

8-2016

Contra Lewisian Naturalness

Dylan Abney

University of Arkansas, Fayetteville

Follow this and additional works at: <http://scholarworks.uark.edu/etd>



Part of the [Metaphysics Commons](#)

Recommended Citation

Abney, Dylan, "Contra Lewisian Naturalness" (2016). *Theses and Dissertations*. 1736.
<http://scholarworks.uark.edu/etd/1736>

This Thesis is brought to you for free and open access by ScholarWorks@UARK. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of ScholarWorks@UARK. For more information, please contact scholar@uark.edu.

Contra Lewisian Naturalness

A thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Arts in Philosophy

by

Dylan Abney
New York University
Bachelor of Arts in Philosophy, 2012

August 2016
University of Arkansas

This thesis is approved for recommendation to the Graduate Council.

Dr. Eric Funkhouser
Thesis Director

Dr. Jack Lyons
Committee Member

Dr. Barry Ward
Committee Member

Abstract

Some ways of talking or thinking about the world are better than others. Most obviously, it is often better to say or believe true things rather than false things. Perhaps less obvious is the notion that our speech and thought ought to, or often in fact *does*, reflect the natural formation or structure of the world. This idea—that we ought to be carving the world at its natural joints—can be found at least as far back as Plato’s *Phaedrus*. More recently, we can see a related idea in the work of David Lewis. In “New Work for a Theory of Universals,” Lewis argued that there is a special set of properties—the *natural properties*—that carve the world at its joints. In Lewis’s sense, the notion of a natural property is a complex one: it plays a number of conceptual and explanatory roles. In this thesis, I first detail some important background on properties in general. I then shift my focus to Lewis’s natural properties and the various roles he had in mind for them. In doing so, I pay particularly close attention to the supposed connection between natural properties and the metasemantic notion of reference magnetism. I consider the question of whether there exists a unique set of properties that can play all of Lewis’s roles. I argue that there is no such set. I then conclude by arguing that there are good reasons to deny that there are natural properties in the Lewisian sense.

Acknowledgments

First, I would like to thank Dr. Eric Funkhouser for his consistent insight and guidance over the past year. He pushed me to make this project into something much better than it otherwise would have been. I am very grateful for that. I would also like to thank Drs. Jack Lyons and Barry Ward for serving on my committee. Thanks are also due to David Cajias, Zach Biondi, and Max Gatyas—many of our conversations over the topics in this thesis were a big help. Finally, I would like to thank my brother, Beck Abney, for his assistance on the images and figures that are featured throughout.

Dedication

For Mom and Dad.

Table of Contents

Section 1: Introduction	1
Section 2: Properties	3
2.1 – Basics	3
2.2 – Traditional roles and historical context.....	6
Section 3: Natural Properties.....	12
3.1 – Understanding natural properties	12
3.2 – Methodological interlude	13
3.3 – Naturalness and its roles	15
3.4 – Carving the world at its joints.....	18
3.4.1 – Two types of joint-carving.....	19
3.4.2 – Interesting cases and other potential issues	22
Section 4: Problems for Natural Properties	28
4.1 – Fundamentality + Necessity + Joint-Carving	28
4.2 – Simplicity + Similarity + Joint-Carving	30
4.3 – Reference magnetism.....	34
4.3.1 – Magnetism + Similarity	38
4.3.2 – Magnetism and other roles.....	39
4.4 – Substantivity and the merely verbal.....	41
4.4.1 – Introduction to substantivity	41
4.4.2 – Substantivity + Joint-Carving	44
Section 5: Conclusion.....	47
Works Cited.....	50

Table of Figures

Figure 1 – Sider’s red-blue world.....	23
Figure 2 – Carving the red-blue world at its joints	23
Figure 3 – A slightly-off carving of the red-blue world.....	24
Figure 4 – Electrons in the actual world and a nearby possible world.....	29
Figure 5 – Joint-carving₁ at the fundamental level	41

Section 1: Introduction

Some ways of talking or thinking about the world are better than others. Most obviously, it is often better to say or believe true things rather than false things. Perhaps less obvious is the notion that we should, or often in fact *do*, talk in such a way that divides (or “carves”) the world up according to its *natural formation*, its *joints*, or its true *structure*. In the *Phaedrus*, Plato’s Socrates is not shy about his preference for this way of talking:

“Now I myself, Phaedrus, am a lover of these processes of division and bringing together, as aids to speech and thought; and if I think any other man is able to see things that can naturally be collected into one and divided into many, him I follow after and ‘walk in his footsteps as if he were a god.’”¹

It is true that when we group things together in a natural way—when we carve the world at its joints, so to speak—we are able to make generalizations that we otherwise could not. Thus, it is good to have a word like ‘tiger’ that groups together all and only tigers—not some haphazard, miscellaneous assortment of objects. This particular word allows us make generalizations like ‘tigers have four legs’. Of course, there may be a few three-legged tigers roaming around the world, but the point is that the generalization remains extremely useful.

More recently, there have been systematic defenses of dividing up the world in a natural way. In particular, the twentieth-century philosopher David Lewis argued that there are special properties—the *natural properties*—that carve the world at its joints. He distinguished the natural properties from the *unnatural properties*, those that “carve reality at the joints—and everywhere else as well.”² In Lewis’s eyes, the natural properties helped not only with things

¹ *Phaedrus*, 266b.

² Lewis (1983, p. 346).

like generalization; he thought they could do a lot of other interesting work as well. For example, Lewis thought the natural properties were the properties whose sharing made for genuine resemblance. The things that have the property *being a tiger* resemble one another more than the things that have the property *being either a tiger or a lion*. This is because the former property is more natural than the latter.

This role of explaining resemblance is just one among many that Lewis had in mind. In this thesis, I consider the question of whether there exists a set of properties that can play all of the various roles that Lewis intended natural properties to play. I argue that there is ultimately no such set.

I proceed in the following way. In Section 2, I lay out some important conceptual and historical background on properties. In Section 3, I detail the various conceptual and explanatory roles that Lewis intended natural properties to play. In Section 4, I raise a number of problems for the joint-satisfaction of these roles. Though I consider a number of problematic sets of roles, I pay particularly close attention to natural properties and their supposed connection to the metasemantic notion of reference magnetism, the idea that natural properties are more “eligible” referents for predicates. I conclude, in Section 5, by arguing that there are good reasons to deny the existence of natural properties as Lewis conceived of them.

Section 2: Properties

2.1 – Basics

Suppose I utter the following sentence: ‘The Empire State Building is tall’. When I say this sentence, it is natural to interpret me as attributing a *property* (or characteristic, or quality) to an object. I am using the expression ‘The Empire State Building’ to refer to, or to *pick out*, a particular object in the world: the building that is located at the southwest corner of Fifth Avenue and 34th Street in New York City. And we might say that I am using the predicate ‘is tall’ to attribute the property *being tall* to this object that I have picked out. A similar phenomenon presumably occurs with the utterance of any declarative sentence.³ So, when one says ‘The Empire State building is in New York City’, for instance, this involves the predication of the property *being in New York City* to a particular object, the Empire State Building.

While some predicates may attribute properties to single objects, others may instead relate multiple objects. For example, suppose I say ‘The Empire State Building is taller than the Chrysler Building’. When I utter this sentence, it is natural to interpret me as relating two objects: the Empire State building and the Chrysler Building. I am using the expressions ‘The Empire State Building’ and ‘the Chrysler Building’ to pick out two distinct objects, and I am using ‘is taller than’ to predicate the *being taller than* relation between these two objects. It is common to call such a relation *dyadic*, as it relates two objects. In principle, however, a relation could relate more than two objects. For example, if I say ‘One World Trade Center is taller than

³ A notable exception would be a sentence like ‘Redness is a color’, which involves the predication of a property *to another property* rather than predication of a property to an object. If one wants to say that properties are themselves objects, then perhaps this is not a genuine exception. There are reasons, however, to distinguish *substances* from *properties*, and I discuss this in more detail later on. At any rate, we should at least recognize this phenomenon, which we might call *second-order predication*.

both the Empire State Building and the Chrysler Building’, I am, of course, relating three objects. I am using a *triadic* relation: *x is taller than y, z*. On a related note, it is also common to think of relations themselves as having properties. For example, the dyadic relation *being taller than* has the property of asymmetry: if *x* is taller than *y*, then it is not the case that *y* is taller than *x*. Compare this to the *being a sibling of* relation, which is symmetric: if *x* is the sibling of *y*, then *y* is the sibling of *x*.

To summarize: properties may be predicated of single objects, and relations may be predicated between more than one object. In what follows, whatever I say of properties may also be said of relations; so I will mostly speak in terms of properties. And whenever I speak of properties, I will typically denote them with an italicized expression like ‘*being F*’, though it is also common practice to denote them with names like ‘*F*-ness’ or ‘*F*-hood’.

It is important to distinguish predicates from properties. A predicate like ‘is tall’ is a linguistic entity—it can be written on paper or spoken aloud. A property, on the other hand, is another kind of entity. It is not obvious exactly *what kind* of entity a property is, as properties seem to be much different in kind from ordinary objects like buildings. Indeed, it is not obvious whether properties truly exist *at all* for this same reason. (I will say more about this topic—the ontological status of properties—shortly.) For now, though, one should at least recognize the distinction between a predicate and a property.

A single predicate may be used to attribute multiple properties. For example, one linguistic community may use ‘is a sport’ to mean *being a competitive outdoor game between two teams*, and another might use the same predicate to mean *being a game featured on ESPN*.⁴

⁴ Both of these uses are more restrictive than our ordinary notion of sporthood, but the point is that a linguistic community just as easily *could* mean these things by ‘is a sport’ if it chose to. Also worth noting: one might claim that we are really dealing with two predicates here—that

Similarly, a single speaker may use the same predicate to attribute different properties in different contexts. For example, one speaker might use ‘is late’ to mean *arriving later than the scheduled time*; but the same speaker, in a more relaxed context, might use ‘is late’ to mean *arriving later than five minutes after the scheduled time*.

Conversely, multiple predicates may be used to attribute a single property. This phenomenon is perhaps most vivid when we think of examples involving more than one language. English speakers use ‘is blue’ to attribute the same property that Spanish speakers attribute with ‘es azul’. The phenomenon can be intralinguistic as well. ‘Is the number two’ and ‘is the immediate successor of one’ could be used to attribute the same property: *being the number two*.

Now that we have established the difference between predicates and properties, we may still be curious about what type of entities properties are and whether they truly exist at all. This question is historically significant. Traditionally speaking, there are two broad answers: Realism and Nominalism.⁵ Realists argue that properties are real, mind-independent entities.⁶ It is often argued, however, that properties cannot be the same type of thing as concrete *particulars*, like rocks, tables, and chairs. You can throw a rock, but you cannot throw a property. Thus, some (but not all) realists take properties to be *abstract* (as opposed to *concrete*) entities, much like one might think of mathematical entities like numbers. This picture of properties is influenced by Plato and his Forms, which were neither located in the mind nor the external, spatio-temporal world. (Philosophers often speak of the Forms as residing in Plato’s

there is a unique linguistic entity for each meaning. If this is the case, then we no longer have a scenario where a single predicate is used to attribute multiple properties.

⁵ For more on this debate, see Quine (1948) and Swoyer (2016).

⁶ I say ‘properties’ instead of ‘universals’ here to count trope theorists (who take properties to be abstract particulars) as realists.

Heaven.) Other realists, like David Armstrong, reject this picture of properties as abstract objects and want to “bring universals down to earth.”⁷ Indeed, some realists may speak of properties as *really existing* out in the physical world, and they may speak of two objects as sharing one and the same property. A realist may even speak of a property’s being *wholly present* wherever it is instantiated by an object. The question of whether a property is wholly present wherever it is instantiated is analogous to other traditional metaphysical questions. For example, there is an interesting question about the persistence of ordinary objects like chairs: are ordinary objects *wholly present* at any given time? Endurantists will answer Yes to this question; perdurantists will answer No, that only a temporal part is present at any given time.

Nominalists, on the other hand, are skeptical of this kind of talk about properties. A nominalist might either be skeptical of universals (and their being wholly present in multiple spatial locations), or the nominalist may simply be skeptical of the existence of abstract objects altogether. I will not go into more detail here, but hopefully this gives the reader a taste of the dispute between realists and nominalists.⁸

2.2 – *Traditional roles and historical context*

Properties are of particular philosophical interest because they can arguably play a wide range of explanatory roles. That is, properties can explain philosophically interesting phenomena. In the philosophy of language, for instance, we are interested in the question of how sentences and sub-sentential expressions (e.g., names and predicates) relate to the world, and we are interested in understanding how sentences can be meaningful. To begin to answer these

⁷ Armstrong (1989, p. 199).

⁸ If you are skeptical of this dispute, you may be in good company. Carnap (1950) famously argued that the question of whether abstract entities exist was a “pseudo-question,” one whose answer depends merely on the linguistic framework we choose to adopt.

questions, philosophers often speak of *compositionality*: the meaning of a sentence has something to do with the meanings of its constituent sub-sentential expressions. Consider the sentence ‘John is happy’. The meaning of this sentence has something to do with the meaning of the name ‘John’ and the predicate ‘is happy’. The meaning (or *semantic value*) of ‘John’ has something to do with the thing it refers to, the actual man who is out there in the world.⁹ But what should we say about ‘is happy’? One answer is that the semantic value of a predicate is a property—in this case *being happy*. As David Lewis (1983, p. 348) puts it: “It is properties we need [...] to provide an adequate supply of semantic values for linguistic expressions.”¹⁰

It is worth noting, however, that this view of properties as semantic values works best if properties are *abundant*. In “New Work for a Theory of Universals,” Lewis treats properties as classes, or sets, of objects. The property *being happy* is just the set of all happy things. In other words, to say that John is happy is to say that he is a member of the set of happy things. On this type of view, it is natural to say there are an infinite number of properties, as there are an infinite number of sets.¹¹ To get a sense of the abundance of properties on this kind of picture, consider the following properties: *being one day old, being two days old, being three days old*, etc. Also consider disjunctive properties: *either being in Arkansas or being one day old, either being in New York or being one day old, either being in Arkansas or being two days old*, etc.

⁹ There is an historically controversial debate in the philosophy of language about whether the meaning of a name is *exhausted* by its referent, but I will not go into detail here. For answers in the affirmative, see Mill (1882), Kripke (1980), and Salmon (1989). For an alternative answer, see Frege (1952).

¹⁰ Though properties can play this semantic role, this is not a particularly good argument for their existence. As Armstrong (1989, p. 199) argued, this so-called “argument from meaning” commits a version of Gilbert Ryle’s ‘Fido’-Fido fallacy: it assumes that since some predicates successfully refer to properties, *all* predicates must successfully refer to some property.

¹¹ Indeed, in the mathematical sense, there are *uncountably* many properties (if properties are merely sets). Lewis (1983, p. 346) amusingly speculates that the number is somewhere around beth-3.

While David Lewis treated properties as abundant, David Armstrong treated them as much more *sparse*.¹² To understand Armstrong's motivations here, we should turn to another role that properties have traditionally played in philosophy.

One might think that properties can explain qualitative similarity or resemblance. After all, surely the fact that one object is similar to (or resembles) another has to do with a property that the objects share. A baseball resembles a softball (at least in part) because both share the property of *having seams*. And so a crude test for degree of qualitative resemblance might be something like the following: the degree of resemblance between any two objects is a function of the number of properties the objects share.

It is natural to think qualitative resemblance has *something* to do with properties, but there is a problem for this crude test: any two entities will share an infinite number of properties, and any two entities will be divided by an infinite number of properties.¹³ This is problematic because it implies that all objects are equally similar (and dissimilar) to one another—an unpalatable result. We want to be able to say that a baseball is more similar to a softball than it is to a basketball, but the crude test will not get us this result.

To see how the crude test fails, consider an example. A baseball has the property *being white*, and a basketball has the property *being orange*. The baseball and basketball thereby share the disjunctive property *being either white or orange*. And they share the property *being white or orange or on Pluto*, even if neither object is on Pluto. And they share the property *being white or orange or a pencil*, even though neither is a pencil. The list goes on, *ad infinitum*. And

¹² See Armstrong (1978).

¹³ As Sider (2011, p. 3) puts it: “[C]onsider two objects x and y . Where F_x and F_y are any features of x and y , respectively, x and y share the feature: *being either F_x or F_y* . And they share the feature *being either F_x or F_y or 1 kg mass*. And they share the feature of *being either F_x or F_y or 2 kg mass*. And so on.”

despite all of these shared properties, the baseball does not thereby become extremely similar to the basketball.¹⁴ The crude test implies that a unicorn and a softball resemble a baseball to an equivalent degree.

So how should we answer this problem? One answer is ontological: just restrict the number of properties that truly exist. In *A Theory of Universals*, David Armstrong argues against the existence of disjunctive properties like *being white or orange or on Pluto*.¹⁵ “Real” or “genuine” properties, in Armstrong’s eyes, are the ones that make for similarity, and disjunctive properties of the form *being F or G* do not make for similarity. Armstrong’s point has not gone uncontested, however. Electrons and cows do not closely resemble each other, but is it not the case that they resemble each other in their sharing of the property *being an electron or a cow*?¹⁶ After all, they both exemplify electron-or-cow-hood, some will argue.

Such examples have driven philosophers like Armstrong to take a step further. Just as there are “real” or “genuine” properties, there is also “real” or “genuine” similarity. Perhaps electrons and cows resemble one another in some respect, but they do not *genuinely* resemble one another—or so some will claim.¹⁷ An interesting consequence of this type of view is that some natural language predicates (like ‘is either an electron or a cow’) will simply not have corresponding properties that can serve as semantic values.¹⁸

¹⁴ One can make a similar move for the properties any two object do *not* share.

¹⁵ See Armstrong (1978, pp. 19-29).

¹⁶ See Sider (2011, p. 2), who mentions this argument but does not endorse it.

¹⁷ Or if they *do* genuinely resemble one another, it must be in some other respect. Perhaps cows and electrons share the property of *being roundish* and resemble one another in that way.

¹⁸ For more on this point, see Armstrong (1978, p. 20). It is also worth mentioning a predicate’s reference to any property will depend on a context of utterance. In principle, a predicate that is grammatically disjunctive could still be used to refer to a property, even if one rejects the existence of disjunctive properties.

David Lewis answered the problem in a different way. His answer was not ontological in the sense that he did not argue against the existence of certain properties—he allowed for an infinite number of properties. Instead, in “New Work for a Theory of Universals” Lewis argued (in response to Armstrong) that there was a special subset of the properties—the *natural properties*—that could take on the role of explaining genuine similarity or resemblance. In some sense, *being a cow* is just *more natural* than a disjunctive property like *being an electron or a cow*.

In fact, Lewis thought that natural properties could play many more roles, too, that properties *simpliciter* could not. For him, natural properties made for similarity, served as a minimally complete supervenience base for the rest of the (classes of) properties, figured into scientific laws, were the properties that carved reality at its joints, and were privileged semantic values that were more eligible for reference.¹⁹ And there were other roles, too. (More on this in Section 3.) Distinctions between types of properties, like Lewis’s between natural and unnatural properties, are common in philosophy, and they can be used to explain many other phenomena. For example, the distinction between *intrinsic* and *extrinsic* properties (which can help explain duplication of objects) is of considerable philosophical interest, as is Locke’s distinction between *primary* and *secondary* properties. Primary properties are objective properties that objects have *in themselves*, while secondary properties are subjective and mind-dependent. The focus of my present discussion, however, will be around Lewis’s notion of a natural property.

More recently, there have been defenses of Armstrong’s and Lewis’s explanations of genuine similarity. For example, in his recent book *Writing the Book of the World*, Ted Sider argues for what he calls a “realism about structure,” which involves a defense of this notion of

¹⁹ See Lewis (1983).

genuine similarity. Lewis argued that natural properties could play many roles, and Sider's book extends this project. First, Sider thought natural properties (or as he often calls it, structure) could play even *more* roles than Lewis indicated. For instance, he thought they could be used to understand the philosophical notion of *substantivity*. There is an intuitive difference between questions like 'Is the Pope a bachelor?' and questions like 'Do electrons repel one another?' The answer to the first question seems to be a matter, perhaps, of what we mean by 'is a bachelor'. If we mean *being an unmarried male*, the Pope counts as a bachelor. If, on the other hand, we mean *being an unmarried male who is eligible for marriage*, then the Pope will not count as a bachelor. Thus, the answer to this question appears to depend on merely *verbal* or *conceptual* matters. The question of 'Do electrons repel one another?' appears to less of a verbal or conceptual matter and more of a *worldly* matter—in this sense, it is a *substantive* matter. What explains the difference between these two questions? Sider says that natural properties can help. (More on this in Section 4.)

Sider argued that natural properties could play other roles, but he also extended Lewis's project in another way. Sider uses his notion of *structure* to apply not only to predicates (which we can say are natural if they refer to natural properties), but also to other grammatical categories: to quantifiers (like the existential quantifier—'∃') and logical operators (like the wedge for disjunction—'∨'). In Sider's words, he goes "beyond the predicate."²⁰

In the next section, I take a deeper look at Lewis's notion of a natural property, and I say much more about the various roles he thought they could play.

²⁰ See Sider (2011, p. 85). Sider applies the idea of structure to many other topics—namely, ontology and logic. In these realms, genuine similarity is not at issue.

Section 3: Natural Properties

In this section, I begin with a brief explanation of what Lewis meant by ‘natural’. I follow that up with some methodological points, and I conclude by discussing the various roles that Lewis thought natural properties were well-suited to play.

3.1 – Understanding natural properties

Obviously there is some difference between *being green* and *being green or blue*. Perhaps we might say that the former is more *natural* than the latter. But what do we mean by this? Is it just that the expression ‘*being green or blue*’ is less simple than ‘*being green*’? Surely not. The way we specify properties in English is merely a contingency. We could just as easily talk about the property of *being green or blue* by using ‘*being grue*’²¹ instead—and even though the word ‘grue’ is shorter than the word ‘green’ this would not make *being grue* more natural than *being green*. Lewis had something different in mind by naturalness. He took it to be mind-independent and objective.

As we saw above, *being green* is more natural, in the Lewisian sense, than *being green or blue* because the sharing of *being green* makes for genuine similarity. But there is more: *being green* also carves the world at its joints better than gerrymandered properties like *being green or blue* or *being-a-person-or-a-desk*. (I will say more on what this means in 3.4.) Lewis also thought there were *perfectly* natural properties that were fundamental, properties that carved *perfectly* at the joints. *Being an electron*, for instance, could turn out to be one of these perfectly natural properties. So while there will be a limited number of perfectly natural properties, Lewis allows for there to be a spectrum of sorts, between those properties that are perfectly natural, and those properties that are highly unnatural and gerrymandered:

²¹ I should note here that I am using ‘grue’ in a different sense from Goodman (1955).

“Probably it would be best to say that the distinction between natural properties and others admits of degree. Some few properties are *perfectly* natural. Others, even though they may be somewhat disjunctive or extrinsic, are at least somewhat natural in a derivative way, to the extent that they can be reached by not-too-complicated chains of definability from the perfectly natural properties.”²²

This should give the reader an idea of what Lewis had in mind when he spoke of natural properties. And after discussing naturalness and its roles below (in Section 3.3), the reader should have an even better understanding.

3.2 – *Methodological interlude*

An interesting thing about naturalness is that it does not have a straightforward analysis or definition. Lewis treats ‘natural’ more like a theoretical term, similar to terms that are featured in physics, like ‘mass’, ‘force’, or ‘electron’. So one should not be overly skeptical about the lack of definition here. When defending his notion of *structure*, an extension of Lewisian naturalness, Sider (2011, p. 9) says, “Let’s be realistic about the extent and value of definitions. Philosophical concepts of interest are rarely reductively defined. Still more rarely does our understanding of such concepts rest on definitions. [...] We generally ‘understand’ philosophical concepts to the extent we know what role they play in our thinking.”

Without a straightforward definition for ‘natural’, our task becomes the following: we must determine whether there is some unique set of properties that satisfies all the roles that Lewis intended naturalness to satisfy. This is the approach we see in Dorr and Hawthorne’s (2013) recent paper, “Naturalness.” Dorr and Hawthorne (D & H) consider all of the various roles Lewis had in mind, and they consider potential problems for the joint satisfaction of these

²² Lewis (1986, p. 61).

roles. D & H's approach is virtuous in one crucial respect: by focusing on the theoretical roles instead of questions involving the term 'naturalness' (or 'natural') itself, they are able to steer clear of any potential merely verbal disputes.²⁴ D & H do not take sides in the debates between what they call 'naturalness enthusiasts' and 'naturalness skeptics', but I think they certainly achieve their ultimate goal of bringing "some structure to the terrain."²⁵

Another methodological point that arises from D & H's paper concerns the treatment of the metasemantics of theoretical terms. Our ultimate question is whether there is a unique set of properties that satisfies the roles Lewis had in mind. If such a set exists, then 'natural' has an extension. But exactly how many roles need to be satisfied in order for us to grant that 'natural' refers? Probably not all of them—only a 'draconian metasemantics' of theoretical terms, as D & H say, would require this. But where do we draw the line?²⁶ While we may want to reject the so-called draconian metasemantics that requires 100% role-satisfaction, an arbitrary cut off of 75% role-satisfaction may be just as objectionable. Surely we will want to be sensitive to the fact that some roles should have greater weight than others. D & H do not offer an answer to this

²⁴ I find D & H's (2013) approach to be very similar Chalmers's (2011) *method of elimination*, in which one eliminates a philosophical term of interest from one's debate in the hopes of staying out of a merely verbal dispute. D & H spend a fair amount of time discussing *Ramsey sentences* and *expanded postulates* instead of the method of elimination, but I think they ultimately make the same point that Chalmers makes.

²⁵ Dorr and Hawthorne (2013, p. 4).

²⁶ Such a problem is reminiscent of Putnam's (1962) 'cat' example, which was also featured in Kripke (1980): Suppose we originally assumed cats were animals—i.e., that they had similar biological makeups as other animals. But suppose, upon further investigation, it turned out that all the things we called 'cats' were actually automata. Would we say that there were not any cats after all? Or would we say that it turned out that cats were automata? This example is analogous to the question at hand, because the automata still play most of the roles that cats played (e.g., we have them as pets, they are inferior to dogs, we think the black ones are bad luck, etc.)—but they no longer play the role of being animals. Is the lack of their playing this role *enough* to say that 'cats' never actually referred to anything?

question in their paper. I will give my own thoughts on this question in Section 5, though these thoughts come nowhere close to constituting a full-fledged theory of theoretical terms.

3.3 – Naturalness and its roles

Dorr and Hawthorne list eleven different roles that Lewis intended natural properties to play, and then they consider different combinations of these roles that might be problematic. My aim in what follows is to more or less extend this approach. I will begin by reproducing their list. This will give the reader a single location to revisit the roles if necessary.

- | | |
|-------------------------------|--|
| <i>Supervenience</i> | Everything supervenes on the perfectly natural properties. |
| <i>Independence</i> | The perfectly natural properties are mutually independent. |
| <i>Duplication</i> | If some bijection from the parts of x to the parts of y maps x to y and preserves all the perfectly natural properties, x and y are duplicates. |
| <i>Non-duplication</i> | If no bijection from the parts of x to the parts of y that maps x to y preserves all perfectly natural properties, x and y are not duplicates. |
| <i>Empiricism</i> | The right method for identifying actually-instantiated perfectly natural properties is empirical. |
| <i>Simplicity</i> | One property is more natural than another iff the former has a definition in terms of perfectly natural properties that is simpler than any definition of the latter in terms of perfectly natural properties. |
| <i>Laws</i> | The conjunction of all the laws of nature can be expressed simply in terms of perfectly natural properties. |

<i>Similarity</i>	The more natural a property is, the more it makes for similarity among things that share it.
<i>Dissimilarity</i>	The more natural a property is, the more it makes for dissimilarity among things that are divided by it.
<i>Magnetism</i>	The more natural a property is, the easier it is to refer to, <i>ceteris paribus</i> .
<i>Necessity</i>	Facts about a property's degree of naturalness are non-contingent.

D & H are thorough in their presentation of this list. They consider a number of interpretations and glosses of each role.²⁷ For present purposes, I do not wish to get bogged down in the exegetical project of determining precisely what Lewis may have meant by the above roles. My reasoning here is twofold. First, D & H are thorough in their presentation of possible interpretations. Second, most of what I have to say about these roles will apply to any reasonable precisification of them.

I will give further explanations of some of these roles in Section 4. Still, I should go ahead and say a number of things about them as they are presented above. Let us begin with Supervenience. To say that a set of properties *A* supervenes on a set of properties *B* is to say that an object cannot have a change in *A*-properties without a change in *B*-properties.²⁸ To consider an example: an interesting question in aesthetics is whether aesthetic properties (properties like *being beautiful*) supervene on microphysical properties. The idea here would be that an object cannot go from being beautiful to being not-beautiful without undergoing some physical change.

²⁷ See Dorr and Hawthorne (2014, pp. 10-41).

²⁸ D & H do present a number of versions of Supervenience, but this gloss will serve present purposes.

Of course, if one thinks beauty is a much more subjective matter and is “in the eye of the beholder,” then one may be compelled to reject such a supervenience thesis.

As D & H make clear, Supervenience combines with Independence to make the claim that the perfectly natural properties serve as a minimally complete supervenience base for all other properties. This is the idea that the perfectly natural properties are the most fundamental properties. For the sake of simplicity, whenever I wish to refer to the conjunction of Supervenience and Independence, I will simply speak of ‘Fundamentality’.

The Duplication role, despite its relatively technical definition above, is not difficult to grasp. Consider what Lewis (1983, p. 355) has to say about it:

“We are familiar with cases of approximate duplication, e.g. when we use copying machines. And we understand that if these machines were more perfect than they are, the copies they made would be perfect duplicates of the original. Copy and original would be alike in size and shape and chemical composition of the ink marks and the paper, alike in temperature and magnetic alignment and electrostatic charge, alike even in the exact arrangement of their electrons and quarks. Such duplicates would be exactly alike, we say. They would match perfectly, they would be qualitatively identical, they would be indiscernible.”

Lewis’s idea is that two objects that are indiscernible in this way will share all of the same perfectly natural properties.

Skipping ahead, the Similarity and Dissimilarity roles should be familiar from Section 2. Lewis thought that the perfectly natural properties were the properties that made for *genuine* similarity or resemblance.

3.4 – Carving the world at its joints

For now, I will leave explanations of some of these roles (in particular, Simplicity, Magnetism, and Necessity) for Section 4. Before moving on, however, I would like to discuss a role that was noticeably absent from D & H’s list:

Joint-Carving The more natural a property is, the better it carves the world at its joints.

D & H may have had a reason for leaving Joint-Carving out of their list of roles. They may think the question ‘Does *F* carve at the joints better than *G*?’ is simply a variant of the question ‘Is *F* more natural than *G*?’²⁹ But I think it is unlikely that this was their rationale—one can understand joint-carving independently of naturalness and the above roles. (I will say more about this shortly.)

Another possibility is that D & H thought joint-carving was simply a matter of fundamentality (which was already addressed by Supervenience and Independence). Indeed, some naturalness enthusiasts appear to treat ‘fundamental’ and ‘joint-carving’ as interchangeable. Sider (2011, p. 5) is explicit about this interchangeability. This usage, however, might strike one as odd—might there be a property that is perfectly joint-carving but not fundamental? When Plato’s Socrates talks about carving at the joints in the *Phaedrus*,³⁰ he uses the example of dividing the body into two, the right and the left. This is a *natural* division,

²⁹ See Dorr and Hawthorne (2014, p. 8). D & H are right to be careful about these kinds of questions—ones like ‘Are all properties on par?’ This is because answers to these kinds of questions tend to include the word ‘natural’, and this takes us back to the original problem.

³⁰ *Phaedrus*, 265e-266a.

given the shape of the human body, but it is not clear that it has anything to do with the fundamental.

3.4.1 – *Two types of joint-carving*

Perhaps D & H take joint-carving to simply be a matter of fundamentality; they might even take joint-carving to be the conjunction of *all* the roles Lewis had in mind. This is possible, but I think there are two distinct understandings of joint-carving that are related to the other naturalness roles in unique ways. In this section, I will try to make these distinct understandings of joint-carving clear, and I will also make the case that both understandings are implicit in the works of Plato, Lewis, and Sider. The upshot is that we will need to be careful to say which understanding of joint-carving we are employing whenever we are relating it to other roles.

Let's begin with the first understanding of joint-carving—I will call it 'joint-carving₁'. To see what I have in mind here, consider a couple examples from Sider (2011). One example has to do with the meaning of 'Europe'. There is a physically prominent line along the Ural Mountains that could help determine the boundary of Europe. Barring the question of whether the physically prominent line *should* determine Europe's border,³¹ what does Sider have in mind when he says that the physically significant line in the Urals is a joint in nature? The takeaway is that we could use 'Europe' to carve at this joint, to have its boundary match up perfectly with the physically prominent line.

Sider later speaks of "a range of equally joint-carving candidates corresponding to the 'fuzziness' in Mars's spatial boundaries."³² I presume the fuzziness Sider has in mind here has

³¹ I should note that Sider in fact thinks the physically prominent line along the Urals should *not* determine the meaning of 'Europe', as it is not joint-carving *in the right way*. For more on this point, see Sider (2011, pp. 48-49).

³² Sider (2011, p. 51).

to do with the fuzziness of Mars's atmosphere. Barring the question of whether Mars may in fact have determinate layers of atmosphere, we can at least see what Sider has in mind here. The question of joint-carving with 'Mars' is again a matter of something like physical prominence. If Mars's atmosphere were less fuzzy, then the joint between its atmosphere and the space beyond it would be more prominent.

For another example, consider the name 'George W. Bush'. We use this name to refer to George W. Bush, not George-W.-Bush-and-the-electrons-within-a-five-meter-radius-of-him (call this entity 'George W. Bush*'). To carve₁ at the joints is to use 'George W. Bush' in the former way. It is clear that Lewis sometimes had something like this in mind:

“Compare Bruce with the cat-shaped chunk of miscellaneous and ever-changing matter that follows him around, always a few steps behind. [...] Bruce, unlike the cat-shaped chunk, has a boundary well demarcated by differences in highly natural properties. Where Bruce ends, there the density of matter, the relative abundance of chemical elements,...abruptly change. Not so for the chunk.”³³

For a rough analogy, think of how a child in grade school is instructed to “stay within the lines” as they are coloring a picture, or to “cut along the dashed lines” with their scissors. Our search for the most joint-carving expressions and corresponding meanings is a similar task. By using 'George W. Bush' to refer to George W. Bush*, we are coloring outside the lines, so to speak.

Now that we have a rough idea of joint-carving₁, let's turn to what I will call 'joint-carving₂'. I will intend joint-carving₂ to be somewhat of a catchall notion, but I can still say a bit about what I have in mind. It is common in philosophy to speak of carving at the joints as

³³ Lewis (1983, p. 372).

dividing objects into natural kinds.³⁴ We have good reasons to distinguish gold from silver, tigers from lions. Such a notion of joint-carving has little to do with carving along physically prominent boundaries; it instead has more to do with *grouping* objects. To take another example from Sider, consider a world that contains two electrons and a cow.³⁵ It is at least in some sense more natural to group the electrons together than it is to group the cow and one of the electrons together.

I am tempted to say that joint-carving₂ is largely a matter of resemblance. That is, when philosophers talk about joint-carving in relation to dividing the world into natural kinds, they may have resemblance in mind. However, if one wishes to argue that dividing up the world into natural kinds is a much more complicated issue that involves more than resemblance, this is fine—I would intend ‘joint-carving₂’ to apply to these more complicated notions as well. The main feature to keep in mind for joint-carving₂ is *grouping*.

From what I have said so far, one might think that the distinction between joint-carving₁ and joint-carving₂ is closely tied to a distinction between grammatical categories. Perhaps singular terms (like names and descriptions) are varyingly joint-carving₁, while general terms and predicates (which usually have extensions with more than one member) are varyingly joint-carving₂.

This is actually *not* what I have in mind. Predicates and properties can be joint-carving₁, and there is also a sense (albeit a trivial one) in which singular terms can be joint-carving₂. The most straightforward case of properties that are joint-carving₁ are haecceities, properties of being particular entities. For example, the property *being George W. Bush* is an haecceity, and it

³⁴ Indeed, such an understanding of joint-carving is present in contemporary literature. See, for instance, Campbell, O’Rourke, and Slater (2011).

³⁵ Sider (2011, p. 1).

arguably carves₁ at the joints in much the same way that the name ‘George W. Bush’ carves at the joints.

Because haecceities only admit of one object in their extensions, it is difficult to see how they might carve₂ at the joints. After all, how is *being George W. Bush* grouping any objects together in the way that *being a tiger* groups all of the tigers together? The answer is not immediately clear. But if joint-carving₂ is a matter of resemblance, and if resemblance is a reflexive relation, then perhaps haecceities will be trivially joint-carving₂. After all, all objects resemble themselves. It is in this sense that we may also say that singular terms are trivially joint-carving₂.

But it is not just haecceitic properties that can carve₁ at the joints. Other properties will be able to as well. Consider, for instance, the property *being a tiger*. When using the predicate ‘is a tiger’, we could refer to *being a tiger minus its left hind foot*. If we used ‘is a tiger’ in such a way, all of the entities in its extension would genuinely resemble one another, but something is obviously off here—we should be including the left hind foot for each tiger. Including the left hind foot, though, is a matter of joint-carving₁, not joint-carving₂. (And oddly enough, leaving out the foot makes for *more* resemblance, which implies that *being a tiger minus its left hind foot* is more joint-carving₂ than *being a tiger*.)

3.4.2 – *Interesting cases and other potential issues*

Sometimes it will be difficult to tell which understanding of joint-carving is involved in a particular case. Consider, for example, Sider’s red-blue world. Sider has us imagine a world that is filled entirely with red and blue fluid.




Figure 1 – Sider’s red-blue world³⁶

How should we “carve up” such a world? We could do it in a number of ways, but to carve this world at its *joints*, we would have to carve along the boundary between the red fluid and the blue fluid, like this:




Figure 2 – Carving the red-blue world at its joints

In Figure 2, our carving is represented by the dashed line, and the shaded areas represent the extensions of *being red* and *being blue*. The area to the left of the dashed line is red, and the area to the right is blue. But consider another possible carving:

³⁶ The figures I use here are very similar to those featured the first chapter of Sider (2011).

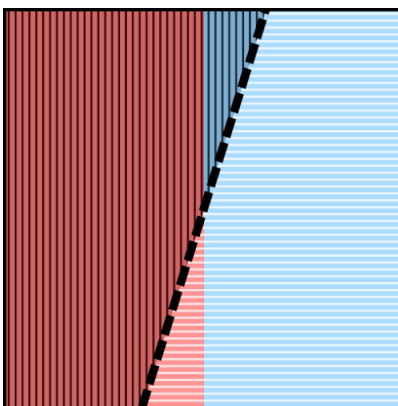


Figure 3 – A slightly-off carving of the red-blue world

This is indeed a way of carving up the red-blue world. As Sider says, we could talk about the area left of the dashed line in Figure 3 as having the property *being bred*, and the area to the right of the line as *being rue*. And if we pointed to some area left of the line as said that it was bred, we would be speaking truthfully. But it is still clear that *being bred* and *being rue* do not carve at the joints.

If *being bred* and *being rue* do not carve perfectly at the joints, then in what sense of joint-carving have they failed? I think this is an interesting case. There is a sense in which the line in Figure 2 carves₁ the red-blue world at its joints, because it carves along the physically prominent line. There is another sense in which the line in Figure 2 carves₂ at the joints, because it divides the fluids into different categories.³⁷

Putting aside overlap cases (where both senses of joint-carving are involved), there may be other problems for these understandings of joint-carving. For example, it is common to talk of carving the mind at its joints.³⁸ Philosophers of mind often ask questions like ‘Is there a joint between cognition and perception?’ or ‘Is there a joint between belief and the imagination?’ In

³⁷ We might even say here that the physically prominent line here is *determined* by facts about similarity.

³⁸ See Lyons (2001).

much the same spirit Plato considered, in Book IV of the *Republic*, whether the soul was divisible into three distinct parts.

If joint-carving₁ is a matter of “coloring inside the lines” and looking for physically prominent features, then it is not easy to see how one could carve₁ at the mental. And if joint-carving₂ is a matter of resemblance, then it is at least more plausible that one could carve₂ the mind at its joints. We commonly speak of some thoughts resembling others.

Given the examples and potential issues we have covered, I will now try to give a more formal characterization of each type of joint-carving. Perhaps one could give much more precise definitions of each, but the hope is that these characterizations will be sufficient for now.

Joint-Carving₁: To carve₁ the world at its joints, one must be worried about *carving* around all and only the right parts of a single object (or of multiple objects in a set).

Joint-Carving₂: To carve₂ the world at its joints, one must be worried about *grouping* objects in a natural way, according to resemblance or some other means.

These characterizations are admittedly vague, but now we can consider, for example, how one might carve the mind at its joints. If we introduce a notion like *functional resemblance*, it will be natural to group together certain mental states and to divide them from others. If, for instance, one thinks some mental states (like beliefs) function to aim at the truth, while others (like imagination) do not function in this way, then there will be more grounds for saying there is a joint between beliefs and imaginative states. This sounds like joint-carving₂.

Not only is it important that we group mental states together in the right way; it is also important that we not *leave out* the right parts or *add* the wrong parts. If, for example, the mental states we call ‘beliefs’ really do function to aim at the truth, then we will want to include that as a feature of theirs. If they do not have this function, then we will not want to include it. This might be a matter of joint-carving₁.

Finally, I want to add a couple more points about these two understandings of joint-carving. The first point is that both understandings will admit of degrees in straightforward ways. An expression’s degree of joint-carving₁ will be directly related to the amount of an object it rightfully includes. Consider the following candidate meanings for the predicate ‘is a tiger’: *being a tiger*, *being a tiger minus its left hind foot*, *being a tiger minus its left hind leg*, *being a tiger minus all four legs*, *being a tiger plus all of the electrons within a five mile radius*. I have arranged this list according to level of joint-carving₁. *Being a tiger* is perfectly joint-carving₁—it includes all and only the right area of tigers. *Being a tiger minus its left hind foot* includes less area than it should, but it includes more of the right area than the two properties that follow it. *Being a tiger plus all of the electrons within a five mile radius* includes all of the area it should, but it also includes far too much—so it is an extremely poor carving. If joint-carving₂ is something like resemblance, then degree of joint-carving₂ will be a matter of how much the objects in a given group resemble one another. This means that even though *being a tiger* carves₁ perfectly at the joints, it does not carve₂ perfectly. There is still a high level of resemblance between all tigers, so *being a tiger* is fairly joint-carving₂. And *being a tiger* is more joint-carving₂ than *being a tiger or a lion*. And *being a tiger or a lion* is more joint-carving₂ than *being a tiger or the Empire State building*.

The second point I wish to make this. One might argue that one of these understandings of joint-carving—more likely joint-carving₂—is more orthodox, or that it is more likely what Plato had in mind in the *Phaedrus*. I do not think such an argument is convincing. I do think Plato was certainly interested in joint-carving₂; but when he speaks of dividing the body up into a right half and a left half, it is at least conceivable that he had something like joint-carving₁ in mind. This is an interesting historical question, but I do not intend to focus on it here. Even if Plato did not have joint-carving₁ in mind, I believe that, given the examples I have shown above, joint-carving₁ is implicit in the work of Sider and Lewis. Joint-carving₁ and joint-carving₂ will have different relations to the other naturalness roles, so it is an important distinction to make.

Section 4: Problems for Natural Properties

There are a number of problems for naturalness as it is typically understood. In this section, I will argue that there are some sets of roles which cannot be jointly-satisfied by any property. Given that naturalness is to be understood in terms of these roles, these arguments then amount to skeptical arguments against naturalness.

4.1 – *Fundamentality + Necessity + Joint-Carving*

I see a problem with a property's satisfying the roles of Fundamentality, Necessity, and Joint-Carving. Recall that Fundamentality is the claim that the perfectly natural properties serve as a minimally complete supervenience base. Necessity is the claim that a property's degree of naturalness is non-contingent. And Joint-Carving is the claim that a property's degree of naturalness corresponds to its level of joint-carving.

Keeping all of these roles in mind, consider an example. Consider the actual world, where electrons are fundamental—they are mereological simples that have no proper parts. Perhaps one day we will discover this to be untrue, but let us stipulate that this is the case for now. So let us stipulate that *being an electron* is fundamental. In contrast, *being an atom* is not fundamental, as atoms are composed of sub-atomic particles like protons, neutrons, and electrons. *Being an atom* is very *close* to being fundamental, but it is not quite there.

Now consider a nearby possible world where electrons *do* have proper parts: suppose these entities can be broken down into halves. Suppose further that this world is the same as the actual world in all other respects.³⁹ In such a world, we might say that *being an electron* is not fundamental, but *being the right half of an electron* is. Another thing is likely, too: *being an electron* is carving₁ at the same joints in both worlds; in this nearby possible world, however,

³⁹ Dorr and Hawthorne (2014, pp. 33-34) consider a very similar example in a footnote.

there are simply more joints to carve at. Recall that joint-carving₁ is a matter of including precisely the right amount of an object in an extension. To make such a carving, we will be moving around the same lines in each world.




Figure 5 – Electrons in the actual world and a nearby possible world

Here we have a situation where a single property—*being an electron*—has different levels of fundamentality in different worlds, but where it maintains its level of joint-carving₁. *Being an electron* carves right along the circles in Figure 5; it does not go outside of the lines, so to speak. But Necessity requires that a property’s degree of naturalness be the same in all possible worlds, so a perfectly natural property cannot be fundamental in one world and non-fundamental in another. Thus, we have found an important property that cannot satisfy all of these roles at once.

One might reply to this scenario by saying that the property of *being an electron* is not really instantiated at the nearby possible world—that the things that I am calling ‘electrons’ in this nearby possible world are not even electrons. Perhaps all there is in this nearby possible world is *being an electron**. This type of reply, as D & H note, could be motivated by a belief in the necessary *a posteriori*.

A proposition is necessarily true iff it is true in all possible worlds. A proposition is *a posteriori* iff it is justified by experience. Traditionally, philosophers (going back at least as far as Hume) treated necessary truths (like ‘All bachelors are male’) as *a priori*, justified

independently of experience. However, Kripke (1980) famously argued that there were some *a posteriori* necessities, some truths that featured natural kind terms. Kripke thought that once we discover the referent of a natural kind term—like ‘water’, whose referent is H₂O—that term rigidly designates that referent. In other words, natural kind terms refer to the same thing in all possible worlds. It follows that *a posteriori* identities like ‘Water is H₂O’ are true in all possible worlds. Similarly, one might think that ‘electron’ rigidly designates the kind of electron we see in the actual world, the kind that is presumably indivisible.

Perhaps we must make a distinction between *being an electron* and *being an electron**. I am actually happy to grant this point. If we have Kripke-like scruples here and want to say that there are no electrons at this nearby world, this is fine; but we are still left with a problem. The problem arises from the assumption that fundamentality is closely tied to joint-carving. Even if we want to say that there is a distinction between *being an electron* and *being an electron**, I will maintain that these two properties carve₁ at the exact same joints in each world. Looking back at Figure 4, both properties carve a perfect circle along the outside of the particles. At the same time, we should all agree that *being an electron** is not fundamental, but *being an electron* is. This allows us to hang on to Necessity, if we so choose, but it is now clear fundamentality and level of joint-carving₁ come apart.

4.2 – Simplicity + Similarity + Joint-Carving

Now consider another set of roles: Simplicity, Similarity, and Joint-Carving. Simplicity is the claim that the naturalness of a property is determined by the length of its definition in a perfectly natural language (a language consisting of only perfectly natural predicates). Although ‘is grue’ is not (from a linguistic perspective, at least) *prima facie* more disjunctive than ‘is blue’, the definition of ‘is grue’ in terms of perfectly natural properties is plausibly longer than the

definition of 'is blue'. This is because 'is grue' (in Goodman's (1955) sense) applies to "all things examined before t just in case they are green but to other things just in case they are blue." So we can say that 'is blue' is more natural than 'is grue'.

As Sider notes, however, this type of approach to comparative naturalness will encounter some difficulties. First, it treats all infinitely-disjunctive definitions as equally natural. Sider's reply to this difficulty is that many of the definitions we are interested in will not be infinite definitions.⁴⁰

Even if it is true that most definitions we need will not be infinite, there is a remaining worry. It is not clear that the number of disjuncts in a perfectly natural language corresponds to naturalness or level of joint-carving. Consider these two disjunctive properties: *being red or orange* and *being red or blue*. Similarity, recall, is the claim that the more natural a property is, the more it makes for similarity amongst the things it applies to. The fact that 'is blue' is more natural than 'is grue' has something to do with the fact that the members of the extension of 'is blue' more closely resemble one another than those of 'is grue'. But what about the extensions of *being red or orange* and *being red or blue*? The members of the extension of *being red or orange* resemble one another more than those of *being red or blue*. If Lewis is right about comparative naturalness, though, we should expect the disjunctive definition of *being red or*

⁴⁰ See Sider (2011, p. 130).

orange to be shorter than the definition of *being red or blue*. But their definitions do not differ in length—both are equally long disjunctions of equally joint-carving properties.^{41 42}

It is clear that Sider may have this type of problem in mind when he says that “A further worry about the Lewisian approach is that mere length of definitions is an inadequate measure.”⁴³ Interestingly, Sider does suggest an alternative strategy that is in the spirit of Lewis’s definitional strategy. If our perfectly natural language is the first-order predicate calculus, and we are only allowing perfectly natural predicates into our language, then we could evaluate definitions in prenex disjunctive normal form; that way we could easily determine exactly how long each definition was.⁴⁴ Additionally, we could come up with new strategies for evaluating various Boolean compounds of perfectly natural predicates. One that Sider mentions: “The number of disjuncts might, for instance, be taken to count against the definition more than the average number of conjuncts per disjuncts.”⁴⁵ After all, *disjunctions* of perfectly natural properties are less natural than *conjunctions* of perfectly natural properties.⁴⁶ We can justify this

⁴¹ These colors (red, blue, and orange) are determinables. There are many ways for an object to be red—it can be scarlet red, mahogany red, or even something as specific as Coca-Cola red. See Funkhouser (2006, pp. 548-549). This means that *being red* does not make for perfect similarity with regard to color. That said, we might consider a similar example to the one above: instead of *being red or orange*, we could take the disjunctions of all the super-determinate shades of red and orange. In such a case, each disjunct would do a better job of making for resemblance, and it would be just as problematic for Lewis.

⁴² Here is another possible example: compare *being an electron or a cow* with *being an electron or a proton*. The latter does a better on making for resemblance, yet we might expect these properties to be equally simple.

⁴³ Sider (2011, p. 131).

⁴⁴ The idea here is that once we have defined a property in a perfectly natural language, we can then ensure that we are comparing the definitions in the fairest way. Prenex disjunctive normal form has all of the quantifiers on the left, with the predicates and variables on the right. Since all first-order logic sentences can be converted to this form, this gives us a way to compare definitions fairly.

⁴⁵ Sider (2011, p. 131).

⁴⁶ This point ultimately derives from Armstrong (1978).

by thinking in terms of resemblance. Conjunctive properties make for genuine resemblance in at least some respect, while disjunctive properties do not guarantee this.

But even if we can come up with a better strategy (like the one we have just considered) for evaluating definitions in a perfectly natural language, this is unlikely to solve the problem raised by the red-or-orange/red-or-blue example. That is, *being red or orange* appears to be more natural than *being red or blue*, but it is not clear that *being red or orange* will have a shorter definition than *being red or blue* in Lewis's perfectly natural language. It does not really matter if we devise a strategy for weighting disjunctions differently than conjunctions if we have no independent reason for thinking that *being red or orange* will have a shorter definition; after all, both properties are disjunctions of equally-natural, equally-joint-carving colors.

I think this problem will generalize, and it can be pressed in at least a couple ways. First, we can come up with a term that is intuitively more natural than another but for which we have no independent reason to think that its definition will be shorter (or otherwise simpler). The red-or-orange/red-or-blue example would fall into this category. Second, we can come up with properties that we might expect to be simpler but that are treated as equally natural or joint-carving. Sider, for instance, claims that *being an unmarried male* and *being an unmarried male eligible for marriage* carve at the joints equally well. But it is pretty plausible that *being an unmarried male eligible for marriage* would have a longer definition in Lewis's perfectly natural language than *being an unmarried male*. *Being an unmarried male eligible form marriage* is adding more constraints on what it is to be a bachelor. It is not as simple.

In response to the example I have just mentioned, one might worry that we are falling prey to the mistake we discussed at the beginning of Section 3: perhaps our reason for claiming that *being an unmarried male eligible for marriage* would have a longer definition just because it

takes more words than *being an unmarried male* to specify the property. It is true that the way we specify properties in English is merely a contingency, so we must address this kind of worry. Let's use '*being F*' to refer to the same property as *being an unmarried male*, and let's use '*being G*' to refer to the same property as *being an unmarried male eligible for marriage*. Even after we evened things out, so to speak, it is *still* plausible that *being G* will have a longer definition in a perfectly natural language.

4.3 – Reference magnetism

Reference magnetism is a *metasemantic* notion. What distinguishes a *metasemantic* notion from a merely *semantic* notion? Generally speaking, *semantics* is concerned with what our words mean; *metasemantics*, we might say, is concerned with questions *about* semantics. One who is interested in metasemantics might ask questions like this: *What makes for a good semantic theory? What is the nature of reference? What is it that makes the words we use have meaning?*⁴⁷ Consider this last question. When philosophers ask this question, they are interested in what ties our words to things in the world. In other words, they are interested in the “semantic glue” that *sticks* words to their referents.

Some philosophers have suggested (in various ways) that it may not be as easy as it looks to get an answer to this question. Following Wittgenstein, Kripke (1982) famously pointed out that our past usage of ‘plus’ may not ensure that we mean *plus* (our normal understanding of the addition function) rather than *quus* (a function that works like normal addition for numbers under 57, but which outputs ‘5’ otherwise). This type of problem generalizes to nonmathematical

⁴⁷ This distinction is closely related to the distinction between a *semantic theory* and a *foundational theory of meaning*. For more on the distinction, see Speaks (2014). There, reference magnetism is treated as a criterion of a foundational theory of meaning. Also see Hawthorne (2007, p. 430), where he speaks of a ‘foundational account of intentionality’.

examples; skeptical worries like this make the search for semantic glue even more important and interesting.

In Section 3.2 of *Writing the Book of the World*, Sider surveys a number of metasemantic theories that attempt to provide semantic glue; but for many of them, as he says, “the glue doesn’t seem to be sticky enough.” One such theory is a simple sort of *descriptivism*, and it works like this: for a given word, there is a set of definitional sentences *S* that contain that word. According to this sort of descriptivism, a satisfactory interpretation of the meaning of a word involves making all of the sentences in *S* true. The problem for such a theory is that we can construct interpretations which make all of the sentences in *S* turn out true, but which also make intuitively *false* sentences come out true, too.

Let’s consider an example. Here is an intuitively false sentence featuring ‘bachelor’: ‘All bachelors are millionaires’. We can assign semantic values that make this sentence true. First, we can assign the set of gold watches to ‘bachelor’. Second, we can assign the set of timepieces to ‘are millionaires’. Since it is true that all gold watches are timepieces, we have made our intuitively false sentence come out true.

But what about the definitional sentences for ‘bachelor’ that would be in *S*? It does not take much work to make these come out true, too. Take ‘All bachelors are male’, for instance. If we assign the set of watches to ‘are male’, then ‘All bachelors are male’ will turn out true because all gold watches are watches. And we could continue to tweak our interpretations so as to make all of the definitional sentences in *S* true.

Given that this simple sort of descriptivism fails, we are going to need a different metasemantic theory. The main failing of the descriptivist account was that it allowed us to assign bizarre semantic values to words. An improved theory should keep this from happening.

That is what Lewis's theory aims to do. Good interpretations, on Lewis's account, involve the assignment of *natural* meanings to words, allowing us to eliminate the bizarre semantic values.⁴⁸ Natural meanings serve as *reference magnets*, pulling words and their referents towards one another.

Importantly, Lewis is *not* claiming that we cannot refer to unnatural properties. We can refer to unnatural properties whenever we like, if that is what we intend to do. For this reason, Dorr and Hawthorne's specification of reference magnetism in terms of 'easiness' of reference is perhaps not ideal; 'eligibility' is the better term to use. Lewis's point is that, generally speaking, we ought to interpret speakers as referring to natural semantic values. This is also why reference magnetism is typically not taken to be full-fledged metasemantic theory; rather, it is typically viewed as an additional constraint on more robust theories. So, for example, one might view reference magnetism as an additional constraint on the simple descriptivist theory above, or one might add it to a causal theory.⁴⁹

It is important to note that reference magnetism is an *externalist* constraint on reference. In the philosophy of language, externalism is the idea that meaning is at least partly determined by features in our environment—or, as Putnam phrased it, meanings just "ain't in the head." In the latter half of the twentieth century, many philosophers (such as Putnam, Kripke, Burge, and Lewis⁵⁰) defended externalism.

⁴⁸ Importantly, other constraints—e.g., a causal constraint—can eliminate (many but not all) bizarre semantic values, too. I say more on this later.

⁴⁹ Both of these options are considered in detail in Sider (2011).

⁵⁰ See Putnam (1975), Kripke (1980), and, more recently, Burge (2010).

Lewis himself favors another view, an “eligibility plus charity” view.⁵¹ On such a view, a good interpretation is one that “strikes the best balance” between eligibility and charity.

Eligibility is a matter of interpreting predicates to pick out more natural properties. *Charity* is a matter of interpreting a speaker’s statements so that they come out to be true.

To see why eligibility on its own is not enough, consider another example. It is fairly well known that ‘jade’ is now typically used to refer to two different elements, nephrite and jadeite. Suppose I pick up a piece of jadeite and I say ‘This rock is jade’. What makes it the case that I am attributing the property *being either nephrite or jadeite* instead of simply *being nephrite*? If one looks only to eligibility, one might interpret me as meaning *being nephrite* as this property is presumably more natural than the disjunctive *being either nephrite or jadeite*. But since I am intending to use the disjunctive property here, something has gone wrong here. If the rock I am referring to is actually jadeite, then the eligibility-only theory interprets my utterance as false. If one adds a constraint of charity, however, one will also interpret my utterances as (for the most part) being true. The charity constraint would thereby eliminate *being nephrite* as a possible semantic value in this situation.

There will be other cases, however, where eligibility is a deciding factor. In a Twin-Earth-style example,⁵² charity may not determine whether ‘water’ refers to the colorless liquid found in lakes and streams or H₂O. In such a case, eligibility would point to the latter candidate.

⁵¹ See Lewis (1984, pp. 227-228). And as Dorr and Hawthorne (2014, p. 28) make clear: “[F]or Lewis, this was just a toy theory. In his considered view, the primary role for naturalness is in the theory of mental content, although naturalness does play a subsidiary role in the story about how semantic facts supervene on mental ones...” Dorr and Hawthorne go on to say much more about Lewis’s theory of mental content.

⁵² See Putnam (1975).

4.3.1 – Magnetism + Similarity

Now that we have a picture of Lewis's metasemantic theory, let's consider how the role of Magnetism might fit in with others. Let's begin with Similarity.

The conjunction of Magnetism and Similarity amounts to the claim that a property's eligibility for reference corresponds to level of similarity (or resemblance) amongst the members of the property's extension. Because *being blue* does a better job at making for similarity than *being blue or green*, the former property is a better candidate for reference. Also consider Twin-Earth-style examples. Does 'water' refer to the colorless liquid found in lakes and streams, or does it refer to H₂O? Because H₂O makes for more resemblance, it is a better candidate for reference.

Magnetism and Similarity will lead us in the right direction in some cases, but there still exist problematic cases. *Being a tiger* does reasonably well on making for similarity. The things that have this property are pretty similar, though there will of course be differences: most notably, differences in sex and size. *Being a lion* does just as well in terms of similarity. The disjunction of these two properties *being either a tiger or a lion* does a little worse on similarity, though we can of course imagine properties that would do even worse. It is plausible that *being a tiger* is a better semantic candidate than *being either a tiger or a lion*. But consider another property. Suppose we know a lion named 'Rex'—and so there is a corresponding property *being Rex the lion*. This property does much better than *being a lion* on similarity, as there is only one member of its extension. Now consider two disjunctive properties: *being either a tiger or a lion* and *being either a tiger or Rex the lion*. Is the latter property a more eligible semantic candidate? Though it does better in making for similarity, this is unlikely.⁵⁴

⁵⁴ It is likely that charity or use would eliminate these candidates in most contexts.

Though I have largely been focusing on the meanings of predicates, it is worth noting that the meanings of singular terms are not straightforwardly determined by similarity. We may, of course, say that any object is similar to itself, but there is no way for us to compare degrees of similarity that different candidates make for. For example, suppose our usage of ‘Bob’ does not favor *Bob* over *Bob-plus-the-atoms-within-a-half-centimeter-from-Bob*. Either candidate trivially makes for perfect resemblance, but *Bob* does better than *Bob-plus-the-atoms-within-a-half-centimeter-from-Bob* on joint-carving₁. Resemblance is not doing any work here.

4.3.2 – Magnetism and other roles

Now let us turn to the question of how Magnetism fits in with other naturalness roles. Let’s start by considering Magnetism and Joint-Carving. The conjunction of Magnetism and Joint-Carving amounts to the claim that a property’s eligibility for reference corresponds to its level of joint-carving.

Recall our distinction between joint-carving₁ and joint-carving₂. Roughly speaking, joint-carving₁ is a matter of “coloring inside the lines.” Joint-carving₂, on the other hand, is a matter of dividing natural kinds, which may just be a matter of similarity (or resemblance).

Philosophers have tended to conflate joint-carving₁ and joint-carving₂; importantly, however, these notions play interestingly distinct roles in conjunction with Magnetism. First consider joint-carving₁. In this sense, carving at the joints rules out bizarre semantic candidates like *being George W. Bush and all the electrons within 5 cm of George W. Bush*. Now consider joint-carving₂. Because joint-carving₂ has more to do with resemblance, carving at the joints in this sense will help eliminate a different kind of bizarre semantic value—like *being either blue or a tiger*.

It is clear that both types of joint-carving can help us eliminate bizarre semantic values, but can a set of properties play both of these joint-carving roles at once? Yes—to an extent. The problem is that while both joint-carving₁ and joint-carving₂ allow for degrees of joint-carving, different properties will have different levels of joint-carving, depending on which version we are looking at. For example, *being an atom* is a reasonably joint-carving₂, but of course not all atoms look alike. A hydrogen atom is different from a helium atom. If we look to joint-carving₁, however, *being an atom* carves perfectly at the joints, even though not all atoms are intrinsic duplicates.

Recall, too, the connections between the two types of joint-carving and other roles. Joint-carving₂ may have some connection to fundamentality. For any fundamental particle, there will be a property, *being that kind of fundamental particle*, that makes for perfect genuine similarity. *Being an electron* may be such a property. But notice, too, that joint-carving₁ is not so closely tied to the fundamental. Just as *being an atom* can carve₁ perfectly at the joints, so too can a property like *being a tiger* (which does not carve₂ perfectly at the joints).

Though the properties that carve₁ at the joints do not need to be fundamental properties themselves, joint-carving₁ is still connected to fundamentality in a more derivative way. This is because the joints, and thus the demarcations, between objects are ultimately determined by fundamental entities. And this actually leads to some interesting questions about joint-carving₁: how are we to “stay inside the lines” at the fundamental level? Here is how I see it working:




Figure 5 – Joint-carving₁ at the fundamental level

Though I think Figure 5 shows the most straightforward way to carve at the fundamental, I am in principle open to another method if it has non-trivial implications.

To summarize: both versions of joint-carving allow us to eliminate a different kind of bizarre semantic value. But both understandings of joint-carving are at odds with other roles. The most joint-carving₁ properties will not necessarily be the most fundamental properties, and a property's level of joint-carving₂ will not necessarily correspond to its level of simplicity in a perfectly natural language.

4.4 – Substantivity and the merely verbal

Though Lewis did not connect natural properties to substantivity, Ted Sider has argued that there is a connection.

4.4.1 – Introduction to substantivity

Some progress is made in philosophy when we discover that a dispute we are engaged in is pointless—we figure out that we need not waste our time on it anymore. One way a dispute can be pointless is if it is nonsubstantive—roughly, if it is dispute merely about the meanings of words rather than the world. When a dispute is merely verbal, it's easy to get the feeling that it is

pointless; if we would just get clear on the meanings of terms, the dispute would evaporate.⁵⁵ A debate about whether the Pope is a bachelor might just involve a disagreement about the meaning of ‘is a bachelor’. To borrow Sider’s example, the parties to this dispute might disagree about whether ‘is a bachelor’ means *being an unmarried male* or *being an unmarried male who is eligible for marriage*. In such a situation it is likely that the root of the disagreement is merely verbal, not substantive.⁵⁶ If we are interested in identifying these types of disputes (and it is not always easy), then it would be helpful to have a good understanding of what makes for a substantive or nonsubstantive dispute.⁵⁷

Sider approaches the problem of substantivity in terms of structure. When asking whether a given question or claim is substantive, we should look to the candidate meanings of all of the constituent expressions. Consider the question of whether online video gaming is a sport. The predicate ‘is a sport’ has a number of candidate meanings: *being an outdoor recreational competition*, *being a recreational competition that involves teams*, etc. Whether online gaming is a sport seems to be a question of which of these candidate meanings we go with, not a question about how the world is.

⁵⁵ As Chalmers (2011) makes clear, not all verbal disputes are pointless. Often we care deeply about the answers to nonsubstantive questions. For example, we care about the judicial interpretation of the term ‘person’.

⁵⁶ We can imagine, however, an alternative context in which the question of whether the Pope is a bachelor *is* intuitively substantive. If, for instance, two parties agree that the meaning of ‘is a bachelor’ is *being an unmarried male who is eligible for marriage*, but one of the parties is just unfamiliar with the commitments of the papacy or does not know who the Pope is, then the question may turn out to be substantive.

⁵⁷ Philosophers often use ‘substantive’ or ‘merely verbal’ to modify different things—e.g., *disputes*, *claims*, or *questions*. Chalmers (2011) speaks mostly in terms of disputes. Sider usually speaks in terms of claims or questions, but he is careful to note that these notions are closely related. See Sider (2011, pp. 47-48).

One might naturally ask what it takes for a particular meaning to count as a semantic candidate. Sider defines candidatehood counterfactually: “If a linguistic community, roughly in our circumstances, *could* have used *E* to mean *m* without seeming ‘semantically alien’—could have used *E* to reach ‘the same semantic goal’ as we use *E* to reach, albeit perhaps by a different route—then *m* is a candidate for *E*.”⁵⁸ I would submit that neither of the meanings I have just considered for ‘is a sport’ is semantically alien in Sider’s sense. In this instance, our semantic goal might just be to determine whether online gaming highlights should be featured on SportsCenter (a program which, for the sake of argument, covers *all* and *only* sports). Both of these meanings will get the job done.

So where does naturalness come into Sider’s account of substantivity? Here is one of the answers that he considers:

“...a nonsubstantive question is one containing an expression *E* whose candidates are such that i) each opposing view about the question comes out true on some candidate; and ii) no candidate carves at the joints *in the right way for E* better than the rest. A candidate *c*₁ carves better ‘in the right way for *E*’ than another candidate *c*₂, to a first approximation anyway, iff *E* is a theoretical term, *c*₁ satisfies enough of the core theory associated with *E*, and *c*₁ carves better than does *c*₂.”⁵⁹

One might think that Sider’s understanding of substantivity admits of degrees. The distinction between the substantive and nonsubstantive might strike one as familiar. It appears to be closely related to other historical divisions—for example, to Hume’s distinction between relations of ideas and matters of fact, or to Kant’s distinction between the analytic and synthetic. If the

⁵⁸ Sider (2011, p. 50).

⁵⁹ Sider (2011, p. 49).

substantive-nonsubstantive distinction is to be understood in the way I have described above, one might think that this suggests a lack of sharp divide between substantive and nonsubstantive claims. Thus, this account of substantivity might be used to motivate a sort of Quinean picture, a picture that blurs the boundaries between the substantive and nonsubstantive.⁶⁰

4.4.2 – *Substantivity + Joint-Carving*

Though Sider is connecting joint-carving to substantivity, it would not be quite right to say that he thinks the substantivity of a question directly corresponds to the joint-carving scores of the question's constituent expressions. But there is still a connection, as we have seen above. And if there is supposed to be this connection, we must yet again be clear on what we mean by 'joint-carving'.

In some places, Sider appears to have joint-carving₂ in mind when he is talking about substantivity. This is the case when he speaks of substantivity and its relation to the special sciences, those sciences other than fundamental physics:

“...I said earlier that being cast in perfectly joint-carving terms normally suffices for substantivity. But being cast in *highly* albeit not perfectly joint-carving terms—a common occurrence in the special sciences—also normally suffices for substantivity. Except for questions that strain the boundaries of taxonomy (a relatively uncommon occurrence), special-science questions normally fall into one of the following two categories: i) each expression has a candidate meaning that carves far better than all other candidates; or ii) each expression has a range of candidates that carve far better than

⁶⁰ See Quine (1951).

do other candidates not in the range, and the question's answer is insensitive to choices of candidates within those ranges. In either case, the question is substantive."⁶¹

If Sider has joint-carving₂ in mind, however, let us consider whether this is plausible. A property like *being a tiger* carves₂ reasonably well, as the members of its extension resemble one another to a high degree. But recall the property we considered in 4.3.1—*being a tiger or Rex the lion*. Such a property is only slightly less joint-carving₂ than *being a tiger*. And now a question like 'Do some tigers have manes?' will come out as nonsubstantive when, intuitively, it is substantive. This is because neither candidate meaning, as Sider's definition calls for, carving *far better* than the other. So understanding substantivity in terms of joint-carving₂ is problematic.

One might say that *being a tiger* in fact *is* carving much better than *being a tiger or Rex the lion*. After all, letting Rex into this group is to let in a completely new animal with different a different evolutionary history, DNA, etc. I might grant this point if pressed, but perhaps this shows the need for further analysis of the 'carves far better than' segment of Sider's definition.

What about joint-carving₁—is it a more plausible guide to substantivity? Unlikely. *Being a tiger* and *being a tiger or Rex the lion* are equally joint-carving₁, because neither property colors outside the lines, so to speak. The objects that each property applies to will be perfectly carved around—Rex's tail, for instance, is not left out of the extension of *being a tiger or Rex the lion*. Thus, joint-carving₁ will not be a good guide to substantivity, at least in the way Sider may have in mind.

⁶¹ Sider (2011, p. 48).

Because both understandings of joint-carving have these problems, I think it is unlikely substantivity can be understood in terms of joint-carving. But perhaps a more refined notion of joint-carving could do the work Sider would like it to do.

Section 5: Conclusion

In Section 3.2, I briefly discussed an interesting point that Dorr and Hawthorne make about the metasemantics of theoretical terms. Suppose one has an incredibly strict (or “draconian”) metasemantics of theoretical terms: a theoretical term *T* has a nonempty extension only if some entity satisfies *all* of the theoretical roles associated with *T*. If my worries in Section 4 are warranted, then such a metasemantics would leave ‘natural’ without an extension. And this would leave us with good reason to be skeptical of Lewisian naturalness. But this metasemantics is probably too strict.

Still, I think that even a more reasonable metasemantics may prevent ‘natural’ from referring. A more reasonable metasemantics might distinguish between roles that are more- and less-central to a given notion. For Lewisian naturalness, distinguishing between more- and less-central roles is no simple task—it is a task that requires a considerable amount of exegetical work. But if we look at Dorr and Hawthorne’s list, we could at least start this project.

Consider, first, a role like Empiricism. Empiricism is less of an explanatory role and more of just a conceptual or epistemic role. This makes some difference. If it turned out that some perfectly natural properties were discoverable by *a priori* means, but they could still do the same explanatory work, then I see no reason to deny the existence of natural properties in that case. This is not to say that all mere conceptual roles will be less central than explanatory roles, but this seems to be the case with a role like Empiricism.

The upshot of Section 4.1 was that we may be forced to drop Fundamentality, Necessity, or Joint-Carving, or hold that none of these *always* obtains. Again, I do not think dropping Necessity would be a major concession on the naturalness enthusiast’s part. But as I stressed in 4.1, dropping necessity will not be enough: it is still the case that level of fundamentality does

not correspond to level of joint-carving. Because of this, I believe a naturalness enthusiast may even be required to drop either Fundamentality or Joint-Carving.

But even if the naturalness enthusiast chooses to drop Fundamentality, Dorr and Hawthorne offer some compelling evidence that Lewis may have intended it to be a central role for naturalness. As they argue, “In ‘New Work’, the main focus is on a claim of *Non-supervenience*: no perfectly natural property is such that the facts about it supervene on the facts about all the other perfectly natural properties. In conjunction with *Supervenience*, this is equivalent to the claim that the perfectly natural properties constitute a *minimal* supervenience base for everything...”⁶³

If Fundamentality is a central role, this leaves open the option of dropping Joint-Carving. But this is probably not a good option for the naturalness enthusiast, as it is plausible that Lewis intended Joint-Carving to be an equally central role. Consider, for instance, what he says in “New Work for a Theory of Universals”: “Properties carve reality at the joints—and everywhere else as well” and, later, “...one thing that makes for the naturalness of a property is that it is a property belonging exclusively to well-demarkated things.”⁶⁴

Given the considerations of 4.1, it appears the naturalness enthusiast is forced to drop a central role—either Fundamentality or Joint-Carving. And given even a somewhat lax metasemantics, one that does not require complete role-satisfaction, we have good reason to say that there are no such things as natural properties.

If one is not convinced by 4.1, then we can turn to Section 4.2, where I considered the roles of Simplicity, Similarity, and (again) Joint-Carving. As we saw in that section, Simplicity

⁶³ Dorr and Hawthorne (2014, p. 13).

⁶⁴ Lewis (1983, pp. 346, 373).

can be explicated in a number of ways: we can come up with different methods for evaluating the simplicity of a definition in a perfectly natural language. If Lewis's length of definition approach could be replaced by a more plausible method of evaluation, then this replacement would be *close enough* to the Simplicity role he had in mind. Even so, I argued in 4.2 that *any* reasonable method of evaluation of definitional simplicity will run into problems when considered in conjunction with Similarity. The role of Similarity is pervasive in "New Work for a Theory of Universals." Lewis says explicitly, "Natural properties would be the ones whose sharing makes for resemblance..." (1983, p. 347). And in *On the Plurality of Worlds*, we see a key passage about Simplicity:

"Probably it would be best to say that the distinction between natural properties and others admits of degree. Some few properties are *perfectly* natural. Others, even though they may be somewhat disjunctive or extrinsic, are at least somewhat natural in a derivative way, to the extent that they can be reached by not-too-complicated chains of definability from the perfectly natural properties."

Thus, I believe the considerations of 4.2 also give us good reason to say there are no natural properties.

Finally, if one is still not convinced, one can turn to Section 4.3 where I pointed out some problems with the role of Magnetism as well. Dropping Magnetism seems to be another less than desirable route for the naturalness enthusiast. Magnetism is an explanatory role that, in Lewis's eyes, was part of the "new work" that natural properties were supposed to do. So if that role cannot be satisfied, then there is yet another reason to deny the existence of natural properties.

Works Cited

- Armstrong, David. *Universals and Scientific Realism*. Cambridge: Cambridge UP, 1798. Print.
- Armstrong, David. "Universals as Attributes." Sosa, Jaegwon Kim and Ernest. *Metaphysics: An Anthology*. Oxford, UK: Blackwell, 1989. 198-208. Print.
- Burge, Tyler. *Origins of Objectivity*. Oxford, UK: Clarendon Press, 2010. Print.
- Campbell, Joseph, Michael O'Rourke and Matthew Slater. *Carving Nature at Its Joints*. Cambridge, MA: MIT Press, 2011. Print.
- Carnap, Rudolf. "Empiricism, Semantics, and Ontology." *Revue Internationale de Philosophie* (1950): 20-40. Print.
- Chalmers, David J. "Verbal Disputes." *Philosophical Review* (2011): 515-66. Web.
- Dorr, Cian and John Hawthorne. "Naturalness." *Oxford Studies in Metaphysics, Volume 8* (2013): 2-77. Web.
- Frege, Gottlob. "On Sense and Reference." *Translations from the Philosophical Writings of Gottlob Frege* (1952): 56-79. Print.
- Funkhouser, Eric. "The Determinable-Determinate Relation." *Nous* (2006): 548-569. Web.
- Goodman, Nelson. *Fact, Fiction, & Forecast*. Cambridge, MA: Harvard UP, 1955. Print.
- Hawthorne, John. "Craziness and Metasemantics." *Philosophical Review* (2007): 427-40. Web.
- Kripke, Saul. *Naming and Necessity*. Cambridge, MA: Harvard UP, 1980. Print.
- . *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard UP, 1982. Print.
- Lewis, David K. "New Work for a Theory of Universals." *Australian Journal of Philosophy* (1983): 343-77. Print.
- . *On the Plurality of Worlds*. Oxford, UK: Blackwell, 1986. Print.
- . "Putnam's Paradox." *Australian Journal for Philosophy* (1984): 221-36. Print.
- Lyons, Jack C. "Carving the Mind at its (Not Necessarily Modular) Joints." *The British Journal for the Philosophy of Science* (2001): 277-302. Print.
- Mill, John Stuart. *A System of Logic*. Harper & Brothers, 1882. Print.
- Plato. *Plato: In Twelve Volumes*. Trans. Harold North. Fowler and Walter R. M. Lamb. Cambridge, MA: Harvard UP, 1975. Print.
- Putnam, Hilary. "It Ain't Necessarily So." *The Journal of Philosophy* (1962): 658. Web.
- . "The Meaning of "Meaning"." *Minnesota Studies in the Philosophy of Science* (1975): 131-193. Print.

Quine, W.V.O. "On What There Is." Quine, W.V.O. *From a Logical Point of View*. Cambridge, MA: Harvard UP, 1961. 21-38. Print.

—. "Two Dogmas of Empiricism." *The Philosophical Review* (1951): 20. Print.

Salmon, Nathan. "How to Become a Millian Heir." *Nous* (1989): 211-20. Print.

Sider, Theodore. *Writing the Book of the World*. Oxford, UK: Clarendon Press, 2011. Print.

Speaks, Jeff. "Theories of Meaning." 26 January 2010. *Stanford Encyclopedia of Philosophy*. Web. 3 June 2016.

Swoyer, Chris. "Properties." 17 February 2016. *Stanford Encyclopedia of Philosophy*. Web. 3 June 2016.