though additional alternatives will be considered in a later section discussing objections. First, analogy might be thought to provide the key to grasping *e*. After all, we use analogies to explain concepts to others all the time in ordinary conversation. Unfortunately, if *e* is not to be analyzed, then *e* is a conceptual primitive, and analogies do not go far enough in explaining what some conceptually primitive phenomenon is to make for genuine understanding. For example, consider a person who is blind and a sighted individual who is attempting to explain facts about color to the blind person. The sighted individual might begin by offering the analogy that color is to the eye what sound is to the ear. They might continue by offering the analogy – playing off the idea that red is a color that involves the passions – that red is to color as drama is to fiction. And they can offer similar analogies to give the blind person some sense of what is involved in the other colors.

___

perceptible constituents of concrete things will deny this point (e.g. David Armstrong [1989]), as will those who agree with platonists that abstracta are not located in space but hold that they causally interact with concreta (e.g. Julian Dodd [2007] and David Friedell [2020]). My opinion is that such views are flawed for independent reasons, but it should be noted that such views might have a way out of the Lewisian argument. However, causal interaction with the abstracta to which *e* pertains is not enough by itself to guarantee that *e* is intelligible: see fn. 3's discussion of Lewis' concerns about van Inwagen's *tu quoque*.

[6] One might think hallucination casts doubt on the principle that relevant perceptions or causal interactions are necessary for grasping the meaning of an internal expression. Take 'color'. Suppose someone blind from birth had their brain stimulated by a scientist in such a way that they seemed to see a red sphere, a blue cone, and a green cube. If the scientist were to tell them 'color' referred to how these apparent objects looked and that there were other colors besides, it's plausible that the blind person would understand what 'color' means even without having seen anything colored or having been causally related to colored things in a way relevant to concept acquisition. But even in this case, the blind person *seemed to* see a colored thing, and no one can *seem to* perceive an abstract object in this sense: abstract objects can no more be imagined or hallucinated than they can be perceived. Modifying the general principle so that it reads "For any *n*-ary expression that applies to some things in virtue of how they are intrinsically, to grasp it one needs to seem to perceive or causally interact with some things to which *e* applies (jointly or individually, as appropriate to the expression), or *e* needs to be analyzable in some way in terms of already familiar notions" would be sufficient to justify premise (2).

This is sufficient to give the blind individual some knowledge, but it would be wrong to say that the blind individual knows what color is. What the blind individual comes to know is that something called 'color' is perceived by the eye, just as sound is perceived by the ear; that among the things falling under its umbrella, one called 'red' is known for stirring emotions; and so on. But to say that the blind individual knows what 'color' means is wrong. They do not know what it signifies, and the same is true for 'red'. What they know, unfortunately, are only existential facts and facts about what some things are called. I cannot find a case where analogy is sufficient to convey understanding of a primitive, and therefore must conclude that analogy won't be adequate to convey understanding of *e*.

Note that I am not saying that there are no cases in which analogy provides sufficient information to yield understanding. It is just that in these cases, the relevant notions are not conceptually primitive. Suppose someone had never heard of Minsk, but did know of Belarus. Someone might offer to such a person the analogy that Minsk is to Belarus as Washington, D.C. is to the United States. This analogy might be enough to provide understanding of the meaning of 'Minsk' to the individual: it is the capital city of Belarus. But this is because all the concepts involved in the definite description that is implied by the analogy are already understood. Even if 'Minsk' is semantically primitive (I take no stance on the semantics of proper names), it isn't conceptually primitive in the same sense as *e*, since familiar concepts can be used to pin down the meaning. Not only can *e* not be defined with familiar notions, but, being a predicate, *e* cannot even be pinned down by a definite description in terms of them.

A second alternative to the three mentioned in the Lewisian argument is that we might grasp *e* by some dependence relations. Take, for example, 'instantiates'. It is plausible that instantiation is grounded in particulars. For example, that a ball instantiates the property of being

red is grounded in the ball, which is red. One might further suppose that properties themselves are grounded in what instantiates them. Familiarity with the grounds of instantiation and properties might be taken to serve as the basis for understanding 'instantiates' and terms for properties. The trouble with this proposal is that it is not in general true that familiarity with grounds is sufficient for gaining understanding of what it is that they ground. To see this, consider the following examples. Someone might know all the parts of a given car and a given roof, but fail to form the concept of a car-roof, their fusion. Likewise, a newcomer to chess might understand all the moves in a given opening but fail to have the concept of an opening. And for a third case, a neuroscientist might understand all the brain states involved in some unusual psychological state while failing to have the concept of that psychological state.

An obvious reply is that even though these cases are possible, the individuals in them are capable of forming these concepts. It would not be difficult to get the individuals to come to understand 'car-roof', 'opening', or a word for the psychological state. Similarly, perhaps someone might be familiar with the grounds for instantiation and properties without grasping 'instantiates' and terms for properties, but they are in a position to understand them. That would undermine premise (2). What is important to notice, however, is what is involved in coming to that further understanding in the three cases I brought up. In the first two, the key is explicit definition in terms of what is already familiar: 'car-roof' and 'opening' can be readily defined. And in the latter case, while it might be difficult to provide a defining description of the psychological state, what is key to naming it is additional experience of and causal interactions with individuals in it. Analysis, perception, and causal interaction are what serve as the conceptual links between the grounds and the grounded. That is exactly what is missing in the case of 'instantiates' and terms for properties, even conceding the facts about grounding. It is what is missing in the case of e.

Moving to the defense of premise (3), remarks apply that are similar to those Lewis makes regarding the possibility of selection being external. Suppose for reductio that *e* is external. One challenge here is simply to explain what *e* means. To understand *any n*-ary expression, whether internal or external, one needs to perceive or causally interact with some things to which *e* applies (jointly or individually), or *e* needs to be analyzable in some way in terms of already familiar notions. Now consider, for example, 'member of'. If we are to believe that 'member of' is external and genuinely understood, it must either be that we perceive or casually interact with some things such that one is a member of the other, or else that 'member of' can be defined using familiar expressions. Since we do not perceive or causally interact with sets, we must rely on familiar expressions.[7] And those familiar expressions must ultimately be defined in terms of expressions that apply to some things that we can perceive or with which we can causally interact. For any such expression, we can ask how it is that we understand *it*; if we understand it not by way of perceiving or causally interacting with some things to which it applies, but in terms of some familiar notions, we can ask how we understand *those* notions; and so on, until we reach the bottom of the chain of familiar notions. My suspicion is that one will find it difficult to define 'member of' in a satisfying way even granting that *e* is external, along with other expressions for distinctive ways abstracta relate to the world such as 'instantiates' or 'bear… to'.

There is, however, an even deeper worry. And that is that any expression like *e* is such that for some way of being and abstractum, an object's being that way necessarily coincides with that object and that abstractum's being such that the expression applies to them. For example, the following sentences are true:

$\Box(\forall x(x$ is red $\leftrightarrow x$ instantiates the property of being red$))$

---

[7] Though some philosophers would only say this is true for pure sets: see fn. 3.

$$\Box(\forall x(x \text{ exists} \leftrightarrow x \text{ is a member of the singleton of } x))$$

An object's being red necessarily coincides with that object and the property of being red's being such that 'instantiates' applies to them, while an object's existence necessarily coincides with that object and its singleton's being such that 'member of' applies to them.[8] The point can be put with e as follows. Whatever e is, at least one instance of one of the schemas

$$\Box(\forall x(x \text{ is blue} \leftrightarrow x \ e \ \mathrm{T}))$$

$$\Box(\forall x(x \text{ is tall} \leftrightarrow x \ e \ \mathrm{T}))$$

$$\Box(\forall x(x \text{ is shaped} \leftrightarrow x \ e \ \mathrm{T}))$$

…

must be true, where 'T' is a schematic letter for terms for the relevant kind of abstracta.[9] Unfortunately, this poses a deep problem. External expressions don't support necessary coincidences of this sort. If an expression does not apply to some things in virtue of how they are intrinsically, there isn't any modal glue to make one of a pair of object's being a certain way necessarily coincide with the expression's applying to that pair. For example, 'far away from' is a paradigmatic external expression. For no A and B is there a necessary coincidence between some property of A and how far away from each other A and B are. Even though some properties entail that A and B are far away if A has them – for example, being 10^{10} lightyears away from whatever is nearest – their being far away does not in turn entail that A has those properties, and even though their being far away entails that A has certain properties – for example, being such that something is far away – A's having those properties does not in turn entail that A and B are far away: the

---

[8] This gloss assumes that 'instantiates' and 'member of' necessarily both exist and have their actual meaning if their actual application conditions are met. This assumption is false, but no matter: the real philosophical problem is posed by the biconditionals, not the gloss.

[9] This assumes e is a binary predicate, but the point and schemas could easily be modified for any n-ary predicate.

entailing and entailed properties are non-overlapping.[10] Other extrinsic expressions include 'near to', 'was sold to', 'existed before', 'graduated after', and 'possesses more wealth than', none of which are such that their application to some pair of objects necessarily coincides with one of the pair being a certain way. So if e is external, no instance of any the schemas mentioned above could be true. Yet some such schema must be true: it seems central to e that it validates one. So e must not be external, which establishes premise (3).

In brief: it seems that e must be internal, but if e is internal, no one understands e. And since no one understands e, no one understands s either. And in general, no one understands common nouns for abstract objects or expressions for the distinctive ways in which they relate to the world. So goes the Lewisian argument.

4. Four Responses

Four responses to the Lewisian argument must be considered: the analysis response, the higher-order response, the neo-Fregean response, and the Moorean response. The first three seek to undermine premise (2) of the argument. These responses focus on providing an account of how e can be understood whether it is internal or external. While an advocate of them might also deny premise (3), the responses are compatible with conceding that e cannot be external without losing its required modal glue. The final response does not target any premise in particular, but merely claims that the denial of the conclusion is more plausible than the conjunction of the argument's premises. I will consider these responses in turn.

The analysis response is that e might be analyzable in terms of other expressions that correspond to how certain kinds of abstracta relate to the world. For example, one might propose that properties are sets. One might then define 'instantiates' in the following way:

---

[10] This is only true for qualitative properties, but the schemas mentioned involve qualitative predicates on the left-hand side of the embedded biconditionals, not non-qualitative ones.

x instantiates *y* := *x* is a member of *y*.

Finally, one might then introduce the following schema for property identifications, using a device from Peter van Inwagen (2006, p. 456) for introducing names of properties:

$$\Pi x\ \Phi x := \{x \mid \Phi x\}$$

Where 'Π*x*' abbreviates 'The property of being an *x* such that' and 'Φ*x*' is a schematic letter for a formula with and only with '*x*' open in it.

Various complaints can be made against the theory of properties this definition of 'instantiates' generates. For instance, the theory it yields is extremely coarse-grained. The issue I want to raise is more fundamental. Once one proposes analyzing 'instantiates' in terms of 'member of', the next question to ask is how we can grasp 'member of'. That expression also describes a distinctive way that a certain kind of abstract object relates to the world. Since sets are abstract, perception and causal interaction cannot explain how we grasp the expression. Supposing 'member of' is internal, if there are any familiar notions in terms of which it can be analyzed, some of them must also be internal expressions for distinctive ways that certain kinds of abstract objects relate to the world (otherwise 'member of' wouldn't be internal and only applicable to some things if at least one is abstract).[11] And then the same question can be asked of them. Eventually, if we are not to get involved in a circular definition or an infinite regress, at least one internal expression describing how a certain kind of abstract object relates to the world will be reached that stands at the bottom of a chain of definitions of other such expressions. There won't be any familiar notions in terms of which to analyze it. That will leave only perception and causal interaction as sources

---

[11] This is a factual point, not a logical one. There is no contradiction in supposing that there are some internal expressions that individually are applicable to pluralities of concrete objects, but some logical combination of which is not applicable to any plurality of concrete objects and is applicable to some plurality containing an abstract object. But it is obvious enough that no such internal expressions could be logically combined to define 'member of'.

for grasping the expression, and they are ruled out by the fact that expression pertains to abstracta. The suggestion, therefore, does not solve the difficulty.

The higher-order response, advocated by Nolan (2020, p. 4728), is to use Ramsey sentences to define *e*. For example, suppose *e* is 'member of', and *e** is 'are the exact members of'. Pretend, for a moment, that for all *x*s, there is a unique set of the *x*s, and the only other principle governing sets is extensionality. (Ignore resulting paradoxes.) One might then introduce the following Ramsey sentence:

> There is an R such that for all *x*s, there is a unique *y* such that the *x*s R *y*, and for all *y*s and for all *z*s, if the *y*s are not the *z*s, then every *w* such that the *y*s R *w* is not identical to any *v* such that the *z*s R *v*.

One could then define *e** and *e* in the following schematic ways:

> the *G*s *e* K* := there is an R such that (i) for all *x*s, there is a unique *y* such that the *x*s R *y*, (ii) for all *y*s and for all *z*s, if the *y*s are not the *z*s, then every *w* such that the *y*s R *w* is not identical to any *v* such that the *z*s R *v*, and (iii) the *G*s R *K*.
>
> …
>
> *A e K* := there are some *x*s such that (i) *A* is one of the *x*s, and (ii) the *x*s *e* K*.

Here, '*A*', 'the *G*s', and '*K*' are schematic letters for terms. Finally, one can say that a set is anything such that some *x*s R it.

The trouble with using Ramsey sentences is that they rely on higher-order quantification, and a good Quinean will suspect that higher-order quantification at best is disguised quantification over relations (and properties and propositions), and at worst is meaningless, strictly speaking. Either result won't do. Clearly, a Ramsey sentence needs to be meaningful for it to be a source of

defining terms. Suppose that instead of being meaningless, it incorporates disguised quantification over relations. Then the full articulation of the above Ramsey sentence looks like the following:

> There is a relation R such that for all *x*s, there is a unique *y* such that the *x*s bear R to *y*, and for all *y*s and for all *z*s, if the *y*s are not the *z*s, then every *w* such that the *y*s bear R to *w* is not identical to any *v* such that the *z*s bear R to *v*.

To understand this sentence, one must understand 'bear… to' and 'relation'. And these are expressions our grasp of which is also in question: 'relation' is a common noun for a specific kind of abstract object, and 'bear… to' is an expression for the distinctive way that kind of abstract object relates to the world. We therefore need an account of how we understand them. Clearly a Ramsey sentence won't do on this understanding of what Ramsey sentences really say, since one would have to use 'bear…. to' (or some equivalent expression) to produce a Ramsey sentence defining 'bear… to' and 'relation'. And that would be circular.

Of course, many philosophers reject the Quinean position on higher-order quantification, and such philosophers may be tempted to embrace the Ramsey sentence solution. Nevertheless, there might still be a problem here. It is plausible that no one can understand higher-order quantification unless they first understand 'relation' and 'bear… to' or some equivalent expressions (and 'property', 'instantiate', 'proposition', and 'represent' as well). When people are first introduced to higher-order quantification, it is explained by using relation-talk: "saying there is an R such that *x* R *y* is like saying *x* and *y* bear some relation to each other, but without the ontology". Even if higher-order quantification should not be understood as disguised relation-talk, practice suggests relation-talk is an indispensable instrument to grasping it. To be sure, this point is controversial: various philosophers have thought that higher-order quantification can be

understood on its own terms.[12] If these philosophers are correct, then Ramsey sentences may provide a solution. The debate cannot be settled here. But if one suspects, as I do, that one needs to understand relation-talk before being able to understand higher-order quantification (if the latter can be understood at all), then one must reject the higher-order response.

The neo-Fregean response is that we grasp e via an abstraction principle. Neo-Fregeans propose that we come to understand some expressions by knowing the truth-conditions for sentences in which those expressions occur. For example, one comes to understand what 'number' means by observing that the truth-condition of any instance of the schema 'the number of Fs is the number of Gs' is given by the corresponding instance of the schema 'there is a one-one correspondence between the Fs and the Gs', and one comes to understand what 'direction' means by observing that the truth-condition of any instance of the schema 'the direction of line a is the direction of line b' is given by the corresponding instance of the schema 'a and b are parallel'.[13] The neo-Fregean idea can readily be extended to expressions pertaining to all sorts of abstract objects. One might propose that one knows what 'instantiates' and terms for properties mean by observing that the truth-condition of any instance of the schema 'a instantiates $\Pi x \, \Phi x$' is given by the corresponding instance of the schema 'Φa', and that one knows what 'member of' and terms for sets mean by observing that every instance of the schema 'a is a member of $\{b_1, \ldots, b_n\}$' is given by the corresponding instance of the schema 'a is among $b_1, \ldots,$ and b_n'.[14] If the neo-Fregeans are correct about what it takes to grasp the meaning of expressions, then abstraction principles such as these make expressions like e intelligible.[15]

---

[12] See, for example, Agustín Rayo and Stephen Yablo (2001) and Crispin Wright (2007).
[13] Bob Hale and Crispin Wright (2002, p. 116). See also Øystein Linnebo (2018).
[14] Setting aside the difficulties that some pluralities do not constitute a set and that some predicates lack a corresponding property as a concession.
[15] A caveat should be made in that with the neo-Fregean's examples, only a predicate is supposed to be grasped by the observation of truth-conditions, whereas in my extrapolations, a new predicate and new terms are proposed to

Unfortunately, grasping truth-conditions is not sufficient for understanding. To explain why, I will introduce some new pieces of language and some principles that make use of them, including an abstraction principle. I will attempt to make it evident that there is no clear meaning to those pieces of language. Consider the novel predicate 'schmoogle'. I offer no definition of 'schmoogle'. Instead, I offer the following principles regarding schmoogling:

$\forall x \forall y[(\neg[x = y] \land [x$ and y are spatiotemporal$]) \to \exists z(z$ schmoogles x and $y)]$

$\forall x s[\exists y(y$ schmoogles the $x s) \to ($the $x s$ are two in number $\land$ the $x s$ are spatiotemporal$)]$

$\forall x \forall y \forall z(z$ schmoogles x and $y \leftrightarrow z$ schmoogles y and $x)$

$\forall x s \forall y s \forall z[(z$ schmoogles the $x s \land z$ schmoogles the $y s) \to ($the $x s = $ the $y s)]$

To put the principles in natural language: all and only pairs of spatiotemporal objects are schmoogled by something, being schmoogled by a given thing is symmetrical, and no two pairs of spatiotemporal objects are schmoogled by the same thing. Say that something is a *schmoogler* just in case there is some x and y that it schmoogles.

These principles license all sorts of inferences about schmooglers. For example, one can easily reason from the above principles to the following conclusions:

Something schmoogles this pair of turtledoves.

Since two schmooglers schmoogle the dogs owned by Adam, on the one hand, and the dogs owned by Sarah, on the other, there is at least one dog that Adam and Sarah do not both own. Moreover, Adam and Sarah both own a pair of dogs each.

Importantly, the principles allow for the following abstraction principle: the truth-condition of any instance of the schema 'a is schmoogled by $|b_1, b_2|$' is given by the corresponding instance of the

---

be grasped by them. A neo-Fregean might say that only one expression can be imbued with meaning at a time given this device. If so, however, the neo-Fregean suggestion I'm making might already be a non-starter. It is hard to see how to use an abstraction principle to explain what 'instantiates' means without making use of terms for properties or how to explain what a term for a property means without invoking instantiation. The same goes for pretty much any pair e and s meeting the conditions mentioned in the exposition of the Lewisian argument.

schema 'a is among b_1 and b_2, b_1 is spatiotemporal, and b_2 is spatiotemporal', taking the '|…|'

notation as a device for introducing terms for schmooglers.

The trouble is that neither the enumerated principles nor the abstraction principle say what

schmoogling is. They in no way provide enough content to answer someone who made the

following remark: "Yes, yes, well and good. I can see how schmooglers relate to spatiotemporal

things. But what is schmoogling, anyhow? What does 'schmoogle' or 'is schmoogled by' mean?

Can you define these expressions for me, either descriptively or ostensively? If they are vague

expressions, can you identify some of their precisifications? I don't know what we're talking about

here." This question seems entirely appropriate and wholly unanswerable from the principles, but

the principles are all one has. Schmoogling is a will-o'-the-wisp. It is not merely that the

phenomenon of schmoogling is not well understood and leaves many questions open. It's that the

principles I've provided aren't enough for there to even be a genuine notion of schmoogling in the

first place.[16] And for that reason, it is doubtful that a neo-Fregean theory of meaning is likely to

succeed, whatever the other merits of abstraction principles might be.[17]

A neo-Fregean might respond that our imagined interlocutor's discomfort with

'schmoogle' only stems from a lack of familiarity. If they can be confident that the schemas

express all the analytic truths about schmoogling, continually using 'schmoogle' might be enough

for their sense that it doesn't have a clear meaning to disappear. Perhaps so. However, that does

---

[16] This connects to the wider literature on incomplete stipulations. For example, Timothy Williamson (1997, pp. 222-227) and David Barnett (2008) offer interesting opinions about the relationship between meaning, incomplete stipulations, and those stipulations' associated expressions. General theories of how they interrelate could bear on the 'schmoogle' case, and intuitions about the 'schmoogle' case might in turn bear on the plausibility of those theories. This merits deeper engagement that ought to be pursued in the future. For now I will just remark that I am sympathetic to views like Barnett's and direct the reader toward my critique of some possible neo-Fregean rejoinders over the next several paragraphs.

[17] Of course, one could stipulate a definition or precisification of 'x schmoogles y and z' if one already has expressions for abstracta in the lexicon: perhaps 'y and z are members of x, y is spatiotemporal, and z is spatiotemporal'. The point is just that the principles given for schmoogling are not enough on their own.

not mean that the schemas are enough for 'schmoogle' to be meaningful. Suppose that after introducing 'schmoogle' to our language, over time we discover two new phenomena, which we can call *boogling* and *doogling*. Both boogling and doogling satisfy the axioms that govern schmoogling. However, boogling and doogling are not identical. We can now ask: does 'schmoogle' refer to boogling, does it refer to doogling, or does it refer to neither? We could not have added as an analytic truth that schmoogling is one or the other beforehand for the simple reason that boogling and doogling were undiscovered when 'schmoogle' was added to our language. And clearly, 'is schmoogling boogling?' and 'is schmoogling doogling?' are meaningful sentences if 'schmoogle' is meaningful (on the pretense that boogling and doogling are real and discovered).[18] The language must provide grounds for there to be some sort of answer to the questions they express. Even if we cannot know the answer, there must be an answer. But the language gives no such grounds. It doesn't even ground the answer that 'schmoogle' is vague: the analytic truths are too incomplete. No causal theory of reference applies to yield answers, since unlike H_2O and 'water', we cannot appeal to some sort of initial baptism.[19] So 'schmoogle' is not meaningful after all, despite familiarity producing comfort with its use. And since we cannot be certain that the schmoogling schemas are uniquely satisfied by some worldly phenomenon, we have no way of knowing that we are not in a boogle-and-doogle-like world. Thus, we have no way of knowing that 'schmoogle' is meaningful after all. Even if grasping truth-conditions is sufficient

---

[18] Note that the question of whether schmoogling is boogling is nominalistic. To see this clearly, it can be put into the idiom of Cian Dorr (2016): is it the case that for any *x* and *y*, for *x* to schmoogle *y* is for *x* is to boogle *y*? Likewise, the question of whether schmoogling is doogling is nominalistic.

[19] Saul Kripke (1980) provides the most well-known articulation and defense of the causal theory of reference.

for understanding some meaningful expression, it isn't sufficient for us to be confident that there is some meaningful expression there to be understood.[20]

At this point the neo-Fregean might respond that the difficulty is posed purely by the need for additional linguistic conventions to be introduced when 'schmoogle' is added to the language. The neo-Fregean might suggest adding the convention that if there is more than one relation that satisfies the schmoogle schemas, 'schmoogle' refers to the most natural one, and that if there is more than one most natural relation among the relations that satisfy the schemas, it doesn't refer to anything at all. Alternatively, the neo-Fregean might suggest adding the convention that if there is more than one such relation, 'schmoogle' has a supervaluationist semantics, and each relation is a precisification of 'schmoogle'. However, in order for these conventions to be helpful, 'relation' must already be understood. As we have already seen, 'relation' is one of the words the intelligibility of which is in question. Therefore, we need an account of how to understand it as well, and a neo-Fregean story that includes conventions that quantify over relations would be patently circular.[21] It is difficult to see how conventions adequate to handle these cases can be added that would not rely on expressions that the Lewisian argument has already put into doubt.[22] For that reason, it seems unlikely that this reply is enough to rescue the neo-Fregean view.

The neo-Fregean might instead respond by suggesting that general truths about how language works will decide whether 'schmoogle' refers to boogling, refers to doogling, has no referent, or is a vague predicate. Perhaps it is just a fact about how language works that in the

---

[20] The boogle and doogle case is reminiscent of the Julius Caesar problem for neo-Fregeanism. For an overview of the Julius Caesar problem, see Richard Kimberly Heck (1997). A fuller comparison of the boogle and doogle case and the Julius Caesar problem must await future work.

[21] This is not to say that if the Lewisian argument is sound, no supervaluationist account of vague expressions can be correct. It is just that how vague predicates get a supervaluationist semantics can't be via some explicit stipulation that involves quantifying over properties and relations.

[22] Higher-order quantifiers can do the trick, but once higher-order quantifiers are permitted the Ramsey sentence solution already becomes available.

absence of an explicit convention, 'schmoogle' will refer to whichever of boogling and doogling is the most natural if no other candidates are more natural than them. Or perhaps it is a general truth that if there is no explicit convention, 'schmoogle' will be vague. Or it might be a general truth that without an explicit convention, 'schmoogle' will refer to boogling or doogling at random. The trouble is that for the neo-Fregean to make their case, they owe us some story about why language works that way: for any of the options, it isn't simply obvious that language does so. But just as there was the difficulty that conventions added to settle questions about the connection between schmoogling, boogling, and doogling would rely on expressions the Lewisian argument puts in doubt, so there is the problem that the new story the neo-Fregean must tell will almost certainly involve those same expressions. It is difficult to see how, for example, a reference magnetist story for language can be told without speaking of properties and relations.[23] And yet in the absence of a story, a proposal regarding the general workings of language seems toothless. The Lewisian argument has radical consequences even for theorizing about semantics: much discussion about meaning in philosophy of language and linguistics cannot survive it. So the neo-Fregean cannot simply borrow from work done in these areas in order to supplement the neo-Fregean narrative. To do so would be to rely on theories of language many key expressions of which are subject to the Lewisian argument's critique, and consequently borrowing from them raises the specter of circularity once again.

The Moorean response to the Lewisian argument is to hold that it is simply an evident fact that language about abstract objects is genuinely meaningful. We use this language all the time, and we are remarkably successful in doing so: there is very little confusion in practice about how to reason about sets, properties, and the like. The Lewisian argument raises a genuine puzzle in

---

[23] See Lewis (1984) for a defense of reference magnetism.

that it makes it difficult to see how we are able to refer to abstract objects and think about them. There is perhaps much interesting philosophical work to be done here in explaining our capacities. But there is no genuine question as to whether we possess the ability to refer to and think about abstract objects: no argument could possibly be strong enough to undercut the obvious fact that we are able to do so. Some premise of the argument must be wrong.

I believe the Moorean position overestimates our epistemic situation vis a vis what we are doing when we seem to think about or refer to abstract objects. Consider again the position we are in with respect to 'schmoogle', 'is schmoogled by', and singular terms for smooglers: we are able to engage in a lot of apparent reasoning about schmoogling and schmooglers, but in the end, expressions for schmoogling lack meaning. My view is that this is the exact position we are in with respect to language about abstract objects. The axioms of Zermelo-Fraenkel set theory (ZF) give a large number of schemas governing sets, but they provide us with no definition of 'set' or 'membership' (one cannot even define 'set' in terms of membership: to say that a set is that which has members excludes the empty set). They enable us to reason in highly systematic ways about sets, but they are insufficient for giving the notions of sethood or membership meaning in the first place. Reasoning about sets is, in the final analysis, not truly reasoning about anything. It is merely apparent reasoning, an activity that looks very much like reasoning about some kind of object but which fails to latch on to any object at all. The same can be said of language governing properties and instantiation, although the axioms there are less clear.[24] Regardless, even if some extremely comprehensive axioms for properties and instantiation were proposed and agreed upon to be adequate to provide a total theory of properties, this would be insufficient to give 'property' and

---

[24] A plausible initial thought is that for every predicate there is a property, but such an approach leads to Russell's paradox. Besides, it is at least an epistemic possibility that there are some properties for which no predicate has ever or will ever be introduced.

'instantiation' meaning absent some definition of those expressions. And by the Lewisian argument, no definition is in the offing.

Our ability to engage in highly systematic apparent reasoning about abstract objects gives us the strong feeling that language about abstract objects is meaningful. This feeling is not irrational in its foundation. Typically, when we can reason about something in a way that is systematic and generates wide agreement across members in our linguistic community, the corresponding language is meaningful. Systematicity is strong evidence of meaning. That evidence, however, is ultimately defeasible when one examines the details of how that language develops and the systematicity and consensus surrounding its use is generated. The schmoogle case undercuts the evidential force that the systematicity and consensus surrounding the use of language about abstract objects provides in favor of that language's meaningfulness. 'Schmoogle' and 'schmooglers' are not meaningful, and common language about abstract objects is plausibly introduced in the same fashion as they are, albeit often implicitly. It is therefore not a Moorean fact that the language is meaningful. It is simply a deeply held feeling that one ought to hold in suspicion once one encounters the Lewisian argument and sees how the language could be used in such systematic ways even absent having a meaning.

5. The Lewisian Argument and Epistemic Arguments Against Platonism

It might be thought that the Lewisian argument does not differ substantially from epistemic arguments against platonism. This would be a mistake. It is important to see how they differ, and how responses that may be satisfactory to epistemic arguments do not address the Lewisian argument.

One epistemic argument raises a justificatory challenge. Given that abstract objects are causally isolated from us and outside of space (and plausibly time), how can we have justification

for believing in them? This challenge to platonism was prominently raised by Paul Benacerraf (1973) and has generated an extensive discussion. Several viable responses are available on behalf of the platonist. For example, a platonist could point out that Benacerraf's argument relies on a causal theory of knowledge, which is now widely rejected.[25] Moreover, it's plausible that belief in abstracta can be justified in the same way belief in any entities can be justified: they form essential parts of simple, explanatory, highly systematic theories.[26] Neither of these points make any headway on the Lewisian argument, however. Consider the point about Benacerraf's argument relying on the causal theory of knowledge. It is one thing to deny that a causal relationship between a subject and a fact is necessary for the subject to know that fact. It is another to deny that some causal or perceptual relationship must hold between a subject and an instance of a certain kind of phenomenon for the subject to be able to think about or refer to examples of that phenomenon or the phenomenon generally. Perhaps I needn't be causally related to the largest pink object to know that it is self-identical – maybe the largest pink object is not even in my spacetime, and so is absolutely causally isolated from me – but surely I need to have been causally or perceptually related to pink objects at some point in my life to know what 'the largest pink object is self-identical' means.[27] Now consider the point that belief in abstract objects can be justified by their essential theoretical roles. In order for a theory to be genuinely explanatory, it has to first be meaningful. And the meaningfulness of those putative theories is exactly what the Lewisian argument seeks to undermine. If a theory utilizes expressions for abstract objects, it is only meaningful if those expressions are meaningful. In fact, the problem runs even deeper: for a

---

[25] See Justin Clarke-Doane (2016) for discussion.
[26] Michael D. Resnik (1997, Ch. 7) pursues this approach, among others.
[27] At least – barring exceptional cases such as the one described in fn. 6 – this is true if pinkness is phenomenal pinkness rather than merely the disposition to reflect such-and-such wavelengths of light. But even if a reductionist account of color is correct, I must surely have been causally or perceptually related to light to know what 'the largest pink object is self-identical' means.

putative theory to even to be a theory, the expressions it utilizes must be meaningful. The Lewisian argument not only threatens the claim of some theories to be explanatory, but their claim to be theories at all.

A second epistemic challenge concerns not how belief in abstract objects can be justified, but how the platonist can explain the reliability of our beliefs about abstract objects. It would seem if our beliefs about abstract objects are correct, it is a giant coincidence: the abstract objects themselves played no role whatsoever in shaping the processes that govern what beliefs we form about them, so it is simply a brute fact that we are so accurate in our judgements concerning them. Accepting such a brute fact is highly undesirable, so we should avoid believing in abstract objects. Hartry Field (1989, pp. 25-30) has articulated this epistemic challenge to platonism.[28] One difficulty with this argument is that it proves too much: it generalizes so far beyond abstract objects that it strains credulity. For example, it is plausible that the modal profiles of various objects – that human beings couldn't be fried eggs, that trees couldn't be clouds, and so on – play no role in shaping the processes that govern our beliefs about what modal profiles various objects have. They certainly don't seem to be relevant to the physical interactions important for how we form beliefs. It then appears that if Field's objection motivates skepticism about abstract objects, a parallel objection motivates skepticism about modality. But there are reasons to doubt that skepticism about modality is ultimately plausible, such as those raised by Timothy Williamson (2007, pp. 134-178).[29] The Lewisian argument does not overgeneralize in this way. The objects which have the modal profiles are, after all, perceptible and causally interact with us. That Tom could not have

---

[28] Note that Field takes himself to be understanding Benacerraf's challenge in a new way. Clarke-Doane (2016) also discusses Field's challenge.

[29] A Field-style argument can also be used to motivate moral skepticism. See David Enoch (2010) for an exposition of this argument. Enoch proposes a solution that strikes me as unsatisfactory for reasons I cannot get into here. To my mind, that Field-style arguments can be generated in so many areas of inquiry is reason to think that something is wrong with the general demand for explanation that Field-style arguments reflect in the first place.

been a fried egg might be irrelevant to how my belief that he couldn't have been was formed, but at least one can perceive and interact with Tom. What motivates the Lewisian argument is not a crass empiricism that recognizes only raw sense experience as a basis for concept formation, but an incredulity at the idea that one can grasp special, kind-determining features of things totally isolated from us.

Regardless of the above concern about the generality of Field's challenge, some philosophers have addressed it. A prominent response is that the fact that abstract objects necessarily both exist and have their distinctive features is sufficient to explain the reliability of our beliefs about them. Given that abstract objects exist and necessarily are how we believe them to be, there is no question as to why our beliefs are reliable: given their content, our beliefs had to be accurate. There is no coincidence here to be explained. To be sure, whether this response is satisfactory is contested.[30] But if it is satisfactory, it is still of no help in addressing the Lewisian argument. The Lewisian argument does not question how it is we come to form reliable beliefs about abstract objects, but whether we have beliefs about them in the first place. If one cannot refer to or think about abstract objects, one can hardly hold any beliefs about them. If the Lewisian argument is sound, what appear to be beliefs about abstract objects are, in the final analysis, pseudo-beliefs. They are psychological states with so much content missing that they do not rise to the level of belief. The fact that abstract objects both exist and are how they are of necessity in no way explains how it is we are able to form beliefs about them. That point holds true even if it might explain how our beliefs about them would be reliable were we able to form such beliefs.

---

[30] For example, Lewis (1986, pp. 111-112) himself takes this line, but Field (1988, pp. 62-67) criticizes his position.

The upshot of the preceding discussion, then, is that various responses to the epistemic arguments against platonism do not undermine the Lewisian argument. If a platonist is to undermine it, they must find some new avenues by which to question it.

6. Nominalism and the Lewisian Argument

The Lewisian argument interacts with nominalism in two ways. First, it strongly motivates it. Second, it simultaneously undermines some nominalist accounts of abstract object discourse. Regarding the first point, as we observed, van Inwagen replied to Lewis' critique of magical ersatzism with a *tu quoque*. Lewis' response only provides a comprehensive account of abstract objects if his widely rejected modal ontology is correct. In contrast, nominalists are not subject to this *tu quoque* in the first place. Nominalists can simply deny that we can talk or think about any kind of abstract object. Magical ersatz possible worlds and sets are just the tip of the iceberg. Strictly speaking, we also can't refer to or think about propositions, properties, relations, numbers, so on and so forth. However useful abstract language is, it is useful nonsense.[31]

The Lewisian argument thus motivates nominalism in a very straightforward way. If the argument is sound, no one can think about or refer to abstract objects. If no one can think about or refer to abstract objects, then no one can use abstract objects to offer genuine theories of phenomena, whether they are explanatory or purely descriptive. And if abstract objects cannot be put to legitimate theoretical uses, there is no reason to believe they exist: indeed, since one can't think about or refer to them, one *can't* believe they exist. Nominalism is not only true – it's the only game in town.

---

[31] The Lewisian argument that I proposed above, of course, makes use of expressions in the sense of expression-types. It thus seems to presuppose some form of platonism about expressions. Happily, this is just convenient shorthand: it would be a simple matter to reframe the argument in terms of token expressions, albeit doing so would make the argument unnecessarily laborious to follow.

However, the Lewisian argument does more than simply motivate nominalism: it also makes the nominalist position much harder to defend against the platonist argument that we need abstract objects to account for the usefulness of the things we say about them. For nominalism to be a truly satisfying philosophical position, nominalists need to provide some account of why utilizing language about abstract objects in our (putative) theories never leads us to falsehoods about concrete objects and is so helpful when reasoning our way to conclusions about the concrete. It is hard to imagine doing physics without applied mathematics, or biology without talking of the anatomical features of various animal species. A particularly attractive way of attempting this is by appealing to what would be the case were abstract objects to exist. Cian Dorr (2008, pp. 36-37) explores this approach in more detail. He proposes paraphrasing any sentence about abstract objects, S, as follows:

(There are numbers, features, etc. $\wedge$ concrete things are how they actually are) $\square\rightarrow$ S

Or, to ensure that abstract objects would behave how we expect, the paraphrase can be put as

([Insert the conjunction of the various axioms of our best theories of abstract objects] $\wedge$ concrete things are how they actually are) $\square\rightarrow$ S

This strategy for paraphrasing claims about abstract objects has the advantage of being amenable to a proof of inferential safety that has been put forward by Lukas Skiba (2019): it can be shown that replacing every sentence about abstract objects with its paraphrase under this strategy never leads to false conclusions about concrete objects. The usefulness of applied mathematics, the abstract portion of biology, and all the rest is assured because their nominalistically-acceptable paraphrases entail the same conclusions about concrete objects.

Unfortunately, if the Lewisian argument is sound, this strategy cannot be embraced by nominalists. The argument entails that language about abstract objects is meaningless. If language

about abstract objects is meaningless, the antecedents of Dorr's paraphrases are meaningless. The paraphrases are, then, meaningless themselves, and it is surely a constraint on paraphrase that a successful paraphrase must be meaningful. Whatever nominalists do, they cannot explain the inferential safety of claims about abstract objects by a superficial adaptation of them.[32]

7. Conclusion

The Lewisian argument is a powerful argument against platonism that proceeds from plausible assumptions about concept acquisition. If it is sound, no one is able to refer to or think about abstract objects, what seem to be successful theories that invoke abstract objects turn out to be collections of meaningless statements, and platonism is metaphysically incoherent. Platonists must respond to this argument not only to show that their view is viable, but to show that they have a view in the first place.

Acknowledgements

---

[32] One might hope to get around this issue by interpreting the claims about abstract objects in the paraphrases through the lens of Lewis: that is, by taking the paraphrases as statements regarding what would obtain were (actual) concrete things how they actually are, modal realism true, Lewis right about properties and the like, and set theory reconstructed in the manner of the Appendix of *Parts of Classes*. Taken this way, no primitive abstract expressions would occur in the paraphrases when fully articulated. But this would raise several problems, not the least of which is that it is hard to evaluate a claim of the form "(actual concrete things are how they actually are ∧ modal realism is true ∧ [insert the ideas of Lewis regarding sets and other abstract objects]) □→ P". What does one make of a claim like that? More simply, how does one evaluate a claim of the form "(actual concrete things are how they actually are ∧ modal realism is true ∧ Q) □→ P", given that modal realism is false? Among other things, it seems at least epistemically possible that there is more than one actual spacetime, a possibility inconsistent with modal realism.

References

Armstrong, D. (1989). *Universals: An Opinionated Introduction*. Boulder, CO: Westview Press.

Barnett, D. (2008). Indeterminacy and Incomplete Definitions. *Journal of Philosophy*, 105(4): 167-191.

Benacerraf, P. (1973). Mathematical Truth. *Journal of Philosophy*, 70(19): 661-679.

Clarke-Doane, J. (2016). What Is the Benacerraf Problem? In Pataut, F (Ed.), *Truth, Objects, Infinity: New Perspectives on the Philosophy of Paul Benacerraf* (pp. 17-43). Paris, France: Springer.

Dodd, J. (2007). *Works of Music: An Essay in Ontology*. New York, NY: Oxford University Press.

Dorr, C. (2008). There Are No Abstract Objects. In J. Hawthorne, T. Sider, and D.W. Zimmerman, *Contemporary Debates in Metaphysics* (pp. 32-63). Malden, MA: Blackwell Publishing.

Dorr, C. (2016). To Be F Is To Be G. *Philosophical Perspectives* 30(1): 39-134.

Enoch, D. (2010). The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope With It. *Philosophical Studies*, 148: 413-438.

Field, H. (1988). Realism, Mathematics, and Modality. *Philosophical Topics*, 16(1): 57-107.

Field, H. (1989). *Realism, Mathematics, and Modality*. Oxford, UK: Blackwell.

Friedell, D. (2020). Abstracta Are Causal. *Philosophia*, 48: 133-142.

Hale, B., and Wright, C. (2002). Benacerraf's Dilemma Revisited. *European Journal of Philosophy*, 10(1): 101-129.

Heck, Richard Kimberly. (1997). The Julius Caesar Objection. In R. Kimberly Heck (Ed.), *Language, Thought, and Logic: Essays in Honour of Michael Dummett* (pp. 273-308). New York and Oxford: Oxford University Press.

Kripke, S. (1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.

Lewis, D. (1984). Putnam's Paradox. *Australasian Journal of Philosophy*, 62(3): 221-236.

Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell Publishing Ltd.

Lewis, D. (1991). *Parts of Classes*. Cambridge, MA: Basil Blackwell Ltd.

Linnebo, Ø. (2018). *Thin Objects*. Oxford, UK: Oxford University Press.

Maddy, P. (1990). *Realism in Mathematics*. Oxford, UK: Oxford University Press.

Nolan, D. (2020). It's a Kind of Magic: Lewis, Magic and Properties. *Synthese*, 197(11): 4717-4741.

Rayo, A., and Yablo, S. (2001). Nominalism Through De-Nominalization. *Noûs*, 35(1): 74–92.

Skiba, L. (2019). Fictionalism, the Safety Result and Counterpossibles. *Analysis*, 79(4): 647-658.

Resnik, M. D. (1997). *Mathematics as a Science of Patterns*. Oxford: Oxford University Press.

Williamson, T. (1997). Imagination, Stipulation and Vagueness. *Philosophical Issues*, 8: 215-228.

Williamson, T. (2007). The Philosophy of Philosophy. Malden, MA: Blackwell.

Wright, C. (2007). On Quantifying Into Predicate Position. In M. Leng, A. Paseau, and M. Potter, *Mathematical Knowledge* (pp. 150-174). Oxford, UK: Oxford University Press.

Van Inwagen, P. (1986). Two Concepts of Possible Worlds. *Midwest Studies in Philosophy*, (11): 185-213.

Van Inwagen, P. (2006). Names for Relations. *Philosophical Perspectives*, 20: 453-477.