

October 2018

EXPLORING THE EASY ROAD TO NOMINALISM

Jordan Kroll
University of Massachusetts Amherst

Follow this and additional works at: https://scholarworks.umass.edu/dissertations_2



Part of the [Metaphysics Commons](#), and the [Philosophy of Language Commons](#)

Recommended Citation

Kroll, Jordan, "EXPLORING THE EASY ROAD TO NOMINALISM" (2018). *Doctoral Dissertations*. 1363.
<https://doi.org/10.7275/12744857> https://scholarworks.umass.edu/dissertations_2/1363

This Open Access Dissertation is brought to you for free and open access by the Dissertations and Theses at ScholarWorks@UMass Amherst. It has been accepted for inclusion in Doctoral Dissertations by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

EXPLORING THE EASY ROAD TO NOMINALISM

A Dissertation Presented

by

JORDAN KROLL

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

September 2018

Philosophy

© Copyright by Jordan Kroll 2018

All Rights Reserved

EXPLORING THE EASY ROAD TO NOMINALISM

A Dissertation Presented

by

JORDAN KROLL

Approved as to style and content by:

Phillip Bricker, Chair

Hilary Kornblith, Member

Alejandro Pérez Carballo, Member

Seth Cable, Member

Joseph Levine, Department Head
Philosophy

DEDICATION

To my parents, and my sister, for all their patience and support.

ACKNOWLEDGMENTS

I would like to thank my advisor Phil Bricker for his careful eye and for his willingness to engage with a project that departs so significantly from his own philosophical convictions. The tone and rigor of my dissertation, if not the views expressed, owe much to Phil. My interest in metaphysics may never have developed were it not for the courses that Phil taught during my first few years at UMass. I would also like to thank my other committee members, Hilary Kornblith, Alejandro Pérez Carballo, and Seth Cable, both for their help on this project and for the courses they've taught during my time at UMass.

I am tremendously grateful to my friends Bob Gruber, Kim Soland, Miles Tucker, and Aaron Washington for philosophical discussions both serious and silly, for pleasant diversions, and for keeping me sane throughout. And of course to Julie Rose, without whom my time at UMass would have been unimaginably different.

I would not have attended UMass at all, let alone written this dissertation, were it not for the philosophical education I received at UBC and UCLA as an undergraduate. I owe a particular debt to Ori Simchen, whose classes on the philosophy of language made me fall in love with technical philosophy.

Lastly, I would like to thank my parents and my sister for their unending patience, their love, and their encouragement. I would not have returned to finish this dissertation without their support.

ABSTRACT

EXPLORING THE EASY ROAD TO NOMINALISM

SEPTEMBER 2018

JORDAN KROLL

B.A., UNIVERSITY OF BRITISH COLUMBIA

PH.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Phillip Bricker

My dissertation is divided into three self-contained chapters, each of which explores some facet of nominalism. The overall aim is to explicate and defend a nominalist approach that recognizes the utility of talking about, or presupposing the existence of, abstract objects even if no such objects exist.

The first chapter begins with a question: why is talk about abstract mathematical entities so useful in describing and explaining the physical world? Here is an answer: talk about such entities is useful for describing and explaining the physical world insofar as there is some appropriate structural similarity between them and the target physical system. But this account leads to a problem: there is no guarantee that the world contains a sufficiently rich ontology for the requisite structures to be instantiated. The primary focus of the first chapter is on exploring ways to resolve this problem as it relates to the metaphysics of quantitative properties.

In the second chapter I present easy road nominalism, a variant of nominalist that accepts that purported reference to abstract objects is an indispensable part of our best

scientific theories. The crucial insight of the easy road strategy is that referential discourse can be useful for reasons that have nothing to do with the existence of the entities purportedly being referred to. While I spend some time explaining the core of the easy road strategy, my focus is on applying the easy road strategy to theories of language. I propose that the presupposition that there are abstract linguistic objects plays a crucial role in explaining linguistic behavior, and provides the basis for a nominalistically acceptable account of content.

In the third chapter I characterize fictionalism, and examine some of the wide variety of fictionalist theories in the literature. Many fictionalist theories depend crucially on the idea that non-literal utterances of a sentence have different content from literal utterances of that sentence. But I argue that the core fictionalist strategy requires no such thing, and that the prevalence of such views has been driven by assumptions about the role of content and truth that are misplaced.

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS.....	v
ABSTRACT.....	vi
CHAPTER	
1. THE PROBLEM OF INSUFFICIENT CONCRETE ONTOLOGY.....	1
1.1 Structural Similarity and Concrete Ontology.....	1
1.2 Quantitative Properties.....	7
1.3 Primitive Ordering and Concatenation.....	10
1.4 Why Platonism Doesn't Help.....	13
1.5 Ideology and Explanation.....	16
1.6 Field's Theory of Quantity.....	17
1.7 Armstrong's Theory of Quantity.....	19
1.8 The Mereology of Properties.....	31
1.9 Properties, Points, and Dimensions.....	41
1.10 The Modal Approach.....	45
1.11 Some Concluding Remarks.....	49
2. LANGUAGE WITHOUT ABSTRACT OBJECTS.....	52
2.1 Introduction.....	52
2.2 Easy Road Nominalism.....	55
2.3 Semantic Pretense-Involving Fictionalism.....	65
2.4 Balaguer's Semantic Fictionalism.....	71
2.5 A Theory of Content.....	77
2.6 Reference and Languages.....	86
2.7 Some Concluding Remarks.....	91
3. FICTIONALISM, CONTENT, AND EXPLANATION.....	97
3.1 Introduction.....	97
3.2 Fictionalism Characterized.....	98
3.3 Modal Fictionalism and Meta-Fictionalism.....	106
3.4 Variants of Content Fictionalism.....	121
3.5 Reference to Abstract Objects as a Descriptive Aid.....	126
3.6 Reference to Abstract Objects as an Explanatory Aid.....	131
3.7 Fictionalism and Content.....	138
3.8 Fictionalism and Truth.....	149
BIBLIOGRAPHY.....	156

CHAPTER 1

THE PROBLEM OF INSUFFICIENT CONCRETE ONTOLOGY

1.1 Structural Similarity and Concrete Ontology

Science aims to describe and explain the physical world around us. It is rather curious, then, that *abstract* objects – scientific models, mathematical objects, and so on – play such a prominent role in current scientific practice; talk about abstract objects is arguably indispensable to science (and if not indispensable, so useful as to be practically indispensable). This is puzzling, or at least it would be if we weren't so accustomed to it. Why is talk about abstract objects so useful in describing and explaining the physical world?

Exactly how one answers this question will depend in large part on one's other philosophical commitments. Platonists accept the existence of abstract objects, and can take talk about models and mathematical objects at face value as truths about a realm of abstract objects. A commitment to platonism does not, in itself, answer our question. Pointing out that mathematics truly describes a realm of abstract objects tells us nothing about why truths about those abstract objects are relevant to describing and explaining the physical world. Truths about one subject matter are not in general relevant to explaining truths about another (e.g. facts about the digestive systems of mice are not in general relevant when explaining the political system in Canada), except when there is some causal connection between the two. But abstract objects are, by definition, causally

inert: abstract objects do not causally interact with physical objects. So if truths about abstract mathematical objects are relevant to scientific practice, and the connection is not causal, we are owed an explanation of just what that connection is.

What ought the platonist say? Here is a sketch of an answer for the platonist: talk about abstract objects is useful for describing and explaining the physical world insofar as there is some *appropriate similarity* between the abstract objects being talked about and the target physical system. Talk about abstract scientific models is useful insofar as the model is sufficiently similar in relevant respects to the physical system being modeled. Mathematics is useful because in employing mathematics we can describe abstract structures that bear an appropriate similarity to the structures instantiated in the physical world.¹ And so on. The details of this sketch can be filled out in various ways. Some might take the only ‘appropriate similarity’ to be identity: physical objects instantiate the very structures being described mathematically. One might instead opt to view ‘appropriate similarity’ as isomorphism: holding that while physical objects instantiate structures of relations isomorphic to the structures of relations that hold among abstract mathematical objects, these structures are nevertheless distinct. Or one might take isomorphism to be too strong a requirement, and require only something looser. I’ve remained deliberately non-committal, since a debate over these details is tangential to the issues I want to explore in this paper.² There is much more to be said here, but for our purposes this broad sketch will suffice.

¹ This general strategy is widely employed by platonists. Some structuralists go further and see numbers as nothing more than the places in certain abstract structures, but it is worth flagging that nothing of the sort is required here. Using structure to explain the applicability of mathematics to science does not require any particular view about the nature of mathematical objects.

² I see no reason to think that the utility of mathematics to science requires something as strong as isomorphism or identity. Exactly how strong the similarity between the abstract and physical structures

Platonists are not the only ones who must find something to say about why talk of abstract objects is so useful in describing and explaining the physical world. Nominalists, in contrast to platonists, reject the existence of abstract objects. So, the nominalist will have to go about resolving our puzzle in a different way. Nominalists must explain the success of science, with all its talk of abstract objects, without positing the existence of abstract objects.

Here is a sketch of an answer for the nominalist: talk about abstract objects is useful for describing and explaining the physical world insofar as there would be some *appropriate similarity* between the abstract objects being talked about and the target physical system, were those abstract objects to exist. The explanation here is much the same as it was for the platonist.³ Talk about abstract scientific models is useful insofar as the model being imagined would be sufficiently similar in relevant respects to the physical system being modeled were it to exist. Mathematics is useful because in employing mathematics we can describe abstract structures that would bear an appropriate similarity to the structures instantiated in the physical world were they to exist. The details differ depending on whether one is a platonist or a nominalist – the platonist sees abstract objects where the nominalist sees fiction – but the general idea is the same: talk about abstract objects is a means of describing and explaining the structure

must be is something that may well vary on a case by case basis, depending on the practical goals at hand. This is obviously so when it comes to scientific models: how closely the model must resemble the target phenomenon, and which simplifying idealizations are appropriate, depends on the particular use for which the model is being employed.

³ The primary difference between the platonist and nominalist here is that the nominalist explanation has a modal character that the platonist explanation does not. One might worry that this modal character is problematic in some way, or at least an additional cost to adopting the nominalist explanation. I agree that it would be if that modal character required additional primitive modality, but I claim that it does not. Unfortunately a full defense of this point would require a paper of its own.

of the physical world.⁴ As Mary Leng puts the point, mathematical explanations “explain a phenomenon by showing it to result from the structural features of the empirical set up.”⁵ Let us call this explanation of the utility of talk about abstract objects to science, whether offered in a platonist or a nominalist variant, the *Structural Similarity Account*.⁶

Can the nominalist version of this strategy really succeed? Can the utility of talk about abstract objects be explained without accepting the existence of abstract objects? I believe it can. And I will assume that it can in what follows. The skeptical reader is invited to read the expansive literature defending the utility of talk about abstract entities in the absence of the existence of any such entities.⁷ My focus is on a different issue, one that has received comparatively little attention. For the Structural Similarity Account to be correct, the physical world must contain structures that are appropriately similar to the abstract structures being described. But this leads to a potential problem: there is no guarantee that the world contains a sufficiently rich ontology for the requisite structures to be instantiated. If the physical world does not contain a sufficiently rich ontology, then

⁴ I should say, more accurately, that *one* of the functions of abstract object talk is to describe and explain the physical world. This is what we might call a ‘first level’ use of abstract object talk, and it is the use that is relevant to our current discussion. But sometimes scientific models are used to model other abstract objects, rather than physical phenomena. And not all mathematics is applied mathematics: pure mathematics makes no attempt at all to describe or explain the physical world. I’ll largely be ignoring these further uses in this paper, although it should be relatively straightforward to see how the picture I’ve sketched extends to these further uses.

⁵ Leng (2012), p. 988.

⁶ It is worth noting that the puzzle I raised is just as pressing for many scientific anti-realists as it is for scientific realists. Anyone who recognizes the incredible productivity of science owes us some explanation of the success of that science, whether scientific theories are taken to be literally true or not. Constructive empiricists, for example, take the aim of science to be to produce empirically adequate (rather than true) theories. Some have argued that only scientific realism can explain the success of science (see Putnam 1975) and hence that anti-realist views such as constructive empiricism are fatally flawed. But constructive empiricists and other anti-realists are able to make use of the Structural Similarity Account to explain the success of scientific theories. For more on constructive empiricism, see van Fraassen (1980) and Monton & Mohler (2017).

⁷ See, among others, Yablo (1998, 2005, 2012), Balaguer (1998, 2009), Melia (2000), Azzouni (2004), and Leng (2005, 2010).

the Structural Similarity Account threatens to collapse. Let us call this *The Problem of Insufficient Concrete Ontology*.

Here is the plan for the rest of this paper. I plan to focus on the Problem of Insufficient Concrete Ontology as it relates to the metaphysics of quantitative properties – properties such as mass, charge, and energy. I will begin by saying something about the general features of quantitative properties that have to be captured by any adequate metaphysics. I'll consider several extant theories of quantity, and show how they succumb to the problem. That is, I'll show how the Problem of Insufficient Concrete Ontology arises in the metaphysics of quantity. The problem arises for the platonist as well as the nominalist, and I'll say something about why platonism does not provide a way out of the problem. I'll then discuss the details of two general strategies available to both the platonist and the nominalist in response to the problem. Although both options have their costs, it is unclear how to explain the success of science without adopting one or the other.

Before pressing on, it is worth making some clarifying remarks about nominalism and abstract objects. Nominalism is sometimes construed as the thesis that there are no properties, and at other times is construed as the thesis that there are no abstract objects. It is often assumed that properties are abstract objects, and hence that nominalists about abstract objects must also be nominalists about properties. And, quite often, nominalists about abstract objects are nominalists about properties. But they needn't be. A philosopher who rejects abstract objects can accept the existence of *in res* properties – properties that are part of the concrete world. Let us reserve the term “nominalism” and its cognates for nominalism about abstract objects.

So what is an abstract object? There is no standard characterization of the distinction between abstract and concrete, but two related distinctions can help illustrate the idea. First, we might characterize abstract objects as those objects that are not part of the spatiotemporal world. This distinction clearly separates paradigmatic abstract objects – mathematical entities such as numbers and pure sets/classes – from paradigmatic concrete objects – electrons, trees, and buildings. But there are abstract objects that arguably do have spatiotemporal properties. Fictional characters are abstract entities, yet it seems plausible to say that Mickey Mouse was invented in the 20th century and didn't exist before that. That said, I will assume that the abstract objects we are primarily concerned with in this paper – abstract quantitative properties – are not (or needn't be) spatiotemporally located. This is generally taken as a feature of abstract properties, since it allows for the existence of uninstantiated properties.

We might, instead, attempt to characterize abstract objects as those that are not causally efficacious. But it has been argued that there are epiphenomenal mental states – mental states that are not causally efficacious – yet mental states are not abstract. So this characterization is not ideal either. But it is close enough for our purposes.⁸ I take it to be a necessary feature of abstract objects that they are not causally efficacious.

As a result, I take it that *in res* properties – causally efficacious properties that are part of the spatiotemporal world – are nominalistically acceptable. The nominalist may still wish to deny the existence of such properties, but they needn't do so. In what follows, I'll be considering nominalist theories both with and without such properties.

⁸ For a more detailed discussion of how to characterize abstract objects see Rosen (2017).

1.2 Quantitative Properties

Quantitative properties – such as mass, charge, and energy – are properties that have magnitudes. An object never merely has mass, but rather has a particular magnitude of mass, such as *2g mass*, or *3g mass*. Likewise for all other quantitative properties; we can always ask how much mass, or charge, or energy, a particular entity has. Quantitative properties feature prominently in our best scientific theories: they factor into the causal story of the concrete world. The trajectories of colliding objects are determined in part by the *mass* properties of those objects. We would not have electricity if particles did not have *charge* properties. The behavior of particles cannot be adequately predicted unless we assume that they have *energy*. And so on. Quantitative properties, like particles, are among the building blocks of the physical world.

A distinction is sometimes drawn between determinable quantitative properties – properties that come in various magnitudes, such as mass, charge, and energy – and determinate quantitative properties – particular magnitudes of those properties, such as *2g mass*, *5C charge*, or *3J energy*. But unless one has a particularly abundant theory of properties, there is no need to posit determinable quantitative properties – determinate quantitative properties do all the work. Saying that an object has mass is just to say that it has some determinate mass property or another.

The distinction between determinable and determinate quantitative properties does draw important attention to a feature of quantitative properties: they are grouped into *families* or kinds. All determinate mass properties belong to the mass family, and all determinate charge properties belong to the charge family. All such families are disjoint:

no determinate mass property belongs to the charge family, and no determinate charge property belongs to the mass family.

Determinate quantitative properties belonging to the same family stand in certain relations to each other.⁹ For one, determinate mass properties in the same family are *incompatible*: no object can have a total mass of 2g and a total mass of 3g, and no object can have a total charge of 1C and a total charge of 2C. Second, all determinate quantitative properties belonging to the same family are *ordered*. So, for example, 5g mass is greater than 2g mass, both of which are greater than 1g mass.¹⁰ Third, determinate quantitative properties belonging to the same family stand in *distance* relations to one another. So, 5g mass is 1 gram of mass distant from 6g mass and 3 grams of mass distant from 2g mass, and hence is closer to 6g mass than to 2g mass. Distance is closely tied to *concatenation*: 2g mass and 3g mass taken together gives you 5g mass. Ordering, distance, and concatenation relations never hold between determinate

⁹ My characterization of the features of quantitative properties follows closely that given by Eddon (2013). See also, among others, Ellis (1966), Field (1980), and Bigelow and Pargetter (1988).

Must all quantities possess these features? Ordering, at the very least, seems to be constitutive of what it is to be a quantity. In any case, I will follow the precedence set in the literature by assuming so. It seems to me (and not just me) that incompatibility, distance, and concatenation are required as well. But I won't insist on it. Someone who thinks that there can be quantities that are only ordered can take my account to focus on those cases where the additional structure is present.

It might also help to distinguish quantities from other properties that we can rank with regards to resemblance. So, for example, we can say that octagons more closely resemble hexagons than they resemble cones. And we can say that red more closely resembles orange than it does blue. But it doesn't make sense to say that one shape is greater than another, or that one color is greater than another. These are cases of what Ellis calls *qualitative equality* (Ellis, 1966, pp. 37-8). These resemblance relations can often be captured spatially: we can, for example, represent the colors in a sphere with three dimensions labeled hue, saturation, and lightness. Because length is a quantitative property, points in this spatial representation of color can be ordered along 3 dimensions, be assigned magnitudes, and stand in distance relations. But the utility of this spatial representation is simply in arranging color properties according to how much they resemble each other. Colors themselves don't have magnitudes.

¹⁰ Multi-dimensional quantities (such as area or volume) can be ordered along each dimension. For such quantities, it may also make sense to construct an absolute ordering among equivalence classes. So, for example, we can order all of the areas of $2m^2$ as being greater than all of the areas of $1m^2$, even though the areas of $2m^2$ differ from each other in length and height.

quantitative properties of different families; it is nonsense to ask whether *2g mass* or *3C charge* is greater than the other.

Mathematically, the relations that hold between quantitative properties are captured using a metric – a function that defines a distance for every pair of quantitative properties in a family. The structure of distance relations holding between quantitative properties in a family is isomorphic to the structure of distance relations holding over an interval of the real numbers. Accordingly, determinate quantitative properties can be assigned numerical values relative to a unit that encodes information about the relations that they stand in to other determinate quantitative properties. The choice of unit is arbitrary: “1kg mass”, “1000g mass” and “2.20462lbs” all refer to the same determinate quantitative property. But a metric tells us nothing about the underlying metaphysics; a metric can only successfully capture facts about quantity if the world holds up its end of the bargain.

Now that we have some sense of what it is that our theory of quantitative properties must capture, what should our metaphysics of quantitative properties look like? Several approaches are possible. One is to take all this talk of properties and relations as literally as possible and adopt a realist stance about properties and relations, taking on an ontological commitment to both. A second approach is to take an antirealist stance and deny that properties and relations exist. This doesn't require denying that there are any objects that have mass. Rather, the antirealist holds that accepting the existence of an object with 3g mass commits one only to the existence of the object with 3g mass, and not to the existence of some further *3g mass* property. Likewise for relations. One could also take a mixed approach, and take a realist stance towards properties and an antirealist

stance towards relations. This sort of mixed view might be motivated by the thought that properties play a substantive causal role in our best scientific theories that relations do not, and hence that we have reason to accept the former but not the latter.

This general debate over the existence of properties and relations is not particular to the metaphysics of *quantitative* properties; a realist about properties more generally will likely opt to be a realist about quantitative properties as well, and likewise for the antirealist. In this paper I will avoid taking a stand in this general debate, although relevant differences in their treatment of quantitative properties will be highlighted. My concern is instead with issues proper to quantitative properties: quantitative properties raise special philosophical difficulties, which emerge regardless of which general framework is adopted. In particular, any adequate theory of quantitative properties should be able to capture the features of quantitative properties outlined above: family, incompatibility, ordering, distance, and concatenation.

1.3 Primitive Ordering and Concatenation

How should the nominalist explain these features of quantitative properties? One option is to take facts about some of these features to be fundamental. For example, one could take facts about ordering and concatenation to be fundamental, and recover facts about family and distance from them. How this theory goes will depend in part on whether one is a realist about properties and relations. A nominalist who rejects the existence of properties and relations altogether can introduce a pair of primitive

predicates for each quantity – “ x is greater or equal (in mass) to y ” and “the (mass) sum of x and y is z ” – that range over objects. The distance (in mass) between two objects e and f , where f is greater than e , is an object g such that the sum of e and g is f .¹¹ There are no properties to group into families on this view, but the primitive predicates are ‘paired-up’ into families of sorts: one ordering relation and one concatenation relation per quantity.

A realist about properties can make do with a single pair of primitive predicates – “ x is greater than or equal to y ” and “the sum of x and y is z ” – that range over determinate quantitative properties. The distance between two determinate quantitative properties a and b , where b is greater than a , is a determinate quantitative property c such that the sum of a and c is b .¹² Properties belong to the same family if and only if one is greater than or equal to the other. Lastly, a realist about both properties and relations can make do with a pair of fundamental relations – x is greater than or equal to y and the sum

¹¹ There is something a bit peculiar about identifying the distance in mass between two objects as yet another object, but there is no other way to make sense of quantities of mass on an anti-realist picture. It is perhaps in part for this sort of reason that anti-realists often opt to capture features of quantity using other fundamental relations that do seem to hold more naturally between objects. See the discussion of Hartry Field’s theory below.

Note also that there is no guarantee that there will be a *unique* object g such that the sum of e and g is f . There may be many such objects, in which case there is more than one object to identify as *the* distance. But shouldn’t there be a unique distance? How can the nominalist get around this uniqueness problem? One option is to deny that we ought to think of ‘the distance’ between two objects as unique. Given that there may be many suitable objects to play the role, we can identify ‘the distance’ with any or all of them. Another option is to identify the distance with the set of all such objects (given a nominalistically acceptable interpretation of the ‘set’). How one opts to resolve this uniqueness problem is not the focus of this paper.

¹² This realist variant of the view, much like the nominalist variant of the view, may suffer from a uniqueness problem (see my discussion in the previous footnote): there may be more than one determinate quantitative property c such that the sum of a and c is b , and hence more than one determinate quantitative property to identify as *the* distance. Whether the realist is faced with this problem will depend on how the realist conceives of properties: as universals or as tropes (more on this distinction later). Here again, as in the previous footnote, I will remain neutral on how best to resolve the uniqueness problem.

of x and y is z – that hold between determinate quantitative properties.¹³ Distance and family can be captured the same way as before. On all such views, there is nothing further in virtue of which facts about ordering and concatenation obtain; facts about ordering and concatenation are taken as fundamental.

There is a serious issue confronting the theories of quantity that I've just presented: they rely on implausible existence assumptions to capture distance facts. To see why, consider again how the anti-realist tries to capture distance facts: the distance (in mass) between two objects e and f , where f is greater than e , is an object g such that the sum of e and g is f . Given this account, the distance (in mass) between two objects e and f can only be captured when there exists an appropriate object g . But what guarantee do we have that there is any such object? Without such an object, the distance in mass between the original two objects cannot be measured. And that's a problem with the theory.

In fact, there are cases where it seems clear that we don't have any object of the appropriate mass. If the standard model of particle physics is correct, then there are a limited number of types of point particles, some of which vary only with regards to mass. For example, the electron, the muon, and the tau are all point particles with $-1e$ charge and $\frac{1}{2}$ spin, which differ in mass. The distance in mass between the muon and the tau is significantly larger than the distance in mass between the electron and muon. But there are no particles with intermediate mass values with which to capture this fact. So the anti-realist runs into the Problem of Insufficient Concrete Ontology.

¹³ For a platonist theory of this sort, that takes facts about ordering and concatenation as fundamental, see Brent Mundy (1987).

The move to a realist account of properties doesn't offer any immediate help to the nominalist: if we are restricting ourselves to instantiated properties, as the nominalist must, then the existence of an appropriate intermediate property with which to capture distance facts is not guaranteed. This sort of difficulty – one that arises from a shortage of appropriate objects or properties – is a persistent problem for nominalistic theories. The scientifically minded nominalist must admit that what exists, and how much of it exists, is a contingent matter. Any adequate nominalistic theory of quantity will have to overcome this difficulty.

1.4 Why Platonism Doesn't Help

What can be said in response? One option is to attempt to resolve the problem by embracing platonism and the existence of uninstantiated determinate quantitative properties. Because the platonist is not restricted by contingent facts about what actually exists, there is no problem with positing the existence of a continuous array of determinate quantitative properties. That guarantees that for any two determinate quantitative properties a and b , there is a determinate quantitative property c that is the distance between a and b .

The platonist strategy is essentially to make up for a lack of concrete ontology with additional abstract ontology. Not enough concrete quantitative properties? No problem! There are abstract quantitative properties to do the work. This strategy is only effective to the extent that the abstract properties *can* do the work the concrete properties

were supposed to do. And abstract properties can do at least some of the requisite work: abstract properties can serve as the missing relata we needed to capture all the ordering, distance and concatenation facts. Unfortunately, they can't do *all* the requisite work, at least not without some unattractive consequences.

It is important to keep in mind why a metaphysics of quantitative properties is needed in the first place: quantitative properties feature prominently in the causal story our best scientific theories tell about the physical world. Without taking into account the *mass* of objects, our best scientific theories would fail to accurately predict the behavior of colliding objects, gravitational attraction, and a host of other observed phenomena. This gives us good reason to believe, whether we are realists about properties or not, that all this talk of "mass" is accurately capturing *something* about the physical world. According to the Structural Similarity Account of the success of scientific explanations involving abstract objects, our use of a metric to model facts about mass will only be successful to the extent that there is a relevant similarity between the structure specified by the metric and a structure in the physical world.

If the platonist finds the Structural Similarity Account appealing, then they must posit an appropriate structure in the physical world. If that's the objective, then the appeal to uninstantiated quantitative properties is theoretically inert: uninstantiated quantitative properties are by definition nowhere to be found in the physical world, and hence they can't be part of physical structure that is being sought. So it seems that the platonist strategy of making up for a deficit in concrete ontology with an expansive abstract ontology is of no help.

Or, at least, the platonist strategy will be of no use if the Structural Similarity Account is on the right track. But what if the platonist rejects the idea that quantitative properties have anything to do with the physical world? Could it not be the case that facts about mass, charge, and energy are at bottom facts about abstract properties and the relations that hold between them? The problem with this strategy is that it leaves mysterious the success of our scientific theories. Why are facts about uninstantiated abstract properties relevant to describing and explaining the physical world around us? They aren't instantiated in the physical world.¹⁴ They have no causal bearing on the physical world. Whence then their relevance? Absent the development of a suitable alternative to the Structural Similarity Account, it appears that positing the existence of abstract quantitative properties does not help the platonist avoid the Problem of Insufficient Concrete Ontology. The platonist, like the nominalist, must find some other strategy. In what follows I will simplify the discussion by focusing on nominalism, but the strategies to be discussed can be adopted (with minor modifications) by the platonist as well.

¹⁴ The focus here on *uninstantiated* properties is important. A proponent of abstract properties could certainly hold that those abstract properties are instantiated in the physical world, and hence that we have a ready explanation of the relevance of the abstract structure: it is instantiated in the physical world! But the advantage that the platonist has over the nominalist is the ability to posit the existence of uninstantiated properties, and uninstantiated properties are (quite obviously) not instantiated and hence not open to this explanation.

1.5 Ideology and Explanation

What alternatives are available? One option is to try and cover for a lack of suitable ontology with an expansion in ideology: the nominalist can simply introduce an infinite number of distance predicates with the form “ x has X grams more mass than y ” for all values of X , and take all such distance facts as fundamental.¹⁵ But given that there are often no instantiated properties to correspond to the “ X grams” of mass indicated, these predicates cannot be given their natural interpretation but must instead be taken as an infinite number of primitive additions to the ideology of the theory. This massive expansion to the primitive ideology of the theory is a rather undesirable consequence, and one to be avoided if possible. This only compounds a difficulty that is already present for those who want to take ordering and concatenation facts as fundamental: such theories leave too much unexplained. Why do the ordering and distance facts so neatly divide the determinate quantitative properties into families? They just do. Why are there no determinate quantitative properties of the same family that fail to stand in ordering and distance relations? There just aren’t. Why does no object ever have more than one mass at a time? They just don’t. In the place of explanations, we are left with stipulations; it is taken as primitive that the facts just are that way.

¹⁵ It is possible, with sufficient technical machinery, to get the same result using a single variable-degree predicate for each quantity, and infinitely many non-logical axioms. Brent Mundy introduces just such a predicate: \leq . The proposition ‘ $a_1, \dots, a_m \leq b_1, \dots, b_n$ ’ is to be intuitively read as saying that the (mass) sum of the objects b_1, \dots, b_n is at least as large as the (mass) sum of the objects a_1, \dots, a_m . The predicate \leq takes a variable number of objects in each position, and may take as few as zero on one side, with the proposition ‘ $\leq a$ ’ meaning that a has nonnegative value. See Mundy (1989), especially pp. 132-4, for details. I’m not convinced that this reduction in ideology really does simplify the theory in any serious way. Distance facts are still taken as fundamental; they are simply described using one variable-degree predicate rather than many fixed-degree predicates.

To be fair, taking certain facts to be fundamental, rather than giving an explanation of those facts in more fundamental terms, is an inevitable part of any theory. Every explanation bottoms out somewhere. But it is generally desirable in theory building to increase explanatory power, and to simplify the fundamental ontology and primitive ideology of our theories as much as possible. So it is a mark against a theory if it takes too many facts as fundamental instead of providing an explanation of those facts. Theories of quantity that rely on primitive ordering and concatenation leave a lot unexplained. Theories of quantity that attempt to respond to the Problem of Insufficient Concrete Ontology with a greatly expanded ideology leave considerably more unexplained. So I take it that such theories are unacceptable. Let's consider some alternatives.

1.6 Field's Theory of Quantity

Not all theories of quantity take facts about ordering and concatenation as primitive. In *Science Without Numbers*, Hartry Field advances a theory that posits a pair of relations for each quantity – between-ness and congruence – that hold between objects.¹⁶ The between-ness relation for mass, $Bet_m(e,f,g)$, can be understood intuitively as saying that e is between f and g with regards to mass. The congruence relation,

¹⁶ This is not quite accurate, in two respects. First, Field talks about “relations” rather than predicates, and I've followed his terminology here, but it is important to keep in mind that Field doesn't mean to posit the existence of relations in any ontologically serious sense. These are a piece of ideology, not of ontology, in his theory.

Second, Field attempts to do away with objects and instead talks about space-time points. Because our focus is on quantitative properties, I will ignore this further element of Field's account. It will simplify our discussion considerably to talk about objects rather than space-time points. However, it should be noted that the objections I raise against Field's theory are equally applicable either way.

$\text{Cong}_m(e,f,g,h)$, can be understood intuitively as saying that the distance in mass between e and f is congruent with the distance in mass between g and h . The mass that an object has is entirely a matter of the $\text{Bet}_m(e,f,g)$ and $\text{Cong}_m(e,f,g,h)$ relations that it stands in: for an object to have *40g mass*, on this view, is not for that object to have a determinate mass property, but rather for that object to stand in particular $\text{Bet}_m(e,f,g)$ and $\text{Cong}_m(e,f,g,h)$ relations. Field then attempts to demonstrate how our ordinary mathematical discourse about quantities is nominalistically acceptable by constructing a representation theorem between the structure of these relations over objects and a connected set of the real numbers.¹⁷

How does Field's theory fare against the Problem of Insufficient Concrete Ontology? Not well. Field's representation theorem can only succeed if the $\text{Bet}_m(e,f,g)$ and $\text{Cong}_m(e,f,g,h)$ relations that hold between objects provide enough structure for there to in fact be an appropriate homomorphism with a connected set of the reals. But there may simply not be enough objects to stand in the requisite relations.¹⁸ I think the force of the difficulty can again be brought out by considering the electron, the muon, and the tau. The masses of these particles stand in ordering and distance relations to each other, but these distance facts cannot be captured by Field's theory. For Field's representation theorem to hold, there must be a continuous array of objects with masses between those of the electron, muon, and tau. But this does not seem to be the case given the standard

¹⁷ The representation theorem will not specify a one-to-one mapping, since different objects can have the same mass. See Field (1980), p. 57.

¹⁸ For objections in this vein, see Armstrong (1988), p. 309, Melia (1998) and Eddon (2013).

model of particle physics: there is no such array of objects with intermediate masses.

Field's theory, then, has difficulty capturing distance facts at the actual world.^{19,20}

1.7 Armstrong's Theory of Quantity

It seems to me that there are only two available options in response to the Problem of Insufficient Concrete Ontology. The first is to insist that the world really does contain a sufficient rich ontology to instantiate the requisite structures, even in cases where prima facie this does not seem to be the case. Call this sort of view an 'Additional Structure' view, since it posits additional concrete structure where intuitively we might see none. The second option is to rely on modal resources to cover for the lack of a sufficiently rich ontology. Let us begin with strategies of the first sort.

One view that posits additional structure is David Armstrong's theory of quantity. To understand Armstrong's views it is important to say something about how he thinks of properties. Realists about properties, like Armstrong, must make a choice about whether to adopt an ontology of universals or an ontology of tropes. *Universals* are properties that

¹⁹ Field may have overlooked this difficulty because the problem does not arise for all quantities. With regards to temperature, Field writes: "temperature and other scalar fields used in physics are assumed to be continuous, and this guarantees that if point x has temperature $\Psi(x)$ and point z has temperature $\Psi(z)$ and r is a real number between $\Psi(x)$ and $\Psi(z)$, then there will be a point y spatio-temporally between x and z such that $\Psi(y)=r$ " (Field, 1980, p. 57). And that seems plausible when it comes to temperature. But it does not seem plausible with regards to other quantities such as mass. If one point has mass $\Psi(x)$ and another point has mass $\Psi(z)$, we have no grounds to conclude that there is another point between them with a mass greater than $\Psi(x)$ and smaller than $\Psi(z)$.

²⁰ Above I raised a secondary problem for theories that take facts about ordering and concatenation as fundamental: such theories require implausible existence assumptions and leave too much unexplained. Field's theory does little better. Field introduces two pieces of primitive ideology, between-ness and congruence, and takes facts about between-ness and congruence to be fundamental. This hardly seems like an improvement over a theory that takes facts about ordering and concatenation to be fundamental.

can be instantiated repeatedly, and are wholly present wherever they are instantiated. So, if there is a *5g mass* universal, it is instantiated by, and wholly present in, each and every object that has 5g of mass. *Tropes*, like universals, are wholly present wherever they are instantiated. However, unlike universals, tropes can be instantiated only once. So for every object that has 5g of mass, there is a *5g mass* trope possessed by that object and no other.

Universals and tropes can be used to do much of the same theoretical work, but they are not entirely interchangeable. Below we will see at least one reason for preferring tropes over universals. But Armstrong prefers universals. Chief among Armstrong's reasons for preferring universals is that they provide a ready explanation of what all objects with 3 grams of mass have in common: every object with 3g mass literally share the *3g mass* universal. Tropes, on the other hand, are not shared. So what does the *3g mass* trope instantiated by one object have in common with the *3g mass* trope instantiated by another? They are both *3g mass* tropes, but in virtue of what? More generally, determinate quantitative tropes divide into equivalence classes: all of the *3g mass* tropes belong to an equivalence class, and all of the *7.5g mass* tropes belong to an equivalence class, and so on. To account for these groupings, the trope theorist has to posit a resemblance primitive.²¹ This resemblance primitive can then be used to group determinate quantitative tropes of the same magnitude. The proponent of universals, on the other hand, has no need for a resemblance primitive to explain what all the objects

²¹ Taking resemblance as primitive is not a feature unique to trope theory. Those who take properties to be classes also require primitive resemblance. And one might well think, as Lewis does, that trope theory is better off here than class nominalism since its primitive of resemblance is "less artificial, hence more acceptable" than the primitives needed by class nominalism (Lewis, 1986a, p. 26).

with 3g mass have in common: the resemblance is explained in terms of the identity of a shared universal.

If this reductive account of resemblance in terms of shared universals can be made to work more generally, then a theory of universals has a significant advantage over a theory of tropes. Whether a theory that includes universals can do without a resemblance primitive is controversial, and depends on whether the proponent of universals can give a reductive account of *partial* resemblance – such as the fact that red resembles orange more than it resembles blue. Armstrong believes that he can, and hence favours universals.²²

With Armstrong’s commitment to universals in mind, how does he propose to deal with quantity? Armstrong takes determinate quantitative properties to be what he calls *non-relational structural properties*. A structural property is constituted by the fact that the proper parts of the object having the structural property have certain properties, and may also require that those parts stand in certain relations to each other. Non-relational structural properties are those structural properties which involves the proper parts having certain properties but don’t require that those parts stand in any particular relations to each other.²³ With this in mind, here is how Armstrong characterizes a determinate mass property:

“*Being two kilograms in mass* appears to be [a non-relational] structural property. For an individual, *X*, (which may be a scattered individual), to be two kilograms in weight there must exist innumerably pairs, triples, etc. of non-overlapping individuals, where the *N*-tuples are such that each member is less than two kilograms in mass, and such that the sum of the members is the individual *X*. The

²² Armstrong thinks that he can capture partial resemblance facts using the partial identity of structural universals (see Armstrong 1989). It isn’t clear that he succeeds. Eddon (2007) argues that Armstrong’s account fails to account for resemblance. Armstrong’s account will also be inadequate if we have reason to reject the existence of structural universals. Lewis (1986a) argues against structural universals.

²³ Armstrong (1988), p. 312.

constituent individuals found in different N -tuples will include individuals having every mass that is less than two kilograms. (Or perhaps not every mass but only every whole number multiple of a certain quantum.) The property of *being two kilograms in mass* is (is identical with) all those non-relational structural properties formed by taking the above-indicated pairs, triples, etc., where the mass of the members of each N -tuple sums to two kilograms.²⁴

On Armstrong's account, if an object e has mass m , then all of the N -tuples of non-overlapping proper parts of e that sum to e will also have a total mass m . And each of those proper parts will have mass less than m . Because there are so many ways of subdividing most objects into parts, Armstrong reasons that for every mass less than m , e will have a proper part with that mass.

Armstrong says remarkably little about how this conception of properties is supposed to help explain features of quantitative properties such as family, ordering, and distance.²⁵ But it is relatively straightforward to see how this might go. Two determinate quantitative properties are in the same family iff one of them, a , is constituted in part by a proper part of any thing instantiating a instantiating the other, b . So, *5kg mass* and *2kg mass* are in the same family because any object that has *5kg mass* will have a proper part that has *2kg mass*. We can account for ordering in a similar way: a determinate quantitative property a is greater than another determinate quantitative property b iff a is constituted in part by a proper part of any thing instantiating a instantiating b .

What about distance? Again Armstrong doesn't say, but here is a natural suggestion. Let the complement of a proper part e for some object f be the unique object g

²⁴ Ibid.

²⁵ For some relevant passages, see Armstrong (1988), pp. 312-16, Armstrong (1989), pp. 106-7, and Armstrong (1978), pp. 116-31. Armstrong spends much more time focusing on whether his account can capture facts about the *resemblance* of quantitative properties in the same family. Armstrong argues that universals are needed to capture facts about resemblance. But a nominalist will reject this requirement, so Armstrong's interest in resemblance is irrelevant to our purposes. And, as noted in a previous footnote, Eddon (2007) argues that Armstrong's theory of quantity can't properly capture the resemblance facts that he is after. I'm inclined to agree.

such that (i) g does not overlap with e , and (ii) the mereological sum of $e + g = f$. Then the distance between two quantitative properties a and b , where a is greater than b , can be determined as follows: take an object f with property a and a proper part of object f with property b , then the distance between a and b is the quantitative property c instantiated by the complement of the part of f instantiating b for f .²⁶ Concatenation can be accounted for as follows: given any two quantitative properties of the same family a and b , a concatenated with b is the quantitative property c that is instantiated by the object f that is the mereological sum of an object instantiating a with a non-overlapping object instantiating b .

There are a couple of attractive features of Armstrong's account. For one, the theory of quantity just outlined is nominalistically acceptable. Second, Armstrong's account effectively offers a reductive account of family, ordering, distance, and concatenation in terms of mereology. Insofar as we must already accept mereological predicates as part of our primitive ideology in order to account for the parts of objects, accounting for quantitative properties requires no new primitives.²⁷ This is a serious theoretical advantage.

But how does Armstrong's theory fare with regards to the Problem of Insufficient Concrete Ontology? We can see how Armstrong's structural properties are supposed to

²⁶ Maya Eddon (2007) considers several other ways that Armstrong might try to capture distance facts, before concluding that Armstrong would have to rely on a metric function that assigns a distance value to every pair of quantitative properties in a family. But without an explanation of why the metric function outputs the values that it does, this amounts to taking distance facts to be fundamental. There is thus a significant advantage to adopting the alternative that I've presented here which grounds distance facts in mereology.

²⁷ I'm giving things a nominalist gloss, of course. A platonist will likely posit fundamental mereological *relations*, rather than primitive mereological predicates. Either way facts about the mereology of objects are taken as fundamental, and an explanation of the features of quantitative properties in mereological terms is desirable.

help: Armstrong reasons that, given an object e with mass m , every mass less than m will be instantiated by some proper part of e . The idea is that once we recognize the existence of structural properties we will see that there *is* sufficient concrete ontology in the world to serve as the relata of the relations that hold between determinate quantitative properties.

Unfortunately, it is not clear that objects have enough parts to guarantee this result. Consider again our familiar example of point-particles. Point-particles are typically taken to be mereological simples: objects with no proper parts. But if point-particles are mereological simples, then it seems that we don't have enough objects to guarantee the existence of determinate mass properties smaller than those instantiated by the point-particles. The problem is not restricted to point-particles. Any view on which objects are composed of mereological simples will place a similar limit on which mass properties can be instantiated; on Armstrong's theory of structural properties, determinate quantitative properties bottom out where the objects do. As a result, we once again have difficulty accounting for distance between quantitative properties: the properties required to ground facts about the distance between the masses of the electron, muon, and tau will simply not exist.

Is there a way out of this problem for Armstrong? There is if he is willing to reject the existence of mereological simples and embrace the thesis that every object has proper parts. There are two ways to go about this. One is to deny that point-particles really exist: perhaps *every* object has spatial volume (no matter how small), and can be further divided into parts that have spatial volume. On this view, the existence of one object

guarantees the existence of an infinitely descending cascade of smaller and smaller objects (with smaller and smaller determinate quantitative properties).

Armstrong favours an alternative option. We can instead reject the idea that point-particles are mereological simples, and hold that they too have proper parts. This thesis comes with some counter-intuitive consequences. For example, it requires that there can be two (or more) indistinguishable objects located at exactly the same place. After all, every point-particle with mass m will have two proper parts with mass $m/2$, three proper parts with mass $m/3$, and so on, all located at exactly the same point. Thus, there is reason to opt for the first option, and hold that every object is spatially extended.²⁸ But nothing important turns on this for our purposes: either option allows one to reject the existence of mereological simples and posit the existence of ‘gunk’ (objects not composed of mereological simples).

That the world is composed of gunk, rather than simples, seems like a genuine possibility for the actual world. If it is, then Armstrong’s theory can account for quantity at the actual world. However, it also seems like a genuine possibility that the actual world is composed of simples, rather than gunk. So the adequacy of Armstrong’s theory depends on a seemingly contingent fact about the world: whether physical objects are composed of mereological simples or of gunk. And if that is so, then Armstrong’s theory is in trouble. After all, we are looking for a *general* theory of quantitative properties – one that will be accurate no matter what the contingent facts turn out to be.

How serious is this difficulty for Armstrong’s theory? If it is truly possible for there to be a world with mereological simples but without gunk, then Armstrong’s theory

²⁸ For more objections to Armstrong’s suggestion that point-particles have parts, see Eddon (2007), pp. 6-8.

is inadequate. But what if worlds with mereological simples but no gunk are metaphysically impossible? If every world is gunky, then Armstrong's theory does guarantee sufficient concrete ontology to account for quantitative properties. I expect that this flat denial of the possibility of worlds without gunk to strike some as a desperate gambit, so let me say something in its defense. It is certainly true that it is *conceivable* that the physical objects in our world are composed of mereological simples rather than gunk. Indeed, it seems plausible that an electron is just such a mereological simple. But it is a matter of some dispute in philosophy whether conceivability entails possibility. It seems to be standard practice in metaphysics to assume that it does not. After all, it is conceivable that there are abstract objects and conceivable that there are no abstract objects. Yet most metaphysicians agree that if abstract objects exist, they exist necessarily.²⁹ So a proponent of Armstrong's theory could hold that we have another case where conceivability is no guide to possibility: the mereological structure of the physical world.

Is this restriction on the structure of the physical world plausible? A platonist might be tempted to take a hard line in response, and claim that there are different standards for the abstract and the physical: abstract structures can be metaphysically necessary, but physical structures cannot. This difference in standards does appear to be reflected to some extent in philosophical practice: philosophers routinely claim that this or that abstract structure is necessarily just so, but are loathe to dispute the contingency of physical structures.

²⁹ With a few exceptions such as fictional characters and the like which are dependent for their existence on how things stand in the physical world.

I think there are several reasons to be skeptical of this hard line. For one, it isn't clear what justifies the difference in standards. In virtue of what is it more reasonable to posit necessary abstract structure than necessary physical structure? Of course, if the nominalist is right that there are no abstract structures, then there might well be good reasons to stipulate that abstract structures are metaphysically necessary when engaging in the fiction of abstract structures. After all, it simplifies our theorizing greatly if we don't have to worry about changes in the abstract realm, or about this or that abstract structure failing to exist. But, assuming that the platonist is right in claiming that there really are abstract structures, why ought conceivability be a better guide to possibility with regards to the physical than it is for the abstract? We aren't free to simply stipulate that it is so.

A second reason for pushing back against the hard line is that there do appear to be cases where metaphysicians are happy to place constraints on the structure of the physical world. Consider the Principle of Universal Composition: given any two or more distinct objects, there is some further object of which those objects are parts. There are those who defend the Principle of Universal Composition, and those who reject it. On the opposite extreme are the nihilists, who believe that there is no object that has another as a part. The debate over the mereological structure of the world is not, generally speaking, taken to be a debate solely about the mereological structure of the actual world. If the Principle of Universal Composition is true, it holds of necessity. If nihilism is true, it holds of necessity. So it seems that there are times when metaphysicians are comfortable attributing necessary structure to the physical world.³⁰

³⁰ It is possible to dispute whether this is indeed an example of necessary metaphysical structure. Bricker (2015) argues that mereological structure is logical structure, and hence would not take this to be an

Whether or not it is feasible to resolve the Problem of Insufficient Concrete Ontology by positing additional concrete ontology is wholly dependent on whether it is feasible to posit metaphysical constraints on physical structure. For if it is not, then there will always be some possible world which lacks the requisite physical structure to account for quantitative properties. However, for the reasons just given, it doesn't strike me as especially implausible to posit metaphysical constraints on physical structure.

Of course, even if that is true, some metaphysical constraints are more plausible than others. I suspect that no one would be willing to endorse the claim that ice cream exists necessarily. So, with regards to Armstrong's theory of quantity, the question becomes whether it is plausible to insist that every world is gunky. We can codify the insistence on gunky worlds into a Principle of Universal *Decomposition*: every object has two or more distinct objects as parts. The principle isn't entirely intuitive – why couldn't there be an object with no parts? – but neither is the Principle of Universal Composition. So, it doesn't seem entirely out of the question to insist on this sort of metaphysical constraint on the structure of the physical world.

I don't want to dwell on this point, as I have nothing decisive to say on either side. The crucial point to keep in mind is that any attempt to resolve the Problem of Insufficient Concrete Ontology by insisting that there is sufficient concrete ontology will have to place some constraints on the structure of the physical world. Exactly what constraints must be placed will depend on the particular theory being put forward, but

example of necessary metaphysical structure. One could respond that the restriction that Armstrong requires is likewise a mereological (and logical, if all mereological structure is logical structure) constraint on the structure of the world. In any case, we have at least a *prima facie* example of a general feature of metaphysics: many metaphysicians attribute necessary structure to the world.

some constraint or another must be placed. This is a cost of such views, and I leave it up to the reader to determine whether the cost is acceptable.

If it is acceptable to insist that every world is gunky, where does this leave Armstrong's theory of quantity? It leaves Armstrong with a sufficient array of objects to capture the ordering, distance, and concatenation relations that hold between determinate mass properties. There remain a couple of difficulties however. One issue is that Armstrong's account doesn't offer up an explanation of the incompatibility of determinate quantitative properties. Armstrong simply stipulates that determinate quantitative properties are non-relational structural properties, and correspondingly stipulates that an object cannot instantiate more than one determinate quantitative property.³¹ But why accept these stipulations? It seems that here Armstrong is forced to rely on a brute metaphysical constraint – it is taken as necessary that quantitative properties are instantiated in the way stipulated.

There is a second difficulty faced by Armstrong's account as well. As we've seen, according to Armstrong if an object e has a determinate quantitative property a , then all of the proper parts of e will have determinate quantitative properties from the same family as a but which are lesser than a . So, for example, the proper parts of an object with $5g$ mass will all have mass properties less than $5g$ mass. This is a reasonable claim about mass. But there are other quantities for which this claim is obviously false.

Consider density. Density is an *intensive* property: the density d of a system s is not additive from the densities of the proper parts of s . If a system s has density d , proper

³¹ See Armstrong (1988), p. 313.

parts of s might also have density d . This is in contrast to *extensive* properties like mass, which are additive in that way. So, the existence of a system s with density d , no matter how mereologically complex, does not guarantee the existence of any less density properties on Armstrong's account.

This is a problem. Consider the account of ordering I offered on behalf of Armstrong: a determinate quantitative property a is greater than another determinate quantitative property b iff a is constituted in part by a proper part of any thing instantiating a instantiating b . For intensive properties, this account is clearly inadequate. It may be that there is a system with a density of 4kg/m^3 that doesn't have any proper part with a density of 2kg/m^3 . It may even be that *no* system with a density of 4kg/m^3 has a proper part with a density of 2kg/m^3 . So it seems that Armstrong's account cannot even capture facts about order for intensive quantities.

Armstrong is aware that intensive properties are problematic for his account. In response he notes that density is not a fundamental property – facts about density are reducible to facts about mass and volume. He gives a similar reductive account of temperature. But the worry remains: the existence of a fundamental intensive quantity seems possible, and it would be better to have a theory of quantity that can account for them.³² Of course, one could again respond by positing a metaphysical constraint on the physical world: all quantitative properties are extensive. But as these constraints add up,

³² Armstrong doesn't explicitly discuss the objection I mention, and it isn't clear that he is aware of *this* problematic aspect of intensive quantities. Armstrong (1988, p. 315) does suggest a way of "getting rid of such quantities": his aforementioned proposal that even point-particles have proper parts. But it isn't clear how that proposal would help with the worry that I'm raising.

Armstrong's account begins to look less and less plausible.³³ I suggest we look elsewhere.

1.8 The Mereology of Properties

There is much to like about Armstrong's theory of quantity, but we've seen that it runs into trouble accounting for quantity in a world composed of mereological simples, or a world that includes fundamental intensive quantities. Is there a nominalistically acceptable account of quantity that can do better? I think there is. Consider a pair of remarks made by Armstrong:

“It may therefore be suggested that lengths are structural properties and that any lesser length stands to any greater length as a proper part to whole.”³⁴

And, in a later paper:

“Could it be, then, that the property of *being two kilograms in mass* actually contains *being a kilogram in mass* as a part?”³⁵

These passages contain a suggestion for the mereology of quantitative properties: every determinate quantitative property has every smaller determinate quantitative property as a proper part.

Armstrong immediately distances himself from these remarks after making them. In the first case, he goes on to clarify that determinate quantitative properties are ‘particularizing’ properties, where this means that they pick out instances that are one

³³ A further difficulty arises for Armstrong with regards to vector properties – properties with both magnitude and direction. These will be discussed further below.

³⁴ Armstrong (1978), p. 122.

³⁵ Armstrong (1988), pp. 311-2.

instance of an object with the property in question. So, the parts of properties are immediately tied again to objects: for a determinate quantitative property *a* to have parts, the parts of an object that instantiates *a* must instantiate the parts of *a*. This should sound reminiscent of my characterization of Armstrong's theory in the last section, and indeed it is an earlier variant of the same view.

In the later paper, Armstrong more clearly distances himself from a mereology of properties. He writes:

“The mereological approach to the complexity of properties is a little too simple. A discussion of *structural* properties will begin to introduce the required sophistication.”³⁶

I disagree. The mereological approach has advantages over Armstrong's approach. The added sophistication of structural properties is what ultimately proves problematic for Armstrong's theory of quantity.

So let us take the mereological approach to quantity seriously. On this approach, *5g mass* has *4g mass* as a part, and *3g mass* as a part, and *2.534g mass* as a part, and *0.00004g mass* as a part, and so on.³⁷ More generally, every determinate quantitative property has every smaller determinate quantitative property as a part. The idea is that *5g mass* just *is* the mereological sum of *3g mass* and *2g mass*, which just is the mereological sum of *1g mass* and *2g mass* and another *2g mass*, etc., and hence that if we have an object that has *5g mass*, then the existence of all smaller mass properties is assured. More generally, we are guaranteed the existence of all mass properties smaller than the largest

³⁶ Ibid.

³⁷ There is a second, unrelated and incompatible, mereological approach to properties available. If properties are classes, as Lewis maintains, then the mereology of classes is also the mereology of properties. Notably, Lewis's theory of the mereology of classes will give quite different results than the mereological account that I'm offering: for Lewis, *4g mass* is *not* a part of *5g mass*. Hence, Lewis's account cannot be used to ground facts about the relations between determinate quantitative properties. See Lewis (1991) for details.

mass property instantiated at a world, since all those smaller mass properties will be parts of the largest mass property being instantiated. So the mereology approach promises the existence of a sufficient array of properties to capture the relations that obtain between determinate mass properties.

So why did Armstrong distance himself from this sort of mereological account of quantity? For a couple of reasons. First, Armstrong noticed that taking the mereology of properties seriously in this way would commit him to certain universals having other universals as parts *twice over*. For example, *4g mass* is the mereological sum of *2g mass* and *2g mass*. So *4g mass* would have the *2g mass* universal as a part twice over. And that's truly a bizarre consequence. What would it mean to have the same thing as a part twice? The worry, of course, generalizes: *4g mass* is also the mereological sum of *1g mass* and *1g mass* and *1g mass* and *1g mass*. So *4g mass* would have the *1g mass* universal as a part four times. This is hard to accept, and it's the reason why Armstrong insisted on quantitative properties being 'particularizing' universals: to say that the *4g mass* universal has the *2g mass* universal as a part twice over is to say that any object with *4g mass* has two non-overlapping proper parts with *2g mass*.

The second reason that Armstrong rejects a mereological account of properties is in response to a battery of objections that Lewis raises against structural universals in his aptly named "Against Structural Universals".³⁸ Lewis attacks the idea that we can make sense the sort of composition of structural universals that Armstrong wants in terms of mereological composition. Lewis focuses in particular on the structural universal *methane*, and its parts *carbon*, *hydrogen*, and *bonded* (a relation). As Lewis notes, carbon

³⁸ Lewis (1986a). Lewis also raises the problem of a universal having another as a part twice over in this paper, but Armstrong is already aware of the problem in his earlier *Universals and Scientific Realism* (1978).

and hydrogen can be bonded in such a way as to compose a compound other than methane. But if that's right, then the sense in which *carbon*, *hydrogen*, and *bonded* compose *methane* cannot be mereological: how, Lewis asks, could two things be composed of exactly the same parts?³⁹ Here it is Armstrong's wider commitments to structural universals such as *methane* that gets him into trouble. But there is no reason for the nominalist to take on these wider (and more problematic) commitments. There is no corresponding difficulty if we restrict our focus to determinate quantitative properties: no two determinate quantitative properties have all of the same parts.

In addition to these problems, there is a general tension between the mereological approach and a commitment to universals. Consider how the mereological approach might be used to capture family, ordering, distance, and concatenation. First, two determinate quantitative properties are in the same family iff one of them has the other as a part. So, *5kg mass* and *2kg mass* are in the same family because *2kg mass* is a part of *5kg mass*. Ordering is easily accounted for: a determinate quantitative property *a* is greater than another determinate quantitative property *b* iff *b* is a proper part of *a*. But what about concatenation? A simple suggestion would be to take concatenation to simply be a matter of taking the mereological sum: *a* concatenated with *b* gives us the mereological sum of *a* and *b*. But here we run into a difficulty. What is the mereological sum of the universal *5g mass* and the universal *3g mass*? We'd like to get the result *8g mass*, but that's not what we get: *3g mass* is a proper part of *5g mass*, and the mereological sum of something *x* with a proper part of *x* just is *x*. So the mereological

³⁹ Lewis (1986a), p. 36.

sum of *5g mass* and *3g mass* is *5g mass*. Not the result we needed. Trying to account for distance runs into similar trouble.⁴⁰

Given this tension between universals and the mereological approach, I propose we drop universals in favour of tropes. While there is only one *5g mass* universal there can be many *5g mass* tropes, and this feature of tropes will prove quite useful. Let us consider how the mereological approach might go. Let's call this the Simple Mereological Theory (SMT). According to SMT, two determinate quantitative properties are in the same family iff one of them has the other, or a trope of the same magnitude as the other, as a part. So, a *5kg mass* trope and a *2kg mass* trope are in the same family because either the *2kg mass* trope is a part of *5kg mass*, or some other *2kg mass* trope (a trope of the same magnitude as our original *2kg mass* trope) is a part of *5kg mass*. Ordering can also be accounted for as follows: a determinate quantitative property *a* is greater than another determinate quantitative property *b* iff *b*, or a trope of the same magnitude as *b*, is a proper part of *a*.

Now, whereas we saw universals run into trouble with regards to concatenation and distance, tropes have no such difficulty. Concatenation is simply a matter of taking the mereological sum of two non-overlapping tropes: *a* concatenated with *b* is the mereological sum of *a* and *b*.⁴¹ So we get the desired result that *2kg mass* and *5kg mass* concatenate to *7kg mass*. Distance requires, as before, the notion of a mereological

⁴⁰ We might try to account for distance, as we did on Armstrong's account, by relying on the idea of a mereological complement. Let the complement of a proper part *a* for property *b* be the property *c* such that (i) *c* does not overlap with *a*, and (ii) the mereological sum of *a* + *c* = *b*. Then the distance between *a* and *b*, where *b* is greater than *a*, is the complement of *a* for *b*. Unfortunately, because every determinate quantitative property overlaps with every other determinate quantitative property of the same family, there simply are no quantitative properties to play the role of complements.

⁴¹ We can also give a more general account of the concatenation of possibly overlapping tropes: *a* concatenated with *b* is the mereological sum of *a* and a trope (i) of the same magnitude as *b* and (ii) which doesn't overlap with *a* (which can be *b* itself if *a* and *b* don't overlap).

complement. Let the complement of a proper part a of property b be the property c such that (i) c does not overlap with a , and (ii) the mereological sum of $a + c = b$. Then the distance between a and b , where b is greater than a , is the complement of a trope a' , which is of the same magnitude as a but which is a proper part of b , with respect to b .⁴² So, the distance between $2kg$ mass and $5kg$ mass is $3kg$ mass, since if we take any $5kg$ mass trope and any proper part of that trope with a magnitude of $2kg$ mass, the complement of the $2kg$ mass trope for the $5kg$ mass trope will be a $3kg$ mass trope.

SMT, like Armstrong's theory of structured universals, offers a reductive account of family, ordering, distance, and concatenation, in terms of mereology. But whereas Armstrong's theory runs into trouble accounting for quantity in a world with mereological simples, SMT has no such problem. That's because SMT doesn't require that the parts of properties are instantiated by some further objects. An object with $5g$ mass also has $2g$ mass. Of course, there is a sense in which it's misleading to say that the object has $2g$ mass. It has *more* than $2g$ mass. When we talk about the mass of an object we almost always talk about the total mass of that object. But that doesn't mean that objects don't also have smaller amounts of mass. According to SMT, they do. So, even in a world consisting of a single point-particle, the mass of that point-particle would have every smaller determinate mass property as a part.

This feature of SMT also allows us to give a reductive explanation of incompatibility. In one sense, SMT simply denies incompatibility: an object with $5g$ mass can (and indeed must!) also have $4g$ mass. But SMT does offer an explanation of our intuition of incompatibility: when we ask how much mass an object has we are asking

⁴² The trope a' is introduced for the case where a is not a proper part of b . In the simpler case that a is a proper part of b , then the distance between a and b , where b is greater than a , is simply the complement of a with respect to b .

about the *total* amount of mass that the object has – that is, the mereological sum of all the mass that the object has – which will be the determinate mass property with the greatest magnitude had by the object. Because determinate mass properties are ordered, only one such property can have the greatest magnitude; it is inconsistent to attribute two different greatest mass properties to the very same object.

SMT also has no problem accounting for fundamental intensive quantities. Suppose, for example, that density was fundamental. According to SMT, $4\text{kg}/\text{m}^3$ has $3\text{kg}/\text{m}^3$ and $2\text{kg}/\text{m}^3$ as proper parts. $2\text{kg}/\text{m}^3$ is $2\text{kg}/\text{m}^3$ distant from $4\text{kg}/\text{m}^3$. Because the proper parts of $4\text{kg}/\text{m}^3$ needn't be instantiated by a proper part of any object that instantiates $4\text{kg}/\text{m}^3$, there isn't a problem account for intensive properties in exactly the same way as extensive properties.

It might be objected that intensive properties don't concatenate, since if one has an object with a density of $2\text{kg}/\text{m}^3$ and another with density $2\text{kg}/\text{m}^3$, one won't end up with a total of $4\text{kg}/\text{m}^3$. But it is important to distinguish how the quantitative properties instantiated by objects determine the quantitative property instantiated by the sum of those objects from the concatenation of the quantitative properties. The distinguishing feature of extensive properties is that if one takes the mereological sum of the objects that instantiate them, the quantitative property instantiated by that sum is the concatenation of the quantitative properties instantiated by the component objects. So, if one has a rock with 2kg mass, and a rock with 3kg mass, their mereological sum of the two rocks will have 5kg mass. But intensive properties don't function this way. No matter how many rocks with $2\text{kg}/\text{m}^3$ that one has, the mereological sum of those rocks will still have a density of $2\text{kg}/\text{m}^3$. But if we take ordering, distance, and concatenation to concern only

the relations between the properties, we can still say that $4\text{kg}/\text{m}^3$ is greater than $2\text{kg}/\text{m}^3$, and that $2\text{kg}/\text{m}^3$ concatenated with $2\text{kg}/\text{m}^3$ is $4\text{kg}/\text{m}^3$. SMT thus has serious advantages over Armstrong's account of quantity in terms of structured properties.

Of course, SMT comes with its own theoretical costs. SMT attempts to resolve the Problem of Insufficient Concrete Ontology by attributing mereological structure to tropes. SMT thus places a metaphysical constraint on the physical world: no quantitative tropes without mereological structure. The details of this mereological structure depend in part on what assumptions we make about quantitative properties. Thus far I've been assuming that quantitative properties are continuous and have no lower bound – that it is possible to have a quantitative property of any magnitude, no matter how small. If that's right, SMT must place a corresponding constraint on the mereology of quantitative properties: determinate quantitative properties are 'gunky' – they have no smallest parts.

But it is possible that, instead, determinate quantitative properties have quanta – smallest discrete magnitudes. If that's right, then SMT needn't require that determinate quantitative properties have no smallest parts; the smallest determinate quantitative properties can be the quanta. On this picture, every determinate mass trope is composed out of quanta of mass. This simplifies the picture somewhat, but it does place a constraint on the nature of quantitative properties.

I have no intention to pick one of these two options over the other. After all, it is an empirical question whether or not there are quanta, and either possibility can be accommodated by SMT. But despite this flexibility in the details, SMT does place a rather substantial constraint on the physical world by requiring that quantitative

properties have mereological structure. And that is a cost. The question is whether it is a cost worth bearing.

There is another difficulty with SMT. There is an important distinction to be made between scalar and vector quantities. *Scalar* quantities, such as mass, have a magnitude but no direction. *Vector* quantities, such as velocity, have a magnitude and a direction. An object can have a velocity of 3m/s in one direction, and another object can have a velocity of 3m/s in an entirely different direction. So far our focus has been on scalar quantities, and I have argued that SMT has serious advantages over the alternative theories we have considered so far when it comes to scalar quantities: SMT resolves the Problem of Insufficient Concrete Ontology in a way compatible with the existence of mereological simples; SMT offers a reductive account of family, ordering, distance, concatenation, and incompatibility, in terms of mereology; SMT can capture these features for both extensive and intensive properties; and SMT can do all of this while placing comparatively limited metaphysical constraints on the physical world.

However, SMT has a much more difficult time with vector properties. Consider a pair of objects moving at 3m/s in the same direction. Intuitively, there is no distance between the velocities of the two objects: they instantiate velocity tropes of both the same magnitude and the same direction. Now consider a second pair of objects moving at 3m/s in opposite directions. Intuitively, there is a distance of 6m/s between the velocities of the second pair of objects: they instantiate velocity tropes of the same magnitude but opposite direction. How can we account for this difference in direction? It isn't clear that any extension of SMT can adequately do so without running foul of the Problem of Insufficient Concrete Ontology. After all, we want to say that the two objects travelling in

opposite direction have a distance of 6m/s, but there is simply no guarantee that any object in the world has a velocity of that magnitude – perhaps we are considering a world where nothing is travelling faster than 3m/s. If so, then the requisite distance trope will simply not exist.

A similar problem arises when we consider objects moving perpendicular to each other. The distance between the two velocity tropes will be a velocity trope with a different direction than either of the two original velocity tropes. But what if nothing is moving in that third direction? There is no guarantee that anything in the world instantiates the appropriate velocity trope. So, again, we run into the Problem of Insufficient Concrete Ontology.

I suspect that the problem here is not merely technical, and points to a serious difficulty that any mereological approach will have with vector properties. Perhaps I am wrong, and SMT can be modified in such a way as to guarantee the existence of the requisite properties. But if not, then it seems that the proponent of the mereological approach will have to reject the existence of fundamental vector properties. This is essentially the tact that Hartry Field takes in *Science Without Numbers*, where he paraphrases away all talk of vectors.⁴³ There is some debate over whether Field's paraphrase strategy is successful, but I want to leave this debate aside. It is one question whether we can paraphrase away talk of vector properties, and another question whether there are fundamental vector properties in the world. It certainly seems possible that there are. But SMT and similar views must reject that possibility, and hold that all quantities

⁴³ Field (1980).

are scalar. So we see another place where SMT must place a metaphysical constraint on the physical world.

1.9 Properties, Points, and Dimensions

Is the mereological approach, despite its limitations, the best option for those who wish to resolve the Problem of Insufficient Concrete Ontology by positing additional structure in the world? Are there any alternatives that look, not to mereological structure, but to some other sort of structure in order to resolve the Problem of Insufficient Concrete Ontology? In this section, I'd like to consider one more approach.

To introduce this alternative, it will be fruitful to begin by considering how quantitative properties are modeled mathematically. The magnitude of a property can be identified with a point on a number line. So an object with a certain mass can be identified with the number 5, and an object with twice that mass can be identified with the number 10. This feature is central to the way we assign names to determinate quantitative properties: we talk about *0 grams* or *5 grams* or *10 grams* of mass. This assignment of numbers to masses works best when structural features of the number line capture structural features of the quantitative property: when, for example, the addition of the numbers mirrors the concatenation of the determinate quantitative properties represented by those numbers.

Some quantitative properties call for a more complex mathematical representation. Vector properties have both magnitude and direction. One way to

represent this additional information is with a vector space. We can represent a velocity, for example, using a three-dimensional space: every point in this space represents a particular velocity (both magnitude and direction). We can then represent the velocity of an object by representing the object as occupying a particular point in the vector space. Vector spaces are thus a means of mathematically representing quantitative properties. In order to do this, vector spaces must have a sufficiently rich mathematical ontology; they must have enough points, arrayed in a space with sufficient dimensions, so that every possible magnitude and direction of the property can be represented as a point. This ontological requirement is not a problem: mathematicians are free to posit as many dimensions as are useful to represent the desired information. The same holds true of the use of number lines to represent scalar properties: mathematicians are free to posit these simple mathematical structures as often as it is fruitful to do so. There is no Problem of Insufficient Mathematical Ontology.⁴⁴ The Problem of Insufficient Concrete Ontology arises because we are not free to simply stipulate the existence of additional concrete structure in the world: we must make do with the concrete world that actually exists.

But what if we take all this mathematical talk of vector spaces as more than a mere stipulation of abstract mathematical structures? What if we instead hold that there are corresponding structures in the world? The proposal is that in addition to the familiar spatio-temporal dimensions (and whatever other dimensions scientists might posit for string theory and so on), there are additional dimensions corresponding to quantitative

⁴⁴ At least, there is no such problem given a reasonably abundant platonist theory. Nor is there a problem for the nominalist who takes a fictionalist approach to abstract objects. But there is room for a problem here: it is possible for one to be a platonist and to also adopt a very restricted account of which abstract mathematical objects exist. I will leave aside this possibility here, and assume that the platonist is happy to accept the existence of whatever mathematical objects the mathematician finds useful to posit.

properties.⁴⁵ With these additional dimensions we also get additional points: we no longer have mere space-time points, but rather points that are located along these additional dimensions as well.⁴⁶ We can then locate objects not just in space-time but in the relevant additional dimension(s) as well. So, for an object to have 10g mass is for it (or perhaps just its mass) to be located at a certain point in the mass-dimension.⁴⁷ Let us call this the *Additional Dimensions Theory* (ADT).

This strategy attempts to resolve the Problem of Insufficient Concrete Ontology in a more direct way than the mereological approach we discussed before. The mathematical structures used to represent quantitative properties do not suffer from insufficient ontology, so if we posit corresponding concrete structures then we won't suffer from insufficient concrete ontology. This approach thus has a serious advantage over the mereological approach we considered earlier: it has no problem accounting for vector quantities. This approach is general and powerful: mathematical structures can be used to represent a wide variety of information, and the existence of corresponding concrete structures would allow us to represent all manner of information about an object by positioning that object in these additional dimensions.

Like all other attempts to resolve the Problem of Insufficient Concrete Ontology by positing sufficient concrete structure, ADT faces a familiar difficulty: there are some

⁴⁵ See Bricker (2017) for a version of this approach.

⁴⁶ There are a number of options here. There is a choice of how to understand these dimensions: is there a single mass-dimension, or a separate mass-dimension at every spacetime point? There is also a choice of how to understand the points that occupy these dimensions: do we introduce a new variety of points to occupy these new dimensions (and do spacetime points remain *mere* spacetime points), or do all points occupy every dimension (having both a spacetime location, and a location in all of these additional dimensions)? I'm going to simply bypass these details in favour of a more general discussion.

⁴⁷ On Bricker's (2017) view, it is determinable quantitative properties, rather than objects, that are located in these additional dimensions (what Bricker calls "vertical" structure in contrast to the more familiar "horizontal" spatio-temporal structure). But these particular details won't matter for our purposes: the relevant feature is the use of additional dimensions to capture the magnitude of quantitative properties.

apparent possibilities that cannot be accommodated. In particular, ADT cannot accommodate the existence of quantitative properties in a world with only spatio-temporal dimensions. Rather, every fundamental quantitative property must be accompanied by at least one additional dimension and an associated increase in the number of points. ADT thus solves the Problem of Insufficient Concrete Ontology with an expanded concrete ontology.

ADT places quite different constraints on the concrete world than the mereological approach we considered previously. Which one is to be preferred? I expect that intuitions about this will vary considerably based on one's other philosophical commitments. Some philosophers might see the close parallel that ADT draws between mathematical models and the concrete world as a positive feature of the view: it allows us to draw an extremely tight connection between mathematical structures and physical structures. Other philosophers are likely to view this close parallel as a defect: the utility of mathematics first led the platonist to accept the existence of abstract mathematical objects (where there are none) and is now being used as a motivation to posit concrete surrogates for mathematical structures. Additional dimensions corresponding to quantitative properties provide a powerful theoretical tool, but they come at a steep cost. After all, mass does not, on the face of it, seem like the sort of phenomenon that calls for the postulation of an additional dimension.

I cannot hope to adjudicate between these two views here. The point I wish to draw attention to is again more general: both views place restrictions on the nature of the concrete world, and neither can accommodate the existence of vector properties in a merely spatio-temporal world. Whether the restrictions placed by one or the other of

these views are an acceptable cost for resolving the Problem of Insufficient Concrete Ontology is another matter. Can we do better? Is there a way of resolving the Problem of Insufficient Concrete Ontology that can accommodate the existence of vector properties in a merely spatio-temporal world? The only way forward that I can see is to abandon the project of trying to resolve the Problem of Insufficient Concrete Ontology by positing additional physical structure, and instead opt for a *modal* approach.

1.10 The Modal Approach

If the concrete world does not have a sufficiently rich ontology and structure then the Structural Similarity Account is in trouble: we cannot explain the utility of talk about abstract objects to science by positing a structural similarity between the abstract objects and some part of the physical world. What to do? If abstract objects aren't being used to capture actual concrete structures, then why are they so useful? The only plausible way forward that I can see, that doesn't leave the success of science entirely mysterious, is to hold that we use talk about abstract objects to capture certain modal features of the world.

It will be instructive to focus again on our example of quantitative properties. The idea now being considered is that we needn't posit a sufficient array of concrete objects or properties to capture all the ordering, distance, and concatenation facts that we want. We can instead rely on how, for example, mass properties *would* concatenate were they to be instantiated. This modal strategy can be put to use in a variety of ways.

One use of the modal strategy is to revive a theory of quantity that relies on primitive ordering and concatenation. A realist about properties can posit primitive facts of the form “*x would be greater than or equal to y*” and “the concatenation of *x* and *y would be z*” where these are taken to capture the relations that would obtain between determinate quantitative properties were they to exist. Using these primitive modal facts, further modal facts about what family determinate quantitative properties would belong to, and what distance relations would obtain between those properties can be specified. Because these are modal facts, and don’t rely on the actual instantiation of the determinate quantitative properties in question, this view does not run afoul of the Problem of Insufficient Concrete Ontology: there is no problem accounting for tricky cases like intensive properties and vector properties, since we can simply stipulate that our ordering and concatenation primitives give us the desired results.

That isn’t to say that this sort of modal approach isn’t without its costs. For one, we lose our reductive explanation of the incompatibility of quantitative properties of the same family. The mereological approach provides us with a ready explanation: there can only be one mereological sum of quantitative properties instantiated by a given object, on pain of contradiction. The additional dimensions approach likewise provides us with a ready explanation: an object (or property) can only occupy one point in these additional dimensions at a time. But the modal approach offers us no such story, and as a result has to take the incompatibility of quantitative properties as primitive.

The modal approach also places metaphysical constraints on the world, just as the non-modal approaches we considered before did. Notably, it is a different sort of constraint: rather than requiring a certain minimum structure to be instantiated, the modal

approach takes there to be primitive facts about how objects and properties *would* be related were they to exist. Which sort of constraint is theoretically preferable I cannot say. And no matter which cost one thinks is worse, there is a cost to be paid either way.

We might try to lighten this cost of the modal view by grounding the primitive modal facts about quantitative properties in some way.⁴⁸ Perhaps these modal facts follow from the nature, or essence, of the quantitative properties. Or perhaps these modal facts fall out of the laws of nature (if we can make sense of that metaphor). But I'm not convinced that these options do anything but reshuffle where the primitive facts lie. If we take these modal facts to follow from the nature of quantitative properties then we've simply packed more into the 'nature' of the quantitative properties. And, moreover, we're left without an explanation of how the modal facts are grounded in the nature of quantitative properties: this will have to be taken as primitive, or constitutive of the nature of the properties, or something of the sort.

Another option is to opt for a hybrid of the modal and mereological approaches, and thereby obtain for a partly reductive account of the relations that quantitative properties bear to one another. We saw that SMT has a difficult time accounting distance and concatenation for vector quantities. We could retain SMT's mereological account of family, ordering, and incompatibility, but introduce a modal concatenation primitive: "the concatenation of x and y *would be* z ". The distance between a and b could then be defined using the concatenation primitive. The general idea is that while magnitude can

⁴⁸ It is worth noting that a general reductive account of modality is not relevant here: it won't help to say that these modal facts reduce to facts about how things stand at a plurality of concrete possible worlds. After all, we can ask why all the possible worlds in which the relevant quantitative properties do exist are worlds in which the quantitative properties are related in the required way. Why are there no worlds, we might ask, where the ordering facts are different? And the proponent of a plurality of possible worlds will have to take this fact about which worlds exist to be primitive.

be captured mereologically, we need to rely on modal facts to capture how vector tropes with different directions concatenate. This sort of hybrid view is something of an unhappy compromise, but it might be the best option available to those who both want to allow for the possibility of vector properties, and want a reductive account of quantitative properties where possible.

One final issue worth reflecting on is a tension between the modal approach, an intuitive understanding of what a ‘structure’ is, and the Structural Similarity Account. Intuitively, we might think that a structure just is some pattern of relations. But if that’s right, and we opt for the modal approach, it turns out that the world doesn’t instantiate the relevant structure: actual concatenation relations have been replaced, at least in some cases, with modal facts about what relations *would* (but don’t actually) obtain. Because we don’t have all the relations needed for the relevant structure, we don’t have the structure. And without the structure, the Structural Similarity Account is in trouble.

There are two ways to resolve this tension on the modal approach. The first is to revise our idea of what a structure is, and to take modal facts about what relations would obtain as relevant to determining what structure is actually instantiated. On this view, a structure is more than a mere pattern of relations: modal facts shape the structure as well. This allows for complex structures to be instantiated even in simple worlds with few objects: we could have the full structure of mass properties even in a world containing a single object, provided that the relevant relations would obtain were the world to contain a sufficiently rich array of objects.

The second option is to retain our intuitive understanding of what a structure is and to instead revise the Structural Similarity Account. Perhaps the success of our talk about abstract objects in science can be explained by the fact that the abstract structures would be relevantly similar to the concrete structures that would exist were the world to contain a sufficiently rich ontology. Thus, our talk of abstract quantitative properties is useful because it captures not only the relations that do obtain between concrete quantitative properties, but also modal facts about what would be the case were more concrete quantitative properties to be instantiated in the world. On this approach, it might still be useful to adopt a fictionalist stance towards the structures, and talk about abstract models fitting the structure of the world – even if the world does not actually have that structure. Either way, the modal approach can be reconciled with some variant of the Structural Similarity Account.

1.11 Some Concluding Remarks

Where does all of this leave us? We began with a puzzle about how to make sense of the utility of abstract objects to the scientific study of the physical world. I suggested a solution – the Structural Similarity Account – and then raised a problem for that solution: the Problem of Insufficient Concrete Ontology. I think the problem here is quite serious, since I know of no plausible alternative to the Structural Similarity Account. If we want to explain the success of science, something must be said about why facts about abstract structures are so useful in describing and explaining the physical world. I've remained

neutral about exactly what sort of similarity relation is required. Some would argue that the connection between the abstract and the concrete needs to be as strong as isomorphism, or even identity, but the problem emerges even if we only insist that some sort of similarity relation is required. And there cannot be an appropriate similarity if the world lacks a sufficiently rich ontology. I hope that I've made the problem salient at least for the metaphysics of quantitative properties, although I hope it is clear that the issue is much more general than that.

We've considered two ways of responding to the problem. The first is to insist that the world does have a sufficiently rich ontology. But this requires placing metaphysical constraints on the structure of the physical world, something that philosophers have been reluctant to do. The second response is to make up for the lack of a sufficiently rich ontology by relying on modal claims about how things would be were the world to have a sufficiently rich ontology. This too requires placing metaphysical constraints on the physical world, although now the constraints are modal in character. Is there another option? One might have thought that the expansive ontological resources available to the platonist would help, but I have argued that this is not the case: given the particular puzzle that the Structural Similarity Account aims to solve, a surplus of uninstantiated entities is not an adequate substitute for the sought after physical structures. So, I see no plausible alternative to the two that I've considered.

If we want to take quantitative properties seriously (in a realist sense or not), we must make some sort of corresponding requirement of the physical world. The cost must be paid somewhere, somehow. Abstractly specifying the structure of something while remaining agnostic about any corresponding structure in the concrete world has its

advantages, but it falls short of a fully developed metaphysics. If the Structural Similarity Account is roughly correct, then abstract structure simply isn't enough.

CHAPTER 2

LANGUAGE WITHOUT ABSTRACT OBJECTS

2.1 Introduction

What is language? To quote David Lewis, it is “a social phenomenon which is part of the natural history of human beings; a sphere of human action, wherein people utter strings of vocal sounds, or inscribe strings of marks, and wherein people respond by thought or action to the sounds or marks which they observe to have been so produced.”⁴⁹

In both philosophy and linguistics, the study of language often proceeds in an explicitly platonistic framework. Lewis immediately contrasts language with *languages*, which he characterizes as particular set-theoretic entities: functions from token utterances to sets of possible worlds.⁵⁰ He then proceeds to explain how languages – these abstract objects – play a crucial role in language. Lewis is far from alone in his platonistic theory of language. One needn’t look very hard to find some explanation or another of how abstract objects factor into language: sentence types, word types, and propositions are among the abstract entities that are taken to play a central role in language.

Nominalists reject the existence of abstract entities. Given the seemingly central role played by abstract linguistic entities in our theories of language, any aspiring nominalist is thus faced with a considerable explanatory burden: how, if there are no abstract objects, are we to make sense of language? One option, of course, is to offer an

⁴⁹ Lewis (1975), p. 3.

⁵⁰ Ibid.

alternative theory of language that makes no mention of abstract entities at all. But I think that nominalists who bet that the best linguistic theories will turn out to eschew reference to abstract entities are doing themselves a disservice. After all, platonistic theories of language still dominate the academic literature, and it is plausible that as the dust settles the best linguistic theories will turn out to refer to abstract entities, just as our best scientific theories refer to abstract mathematical entities. Nominalism is a much stronger ontological thesis if it can account for that possibility. So, while the strategy of eliminating reference to abstract entities is not entirely implausible, I will set it aside here. My project in this paper is conditional. Let us suppose that our best theories of language *do* turn out to involve reference to abstract linguistic objects. What then for the nominalist?

The question that the nominalist faces here is quite similar to questions that nominalists face elsewhere: abstract objects feature not just in our best theories of language, but in many of our best philosophical and scientific theories more generally. The case that has drawn by far the most attention in the literature is mathematics. Mathematics appears to be indispensable to our best scientific theories. Given that mathematical sentences frequently refer to or quantify over mathematical entities, it appears that nominalism is at odds with our best scientific theories. If nominalism is to be at all plausible, then this apparent tension must be dealt with. Some story must be told about why nominalism and our best scientific theories are not in conflict despite first appearances.

There are a number of responses a nominalist can make. One strategy which has received considerable attention recently, and which I will focus on in this paper, is

sometimes called ‘easy road’ nominalism.⁵¹ Easy road nominalists accept that mathematics plays an indispensable role in our best scientific theories, but deny that the contribution that mathematics makes to the success of science gives us any reason to believe in the existence of abstract mathematical entities.⁵² The key to this strategy is to recognize that there is a (rather significant) difference between taking mathematical *discourse* to be indispensable to our best scientific theories and taking mathematical *entities* to be indispensable to our best scientific theories. Mathematical discourse may well be indispensable to our scientific theories for reasons quite unconnected to the actual existence of mathematical entities. If so, then the use of mathematical discourse would give us no reason to posit the existence of abstract mathematical entities.

While much of the literature has focused specifically on applying the easy road strategy to mathematical entities, there is nothing that prevents us from applying it elsewhere. Comparatively little has been said in defense of nominalism with regards to abstract linguistic entities. This is understandable: because mathematics seems to be indispensable to our best scientific theories, it has long been regarded as a hard case for the nominalist. And there is good reason to suppose that the strategies that the nominalist

⁵¹ There are a number of alternative nominalist strategies that I will not be discussing in this paper. ‘Easy road’ nominalism is most frequently contrasted with the ‘hard road’ nominalist strategy of trying to show that mathematical discourse is not indispensable to science but is rather a practical convenience that could be eliminated if we were so inclined. See Field (1980). I do not argue against hard road nominalism in this paper, but simply concede to the platonist that mathematics plays an indispensable role in our best scientific theories.

Another nominalist strategy is to argue that mathematical entities are *not* abstract. Yet another nominalist strategy is to deny that mathematical sentences are factive – that is, to deny that mathematics attempts to tell us anything about how the world is. See Pérez Carballo (2014). If this were right, then mathematical sentences would not be the sort of sentence that carries ontological commitments. In this paper I will have little to say about any of these alternatives to easy road nominalism.

⁵² This type of strategy has been defended by Melia (2000), Azzouni (2004), and Leng (2005), among others. Balaguer (1998, 2009) defends the very similar easy road to agnosticism (he employs the easy road strategy, but ultimately remains agnostic about the existence of abstract entities). Yablo (1998, 2000, 2005, 2012) defends the easy road to there being no fact of the matter about whether mathematical objects exist (he too employs the easy road strategy without ultimately endorsing nominalism). In what follows, I will lump Balaguer and Yablo in with the nominalists.

can employ with regards to mathematics can likewise be applied to other domains. But it would be a mistake to think that the same strategies can be carried over to a nominalistic treatment of language without modification. Abstract objects play a role in language that they do not play in mathematics, and as a result applying the easy road strategy to language requires special consideration.

Here is a brief roadmap for the rest of this paper. My primary aim is to assess the applicability of the easy road strategy to abstract linguistic entities. In the next section I will introduce easy road nominalism in more detail, and note a special challenge that language presents. I will then discuss two nominalist treatments of proposition found in the literature, and some shortcomings of those accounts. One lesson that becomes clear in this discussion is that for an easy road nominalist treatment of language to be plausible, it must be paired with a nominalistically acceptable theory of content. In §5 I offer such an account. In section §6 I widen my focus from propositions to reference and to abstract linguistic entities more generally. I conclude with some remarks on semantics, and a defense of the idea that it is acceptable for there to be a significant mismatch between the sentences that language users take to be true and the sentences that come out true on our semantic theory.

2.2 Easy Road Nominalism

Easy road nominalism begins by making a major concession to the platonist: reference to, and quantification over, abstract entities is an indispensable part of our best

theories (or, at least, is so useful as to be practically indispensable).⁵³ This is a serious concession, but it is important to recognize that it doesn't amount to an endorsement of abstract entities. As I noted in the introduction, it is one thing for discourse about certain entities to be indispensable and quite another thing for those entities to themselves be indispensable.

Let us focus, for now, on the case of mathematical entities. The question we are faced with is this: what, if anything, can be said in favour of moving from the indispensability of mathematical discourse to the existence of mathematical entities? Platonists are sometimes inclined to think that no argument is needed: mathematical discourse is *obviously* true, and hence mathematical entities exist. But this is too quick. To bridge the gap between the indispensability of the discourse and the existence of the entities requires that the platonist establish two claims: (1) that mathematical discourse carries an ontological commitment to mathematical entities, and (2) that mathematical discourse is true.⁵⁴

But *is* mathematical discourse obviously true? Not if mathematical discourse carries an ontological commitment to abstract entities. The ontological status of abstract

⁵³ This statement will undoubtedly strike some readers as self-undermining: if reference is understood as a relation, then you *can't* refer to something that doesn't exist (since the relation cannot hold without the relata). Throughout this paper, I use 'reference' in a looser sense to also cover cases of empty reference. It is the referential behavior of language users that is of interest to the easy road nominalist, regardless of whether these attempts at reference manage to genuinely refer to anything at all. Likewise for talk of quantification: it is the attempts at quantification over abstract objects that are of interest to the easy road nominalist, whether or not such attempts succeed at genuinely quantifying over anything. The easy road nominalist doesn't need to explain the presence of genuine reference and quantification over abstract objects, but rather the presence (and utility) of referential and quantificational language in our best theory. I discuss this issue in more depth in §6 below.

⁵⁴ In phrasing things this way, I'm glossing over some complications. It is not necessary that all mathematical discourse be ontologically committing. Nor will the truth of any mathematical sentence do – to take an obvious example, the truth requirement would not be met by platonists showing that “there are no numbers” is true. It might be more accurate to say that what is required is that there is some sentence that is true, indispensable, and that carries an ontological commitment to a mathematical entity. But this won't quite do either – it is unlikely that any single sentence is indispensable.

entities is not a trivial or obvious matter – philosophers have spent a considerable amount of time arguing over precisely this point (this paper included). If the truth of mathematical discourse turns on whether or not abstract entities exist, then the truth of mathematical discourse is no trivial matter.⁵⁵ This is not to say that it's entirely unreasonable to hold that mathematical discourse is obviously true. Rather, it's to say that it is unreasonable to hold that mathematical discourse is obviously true *while also* maintaining that the truth of mathematical discourse requires the existence of abstracta. Deep ontological questions are not so simply solved.

Now, does mathematical discourse carry an ontological commitment to abstract mathematical entities? Platonists say yes. So do some easy road nominalists. For such nominalists, mathematical discourse carries an ontological commitment to abstract mathematical entities, and hence mathematical discourse is strictly speaking false (given that there are no such entities).⁵⁶ Other easy road nominalists disagree. For such nominalists, the truth of a typical mathematical utterance does not require the existence of mathematical entities.⁵⁷ This debate turns on subtle issues that I intend to bypass here: let us grant that the platonist is correct in thinking that mathematical discourse carries an ontological commitment to abstract mathematical entities. This has the advantage of allowing us to see that the plausibility of easy road nominalism does not depend on

⁵⁵ For a discussion of this point, see Balaguer (1998).

⁵⁶ See Balaguer (1998), Melia (2000), Leng (2005).

⁵⁷ There are important differences among easy road nominalists who deny that the truth of a typical mathematical utterance relies on the existence of mathematical entities. Easy road nominalists that would self-identify as 'hermeneutic fictionalists' hold that typical mathematical utterances are *non-literal* (where this may or may not amount to being uttered partly on the basis of pretense), and that non-literal utterances of mathematical sentences do not carry an ontological commitment to mathematical entities. See Yablo (2001, 2005). Jody Azzouni (2004), on the other hand, argues that a literal utterance of "the number 2 exists" is true (and about the number 2), despite the fact that numbers don't exist.

rejecting standard platonist assumptions about the content and ontological commitments of mathematical discourse.⁵⁸

If bridging the gap between the indispensability of mathematical discourse and the existence of mathematical entities requires an argument, what should the argument be? It seems to me that the most promising platonist strategy here is an inference to the best explanation: the best explanation of the indispensability of mathematical discourse is that there are mathematical entities, and those entities have the properties that the discourse tells us they do. After all, one might suppose, isn't the indispensability of mathematical discourse entirely mysterious unless there are in fact mathematical entities?

This explanation is clearly unsuitable for the easy road nominalist, and hence an alternative explanation of the indispensability of mathematical discourse must be provided. Let us consider how this alternative might go. Consider the sentence (S₁): "The number of objects in the room = 2". S₁ is a claim about the identity of two numbers – the number of objects in the room and the number 2. But the purpose of uttering S₁ is not to make a claim about the platonic realm. Rather, an utterance of S₁ is used to convey a fact about the physical world: there is an object x in the room, and a distinct object y in the room, and every object in the room is identical to x or y. Here, mathematical discourse is

⁵⁸ Of course, if mathematical sentences carry ontological commitments to abstract entities, and there are no such entities (as the nominalist maintains), then such sentences must, strictly speaking, be false. Some platonists have taken this to be a serious problem: David Lewis colorfully remarked that that he would be "moved to laughter at the thought of how *presumptuous* it would be to reject mathematics for philosophical reasons" (1991, p. 59). As a number of philosophers have emphasized however, easy road nominalism does not require the rejection of mathematics in any serious way: the nominalist is (quite rightly) not suggesting any change in scientific practice. A proper assessment of science and mathematics shows that the ontological issue at stake between nominalists and platonists is entirely mathematically unimportant and uninteresting, and makes no difference to the empirical success of those theories. Even if a scientific theory is made literally false by the non-existence of numbers, that theory will still, as Mark Balaguer puts it, "paint an essentially accurate picture of the physical world" (2009, p. 134). For a discussion of these issues see Balaguer (1998 and 2009), Leng (2005), and Daly (2006). See also my discussion in "Language Without Abstract Objects".

being used to make a nominalistically acceptable claim about the world. The mathematical discourse is dispensable, but it has a certain descriptive convenience. To highlight this point, consider sentence (S₂): “The number of objects in the room = 2317”. As with S₁, what S₂ tells us about the physical world can be specified in a nominalistically acceptable vocabulary (exercise left to the reader). But S₂ allows us to express what we want to say about the physical world much more simply than we could without mathematical discourse.

In other cases, what a mathematical sentence requires of the world can only be specified in a mathematics-free vocabulary by an infinite disjunction. Consider the sentence (S₃): “The average star has 2.4 planets”.⁵⁹ For S₃ to be true, it must be the case that: either there exists 5 stars and 12 planets, or there exists 10 stars and 24 planets, or there exists 15 stars and 36 planets, etc. The sentence S₃ thus has an indispensable expressive brevity – it can be used to express nominalistically-acceptable content that otherwise couldn’t be specified without an infinite disjunction.

Now consider a case discussed by Yablo. Yablo has us imagine a physicist studying escape velocity. As a result of her experiments, the physicist knows many facts of the following form: “(A) A projectile fired at so many meters per second from the surface of a planetary sphere so many kilograms in mass and so many meters in diameter will (will not) escape its gravitational field.”⁶⁰ Now, this is not a sentence about mathematical entities – it is a sentence about projectiles, and an interesting connection between their velocities and their masses when subject to the effects of gravity. But it is close to impossible to talk about velocities or masses without bringing in mathematical

⁵⁹ This example is drawn from Melia (1995), p. 226.

⁶⁰ Yablo (2005), pp. 94-5.

language. The physicist will have difficulty recording the facts of form (A) without quantifying over numbers. Facts of form (A) even *feel* mathematical in a sense, since we are so familiar with the type of mathematical language normally used to express such facts.

Let us consider how the physicist might try to record the facts of form (A) without quantifying over numbers. She might try to do so using numerical adjectives: “A projectile fired at 1 meter per second..., a projectile fired at 2 meters per second..., etc.”⁶¹ Unfortunately, this is simply impossible to do: given that velocities range along a continuum, our physicist would have to write uncountably many sentences using uncountably many primitive numerical adjectives. What alternatives does the physicist have? As Yablo notes, the obvious alternative is to quantify over numbers. The physicist can drop the idea of using numerical adjectives and say instead: “(B) For all positive real numbers M and R , the escape velocity from a sphere of mass M and diameter $2R$ is the square root of $2GM/R$, where G is the gravitational constant.”⁶²

Here the physicist is making use of numbers as *descriptive aids*. The physicist has no interest in numbers, and is trying to convey something that has nothing to do with numbers – facts about the escape velocity of physical objects. Yet, mathematical language appears to be indispensable to the physicist. Without using mathematical discourse, the physicist is simply unable to report her findings.

An important feature of the utility of mathematics as a descriptive aid is that it is nominalistically acceptable; the descriptive advantages conferred by mathematical

⁶¹ Yablo is here supposing that numerical adjectives do not carry an ontological commitment to numbers. If one thinks that the use of numerical adjectives does carry an ontological commitment to numbers, this will only reinforce the overall point that we cannot express the facts of form (A) without relying on mathematical language.

⁶² Ibid.

discourse don't depend on the actual existence of numbers. The physicist is merely attempting to convey the physical facts of form (A) in a concise way. Mathematical language is useful here because of what it can be used to express about the physical world. The truth of sentence (B) requires, in addition to the existence of real numbers, that the physical world behaves in a certain way: that the facts are of form (A), rather than some alternative (A)*. Using mathematical discourse, and talking as if there are numbers, allows us to express what we want to express about the physical facts whether or not there really are numbers. As Yablo puts the point: "That (B) succeeds in gathering together into a single content infinitely many facts of form (A) owes nothing whatever to the real existence of numbers. It is enough that *we understand what (B) asks of the non-numerical world*, the numerical world taken momentarily for granted. How the real existence of numbers could help or hinder that understanding is difficult to imagine."⁶³

These considerations point to a nominalistically acceptable explanation of the indispensability of mathematical discourse: mathematical discourse is an indispensable descriptive aid. If this is right, then the prevalence of mathematical discourse in our best scientific theories is no evidence for the existence of abstract mathematical entities.⁶⁴

One important question we are left with is *why* mathematics should turn out to be so useful; there remains something mysterious about the utility of mathematical language as a means of expressing nominalistically acceptable facts. If there really are no

⁶³ Yablo (2005), p. 95. The italics are Yablo's. Yablo makes specific reference to the *content* of (B), but it should be noted that there is no specific requirement that the expressive power of (B) comes from it carrying the facts of form (A) as part of its content. It suffices that the facts of form (A) can in some way or another be easily understood or inferred from an utterance of (B).

⁶⁴ This is, of course, not the end of the debate. The platonist can argue that there are *other* uses of mathematics that are not nominalistically acceptable. In particular, they can point to the fact that mathematics not only plays a descriptive role in science, but also an *explanatory* role. I have argued elsewhere that the explanatory role of mathematics in science is also nominalistically acceptable. See my "Fictionalism, Content, and Explanation". See also Leng (2012).

mathematical entities, then what is it about mathematics that makes it so *uniquely* well suited to playing this role? Why is no alternative, non-mathematical, language available instead? Easy road nominalists have often neglected to say much in response to this charge, which can contribute to the perception that some sleight of hand is going on. So let me say something about it here, in hopes of dispelling the mystery.

Mathematical language is particularly well suited for describing and reasoning about purely *structural* features of the world, for several reasons.⁶⁵ First of all, we are free to stipulate the existence of mathematical entities with the desired structural properties.⁶⁶ Rather than doing the difficult and unnecessary empirical work of finding objects that stand in the structural relations we want to study, we can simply focus on the structure being described mathematically. Moreover, we can follow the standard practice of taking mathematical entities to have their structural properties necessarily, thus avoiding the need to worry about changing structural features. Furthermore, because our interest is solely in structural relations, mathematical entities are ideal: we can simply stipulate that they lack the properties we are uninterested in. These are serious advantages. By divorcing our focus from particular physical systems and the structures they instantiate, we are able to say something more general about any system that instantiate the relevant structure. Thus, if our interest is in talking generally about structure, there are considerable benefits to focusing our attention on (non-existent)

⁶⁵ The suggestion that mathematics is concerned with structure is certainly not a novel one. Mathematical structuralists have long held that mathematics is the study of structure. Of course, they typically take structures to be abstract entities of some kind. That obviously will not do for the nominalist. There are various options open to a nominalist who wishes to recognize that the world has structure. One might, for example, opt to view structural properties as *in res* properties of some variety (either as universals, or as tropes). Or, one might eschew properties altogether, and hold that the truth of statements about the world's structure do not require the existence of properties.

⁶⁶ There are some limits on this stipulation of course. We can't make inconsistent stipulations (at least without adopting some appropriate paraconsistent logic). But then again, logically inconsistent structures would not be of much use in studying the structures actually instantiated in the physical worlds.

mathematical entities. Similar benefits can be gained by stipulating appropriate non-mathematical abstract entities: abstract models are of use in scientific reasoning for similar reasons.

A second advantage of mathematical language is that the names of many mathematical entities are generated algorithmically. There are an infinite number of structural properties that a system might instantiate. Two objects might be 1 meter apart, or 2 meters apart, or so on. Because we cannot learn an infinite number of names for different distance relations, the names for various distance relations must be generated in an algorithmic way that transparently encodes information about how the distance relation referred to by any given name compares to other distance relations. Numerals are suitable for this purpose, since they are generated in an algorithmic way that syntactically encodes information about the mathematical entity in question. We can readily determine various features purportedly had by the number referred to by any numeral, even if we've never seen that numeral before. Together, these two features of mathematical language make it extremely useful for describing and reasoning about the structure of the world.⁶⁷

There is of course much more to be said, but this should suffice as an introduction to easy road nominalism. My aim in this paper is not to give a thorough defense of the easy road strategy.⁶⁸ Rather, I will assume that the general strategy is successful with regards to

⁶⁷ For an extended discussion of some other features of mathematical objects that suggest that play a merely descriptive role in our discourse, see Yablo (2000), pp. 301-4, and Yablo (2001), pp. 22-3. Yablo (2000) doesn't restrict his claim to mathematical objects, but applies it to all 'Platonic Objects' – those objects "with a tendency to turn up unexpectedly in truth-conditions" (2000, p. 277). These include not just abstract objects but also entities like possible worlds and events.

⁶⁸ The skeptical reader is referred to the works cited in footnote 52 above.

abstract mathematical entities, and evaluate whether it can be extended to deal with abstract linguistic entities as well.

The first thing to note is that there is nothing about easy road nominalism that restricts its applicability solely to mathematical discourse. The general insight behind easy road nominalism is that reference to, and quantification over, certain entities can have utility (even indispensable utility) altogether unrelated to the existence of any such entities. One cannot infer from this, of course, that there are no such entities. But if the presence of referential and quantification language in our theories can be explained independently of the existence of the entities purportedly referred to or quantified over, then the presence of this language in our theories is no evidence for the existence of those entities either.⁶⁹

That said, there is a complication that arises when the easy road strategy is applied to abstract linguistic entities that doesn't arise in the case of mathematics. In providing a nominalistically acceptable explanation of the utility of mathematical discourse I leaned heavily on the idea that mathematical discourse can serve a critical descriptive role: it can be used to communicate *content* about the physical world. We've thus far been taking it for granted that we can utter sentences, and that those sentences have content. This seems like a reasonable assumption in general, but it looks rather suspect when applied to abstract linguistic entities. Suppose that I were to simply claim that talk about propositions plays a useful descriptive role (by communicating

⁶⁹ Easy road nominalism thus involves a rejection of (at least some part of) the Quinean account of ontological commitment. It does not, however, require a direct rejection of Quine's (1948) criterion of ontological commitment, formulated as follows: "A theory is committed to those and only those entities to which the bound variables of the theory must be capable of referring in order that the affirmations made in the theory be true" (p. 33). This is because it is open to the easy road nominalist to hold that our best theories aren't strictly speaking true – and hence that we needn't (and oughtn't) accept the ontological commitments of our best theories. Those easy road nominalists who do think that our best theories are true reject Quine's criterion of ontological commitment.

nominalistically acceptable content) and hence that the nominalist can safely explain the utility of talking about propositions. This is unlikely to satisfy a platonist: if someone thinks that propositions play a role in the semantics of our utterances, then we can't simply declare that propositions don't exist and expect that our utterances are able to carry the content that they do. If content depends on the existence of propositions, then explaining away propositions by appealing to content introduces a vicious circularity into our account. The easy road nominalist must take care here, and offer an account of language that doesn't involve an unexplained reliance on content. In the following section we'll look at a theory of propositions that fails to meet this standard.

2.3 Semantic Pretense-Involving Fictionalism

In "The Story About Propositions", Armour-Garb and Woodbridge defend a view that they call *Semantic Pretense-Involving Fictionalism* (SPIF).⁷⁰ According to SPIF, language users use 'semantic pretense' to make claims about real aspects of the world "by making *as if* to talk about other (pretend) things and situations."⁷¹ Semantic pretense isn't a form of explicit pretense: language users aren't, in general, aware that they are

⁷⁰ Armour-Garb & Woodbridge (2012). Thus far I have avoided use of the term 'fictionalism', although a number of authors employing the easy road strategy would describe their view as a version of fictionalism. Unfortunately, as a result of the breadth of fictionalist theories, the term "fictionalism" has become bloated with unhelpful connotations. The easy road strategy does not require an "according to the fiction" operator or require that language users are engaged in explicit pretense of any sort. The analogy to fiction is just the following: as with fiction, we talk about objects that don't really exist. (Of course, there are platonist accounts of fiction according to which fictional objects exist, in which case the analogy is of no use at all.)

⁷¹ Ibid. p. 637.

engaged in semantic pretense.⁷² This is pretense that is supposed to occur at a subconscious level in our semantic processing, and which impacts the content of our utterances. The focus of their paper is on developing SPIF as applied to propositions: language users engage in semantic pretense when talking about propositions, as a means of making other (nominalistically acceptable) assertions about the world.

Their theory is a variant of easy road nominalism. As they put the point: “Propositions may not exist, but talking *as if* they do extends the expressive capacity of our language in useful and important ways.”⁷³ However, the way that talking about propositions extends the expressive capacity of our language on their view is quite different than the easy road proposal outlined in §2 above. The easy road strategy outlined in §2 makes no revisionist claims about the semantics of our assertions: claims that refer to abstract objects are (false) claims that have platonistic content. The value of those claims, it was contended, comes from the fact that they also make requirements of the physical world: part of the content of those sentences is nominalistically acceptable, and that is enough to explain their inclusion in our best scientific theories. Armour-Garb and Woodbridge make a much stronger claim about the content of assertions that refer to propositions: through the mechanism of semantic pretense, such claims are assertions of nominalistically acceptable content.

What sort of nominalistically acceptable content is attributed to discourse about propositions on their SPIF account? Let us focus on their treatment of that-clauses. They claim that proposition-talk relates speakers “to (pretend) Russellian *complexes* of

⁷² Ibid. p. 656.

⁷³ Ibid. p. 646.

(pretend) objects, properties, and relations”.⁷⁴ This allows speakers to make assertions about complex use-features of linguistic items that would otherwise be extremely long and technical to specify. Consider the following example:

...the serious assertion made by “(8) Goldbach’s Conjecture is that every even number greater than 2 is the sum of two primes” is that the name ‘Goldbach’s Conjecture’ has the same (singular-term) long-arm conceptual role that the speaker of (8) attaches to ‘that every even number greater than 2 is the sum of two primes’.⁷⁵

It should be clear from this example that the content assigned to utterances involving that-clauses on the SPIF account isn’t something that is introspectively accessible to the speaker. The notion of a ‘singular-term long-arm conceptual role’ is a piece of technical jargon. To drive home this point, here is how Armour-Garb and Woodbridge explain what they have in mind when they introduce the notion of a long-arm conceptual role:

As we see it, the primary purpose of [the] pretense is to provide a practical means for speakers to talk about certain complex use-features of linguistic items and cognitive states. The kinds of features we have in mind are those included in what are sometimes called *long-arm conceptual roles*. In addition to narrow computational or inferential role (that is, the position an utterance of thought-state occupies in an inferential network, determining what claims it follows from and what claims follow from it) these use-features include certain social factors and inferential and causal connections to actions, as well as... “indication relations” (certain causal relations, including perceptual connections to the world). Direct specifications of attributions of these use-features would inevitable be extremely long, complicated, and technical.⁷⁶

I will not attempt to unpack this description, since it is not important for our purposes precisely what a singular-term long-arm conceptual role is. The important point is simply that the phrase “long-arm conceptual role” is not merely a way of denoting a familiar

⁷⁴ Ibid. p. 650.

⁷⁵ Ibid. p. 660.

⁷⁶ Ibid. p. 650, their italics.

concept that we should expect ordinary language users to have a grasp of. Rather, it picks out a complex feature of language use.

The SPIF account of the content of utterances involving that-clauses strikes me as implausible. Is an ordinary speaker who utters sentence (8) really making an assertion about ‘long-arm conceptual roles’? Even on reflection, it doesn’t seem to me that what I would be saying by uttering (8) is something about long-arm conceptual roles. And this is a problem. In assessing an account of content, we should take the phenomenology of language seriously: a semantic theory becomes implausible when it assigns contents to sentences that differ significantly from what language users believe the contents of those sentences to be.

The objection here isn’t that Armour-Garb and Woodbridge rely on technical jargon. Rather, it is that Armour-Garb and Woodbridge have misidentified the content of sentences involving that-clauses. Sentence (8), on the face of it, is a sentence about Goldbach’s Conjecture. It is not a sentence about the *conceptual role* of the name ‘Goldbach’s Conjecture’. Of course, we use the concept of Goldbach’s Conjecture in thinking about Goldbach’s Conjecture. And we likewise use the name “Goldbach’s Conjecture” in making utterances about Goldbach’s Conjecture. But using a concept or a name is one thing, and talking about that concept or name is quite another. The SPIF account of the content of utterances involving that-clauses conflates the use and mention of conceptual roles.⁷⁷

⁷⁷ Armour-Garb and Woodbridge are explicit about the fact that they are collapsing the use/mention distinction and take this to be a benefit of their view (2012, p. 650). As they see it, the whole point of relying on that-clauses is that speakers are thereby able to communicate about complex use-features of linguistic items that would otherwise be extremely long and technical to specify.

This sort of move is not without precedence. There is a long tradition in philosophy of taking claims that do not seem at first glance to be about language use to really be telling us something about

The problem is that SPIF is intended as a descriptive theory – it attempts to give an account of the content of actual utterances made by ordinary speakers. But sentence (8) doesn't seem to be expressing anything especially complex or technical. Rather, it seems to express something quite straightforward about what Goldbach's Conjecture is. This significant mismatch between the content that utterances involving that-clauses seem to have and the content they are purported to have according to SPIF is a cost of the theory.⁷⁸

A related difficulty arises when we consider sentences that are explicitly about propositions. Armour-Garb and Woodbridge consider the following sentence as an example: "There are numerous propositions about the distribution of physical properties that no one will, or could, ever assert or entertain."⁷⁹ Armour-Garb and Woodbridge recognize that this sort of sentence requires a different treatment. Their strategy is to say that in such cases the assertion being made is really "only about the content of the make-

language use. Consider the following example. Suppose that someone says that the content of an utterance of (T) "Trump is the president of the US" is that the terms "Trump" and "the president of the US" have the same referent. I grant that there is something seemingly reasonable about this proposed analysis. After all, there is a close connection between identity claims and claims about co-referential terms. That said, this is not my preferred analysis of the content of an utterance of (T). It seems to me that although the terms "Trump" and "the president of the US" are being used in the sentence (T) to make an identity claim, the terms are not themselves the topic of conversation. But I recognize that my position here is at least somewhat controversial, and that this sort of analysis will strike many philosophers as reasonable.

Armour-Garb and Woodbridge are relying on essentially the same sort of strategy: utterances involving that-clauses are understood to really be about complex use-features of linguistic terms. But I think their position is substantially more difficult to defend than the analogous analysis of sentences like (T). After all, ordinary language users are able to recognize the connection between two objects being identical and the terms we use to refer to them being co-referential. But the connection is much less obvious, and much more complicated, when it comes to the analysis that Armour-Garb and Woodbridge offer of that-clauses.

⁷⁸ I don't want to go so far as to suggest that language users have perfect introspective access to the content of their utterances. And perhaps in some cases language users really are saying something about the terms or concepts they are using without recognizing that that is so. But if language users are radically mistaken about what they are saying when they make utterances then the communicative function of language becomes impaired. And so any mismatch between the apparent content of an utterance and the content assigned to that utterance according to a theory of content should at least be treated with caution. The greater the mismatch, the more cause there is for concern.

⁷⁹ Ibid. p. 653.

believe that is behind proposition-talk” – in particular about the pretense that is stipulated as background for the make-believe.⁸⁰ This response only compounds my previous worry. Recall that semantic pretense is supposed to happen subconsciously, and be disconnected from the attitudes of speakers. Here again we have a case where the content of utterances departs, in a rather significant way, from the content that our utterances seem to have: utterances about propositions are supposed to really be utterances about the subconscious make-believe behind proposition talk. But again, this seems implausible as a descriptive claim.

Armour-Garb and Woodbridge’s treatment of explicit proposition talk also exposes a worry about the notion of semantic pretense. Armour-Garb and Woodbridge repeatedly talk about the *stipulations* of the make-believe. But who or what is making these stipulations? It only seems plausible that speakers could be making these stipulations if the pretense involved was explicit. But the idea that people uttering sentences including that-clauses are engaged in explicit pretense as a means of making assertions about long-arm conceptual roles doesn’t fit with the phenomenology of language use. While this is an empirical matter, it would be a shocking result indeed if most speakers were engaged in this sort of explicit pretense when using that-clauses. So Armour-Garb and Woodbridge are caught in a dilemma. They must either hold that speakers are aware enough of the semantic pretense to make stipulations about it, or they must hold that stipulations about the make-believe of propositions are being made at the subconscious level. But how exactly we are supposed to understand stipulation, and ‘making as if to talk about pretend things’, at a purely subconscious level is not clear.

⁸⁰ Ibid. p. 653.

Before moving on, I would like to call attention to a further difficulty for SPIF. At the end of the last section I specified a desideratum for any adequate nominalist theory of language: it must not involve an unexplained reliance on content. Armour-Garb and Woodbridge’s theory fail this test. Even if the claims they made about the content of sentences involving that-clauses weren’t problematic for the reasons outlined above, their theory would still rely on an unexplained notion of content. They reject the existence of propositions, but offer no account of the content of sentences. So the platonist can respond to SPIF by noting that even if Armour-Garb and Woodbridge are correct in thinking that speakers do not commit themselves to the existence of propositions, the existence of propositions is nevertheless needed to explain how it is that sentences have content. Something more must be said in response to the platonist.

2.4 Balaguer’s Semantic Fictionalism

In his “Attitudes Without Propositions”, Mark Balaguer advances a theory of propositions that closely follows the easy road strategy described in §2. His *Semantic Fictionalism*: “(a) “admits that platonists are *right* about the truth conditions of ‘that’-clause-containing sentences (i.e., admits that such sentences really are “about” propositions), but nonetheless, (b) maintains that there are no such things as propositions, and thus (c) concludes that the truth-conditions of ‘that’-clause-containing sentences are never *satisfied* – i.e., *that there are no true ‘that’-clause-containing sentences* (except for

those which are *vacuously* true).”⁸¹ Balaguer defends his semantic fictionalism by arguing for two claims.⁸²

- (NC) Empirical Science has a purely nominalistic content that captures its “complete picture” of the physical world, and
- (COH) It is coherent and sensible to maintain that the nominalistic content of empirical science is true and the platonistic content of empirical science is fictional.

The purpose of abstract object talk, according to Balaguer, is to help us communicate the nominalistic content of our theories. This should sound familiar; Balaguer’s Semantic Fictionalism is one of the earliest defenses of easy road nominalism.

While I ultimately want to endorse a theory very close to that put forward by Balaguer, there are a couple of issues in his account that are worth drawing attention to. Balaguer’s argument for (NC) rests on three premises:⁸³

- i. Abstract objects – if there are such things – are not causally relevant to the operation or state of the physical world;
- ii. If (i) is true, then the mixed facts expressed by empirical science supervene on more basic facts that aren’t mixed, i.e., that are either purely nominalistic (i.e., purely physical, neurological, or whatever) or else purely platonistic; and
- iii. If the consequent of (ii) is true, then (NC) is true, i.e., empirical science has a nominalistic content that captures its complete picture of the physical world.

Both premise (ii) and premise (iii) are too strong; in both cases one can accept the antecedent of the conditional without accepting the consequent. Let us start with premise (ii). As Balaguer stresses, abstract objects are not causally relevant to the physical world. From this it follows that the basic *causal* facts aren’t mixed facts. But the more general claim that mixed facts (facts involving both physical and abstract objects) supervene on

⁸¹ Balaguer (1998), p. 806, his italics.

⁸² Ibid. p. 812.

⁸³ Ibid. p. 815-6.

more basic facts that aren't mixed doesn't follow. The fact that abstract objects are not causally related to physical objects doesn't mean that there are no basic non-causal facts connecting abstract and physical objects; it is open to the platonist to posit fundamental non-causal relations holding between abstract and physical objects.

Is there a way to fix premise (ii)? Not unless one accepts the principle that all basic facts are basic causal facts. But even the nominalist is unlikely to accept this principle: facts about the instantiation of properties, about the location of objects in space, and so on, are not causal facts and do not supervene solely on causal facts. So premise (ii) should be rejected.

Premise (iii) is likewise problematic. Suppose that our scientific theories express mixed facts, and that there are no basic mixed facts. What does this tell us about the *content* of our theories? Balaguer thinks that it tells us quite a lot: if the mixed facts supervene on basic facts that aren't mixed, then the mixed content of our scientific theories can be neatly divided into its nominalistic content and its platonistic content. But why should this follow? Absent a worked out theory of content, this is too bold an assumption. Here Balaguer runs afoul of the desideratum on a nominalist theory of language I specified in §2: it must not involve an unexplained reliance on content.

There is a third, and related, difficulty with Balaguer's reasoning. Balaguer supposes that once he has established (NC) and (COH) the adequacy of nominalism follows. And it is easy to see why he would think so. According to (COH), it is sensible to hold that the nominalistic content of empirical science is true but the platonistic content of empirical science is false. But if it is sensible to hold that the platonistic content of empirical science is false, then doesn't it follow that it is sensible to hold that

abstract objects don't exist? No. This line of reasoning fails to take into account the possible role of abstract objects in a theory of content. For suppose that the platonist is right in identifying the contents of sentences with propositions.⁸⁴ This is entirely compatible with (NC) and (COH). After all, the relevant propositions are neatly divided into those that capture nominalistic content and those that capture platonistic content as (NC) requires. And it is further possible that all the propositions with platonistic content are false, as (COH) requires. But despite the compatibility of this view with (NC) and (COH), this is not a nominalistically acceptable theory: we are relying on abstract objects as contents. Balaguer considers abstract objects only in their role of making-true the platonistic content of scientific theories. But because theories are linguistic in nature, abstract linguistic objects may have an additional role to play.

Despite these flaws in his arguments, I think that Balaguer's easy road nominalism is essentially on the right track. The appeal to unmixed basic facts isn't crucial to the easy road strategy, and can be dropped. There are other, better, reasons for thinking that our scientific theories have nominalistic content. Exactly what form the nominalistic content of our theories take is a matter of some dispute. The literature on this topic has revolved around attempts to specify the nominalistically acceptable content **N** of a theory **T**, where theory **T** contains mathematical discourse.⁸⁵ Joseph Melia, for example, favours the use

⁸⁴ To keep things simple, I will focus on the platonist view according to which propositions are the contents of sentences. Platonists are not required to endorse this view – they may reject propositions, or draw a distinction between the proposition expressed by a sentence and the content of that sentence. But the arguments in this paper will apply just as well to whatever alternative contents that the platonist posits.

⁸⁵ See, for example, Melia (2000), Colyvan (2010, 2012), Yablo (2012), and Liggins (2012).

of retraction: one first states T, and then ‘takes back’ the platonic elements of that theory to arrive at N.⁸⁶

Rather than delve into the specifics of Melia’s strategy, or of any other strategy for extracting N from T, I would like to approach the issue from a different angle. I think there is a compelling reason – independent of any particular account – for the platonist to admit that the nominalistic content of scientific theories is, in general, extractable.

Easy road nominalism is designed to reconcile nominalism with realism about science: the idea is to show that nominalism is compatible with our best scientific theories being essentially correct about the physical world. If the nominalist were happy to reject current scientific practice, then easy road nominalism would be unnecessarily complicated. So, a certain deference to scientific practice is being taken for granted. Now, if we pay attention to scientific practice it is clear that scientists *can* and *do* extract nominalistic implications from scientific theories quite frequently. Consider again Yablo’s sentence (B): “For all positive real numbers M and R, the escape velocity from a sphere of mass M and diameter 2R is the square root of $2GM/R$, where G is the gravitational constant”. Scientists understand, without much difficulty, what this says about the physical world. A regularity regarding the behavior of matter is being recorded. Scientists are able to design experiments, and use instruments, to test whether there is such a regularity in the behavior of physical objects.

It is important to emphasize that abstract entities cannot play a role in governing the behavior of physical objects. Abstract entities like numbers are wholly causally inert, and *cannot* play any causal role in physical systems. Whatever regularities exist between

⁸⁶ Melia (2000), pp. 466-71.

mass and escape velocity, they are not mediated by abstract entities. Nor can scientists test or measure the platonic implications of scientific theories – again, abstract entities are causally inert and cannot be accessed using empirical means. If a theory T is confirmed, it is only ever confirmed on the basis of what it requires of the empirical world. If scientists couldn't easily extract the nominalistic implications of sentences like (B), then sentences like (B) would lose their scientific utility. In general, the very applicability of mathematics and other talk of abstract entities *depends* on the ability of scientists to determine what their theories require of the physical world.

If this is right, then we have good reason to accept that our theories carry nominalistic content. The fact that scientists *do* extract nominalistic content from scientific theories is proof enough that they can, regardless of the psychological processes they use to extract that content. The fact that abstract entities are causally inert is enough to guarantee that no matter how scientists extract the nominalistic content of scientific theories, they are doing so in a nominalistically acceptable way.

Now, it's true that the practice of science doesn't show that scientists can extract *all* of the nominalistic implications of their theories, and hence doesn't show that they can fully specify **N** for some theory T. But so what? The utility of abstract object talk as a descriptive aid doesn't require any such thing. Rather, it simply requires that such discourse be able to express *some* nominalistically acceptable facts that are otherwise inexpressible (short of using infinite disjunctions, infinite primitive predicates, or the like). And it is clear that scientists do make use of abstract object talk in precisely this way.

2.5 A Theory of Content

I've argued that we have good reason to accept that our theories have nominalistic content, but I've yet to say anything about why we should think that our best theory of linguistic content will be nominalistically acceptable. In fact, I think what the nominalist should say here is fairly straightforward, and begins with a remark that Balaguer makes in his paper.⁸⁷

Balaguer points out that, for the platonist, any explanation of the fact that a sentence *means* anything to language users will have to rely on the fact that language users form a mental representation of the proposition expressed by that sentence. After all, abstract objects can play no causal role in language use – if propositions are to have any role at all it has to be in virtue of language users having some sort of mental connection of them. And this, Balaguer goes on to argue, commits the platonist to nominalistic content-determining properties for since “propositions (if there are such things) are causally inert, our brains are going to have to form their “representations” *on their own*, so to speak, i.e., without receiving any “help” from any abstract objects.”⁸⁸ The idea is simply this: whatever the causal story is behind the mental representations that language users form of propositions, it will be nominalistically acceptable.

A couple of remarks are in order. First, words such as “represent” (and “about”) are sometimes used in such a way that representation fails if there is no such object as the object that is supposedly being represented. On this use, it would be impossible to form a mental representation of a proposition if there are in fact no propositions. Accordingly, it

⁸⁷ Balaguer (1998), pp. 822-4.

⁸⁸ Ibid. p. 824.

would be illegitimate for the nominalist to rely on talk of mental representation. Now suppose that a platonist uses “represent” in this way in explaining the causal behavior of language users. Does this pose a problem for my claim that the causal story behind the mental representations that language users form of propositions must be nominalistically acceptable? It does not. It is important to recognize that there is something sloppy about using representation in this way: because abstract objects are causally inert, the difference between a successful representation of an abstract object and an otherwise identical mental state that fails to represent due to the non-existence of that abstract object *can play no causal role*. The ‘failed’ representation will serve just as well as the ‘successful’ representation in any explanation of the behavior of the relevant agent. Platonists sometimes treat relations to abstract objects as if they have a place in the causal story, but this can’t be right: whatever the physical world needs to do to hold up its end of the bargain is sufficient for the causal story. And the platonist is committed to there being something the physical world needs to do to hold up its end of the bargain. For our purposes, then, it is acceptable to always use talk of “representation” in a broader sense, to cover both cases of successful representation and representation that fails due to the non-existence of the object being represented.

The second note I want to make here is that the focus on “mental representation” is not crucial. While Balaguer takes the platonist to be committed to the mental representation of propositions, the general argument being made can be applied more broadly to platonist theories that reject mental representation. If language users relate to propositions in any way, the platonist will need an account of what it takes for language users to so relate. Some theory is needed of how the physical world must be in order for

language users to stand in the appropriate relations to propositions. Whatever this story is, it will be nominalistically acceptable.⁸⁹

Of course, the causal story that the platonist tells is not the whole story. And so the argument so far doesn't, by itself, guarantee that a theory of content will be nominalistically acceptable. But it is now possible to see how the nominalist can adapt, with fairly minimal modification, any theory that the platonist gives. Take any complete platonistic theory of content. This will include the nominalistically acceptable causal story, some abstract linguistic entities, and some relations between language users and those abstract linguistic entities. Now, consider the theory one gets by simply removing the abstract linguistic entities, and relations to those entities, from the picture. One is left with the same causal story. So if the original platonist theory was an adequate causal account of language use, the new nominalist theory will be an adequate causal account as well. Nothing further is required to guarantee that language use is nominalistically acceptable. For one of the very features of abstract objects that make them so useful – their abstractness – guarantees that they play no causal role in language use. And hence any platonist theory can be so adapted.

Leaving things at this point would result in a very unsatisfactory theory of language, for a couple of reasons. The first is that platonist theories of language don't typically come packaged with a thorough causal story of language. We are still relatively ignorant of the

⁸⁹ Strictly speaking, this isn't quite right: there is room for a platonist to hold that language users stand in relations to propositions, but that there is nothing in common (in the physical world) among language users who stand in such relations. But in this case the appeal to content loses its explanatory power. The purpose of talking about content was to provide a higher-level explanation of language use – one that abstracts away from specific causal details. But if sameness of content is not supported by some regularity in the world, then an appeal to content can't be used to provide such a higher-level explanation: there is nothing in common in the causal story to abstract away from, and hence content cannot help explain linguistic behavior.

neuro-physiological mechanisms involved in language use; a great deal of further empirical research is required. So, while the platonist is just as committed as the nominalist to their being such a story, it has yet to be specified. So the nominalist who just claims that this causal story – whatever it may be – will be nominalistically acceptable is left without much of a theory of language at all. At best, the nominalist is offering the promise of a theory.

What platonist theories of language do instead is to provide a higher-level explanation of language use – one that glosses over the causal details. And here, language use is typically explained in part by talking about the very abstract entities that the nominalist wants to eliminate. Consider, for example, a platonist theory on which what it is for an utterance to have a particular content is for it to bear an appropriate relation to a proposition *p*. If we simply eliminate propositions and relations to propositions from this theory, then the result will be a nominalist theory on which utterances don't have content: since no utterances stand in any relations to propositions, no utterances have content. In the absence of a complete causal explanation of language, the platonist has at least provided a higher-level explanation of language. The nominalist strategy eliminates even this explanation, leaving us without much of a theory at all.

It is important to be clear about exactly what is at stake. The problem is not that the nominalistically acceptable causal explanation will be low-level, and that abstract objects can play a role in any higher-level explanation.⁹⁰ The nominalistically acceptable features of the world are not so constrained: it is not just the complicated causal histories of particular instances of language use that are nominalistically acceptable, but also

⁹⁰ Mark Colyvan makes this sort of objection to nominalism in a different context – regarding mathematical explanation. See Colyvan (2010) pp. 302-3 and Colyvan (2012) pp. 1034-6.

higher-level regularities and structural features of the world. Talk about abstract objects may function as a useful descriptive aid in our characterization of these higher-level regularities, but abstract objects cannot play a more substantive role than that. Abstract objects, it bears repeating, are causally inert: they *cannot* influence the behavior of physical objects. This is as true of higher-level features of the physical world as it is of lower-level features of the physical world. So we should not think that the platonist has an advantage in offering an explanation of language at the desired level of generality.

Where the existence of abstract objects can play a role is in other, non-causal, relations. Platonists are free to offer theories of content, or of reference, or of truth, that depend in some way on the existence of abstract objects. The nominalist rejects the existence of abstract objects, and hence should reject these platonist theories. And here it is important for the nominalist to have an alternative theory on offer.

So, what should our nominalist theory of content look like? One major constraint on our theory is that we will have to make do without the existence of anything to serve as the contents of our utterances. I have been assuming, in this paper, that if propositions exist then they are abstract – there is nothing in the world that is suitable to play the role of propositions. But if we do not have a nominalistically acceptable replacement for propositions, then we have no entities to serve as the contents of our utterances. Our theory of content, then, will have to do without contents.

Here is my proposal. Language users *presuppose* the existence of propositions, and interpret sentences as expressing propositions. This is done systematically so that there are strict regularities in how speakers interpret sentences. Speakers in the same linguistic community will attribute the same (or very similar) contents to the same

sentences. A sentence s can be said to have content p in context c if the relevant linguistic community in context c interprets s as expressing the proposition p . Note that there is nothing about this that requires that propositions actually exist. It is the mental representations of language users – the attributions of content to token sentences – that do all the work. And we've already seen that the platonist is just as committed as the nominalist to language users having this sort of mental capacity. Where the platonist sees language users as correctly identifying the propositions that sentences express, the nominalist sees failure: there are no propositions, so language users cannot correctly identify the propositions that sentences express. Language users are making attributions based on a false presupposition.

The core of the proposal is that language users make use of a false presupposition – that there are contents – as a means of attributing meanings to sentences. By (falsely) presupposing that there are contents, languages users are able to give sentences content. While in some contexts presupposition failure is a serious problem, here it can be safely ignored. The presupposition is a mere means of facilitating communication, and the failure of the presupposition is pragmatically irrelevant. The falsity of the presupposition does not collapse the content that is communicated: by making such presuppositions, rich truth-conditions can be specified and communicated. Reference to abstract objects thus plays a critical role in language quite distinct from the use of reference to abstract objects as a descriptive aid.

Here we must be careful to distinguish the claim that a sentence s has content p from the claim that there exists a content p expressed by sentence s ; the former claim is true but the latter claim is false. Sentences are represented as having content, even though

there are no entities to serve as *contents*. It is the presupposition that there are propositions, rather than the actual existence of any such entities, that grounds the content and the truth-conditions of sentences.

The proposal I am making should not be taken to require any sort of explicit presupposition on the part of the language user: the presupposition happens at the level of semantic processing.⁹¹ A devout nominalist can explicitly reject the existence of abstract objects while at the same time making use of such a presupposition on a semantic level.⁹² Nor does the presupposition need to involve in any way the idea that propositions are abstract; what is needed is just the presupposition that there are contents that can be transmitted from speaker to listener via the use of language. Language users needn't waste any energy thinking about the nature of these propositions – they simply act as if there are such propositions, and process language accordingly.⁹³

⁹¹ While explicit presupposition is not required, I don't mean to suggest that the presupposition is entirely hidden either. Generally speaking, people treat talk of abstract objects quite differently than they treat talk of concrete objects. Few people take the time to consider whether abstract objects exist. And, when asked, a surprisingly common sentiment is that it is a category mistake to even ask whether abstract objects exist. (In my experience, this is the second most common response, following simple befuddlement at the question). This attitude makes sense if people are simply presupposing the existence of abstract objects when it suits them. For more on the differences between the way that people treat talk of abstract objects and the way that people treat talk of concrete objects, see Yablo (2000), pp. 301-4. I don't want to put much stress on this point though, since the explicit beliefs of language users are largely irrelevant: the semantic presupposition can function equally well regardless of the explicit beliefs being held.

⁹² These remarks are fairly broad, but I would like to avoid taking a stand on just how modular the mind is or on which modules are involved in semantic processing. The crucial point is just that one part of the mind can rely on a presupposition that isn't made elsewhere (and would be rejected elsewhere).

⁹³ I am not entirely happy with the choice of the word "presupposition" here. I would like the account I give to be as general as possible. The more general point is this: the platonist, just like the nominalist, owes us a story of the psychological underpinnings of language. The platonist story will involve certain relations between language users and abstract objects (s recognizes x, etc.). The nominalist can give *precisely the same* psychological story, apart from holding that language users fail to enter into relations to abstract objects in virtue of the fact that there are no abstract objects. But how things are physically, and how things are 'in the head' – how things are in every way that is causally relevant to how language users process language – is exactly the same. On the platonist picture, language users (correctly) *believe*, on some level, that propositions exist. But whatever this belief is, it needn't be explicit – people are able to use language without forming explicit beliefs about the existence of propositions. It is for this reason that I chose the word "presuppose" rather than "think" or "believe": there is nothing deliberate and thoughtful about the choice – it is a built-in part of semantic processing.

While this is not a detailed causal story of language use, it provides an alternative to the high-level explanation that platonists provide in terms abstract objects, and shows how the nominalist can recover a meaningful notion of content. And it does so without abandoning the idea that the role of a semantic theory is to offer an account of what our utterances say about the external world. The central idea is just that we can fruitfully talk about the content of a sentence without positing the existence of ‘contents’.

The proposal I have put forward differs in a small but significant way from the easy road strategy I outlined earlier. It isn’t discourse about propositions that plays the crucial role in language. Rather, it is the presupposition that there are propositions. Whereas discourse about propositions may be a descriptive aid in our theorizing about language, the presupposition that there are abstract objects is, on my proposal, a critical part of the *use* of language. This presupposition comes with serious advantages.

For one, the presupposition that there are propositions simply bypasses the question of whether or not propositions exist; language users can form, and make use of, the same mental representations of propositions regardless of the ontological status of propositions. If this nominalist is right in thinking that there are no propositions, then this benefit is gained without any ontological cost – the actual existence of propositions plays no important role in language use, and can be safely eliminated from our theories. The nominalist theory is thus more parsimonious than the platonist alternative.

The presupposition account also allows for considerable flexibility: the propositions presupposed can take whatever form is most useful to the task at hand. Semantic modules can employ whatever presupposition results in reliable and efficient semantic processing, thereby reducing the amount of mental processing required.

Language users are, in general, not concerned with matters of metaphysics. Language users can rely on a shifting ontology – taking on a commitment to some entities, for certain purposes, and rejecting those entities for other purposes. Indeed, I see no reason why language users could not presuppose the existence of inconsistent entities (impossible worlds, etc.) when it suits their purposes. We should expect them to make whatever presuppositions are practically useful for them to take on, even if on rational reflection they would deny the existence of the presupposed entities.

In this paper I have been assuming that propositions, if they exist, are abstract objects. Would matters be any different if the nominalist were able to find a concrete surrogate for propositions – something nominalistically acceptable that could serve as the contents of utterances? I don't believe so. If what I've said here is on the right track, there would be no reason to suspect that actual language use involves any such objects. The presupposition that there are contents is sufficient. Speakers can focus on the task of making an utterance that will communicate the correct content to those listening. Our linguistic capacities, like all mental capacities, have developed for practical purposes – to allow for communication. The developmental pressures on semantic mechanisms would emphasize reliability and efficiency. Getting the ontological status of propositions right is of no practical importance, and hence we should not expect our semantic mechanisms to be sensitive to such issues. If making false presuppositions improves reliability or efficiency, then a semantic module that makes those presuppositions is to be preferred.

Of course, these remarks are highly speculative. While the developmental pressures on semantic mechanisms would emphasize reliability and efficiency, evolutionary and social pressures don't always result in optimal processes. And the actual

neuro-physiological mechanisms involved in semantic processing are an empirical matter that is not yet well understood. Indeed, our best semantic theories may in the end make no mention of propositions at all. But the nominalist ought not dismiss propositions out of hand: various philosophical arguments have been advanced for the conclusion that propositions play an essential role in language.⁹⁴ My aim here has been to argue that any role for propositions in a theory of content can be played equally well (if not better) by the mere presupposition that there are propositions.

2.6 Reference and Languages

Thus far I have been focusing on the role of propositions in a theory of content. But propositions are also frequently taken to play quite a different role as well: as the referents of that-clauses. Let us suppose once again that the platonist is on the right track, and that that-clauses refer to propositions. If there are no propositions, as the nominalist contends, then every occurrence of a that-clause will be a case of empty reference: there are no propositions to serve as the referents of that-clauses.

There is nothing particularly special about propositions in this role. So rather than dealing with propositions in particular, I would like to deal with reference more generally. A semantic theory must not just assign semantic values to sentences but also to sub-sentential linguistic elements. One of the important semantic roles of words and phrases is to *refer* to objects in the world; referential terms function to connect what is

⁹⁴ See, for example, Balaguer's (1998) discussion of the Frege-Bealer-Schiffer argument (we need propositions to account for the truth of sentences with a 'that'-clause) and of the Frege-Katz argument (we need propositions to account for the content of public-language sentences).

being said to particular objects in the world. So far our focus has been on propositions, but words and phrases can be used to refer to all manner of abstract object. The challenge in every case is the same: if abstract objects do not exist, then all purported reference to them is empty reference. How to deal with referential terms that lack a referent is a longstanding issue in philosophy. A fully developed nominalist account of reference would require a paper of its own, but a few general remarks are in order here.

The first is that we should not, in general, expect the referents of referential terms to be a relevant causal factor in semantic processing. There are exceptions of course – language users can refer to entities that are so relevant – but not for abstract entities: abstract objects are causally inert. And so we can draw on considerations similar to those in the previous section. For one, both the platonist and the nominalist are committed to there being a complete nominalistically-acceptable causal story of language use. And second, we can employ a similar presupposition-based strategy as before to explain referential behavior. Where the platonist claims that we need abstract referents, the nominalist can claim that the mere presupposition that we have such referents will do: if language users presuppose that the requisite referents exist, they will freely form mental states about those referents without concern for whether or not the referents do in fact exist.

We should not be misled here by the fact that reference is a relation between a linguistic item and an object. In some cases the presupposition of reference is supported by the actual existence of a referent, and accordingly we have genuine reference. But in other cases we merely have the presupposition of reference, and this presupposition of reference is sufficient for semantic processing involving empty-referential terms. For

even if we adopt an externalist account of the individuation of mental states – an account on which mental states are identified in part based on how things stand ‘outside the head’ – it is only the neuro-physiological workings of the brain that are causally involved in semantic processing. Successful and unsuccessful reference play the same role in semantic processing, and hence the non-existence of a purported referent can be safely ignored when our focus is on linguistic behavior.

This is not to say that the existence of abstract referents is of no significant at all. The existence (or non-existence) of abstract referents plays an important role in determining the truth-value of utterances which refer to those entities. How, exactly, the non-existence of a referent impacts the truth-value of a sentence would be straightforward to deal with if all sentences containing empty referents were false. But things are not so simple. When introducing easy road nominalism I remarked that some easy road nominalists want to take standard mathematical sentences referring to numbers to be false, while others want to take such sentences as true. And philosophers disagree over whether the sentence “Voldemort is evil” is (i) strictly speaking false but good enough for communicative purposes (since Voldemort doesn’t exist, but the sentence can be used to communicate something true about how the character Voldemort is portrayed in the canonical Harry Potter fiction), or (ii) literally true (either because the literal truth of the sentence doesn’t require the existence of Voldemort, or because Voldemort does exist as an abstract fictional entity). And there are some sentences involving empty reference that nearly all nominalists agree are true, such as “Voldemort doesn’t exist”.

A variety of stances are thus available to the nominalist. On one extreme, the nominalist could simply take *all* sentences involving empty reference to be false – even

sentences such as “The number 2 doesn’t exist”. While this is an extreme position, I don’t find it entirely outlandish. I’ve already argued that language users should not shy away from literal falsehood when it suits their communicative purposes. The literal content of an utterance only takes us part way to what is actually communicated; pragmatic factors play a role as well. Regardless of whether the sentence “The number 2 doesn’t exist” comes out as literally true or false, it is clear enough what this sentence communicates to us about the ontology of the world.⁹⁵

More moderate positions are also available, on which (at the very least) direct existential claims such as “The number 2 doesn’t exist” come out literally true. There are various ways such a position can be developed, and I won’t commit to any particular strategy here.⁹⁶ What I would like to point out is that the existential status of abstract referents presents no special difficulty. The story we tell for the sentence “The number 2 doesn’t exist” will not be substantially different than the story we tell for the sentence “The Loch Ness monster doesn’t exist” – whether or not the empty reference is to a

⁹⁵ While this view results in the simplest semantics, it comes with complications. If “The number 2 doesn’t exist” is strictly speaking false, one might worry that nominalism is self-undermining. After all, the nominalist wants to maintain that the number 2 doesn’t exist (since no abstract objects exist). This consequence looks much worse than it is – the view being presented is not self-undermining. It is important to recognize that the falsity of nominalist claims doesn’t mean that there *are* abstract objects. That is only one way that such claims could be false. Nominalist claim might also come out false, on this view, because expressing the core ontological claim of nominalism depends on the use of empty reference (which results in any statement of nominalism coming out false). And while this is not ideal, we nevertheless are able to understand what the core ontological claim of nominalism is. And related, true, claims can still be expressed: “Every purported reference to abstract objects is a case of empty reference”.

⁹⁶ The simplest option, perhaps, is to take a claim of the form “The number 2 doesn’t exist” as really expressing that the “The number 2” fails to refer to anything. But, as I stated earlier in the paper, I think this sort of strategy relies on a use/mention conflation that ought to be avoided in general. I’m more partial to a pretense-based account of reference. Those interested in a developed pretense-based account of reference are referred to Crimmins (1998). Although I don’t agree with all of the details of Crimmins account, he points to various ways in which pretense surrounding reference can play a useful role in language.

purported abstract object or to a purported physical object makes no difference.

Consequently, I see no special role for abstract objects to play in a theory of reference.⁹⁷

We've now seen how the presupposition-based strategy can be generalized to cover any abstract objects that appear as semantic values. We can also generalize the presupposition-based account in a second way: to abstract linguistic entities that do not serve as semantic values. Propositions are not the only abstract linguistic entities that appear in theories of language. Other candidate entities for elimination include public languages, word types, and sentence types. It is easy to see that we can extend a presupposition-based account to cover these additional abstract entities as well: language users presuppose that there are languages, that there are words and sentences belonging to those languages, and that those words and sentences have content. These presuppositions can do all the same important linguistic work whether or not there really are any such entities. Just as the presupposition that there are propositions allows us to think of ourselves as transmitting or sharing content, the presupposition that there are word types allows us to think of ourselves as acquiring, and sharing, words – if I teach you a new word, you can use that *same* word. Words and sentences can be seen as public resources, to be freely shared among language users.

I opened this paper by remarking on Lewis's distinction between language and languages. Public languages seemingly play a crucial role in language: they enable coordination between language users by providing an external standard that language

⁹⁷ This claim is complicated somewhat by views on which there is no truly empty reference: wherever we have reference to something that 'doesn't exist', what we really have is reference to an abstract *fictional* entity, or to an entity that doesn't exist at our world but that exists at other possible worlds, or so on. So a nominalist cannot be entirely neutral here: some account of empty reference must be adopted according to which empty referential terms really do fail to refer to anything at all.

users learn and cohere to. But it is a mistake to think that a public language is required for this role. Indeed, it is a mistake to think that a public language *can* play this role. Abstract entities cannot figure into the causal story of how agents use utterances and inscriptions to communicate. Language users can't be relying on the existence of abstract public languages in order to coordinate their linguistic behavior. This point is easily obscured by the way we often talk about language. Lewis, for example, talks about language users *using* an (abstract) language, where this amounts to them having conventions of truthfulness and trust *in L*. But what does it mean to 'use' a language? If this requires some sort of relation between language users and abstract languages, then 'using a language' is causally idle. It is the mental states of the language users that do the real work. The shared presupposition that there is a public language that both speaker and listener are employing is enough to solve the coordination problem and enable linguistic communication. And so we can give a Lewisian theory of language – a theory in which linguistic communication depends on language users following certain conventions of truthfulness and trust – without positing the existence of any languages.⁹⁸

2.7 Some Concluding Remarks

I would like to close with some remarks on semantics. It is worth asking *why* a theory of content is important in explaining language use. If the nominalist is right in thinking that there are no propositions to serve as contents, what value is there in a theory

⁹⁸ Of course, if we opt for the version of easy road nominalism on which ordinary mathematical claims are strictly speaking false then we should reject the idea that language use is governed by a convention of truthfulness (or at least be very careful in how we apply this convention).

of content? It seems to me that semantics has two primary roles to play in language. For one, it connects language to the world. Language users intend to communicate about the world around them. They cannot do this unless words or phrases refer to, or are about, particular things or types of things in the world. So it is important that words and phrases can be used to successfully refer to things in the world. And second, it provides compositional rules whereby language users can construct and understand novel sentences from a comparatively small set of words. Human beings have limited mental capacities, and I can think of no plausible alternative for the assignment of contents to sentences than that we rely on a compositional semantics. In combination these two features of semantics enable efficient and reliable communication about the world.

It is a widespread idea that one desideratum on a good semantic theory is that it fits with the truth-claims of language users: there should be a close match between the sentences that language users take to be true and the sentences that come out true on our semantic theory (with some allowance for false beliefs). I think this is a mistake. The primary purpose of semantics is to enable efficient and reliable communication about the world. This would require that our semantic theory match the truth-claims of language users *if* the most efficient and reliable means of communicating about the world was to always to utter truths. But this is not necessarily the case, for two reasons.

First, it may be more efficient to employ semantic rules that result in false content and to rely on pragmatics to ensure successful communication than to employ more complicated semantic rules that result in true contents. A literal falsehood can be used to communicate something true. This is fairly uncontroversial in a limited range of cases: cases of metaphor, hyperbole, and other figurative language. But the easy road nominalist

can argue, as I have, that the same is true of discourse more generally and can be applied to great effect when talking about abstract objects.

Second, it may be more efficient to rely on false presuppositions in semantic processing than to ensure that no reference is made to non-existent entities. If a widespread error of this sort is part of the most efficient and reliable way to communicate, then we should expect and desire widespread error in our semantic processing. If the presupposition that abstract objects exist is of pragmatic utility, then we should expect language users to persist in making that presupposition regardless of whether abstract objects exist. And if the best way of integrating that presupposition into the regular semantic processing of language users results in false utterances, then we should expect language users to repeatedly make such false utterances. Of course, if language users are presupposing the existence of abstract objects, then we should not expect them to judge utterances that refer to such objects as false. And as a result, we should not be surprised to see a rather significant mismatch between the sentences that language users take to be true and the sentences that come out true on our semantic theory. As paradoxical as it sounds, perhaps the best way to accurately describe the world involves some falsehood. Perhaps it relies on the use of talk about abstract objects to serve as a descriptive aid. And perhaps it relies on the presupposition that there are abstract linguistic objects as a means of enabling efficient and reliable communication.

In §2 I distinguished two versions of easy road nominalism. One, which I have been working with in this paper, opts for a simple semantics on which apparent reference to and quantification over abstract objects is treated as genuine reference to and

quantification over those objects. The result is that many of our ordinary claims that refer to abstract objects come out false. I have just said something about why this does not strike me as problematic. But it is important to note that the presupposition-based nominalist strategy I outlined in this paper doesn't depend on the acceptability of the use of widespread falsehood in linguistic communication. A nominalist can instead hold that many of our ordinary claims that appear to refer to abstract objects come out true: a nominalist can accept that the presupposition that there abstract linguistic objects plays an important role in language use, but hold that the contents that speakers attribute to utterances are more complicated than I've allowed for here.⁹⁹

There is also a third option. I've claimed that language users are largely unconcerned about the ontological status of abstract objects, and that this leaves room for language users to be widely mistaken about the truth-values of sentences that presuppose the existence of abstract objects. It likewise leaves room for a significant degree of semantic imprecision; perhaps there simply is no fact of the matter about the truth-value of utterances that rely on a false presupposition that abstract objects exist. An analogy might be helpful here. Consider the sentence "there are solid objects". Modern science has shown us that our intuitive notion of a 'solid' breaks down at the microscopic level: even the densest of objects are composed of particles separated by comparatively vast areas of empty space. In light of this discovery, is the sentence "there are solid objects" true or false? One response is to hold that there is no correct answer to this question: the content of the sentence is too imprecise to render a verdict here. A similar claim can be

⁹⁹ The general idea can be motivated as follows: in cases where talk of abstract objects is serving as a mere descriptive aid, all the speakers care to be communicating is whatever is entailed for the physical world. So that is all that we should take to be a part of the content of what is being said. As a result, many of our utterances that involve abstract object talk turn out to be true. The complication here is, of course, to give a plausible and consistent account of just what these semantic contents are.

made in the case of abstract objects: generally speaking, the content we assign to sentences that refer to abstract objects is silent about the case in which our presupposition that abstract objects exist fails. The analogy is not perfect. It was an empirical discovery that our intuitive understanding of solidity did not in fact apply to paradigm examples of solid objects, but there is no risk of an analogous empirical discovery that our presupposition that abstract objects exist is false. Nevertheless, I think that the ontological status of abstract objects is a matter of fact, and something that we can be mistaken about. And so, if semantic imprecision occurs in one case I do not see in principle why it could not occur in the other.¹⁰⁰

It is also worth noting that successful communication can allow for quite significant differences in the semantic rules employed by language users. For suppose that we have three language users who hear the utterance “The number of bridges in town is two” – one who takes this sentence to require the existence of the number 2, one who doesn’t, and one who takes the matter to be semantically ambiguous. What the three language users think this sentence tells us about abstract mathematical entities is, of course, different. But the three language users will agree about what the sentence requires of the physical world: there must be a bridge in town, and another (distinct) bridge in town. If the communicative purpose at hand is simply to communicate about the bridges in town, then the difference in their semantic views is irrelevant. The general result I am pointing to here is this: successful communication about the physical world tolerates widely divergent semantics for abstract objects. So, it does not seem implausible to me

¹⁰⁰ The possibility of semantic imprecision here is no threat to my ability to make ontological claims. It is always open to language users to be more precise in their claims than is typical. And it is clear enough that I intend to be making utterances of the form “abstract objects don’t exist” in an ontologically weighty sense.

that there could be a wide degree of variability in the semantics that language users employ to deal with abstract objects. How much, if any, semantic variability there is in practice is of course an empirical matter.

I opened this paper by claiming that nominalists are faced with an explanatory burden: how, if there are no abstract objects, are we to make sense of language? If what I have said in this paper is correct, there is no special explanatory burden on the nominalist. The supposed role that abstract objects play in language can always be captured just as well by presupposition and empty reference. This, of course, tells us little about the actual causal mechanisms that are employed in language use. But here, at least, the platonist and the nominalist are on an equal footing.

CHAPTER 3

FICTIONALISM, CONTENT, AND EXPLANATION

3.1 Introduction

There is an apparent conflict between nominalism and ordinary language. According to nominalism, abstract objects don't exist. And yet speakers frequently make utterances that refer to abstract objects. The nominalist might have been able to brush aside this apparent conflict if talk about abstract objects was restricted to throwaway comments, but that is not the case. Reference to abstract objects appears with great frequency in contemporary scientific and philosophical theories. If nominalism is to be plausible, the nominalist thus owes us some story about why their view does not require a widespread rejection of contemporary scientific and philosophical theories.

There are various strategies available to the nominalist. A number of these strategies go under the name of *fictionalism*. Unfortunately, there is no standard characterization of fictionalism in the literature. The label 'fictionalism' is used in a variety of incompatible ways, generally to pick out a restricted range of fictionalist theories.¹⁰¹ Broadly speaking, fictionalist theories attempt to account for the acceptability of discourse that refers to unwanted entities by drawing an analogy to *fiction*.

¹⁰¹ Compare the characterizations of fictionalism offered in the Stanford Encyclopedia of Philosophy entries on fictionalism (Eklund 2015), fictionalism in mathematics (Balaguer 2011), and modal fictionalism (Nolan 2016).

Consider the sentence “Santa Claus lives at the North Pole.” In uttering that sentence I do not intend to make the assertion that Santa Clause lives at the North Pole. I do not believe that, strictly speaking, Santa Claus exists and resides at the North Pole. In uttering the sentence I do not commit myself to the existence of Santa Claus, even though a speaker could assert the sentence to take on such a commitment. The idea that we utter sentences that we are nevertheless unwilling to accept as literally true is thus familiar from everyday contexts in which we utter fictional claims. The fictionalist wants to appropriate this idea and apply it to other areas of discourse; sentences of a discourse are taken to be analogous, in this way, to fictional claims.

Spelling out what this analogy with fiction amounts to is neither straightforward nor unproblematic. Fictionalist theories vary widely in their details, and many face serious objections. My aim in this paper is to explore several versions of fictionalism as a means of elucidating, and ultimately defending, the viability of the fictionalist strategy. After characterizing fictionalism in slightly more detail, I will begin by considering early modal meta-fictionalism. Although modal meta-fictionalism faces some serious difficulties, it provides a useful introduction to some of the main issues that fictionalism must address.

3.2 Fictionalism Characterized

It is worth emphasizing that fictionalism is not committed to the existence of fictions, or to the claim that the entities characteristic of the discourse exist as fictional

entities. Indeed, any version of fictionalism that the nominalist is interested in had better not rely on the existence of fictional entities. The analogy with fictional entities is supposed to be the following: abstract objects, like Santa Claus, don't exist. The nominalist cannot merely replace one purported abstract object with another. Platonists who believe that Santa Claus exists as some sort of abstract object won't find much use in the analogy with fiction.

In contrast to the eliminative anti-realist, the fictionalist recognizes that there can be value in referring to or quantifying over entities that don't exist.¹⁰² The modal fictionalist thinks that we should continue to utter sentences that refer to possible worlds, despite wanting to avoid an ontological commitment to possible worlds. The strategy here is not to claim that sincere assertions don't carry the ontological commitments that we think that do; the fictionalist allows that we *can* utter such sentences in such a way as to take on ontological commitments.¹⁰³ Someone could sincerely assert "Santa Claus lives at the North Pole" and thereby take on a commitment to Santa Claus. The idea is simply that it is also possible to assert the sentence in such a way as to not take on such a commitment. Since sentences can be used in these two ways, it would be helpful to have

¹⁰² Of course, this characterization is self-undermining if 'reference' and 'quantification' are understood as success terms. On this understanding of the terms, if the entities purportedly referred to don't exist then we don't have genuine reference. We nevertheless still have (empty) referential terms. And if the entities purportedly quantified over don't exist, then we don't have genuine quantification. But we still have quantification language. The fictionalist sees value in retaining that referential and quantificational language even when there are no entities being referred to or quantified over.

¹⁰³ There are nominalists (and agnostics) who reject the idea that our sincere utterances do carry ontological commitments. See Azzouni (2004), who argues that a literal utterance of "the number 2 exists" is true and about the number 2, but nevertheless doesn't carry an ontological commitment to numbers. See also Yablo's brief remarks in Yablo (1998), pp. 258-60. This claim should not be taken as representative of Yablo's position. Yablo's fictionalism evolves considerably from Yablo (1998) to Yablo (2001) to Yablo (2005), and no similar claim appears in his later papers.

some terminology to distinguish between those utterances that carry such a commitment and those utterances that don't.

Let us call an utterance that expresses precisely the content it seems to have, and carries an ontological commitment to the entities it refers to or quantifies over, a *literal assertion*. A literal assertion of "Santa Claus lives at the North Pole" is true if and only if Santa Claus really does exist, and lives at the North Pole. Likewise, a literal assertion of "There is a possible world at which all elephants are pink" is true if and only if there are possible worlds, and at least one of them contains no elephants that are not pink. Not all utterances are literal assertions. Let us call an utterance that is not a literal assertion a *non-literal utterance*. If I were to utter "Santa Claus lives at the North Pole", I would be making a non-literal utterance. One could not accuse me of inconsistency if I later claimed that Santa Claus does not exist.

A major division among fictionalists concerns how best to understand these non-literal utterances. All fictionalists about a particular domain accept that we should utter, but not literally assert, sentences of the target discourse. But, while fictionalists agree that not all utterances are literal assertions, they disagree about whether all utterances are assertions of some sort or another. A literal assertion requires that the speaker endorses, or accepts, the literal content of the utterance being made. If utterances can have content other than their literal content, then there can be assertions that are not literal assertions.¹⁰⁴

¹⁰⁴ The terminology of 'utterance', 'assertion', and 'literal assertion' are not standard in the literature. However, it is clear that something like these distinctions are implicit in all fictionalist theories. Eklund (2015), for example, talks about 'pretend-assertions', rather than about utterances that fall short of literal assertion. Yablo talks about sentences being "advanced in a fictional or make-believe spirit" (2001 p. 3), rather than about utterances that fall short of literal assertion. Both formulations have the unfortunate effect of contrasting literal assertion to pretense, or make-believe. It is clear, however, that many fictionalists have

According to *content fictionalism*,¹⁰⁵ the difference between a non-literal utterance and a literal assertion is a difference in content. Non-literal utterances of sentences have *different* content from literal assertions of those same sentences. The fact that a speaker is not asserting a sentence literally has an impact on what is said. For the content fictionalist, a non-literal utterance is still an assertion. It is simply an assertion of some content other than the literal content.

Consider again the sentence “Santa Claus lives at the North Pole.” According to the content fictionalist, when a speaker utters this non-literally the content of the speaker’s utterance is not that Santa Claus lives at the North Pole. Just what content the sentence does have when uttered non-literally is a matter of dispute among content fictionalists, but one proposal is that the content of a non-literal utterance of the sentence is that according to the Santa Claus fiction, Santa Claus lives at the North Pole. This is why, according to the content fictionalist, a non-literal utterance of “Santa Claus lives at the North Pole” does not commit the speaker to the existence of Santa Claus: the content being asserted doesn’t carry an ontological commitment to Santa Claus. Likewise, according to the mathematical content fictionalist, the content of a non-literal utterance of the sentence “2 is the only even prime number” does not carry an ontological commitment to numbers.

no interest in pretense or make-believe. Despite his focus on pretend-assertion, Eklund is aware that there are alternatives. See Eklund (2015) section 2, and Eklund (2005).

Other writers do not focus on pretense or make-believe. For example, while characterizing mathematical fictionalism, Balaguer (2011) explicitly states that “there is no claim here that mathematical discourse is a kind of fictional discourse”. Although he never says so explicitly, it is clear that Balaguer does take the mathematical fictionalist to be committed to the claim that we should utter (but not assert) mathematical sentences. By using the terminology of ‘utterance’, ‘assertion’, and ‘literal assertion’ we can make these commitments explicit, which will help us in categorizing and evaluating various fictionalist theses.

¹⁰⁵ The terminology here is my own, although Eklund (2015) draws a similar distinction between ‘content fictionalism’ and ‘force fictionalism’. See Eklund (2015) section 2.3. Because Eklund restricts his focus to descriptive fictionalism (which I will briefly discuss below), it is unclear whether he would view the distinction I draw as a natural extension of his distinction or as a separate but related distinction.

Content fictionalism is not the only option for explaining how non-literal utterances avoid ontological commitment. According to *attitude fictionalism*, non-literal utterances of sentences have the same content as literal assertions of those sentences. A non-literal utterance is simply an utterance where the speaker is not fully endorsing the content of that utterance; the fact that the utterance falls short of literal assertion is solely a feature of the speaker's attitude towards the literal content of that utterance. For the attitude fictionalist, a non-literal utterance is not an assertion at all; the literal content is being uttered, but not asserted.

According to the attitude fictionalist, when a speaker utters "Santa Claus lives at the North Pole" non-literally the content of the speaker's utterance *is* that Santa Claus lives at the North Pole. The speaker is simply not asserting that content. In this vein, Richard Joyce writes that a sentence about goblins "has exactly the same content whether it is used as part of a fairy tale or to foolishly assert something false. What changes is the 'force' with which it is uttered."¹⁰⁶

According to the attitude fictionalist, a speaker avoids taking on an ontological commitment to Santa Claus, not because the content expressed by a non-literal utterance of "Santa Claus lives at the North Pole" doesn't carry an ontological commitment to Santa Claus, but because a speaker doesn't take on an ontological commitment merely by uttering a sentence. A speaker is only committed to the ontological commitments of the sentences she asserts. So, according to the mathematical attitude fictionalist, the reason why a non-literal utterance of "2 is the only even prime number" does not commit the

¹⁰⁶ Joyce (2005), p. 293.

speaker to the existence of numbers is simply that the speaker is not asserting the content of her utterance.

We've now seen two general strategies for uttering sentences without taking on the ontological commitments of those sentences. According to content fictionalism, we can utter a sentence as a means of asserting something other than its literal content. And according to attitude fictionalism, we can utter sentences without really asserting anything at all. Which strategy is to be preferred? We are not yet in a position to properly evaluate this question, but we will return to it later in the paper. For now, let us suppose that either strategy will do.

What work does fictionalism do for the nominalist? By itself, not much. That we *can* avoid ontological commitments isn't a reason to think that we *should* avoid ontological commitments. The fictionalist strategy outlined above is not in any way restricted to abstract objects. We can just as well avoid making literal assertions about cars, or people, or houses, or electrons. Yet the nominalist isn't looking to reject such entities, and hence has no desire to avoid literal assertions about such entities. The ability to avoid ontological commitments is only useful if we have reason to avoid ontological commitments.

But here we have a puzzle. If reference to abstract entities is theoretically useful, then isn't that reason to take on an ontological commitment to abstract entities? And if the nominalist is right in thinking that abstract entities don't exist, then shouldn't we simply eliminate reference to abstract entities from our discourse? For fictionalism to be of any use, the answers to these questions must be "no" and "no". For if the answer to the first question were "yes", then the nominalist would have reason to abandon nominalism.

And if the answer to the second question were “yes”, then the nominalist would have reason to eliminate reference to the offending entities rather than fall back on non-literal utterances. So, a major explanatory goal of any fictionalist theory must be to provide a story about *why* reference to and quantification over certain entities is useful even if those entities don’t exist. In fact, I take this to be the central explanatory goal of a fictionalist theory, and it is a theme that we will return to again and again in this paper.

How, exactly, this story will go differs depending on the entities in question and the role that they play in the theories in which they appear. But it is crucial from the outset to keep in mind that there is a rather significant difference between reference to abstract entities and the abstract entities themselves. If reference to abstract entities can play a useful theoretical role even if those entities don’t exist, then the nominalist has a reason to continue to refer to those entities without accepting their existence.

Another important distinction between varieties of fictionalism concerns whether or not the theory is advanced in a descriptive spirit. As I’ve characterized fictionalism thus far, it provides an account of how language users can avoid taking on the ontological commitments of the sentences they utter by uttering those sentences non-literally. And, as I’ve just noted, this ability to avoid taking on ontological commitments is itself not of much use unless it is paired with some explanation of why we *should* avoid such commitments. So, the core fictionalist strategy I’ve outlined can be thought of as a prescriptive thesis: it is concerned with the attitude that we should adopt towards sentences of the target discourse.

But some fictionalists want to make a further claim about the attitude speakers actually do adopt towards sentences of the target discourse. According to these descriptive fictionalists, we should interpret the typical utterances made by speakers about some target domain as being non-literal. So, while all mathematical fictionalists hold that we should refrain from literal assertions of mathematical sentences, a descriptive mathematical fictionalist makes the further claim that speakers typically do in fact utter mathematical sentences non-literally.

The descriptive fictionalist is also committed to the prescriptive thesis. A descriptive fictionalist thinks that we both do and should utter the sentences of the target discourse non-literally.¹⁰⁷ Because descriptive fictionalists are committed to both the descriptive and the prescriptive theses, I will use the term “prescriptive fictionalism” in this paper to pick out those varieties of fictionalism that endorse only the prescriptive thesis, and “descriptive fictionalism” when a further descriptive claim is being made.¹⁰⁸

The distinction between descriptive and prescriptive fictionalism is independent of the distinction between content and attitude fictionalism. So we can distinguish four major varieties of fictionalism: descriptive content fictionalism, descriptive attitude fictionalism, prescriptive content fictionalism, and prescriptive attitude fictionalism.

¹⁰⁷ This is not to deny that someone could endorse the descriptive thesis but not the prescriptive thesis, and hold that although we do utter the sentences of discourse non-literally, we should assert those sentences literally and accept their ontological commitments. A platonist who was convinced by the arguments of the descriptive fictionalist about our actual language use could adopt such a position. But this position would not be a version of fictionalism.

¹⁰⁸ This distinction is made in the literature using a variety of different terms. I purposefully avoid the more common terminology of “hermeneutic fictionalism” (descriptive) and “revolutionary fictionalism” (prescriptive), because these terms are often used to distinguish exclusively between descriptive and prescriptive varieties of content fictionalism. See, for example, Burgess (2004) p. 23. The “prescriptive” and “descriptive” terminology originates in Joyce (2001), p. 14. Nolan draws the same distinction using the terms “normative” (prescriptive) and “descriptive”. See Nolan (2016), section 1.3.

3.3 Modal Fictionalism and Meta-Fictionalism

Modal fictionalism is fictionalism not about modality in general, but about possible worlds. Modal fictionalists want to retain talk about possible worlds, but reject the existence of any such entities. One can be a modal fictionalist and also believe in the existence of abstract objects – perhaps even leaving room for the existence of possible worlds as abstract fictional entities. But my interest is in fictionalism as a nominalist strategy, so we'll leave such views aside. The modal fictionalist I have in mind rejects the existence of both possible worlds and abstract objects.¹⁰⁹

A modal fictionalist wants to be able to utter the sentence “There is a possible world at which pigs fly”. How should we understand such utterances? Perhaps the simplest, and most commonly discussed, form of fictionalism is meta-fictionalism. According to meta-fictionalism, if the content of a literal utterance of a sentence is P, then the content of a non-literal utterance of that sentence is that according to the fiction, P. Accordingly, a non-literal utterance of a sentence of the discourse is true if and only if the literal content of that sentence is true *according to the fiction* that governs that discourse.¹¹⁰ So, an utterance of “There is a possible world at which pigs fly” by the

¹⁰⁹ There is also room to be a nominalist and believe in the existence of possible worlds. David Lewis, while not himself a nominalist, advances a theory of *concrete* possible worlds. See Lewis (1986b). I don't expect that there are many nominalists who find the notion of concrete possible worlds to be appealing. Generally speaking, the sorts of arguments that Lewis advances in favor of possible worlds are the sorts of arguments that are also advanced in favor of abstract objects. If nominalists are not moved by these arguments for abstract objects, they are unlikely to be moved by similar argument for concrete possible worlds. Concrete possible worlds are exactly the sort of ontological excess that most nominalists intend to do without when they reject abstract objects. So, in what follows, I will ignore the possibility that a nominalist might accept the existence of concrete possible worlds.

¹¹⁰ What is it for a sentence to be true ‘according to the fiction’? There are several options available to the fictionalist here, but I think that the most promising proposal is one put forward by Mark Balaguer. On this view, a sentence is true according to the fiction if and only if it would have been true if there had actually existed abstract objects of the kind being described. See Balaguer (2009) and (2011).

nominalist should be interpreted as having an implicit “according to the fiction” operator, and as expressing the content that according to the fiction of possible worlds, there is a possible world at which pigs fly. A nominalist’s utterance of “There is a possible world at which pigs fly” thus comes out true even if there are no possible worlds, since the content it expresses has to do with how things stand in the fiction of possible worlds rather than with how things stand with possible worlds.¹¹¹

We have yet to say anything about the utility of meta-fictionalism: why might the nominalist be interested in making utterances about how things stand in the fiction of possible worlds? There is an obvious answer: because the nominalist is interested in talking about the modal facts. Rejecting the existence of possible worlds is not tantamount to rejecting modality. The nominalist wants to hold that it is still possible for pigs to fly, even if there is no possible world at which pigs fly. Whereas the realist about possible worlds typically wants to ground modal facts – facts about what is possible and what is necessary – in facts about possible worlds, the nominalist does not.¹¹² The modal fictionalist rejects the idea that modal facts are grounded by facts about possible worlds, yet wants to retain talk about possible worlds. What, then, is the connection between how things stand in the fiction of possible worlds and the modal facts?

¹¹¹ Meta-fictionalism must be filled out in various ways. For one, there is no single fiction about possible worlds to count as “the fiction”. Some particular story must be chosen, such as Lewis’ (1986b) modal realism. The task of specifying the appropriate fiction is a matter of pragmatic preference; not all fictions will be equally useful. The meta-fictionalist must choose a fiction that suits her purposes. Rosen (1990), for example, constructs a fiction based on Lewis’ modal realism. See Rosen (1990), pp. 333-5. More than one theory might do equally well, in which case the prescriptive fictionalist must simply choose one or the other to adopt for the discourse. This becomes problematic if meta-fictionalism is being advanced as a descriptive theory and we want to ascribe the same content to different nominalists making utterances about possible worlds. After all, it is not plausible that speakers have coordinated and settled on one particular fiction. If we are not concerned with making a descriptive claim then we can set this worry aside: speakers *could* coordinate on a particular fiction if need be.

¹¹² I talk about metaphysical “grounding” here, but I don’t mean to be making any strong commitments to a metaphysical framework here. The unhappy reader is invited to paraphrase my talk of grounding as talk about supervenience and fundamentality.

One response, sometimes called ‘strong’ modal fictionalism, is that the fiction of possible worlds grounds the modal facts. The strong modal fictionalism simply replaces possible worlds with the fiction of possible worlds as the grounds for modal facts. This is a bold claim: how things stand modally is a matter of how things stand in the fiction of possible worlds. Strong modal fictionalism strikes me as highly implausible. For those who do accept the existence of abstract objects, it is hard to see how construing possible worlds as abstract fictional objects would be any improvement over alternative ersatz theories on which possible worlds are construed as abstract objects. For the nominalist, the problem is more acute: the existence of abstract fictions are rejected altogether. So the only way I can see to interpret strong modal fictionalism in a way consistent with nominalism is as the claim that modality is ‘all in the head’ of those who are engaged in the fiction of possible worlds. This is not a claim I find plausible, and although I won’t offer any arguments against it, I will set it aside here.¹¹³

If strong modal fictionalism is rejected, then it becomes quite apparent that modal fictionalism is *not* a theory of modality. No claim is being made that the modal facts are grounded in facts about possible worlds or the fiction of possible worlds. Modal fictionalism is thus not tenable unless it is paired with an independent, nominalistically acceptable, theory of modality.

What should a nominalistically acceptable theory of modality look like? I know of no real option but to rely, in some way or another, on *primitive* modality. A theory of

¹¹³ I’m not entirely unsympathetic to the idea that much of *metaphysical* modality is a conceptual matter. See Sidelle (1989). However, it seems to me that any such view still depends on some primitive logical modality – on the notion of logical possibility, or consistency. And it seems to me that this sort of view loses any plausibility it has on the strong modal fictionalist variant where modality is a matter of how things stand with our fiction of possible worlds. Modality plays too important a role in too many philosophical theories to be dismissed as an artifact of our story telling.

possible worlds, on the other hand, promises to deliver modality without primitive modality. For this reason, the inability of nominalists to rely on possible worlds in providing an account of modality is often seen as a serious cost of adopting nominalism. Because this paper is part of a general defense of nominalism, I want to make a few remarks here before returning to our question of why nominalists might be interested in how things stand in the fiction of possible worlds.

The first point to emphasize is that relying on primitive modality is not the same thing as taking all modal facts as primitive. Consider the sentence “It is (metaphysically) possible for pigs to fly”. For the realist about possible worlds, the truth of this sentence is grounded by the fact that there is a possible world in which pigs fly. What about for the nominalist? My preferred strategy is to ground facts about possibility in facts about logic. On this view, facts about modality more generally are a matter of logical consistency. Facts about consistency are irreducibly modal: consistency concerns not just what is but what *can* be.¹¹⁴ Once we have logical modality, we can get metaphysical modality: the (metaphysical) possibilities are those that are consistent with whatever further metaphysical constraints apply.¹¹⁵ The concepts of *pig* and *flying* can be consistency

¹¹⁴ My claim that facts about consistency are irreducibly modal is controversial. Phillip Bricker has suggested to me that although he takes consistency as ideologically basic, he does not consider it a modal notion. Rather, on his account, every logically consistent theory is true at some world, and it is those worlds that ground the modal facts. But I’m not convinced that consistency ought to be regarded as a non-modal notion. Insofar as consistency constrains the space of possible worlds, it strikes me as a modal notion: it places a constraint not only on which worlds do exist, but on which worlds *can* exist. Can the realist take consistency to constrain only which worlds do exist, but not to constrain which worlds can exist? I must admit that I have difficulty understanding this proposal. But suppose that we grant that consistency ought to be regarded as a non-modal notion. Can the nominalist then ground modal facts directly in (non-modal) facts about consistency, and similarly bypass any primitive modality? Perhaps. Logical possibility simply reduces to logical consistency on this view, and if consistency is not a modal notion then the modal facts are grounded in non-modal logical facts.

¹¹⁵ I’m sympathetic to the idea that what is metaphysically possible (rather than logically possible) can vary contextually, and that the metaphysical constraints relevant to determining what is metaphysically possible can vary based on what is contextually salient. But the general nominalist account I’m sketching here does not depend on that being the case.

applied to the same entity, and accordingly it is metaphysically possible for there to be a pig that can fly.¹¹⁶ While this is only a bare sketch of a theory, its aim is illustrative: the nominalists' commitment to primitive modality is not a commitment to widespread primitive modality. And indeed, it is not obvious to me that the nominalist is worse off than the realist about possible worlds here.

Realists about possible worlds see the space of possible worlds as saturated: for any given theory of possible worlds, every possibility that is consistent with the metaphysical constraints of that theory will be captured by some possible world. This allows for an analysis of possibility and necessity in terms of the possible worlds that exist. But what guarantees the existence of all, and only, these possible worlds? How can we be so certain that there is a possible world at which pigs fly? And how can we be so certain that there are no logically impossible worlds, at which, for example, it is true that pigs fly and also true that it is not the case that pigs fly? The tight connection between possible worlds and possibility breaks down in either scenario.

There are two answers that the realist can give. The first is that the fact about which worlds exist is simply a brute fact, in need of no explanation. If this approach is taken, it isn't clear that the realist has an advantage over the nominalist: where the nominalist posits brute modal facts, the realist posits brute facts about the existence of

¹¹⁶ One of the more serious challenges to the idea that we can recover metaphysical modality from logic and non-modal facts (both conceptual and empirical) comes from *a posteriori* necessities, which intuitively seem more resistant to such a treatment. The interested reader is referred to Sidelle (1989) for an extended response to this challenge.

A second worry concerns nomological modality. Can it be given a similar treatment? We might try the following: nomological possibility is a matter of consistency with the laws of nature. But the metaphysics of the 'laws of nature' is a notoriously difficult issue, so just what this proposal amounts to isn't entirely clear.

possible worlds. There is a difference in the brute facts, but not one that makes for a clear theoretical advantage.¹¹⁷

The second answer that the realist can give is that what worlds there are is determined by facts about consistency. A world in which pigs fly exists because it is consistent with the constraints placed on the space of possible worlds. A world in which it is true that pigs fly and also true that it is not the case that pigs fly does not exist, because it is logically inconsistent. This is an appealing explanation, but can only be bought at the price of introducing some primitive modality: consistency, or some suitably similar notion, is playing an important explanatory role.¹¹⁸ Here too the realist does not have a clear theoretical advantage over the nominalist: both rely on primitive modality.

I cannot defend a fully developed, nominalistically acceptable, theory of modality here. So, let us suppose that the nominalist has such an account and return to the question at hand: why might the nominalist be interested in talking about possible worlds?

Modal meta-fictionalism answers this question in two steps. First, talk of possible worlds is understood as talk about the fiction of possible worlds. Second, biconditionals are introduced to connect talk about the fiction of possible worlds with modal facts.¹¹⁹

Possibly *P* if and only if according to the fiction of possible worlds, *P* is true at some possible world.

¹¹⁷ Indeed, it seems to me that this answer gets the order of explanation backwards. Why does the realist believe that there is a world at which pigs fly? Because it is consistent. If the space of possible worlds is not as saturated as the realist supposes, the correct analysis would be that there are some possibilities unaccounted for, not that possibility is more constrained than we realized. See Balaguer (2010), p. 184, for a similar line of thought.

¹¹⁸ Although see my remarks in footnote 14 above. The conclusion that the realist must rely on primitive modality can be avoided if consistency is not regarded as a modal notion.

¹¹⁹ I follow the formulation given in Nolan (2016). This simple formulation has technical difficulties that we will not concern ourselves with.

Necessarily P if and only if according to the fiction of possible worlds, P is true at all possible worlds.

It is important to note that the first of these two steps is sufficient to guarantee that the content of non-literal utterances about possible worlds carry nominalistically acceptable content: sentences that purport to be about possible worlds are really about the fiction of possible worlds. But we still need an explanation of why a fictionalist would make utterances about possible worlds. After all, we want to know why fictionalists would care about what's true in the fiction of possible worlds. We've said that the fictionalist is ultimately interested in the modal facts – facts about what is possible and necessary – but we've also rejected strong modal fictionalism – the view that the modal facts are grounded in by facts about the fiction of possible worlds. So we want to know exactly what the connection between the modal facts and the fiction of possible worlds is. The biconditionals are introduced to provide just such an explanation: we can see how some of the claims that are made about what is true in the fiction of possible world are connected to modal facts. There is thus some utility in talk about possible worlds: it provides a convenient means of talking about the modal facts. Unfortunately, the meta-fictionalist account faces some serious difficulties.

Meta-fictionalism is a version of content fictionalism: non-literal utterances about possible worlds are understood to really be claims about how things stand according to the relevant fiction. To explain why we should care about how things stand according to the fiction, the meta-fictionalist introduces biconditionals that connect some facts about the fiction to some further facts that we do care about. In the case of modal meta-

fictionalism, the relevant fiction is the fiction of possible worlds, and the further facts are facts about what is possible and necessary.

In order to properly evaluate meta-fictionalism, it is important to say a bit more about fictional discourse more generally. We cannot treat all utterances of the sentence “Santa Claus lives at the North Pole” in the same way. There is an obvious difference between a literal utterance made by a child who believes in Santa Claus, and a non-literal utterance made by an adult who is playing along with the fiction. But it is also important to keep in mind that the adult doesn’t have to play along with the fiction. When playing along with the fiction, the adult is correct in uttering the sentence “Santa Claus lives at the North Pole”.¹²⁰ However, when not playing along with the fiction, uttering the sentence “Santa Claus lives at the North Pole” would be incorrect.

Likewise, when a fictionalist makes a non-literal utterance of the sentence “There is a possible world in which pigs fly”, they are *engaged* in the fiction – they are playing along with the fiction that there are possible worlds, without accepting that possible worlds really exist. But there are contexts when the fictionalist will not be willing to utter the sentence “There is a possible world in which pigs fly”. When fictionalists enter the philosophy room and drop the fiction, they will generally want to deny, or overturn, what they said earlier. For example, a fictionalist about numbers might assert “2 is the only even prime number” when in a regular conversational context, but then later hold that this

¹²⁰ The content fictionalist has an easy account of correctness-conditions: a sentence of the discourse is correct if and only if its non-literal content is true. For the content fictionalist, what is being asserted by a non-literal utterance of a sentence of the discourse *is* true (and hence correct), it’s just that what is being asserted is not the literal content of those sentences.

The attitude fictionalist, on the other hand, must distinguish between correctness-conditions and truth-conditions. After all, for the attitude fictionalist, a non-literal utterance of “Santa Claus lives at the North Pole” is correct in the relevant fiction, but it is not true. One option for the attitude fictionalist is to hold that a sentence of the discourse is correct if and only if the literal content of the sentence is true according to the relevant fiction.

earlier assertion is false since there are no numbers at all. In this later context, we can say that the fictionalist is *disengaged* from the fiction.¹²¹ Any reasonable variant of fictionalism must allow for both these types of utterances, just as we do for cases of fiction: when I am not engaged in the Santa Claus fiction, I will deny that Santa Claus lives at the North Pole. This gives us a new way to frame the work that modal meta-fictionalism is trying to do when it draws a connection between claims about possible worlds and claims about how things stand in the fiction of possible worlds: it connects the sentences that *engaged* modal fictionalists are willing to utter to sentences that *disengaged* modal fictionalists are willing to assert.

One objection to meta-fictionalism is that, as a version of content fictionalism, it does not adequately distinguish between describing a story and telling a story. That is to say, it fails to give an adequate account of the content of utterances made by engaged speakers. When a speaker engages with a fiction, they participate in a bit of pretense or make-believe – they are taking onboard the fiction of numbers, or of possible worlds, and talking accordingly. But it is one thing to engage in a bit of make-believe, and quite another to talk *about* the make-believe. As Richard Joyce notes, “if every sentence of the story uttered contained an unpronounced fiction operator, then there is no sense to be made of the claim that the storyteller is pretending.”¹²² Yablo compares this to the rules of English: speakers rely on the rules of English when uttering English sentences, but they are *not* talking about those rules in doing so. Likewise, when fictionalists utter “The number of apostles is twelve”, they are relying on the number fiction, but not talking

¹²¹ The terminology of “engaged” and “disengaged” is drawn from Yablo (2001). See also Joyce (2005) p. 292.

¹²² Joyce (2005), p. 291.

about it.¹²³ For this reason, few content fictionalists are now attracted to the idea that the content of the relevant non-literal utterances concerns how things stand in the appropriate fiction.

The force of the objection is that meta-fictionalism is motivated by an overly naïve understanding of the content of typical fictional utterances; speakers who are engaged in a fiction should not be interpreted as talking about the fiction. The extent to which this should worry the meta-fictionalist depends in large part on whether or not meta-fictionalism is advanced as a version of descriptive content fictionalism. As a descriptive claim, meta-fictionalism is problematic; if non-literal utterances about abstract objects are to be treated in the same way as paradigmatic utterances by speakers engaged in a fiction, then meta-fictionalism is not a plausible theory.

But there is still room for the core of the meta-fictionalist approach to be adapted as a version of prescriptive content fictionalism. Meta-fictionalism may not stand up as a plausible descriptive theory of the content of typical fictional utterances, but it is open to the meta-fictionalist to reject a close analogy with fiction: it is still the case that a modal meta-fictionalist could make non-literal utterances about possible worlds as a means of talking about how things stand according to the fiction of possible worlds. That said, rather than giving up on the analogy with fiction and retaining the meta-fictionalist account of the content of utterances, the fictionalist may well prefer to retain the analogy with fiction and give up on the meta-fictionalist account of the content of utterances. After all, it was precisely this analogy with fiction that motivated the meta-fictionalist approach in the first place.

¹²³ Yablo (2001), p. 6.

An attitude fictionalist will likely see the difficulty here as a reason to reject content fictionalism altogether: if we consider the content of an utterance of “Santa Claus lives at the North Pole” made by a speaker engaged in the fiction, the simplest explanation is that the engaged speaker has uttered the sentence with exactly the same content as it normally has. But because this utterance was made in the context of being engaged in the fiction, it should not be taken literally and the speaker needn’t take on the ontological commitments carried by that sentence. There are, however, alternatives available to the content fictionalism. We’ll consider some of those alternatives in the next section. But for now, let us turn our attention to another serious difficulty with meta-fictionalism.

A second difficulty with modal meta-fictionalism is with the biconditionals that are crucial to explaining why the nominalist should care about what is true in the fiction of possible worlds. The objection, known as the Brock-Rosen objection, can be formulated as follows:¹²⁴

Consider the following sentence: “Necessarily there exists many worlds”. The meta-fictionalist does not want to be forced to endorse this sentence as true. Yet, given the modal meta-fictionalists’ biconditionals, the sentence will be true if according to the fiction of possible worlds it is true at all possible worlds that there exist many worlds. But that *is* true according to (the most common variants of) the fiction of possible worlds. The meta-fictionalist must thus take on a commitment to the existence of possible worlds.

¹²⁴ This formulation is based on the presentation of the problem in Nolan and O’Leary-Hawthorne (1996) p. 24, although I present the objection against sentential formulations of meta-fictionalism, whereas Nolan and O’Leary-Hawthorne formulate their objection against Rosen’s propositional formulation. The Brock-Rosen objection was originally presented by Brock (1993) and Rosen (1993).

Originally, some thought that the lesson of the Brock-Rosen objection was that a bit of technical sophistication is needed – that to avoid such embarrassing results, the meta-fictionalist must be very careful about which fiction of possible worlds they employ and about the exact specification of the biconditionals. But the objection is not merely a technical one. In fact, Nolan and O’Leary-Hawthorne have shown that similar arguments can be generated not just for modal meta-fictionalism but for meta-fictionalism about other entities, such as numbers, as well.¹²⁵

Imagine a mathematical meta-fictionalist who endorses the following biconditional: “For any sort of entity Φ , there are n Φ s iff according to the mathematical fiction, the number N numbers the Φ s”, where n stands for a numerical adjective and N names the number corresponding to the numerical adjective. Given this biconditional, the mathematical meta-fictionalist holds that there are two moons around Mars if and only if according to the mathematical fiction, the number 2 numbers the moons around Mars. Now consider the sentence: “There are (at least) three numbers”. Given the biconditional, this sentence is true if according to the mathematical fiction, the number 3 numbers the numbers (or at least numbers some subset of the numbers). But this is true according to the mathematical fiction, so our mathematical meta-fictionalist is thereby seemingly committed to numbers.¹²⁶

¹²⁵ Nolan and O’Leary-Hawthorne (1996). Note that my claim that the Brock-Rosen objection is not merely a technical one is controversial. See Nolan (2016) for more. The extent to which the Brock-Rosen objection can be considered a technical one depends a great deal on what *role* the biconditionals are designed to be playing. If the biconditionals are being used to connect claims about the fiction to further claims of interest to the nominalist, as I suggest they must be, then we want the biconditionals to be quite general. See the discussion of this issue in the remainder of the section. Any technical solution to the Brock-Rosen objects functions to restrict the applicability of the biconditionals.

¹²⁶ Nolan and O’Leary-Hawthorne (1996), pp. 25-6.

How serious is the Brock-Rosen objection? It is undeniable that, at least in some cases, the biconditionals lead to problematic results for the meta-fictionalist. There is no other option but to restrict the use of the biconditionals in some way or another. One suggestion, made by Matti Eklund, is that we restrict the use of the biconditionals to speakers who are engaged in the fiction: “it is only when speaking *within the fiction* that we rely on them”.¹²⁷ Eklund is correct that the biconditionals are rendered harmless when speaking within the fiction. But recall the role that the biconditionals were meant to be playing. Meta-fictionalism involves two steps. First, non-literal utterances (by the engaged fictionalist) are understood as claims about how things stand ‘according to the fiction’ (claims that the disengaged fictionalist is able to accept). The biconditionals were then used to connect claims about how things stand according to the fiction, to some further claims that the fictionalist is actually interested in – claims about how things stand in the fiction of possible worlds are connected to bare modal claims. Note that the biconditionals are not intended to be employed by fictionalists engaged in the fiction at all. The biconditionals connect two sets of claims that the disengaged fictionalist will regard as true: claims about the fiction of possible worlds and claims about what is possible and necessary. So, if we follow Eklund’s suggestion of restricting the use of the biconditionals to engaged speakers, then the biconditionals cease to play the role they were designed for. This obviously will not do.

Another suggestion is to restrict the cases in which the biconditional is applicable.¹²⁸ After all, not all uses of the biconditional prove problematic. Consider

¹²⁷ See Eklund (2015).

¹²⁸ See Yablo (2001), pp. 8-12, for an example of this sort of move. He speaks of “principles of quasi-assertion” rather than biconditionals, and restricts them to cases of “applied X-talk” – cases where numbers are only “representational aids” rather than “things represented”. Terminology aside, his suggestion is

again one the biconditionals employed by the modal meta-fictionalist: Necessarily P if and only if according to the fiction of possible worlds, P is true at all possible worlds. This biconditional runs into trouble when P is itself about possible worlds. Likewise, the mathematical meta-fictionalist's biconditional "For any sort of entity ϕ , There are n ϕ s iff according to the mathematical fiction, the number N numbers the ϕ s" runs into trouble when the entity ϕ is itself an abstract mathematical entity. What if we were to simply restrict the applicability of the biconditionals by explicitly ruling out this type of case?

The result is that we have some statements about how things stand according to the fiction that the disengaged meta-fictionalist will accept as true, but that do not correspond to any statement about what is possible or necessary that the disengaged meta-fictionalist will accept as true. If we restrict the applicability of the biconditionals, we likewise restrict the meta-fictionalist explanation of why the nominalist would spend time making non-literal utterances about possible worlds.

How damaging is this restriction of the biconditionals? It would be damaging if the meta-fictionalist otherwise had a solid explanation for the utility of talk about what is true according to the fiction of possible worlds. But suppose for the moment that the Brock-Rosen objection could be resolved in some other way, and that the biconditionals did not have to be restricted. It would still be the case that there are great many sentences about what is true in the fiction that are simply not covered by the biconditionals. After all, the biconditionals only provide a connection to the modal facts for claims of the form " P is true at some possible world" and " P is true at all possible worlds". But the

roughly the one I'm considering here: allow the biconditionals to be applicable in some cases but not in others.

statements made when talking about possible worlds are not so constricted. Consider, for example, the sentence “Qualitatively indiscernible possible worlds are identical”. The meta-fictionalist will interpret a non-literal utterance of this sentence as having the content that according to the fiction of possible worlds, qualitatively indiscernible possible worlds are identical. But why should the nominalist care about this fact? This fact about the fiction of possible worlds does not seem to be directly connected in any way to a modal claim that the nominalist would care about.

And this is a serious problem. It won't do for the fictionalist to provide an explanation of the utility of some subset of the non-literal utterances that they make. If only a subset of the sentences of the discourse are useful to utter, then we would expect the fictionalist to restrict the utterances that they make to those that are useful. But in practice, the non-literal utterances that fictionalists make are not so restricted. We need an explanation of why these sentences, which the fictionalist find useful to utter but to which the biconditionals are not applicable, are useful to utter. The explanation we have been working with so far – that each non-literal utterance is useful because it tells us about some fact that we do care about – collapses.

This problem is acute. Even if the modal meta-fictionalist is able to restrict their utterances in such a way so as to avoid this problem, the same can certainly not be done in the case of mathematical entities. The problem is that claims about mathematical entities are extremely widespread and are a crucial component of mathematical practice. But those are precisely the sorts of claims that prove to be problematic for the biconditionals that the meta-fictionalist relies on. Such utterances cannot be eliminated in practice, and as a result the meta-fictionalist is left without an explanation of why the

nominalist should care how things stand in the fiction of numbers. Unless some alternative explanation of the utility of non-literal mathematical utterances can be provided, mathematical practice is left without a nominalistically acceptable explanation.

If meta-fictionalism cannot offer an explanation of why a nominalist would be interested in making mathematical utterances then meta-fictionalism must be rejected. What alternatives are available? In the following section we will consider several variants of content fictionalism, and see if they fare any better than meta-fictionalism.

3.4 Variants of Content Fictionalism

Meta-fictionalism, it was claimed, made the mistake of conflating engaging in a fiction with talking about that fiction. For the content fictionalist, then, the question is: if non-literal utterances don't have content that is about the fiction, what is the content of non-literal utterances. One fairly natural option that comes to mind is a view that Yablo calls *object fictionalism*. According to object fictionalism, non-literal utterances are really assertions about the world – namely the way the world needs to be in order to make the literal content of those utterances true according to the relevant fiction. So, a mathematical object fictionalist might hold that the real content of a non-literal utterance of a sentence of form “The number of E's = n” is that there are n E's. For example, the content of a non-literal utterance of the sentence “The number of planets = 8” is that there

are 8 planets; there being 8 planets is what is required of the world in order for it to be true according to the fiction of mathematics that the number of planets = 8.¹²⁹

Object fictionalism does not conflate engaging in a fiction with talking about that fiction; the engaged fictionalist uses the fiction of possible worlds, or the fiction of numbers, to make claims directly about the world. In a sense, object fictionalism collapses the two steps of meta-fictionalism into a single step: non-literal utterances in the fiction are taken to be assertions about the fiction-independent world.

Object fictionalism is not a tenable view.¹³⁰ While there are no biconditionals to deal with, something akin to the Brock-Rosen objection is still applicable. To see the problem, consider the following sentence: “The number of even prime numbers = 1”. What is required of the world in order for this sentence to be true according to the mathematical fiction? What is required is that there is one even prime number. Hence, the real content of the sentence, according to the object fictionalist, is that there is one even prime number. Here we see that the same sort of sentence that proved problematic for meta-fictionalism likewise proves problematic for object fictionalism. The fictionalist engaged in the fiction of mathematics wants to accept the sentence “The number of even prime numbers = 1”, since it is true according to the fiction of mathematics, yet according to object fictionalism this commits him to the existence of an even prime number.

The problem here is even worse for the object fictionalist than it was for the meta-fictionalist. We saw how the meta-fictionalist could restrict the applicability of their biconditionals to avoid unwanted commitments. When doing so, the meta-fictionalist still provided us with an account of the content of every non-literal utterance. What was

¹²⁹ Yablo (2001), pp. 6-7.

¹³⁰ Yablo (2001), pp. 8-12.

lacking, in the absence of the biconditionals, was an explanation of why we would care about that content. Object fictionalism, on the other hand, cannot even provide us with an account of the content of non-literal utterances without running into trouble.

A more sophisticated variant of content fictionalism, proposed by Yablo in response to this defect with object fictionalism, is his *relative reflexive fictionalism*.¹³¹ Yablo argues that reference to abstract objects plays two distinct roles in our discourse, and that we must treat the cases separately; if we focus on one of the roles, and not the other, we will inevitably end up with an inadequate account of content. Here is how Yablo characterizes the distinction:

There are actually two roles X's can play. Sometimes they function as representational aids. This is how butterflies function in "I had butterflies in my stomach," and numbers function in "the number of Martian moons is 2." Other times they function as things-represented. This is how butterflies function in "the butterflies were splattered all over the windscreen," and how numbers function in "there are no numbers, that's just a way of talking."

Here Yablo draws an analogy between the way we talk about numbers and the way we use instances of figurative language. A speaker familiar with the figurative use of the sentence "I had butterflies in my stomach" will understand that, despite the presence of the word "butterflies", the sentence isn't about butterflies at all. Someone who utters that sentence is instead indicating that they are experiencing a particular sort of sensation associated with being anxious or nervous. Someone who insists on interpreting the sentence literally – as being about the actual presence of butterflies in the speaker's stomach – has missed the point entirely. Yablo's suggestion is that numbers often play a

¹³¹ Yablo (2001), pp. 11-15. It should be noted that Yablo's views on fictionalism evolve significantly over time. So when I talk about Yablo's views, they should be understood as his views at the time of the article I cite.

similar role: when someone utters the sentence “the number of Martian moons is 2” they are talking about the Martian moons, not about the number 2. To interpret the speaker as saying something about numbers is to interpret them too literally.

Notice that the analogy to fiction that Yablo is working with in this case has shifted substantially from the analogy to fiction that we originally had in mind for the modal meta-fictionalist. The modal meta-fictionalist was trying to give an account of the utterances of speakers who continued to speak about possible worlds despite disbelieving in, or being agnostic about, such entities. A connection was drawn to similar cases of fiction: an adult who makes utterances about Santa Claus despite disbelieving in Santa Claus. In the case of numbers as representational aids, Yablo is instead drawing a connection to figurative language. Someone who utters the sentence “I had butterflies in my stomach” is doing something quite different from someone who utters “Santa Claus lives at the North Pole”. The use of the former sentence doesn’t require that the speaker engage in the fiction of butterflies, or anything of the sort.¹³²

Numbers do not *only* function as representational aids. In other cases, speakers *are* talking about numbers, or butterflies. It won’t do to simply brush aside all talk of numbers as really being about something else. In Yablo’s diagnosis, that’s precisely the problem with object fictionalism. According to relative reflexive fictionalism, the account of content given by object fictionalists only applies to sentences in which the entities of the fiction (numbers, for example) function only as representational aids. When numbers function as things-represented, on the other hand, the engaged fictionalist should be

¹³² An interesting feature of Yablo’s account is that it is intended as a version of descriptive content fictionalism. That is to say, Yablo isn’t suggesting that speakers *could* treat mathematical sentences analogously to figurative language, but rather that speakers typically *do* treat mathematical sentences analogously to figurative language. Yablo spends considerable effort trying to make this descriptive claim plausible. See Yablo (2000) and (2001).

understood as trying to portray the numbers as they are supposed to be imagined by those who are playing along with the mathematical fiction. So, when the engaged fictionalist utters the sentence “The number of even prime numbers = 1” this is not, as the object fictionalist holds, an assertion that there is one even prime number. Rather, the engaged fictionalist is playing along with the mathematical fiction, and evaluating mathematical sentences based on “how things are imagined to be” according to that fiction.¹³³

How is the distinction between representational aids and things-represented supposed to help relative reflexive fictionalism avoid a variant of the Brock-Rosen objection? The distinction is supposed to help because all of the problematic sentences – those that are susceptible to some variant of the Brock-Rosen objection – are sentences in which abstract objects appear as things-represented. This allows us to treat sentences in which abstract objects appear only as representational aids in exactly the way that the object fictionalist does. It simply remains to give a suitable account of sentences in which abstract objects appear as things-represented. Unfortunately, Yablo does not fully fill out the details of this proposal. We are told that the engaged fictionalist evaluates sentences in which numbers appear as things-represented relative to how things are imagined to be in the fiction of numbers. But we aren’t given a story of the content of those utterances. Indeed, nothing that Yablo says commits us to thinking that relative reflexive fictionalism is a version of content fictionalism at all: evaluating the acceptability of mathematical sentences based on how things are imagined to be according to the fiction is entirely compatible with an attitude fictionalist approach on which the content of such utterances just is its literal (but false) content.

¹³³ Yablo (2001), p. 14.

But let us suppose that relative reflexive fictionalism is intended as a variant of content fictionalism. One proposal for filling out the view is that, in cases where numbers function as things represented, meta-fictionalism was on the right track. As a descriptive thesis, this is difficult to accept: we've already seen that meta-fictionalism conflates engaging in a fiction with talking about that fiction. But if we understand relative reflexive fictionalism in a more prescriptive spirit, then it seems plausible enough that nominalists could utter sentences such as "The number of even prime numbers = 1" as a means of talking about what is true according to the fiction. This is a bit of an unhappy compromise, since Yablo originally intended relative reflexive fictionalism as version of descriptive content fictionalism. But I'm not sure how else to fill out the missing details in his proposal. Moreover, we are still left without an explanation of *why* nominalists would be interested in making utterances about how things stand in the fiction. So we have not made much progress over the original meta-fictionalist proposal we began with.

3.5 Reference to Abstract Objects as a Descriptive Aid

Let us take stock of where this leaves us. We're in search of an explanation for why nominalists should be interested in making non-literal utterances about abstract objects. And we have a proposal: non-literal utterances are useful because they describe (or entail) further facts of interest to the nominalist. We've now considered several suggestions for how that might go, all of which have come up short. The point I hope to

drive home here is that the difficulties we've seen are not the result of insufficient technical sophistication. No alternative theory of content will fare better.

The fundamental problem is that many sentences in which abstract objects appear as things-represented simply lack much nominalistically acceptable content. Consider again the sentence "The number of even prime numbers = 1". If we try to look past what this is telling us about numbers, or about how things stand according to the fiction of numbers, it's hard to see that anything much is left. What else does this sentence tell us? Very little, if anything at all. And it is hard to see how so very little could be so crucial to scientific practice. The only reasonable conclusion, it seems to me, is that such sentences are important precisely because of what they tell us about numbers (or, from the nominalist's perspective, about the fiction of numbers). If we are looking for an explanation of why such utterances are useful to the nominalist, it won't do to say that it is because of what such utterances tell us about the nominalistically acceptable world. They tell us very little apart from how things stand in the fiction of abstract objects.

Note that the problem here is not with content fictionalism. The idea that non-literal utterances are useful because of what they tell us about the nominalistically acceptable world is entirely compatible with attitude fictionalism. Whereas the content fictionalist focuses on content, the attitude fictionalist does not. But the explanation is essentially the same: non-literal utterances are useful because they can be used to convey (in some form or another) information of interest to the nominalist. The problem is that, considered in isolation, the sentence "The number of even prime numbers = 1" doesn't seem to tell us much at all about the physical world. The proposal that non-literal utterances are useful because they inform us about further facts of interest to the

nominalist is inadequate because it cannot account for the utility of uttering a sentence that doesn't tell us anything of interest about the physical world.

It wouldn't do to overstate the problem here. What I've pointed out is that the proposal that non-literal utterances are useful because they describe further facts of interest to the nominalist can't be the *whole* story. But that doesn't stop it from being a big part of the story. After all, in many cases this proposal seems to get things exactly right.¹³⁴

Consider the sentence (S_1): "The number of objects in the room = 2". S_1 is a claim about the identity of two numbers – the number of objects in the room and the number 2. But the purpose of uttering S_1 is not to make a claim about the platonic realm. Rather, an utterance of S_1 is used to convey a fact about the physical world: there is an object x in the room, and a distinct object y in the room, and every object in the room is identical to x or y . Here, mathematical discourse is being used to make a nominalistically acceptable claim about the world. The mathematical discourse is dispensable, but it has a certain representational convenience. To highlight this point, consider sentence (S_2): "The number of objects in the room = 2317". As with S_1 , what S_2 tells us about the physical world can be specified in a nominalistically acceptable vocabulary (exercise left to the reader). But S_2 allows us to express what we want to say about the physical world much more simply than we could without mathematical discourse.

In other cases, what a mathematical sentence requires of the world can only be specified in a mathematics-free vocabulary by an infinite disjunction. Consider the

¹³⁴ The remainder of this section repeats material that is also covered in my paper "Language Without Abstract Objects", with some minor modifications. In both papers it serves to illustrate the descriptive power of language that involves reference to abstract objects for the nominalist.

sentence (S_3): “The average star has 2.4 planets”.¹³⁵ For S_3 to be true, it must be the case that: either there exists 5 stars and 12 planets, or there exists 10 stars and 24 planets, or there exists 15 stars and 36 planets, etc. The sentence S_3 thus has an indispensable expressive brevity – it can be used to express nominalistically acceptable content that otherwise couldn’t be specified without an infinite disjunction.

Now consider a case discussed by Yablo. Yablo has us imagine a physicist studying escape velocity. As a result of her experiments, the physicist knows many facts of the following form: “(A) A projectile fired at so many meters per second from the surface of a planetary sphere so many kilograms in mass and so many meters in diameter will (will not) escape its gravitational field.”¹³⁶ Now, this is not a sentence about mathematical entities – it is a sentence about projectiles, and an interesting connection between their velocities and their masses when subject to the effects of gravity. But it is close to impossible to talk about velocities or masses without bringing in mathematical language. The physicist will have difficulty recording the facts of form (A) without quantifying over numbers. Facts of form (A) even *feel* mathematical in a sense, since we are so familiar with the type of mathematical language normally used to express such facts.

Let us consider how the physicist might try to record the facts of form (A) without quantifying over numbers. She might try to do so using numerical adjectives: “A projectile fired at 1 meter per second..., a projectile fired at 2 meters per second..., etc.”¹³⁷

¹³⁵ This example is drawn from Melia (1995), p. 226.

¹³⁶ Yablo (2005), pp. 94-5.

¹³⁷ Yablo is here supposing that numerical adjectives do not carry an ontological commitment to numbers. If one thinks that the use of numerical adjectives does carry an ontological commitment to numbers, this will only reinforce the overall point that we cannot express the facts of form (A) without relying on mathematical language.

Unfortunately, this is simply impossible to do: given that velocities range along a continuum, our physicist would have to write uncountably many sentences using uncountably many primitive numerical adjectives. What alternatives does the physicist have? As Yablo notes, the obvious alternative is to quantify over numbers. The physicist can drop the idea of using numerical adjectives and say instead: “(B) For all positive real numbers M and R , the escape velocity from a sphere of mass M and diameter $2R$ is the square root of $2GM/R$, where G is the gravitational constant.”¹³⁸

Here the physicist is making use of numbers as representational aids. The physicist has no interest in numbers, and is trying to convey something that has nothing to do with numbers – facts about the escape velocity of physical objects. Yet, mathematical language appears to be indispensable to the physicist. Without using mathematical discourse, the physicist is simply unable to report her findings.

An important feature of the utility of mathematics as a descriptive aid is that it is nominalistically acceptable; the descriptive advantages conferred by mathematical discourse don’t depend on the actual existence of numbers. The physicist is merely attempting to convey the physical facts of form (A) in a concise way. Mathematical language is useful here because of what it can be used to express about the physical world. The truth of sentence (B) requires, in addition to the existence of real numbers, that the physical world behaves in a certain way: that the facts are of form (A), rather than some alternative (A)*. Using mathematical discourse, and talking as if there are numbers, allows us to express what we want to express about the physical facts whether or not there really are numbers. As Yablo puts the point: “That (B) succeeds in gathering

¹³⁸ Ibid.

together into a single content infinitely many facts of form (A) owes nothing whatever to the real existence of numbers. It is enough that *we understand what (B) asks of the non-numerical world*, the numerical world taken momentarily for granted. How the real existence of numbers could help or hinder that understanding is difficult to imagine.”¹³⁹ Mathematical language is a powerful communicative tool, helpful for describing a great many nominalistically acceptable facts.

3.6 Reference to Abstract Objects as an Explanatory Aid

The descriptive benefit to the nominalist of making reference to abstract objects is very real. But we’ve seen that the descriptive explanation falters when we consider sentences that have as their primary subject matter the fiction itself – sentences that are about the properties of possible worlds, or about the properties of numbers. And such sentences are a central, arguably ineliminable, component of scientific practice. So the descriptive explanation cannot be the whole story.

The focus in the literature on the descriptive benefit of reference to abstract objects has sometimes led to the impression that this is the *only* explanation available to the nominalist. Mark Colyvan has objected to fictionalism¹⁴⁰ on the basis that

¹³⁹ Yablo (2005), p. 95. The italics are Yablo’s.

¹⁴⁰ Mark Colyvan does not object directly to fictionalism, but to ‘easy road’ nominalism. Easy road nominalism and fictionalism cover much of the same ground, and due to inconsistencies in the use of the term ‘fictionalism’ it is often avoided altogether. But there are differences. Fictionalism refers broadly to theories that attempt to explain the acceptability of reference to certain entities without taking on a commitment to such entities, normally by drawing an analogy to fictional or figurative language in which we do something similar. There is no requirement in fictionalism that the reference to the entities be ineliminable from our discourse. However, one of the central areas of discourse for the fictionalist to

mathematics not only plays a descriptive role in science, but also an *explanatory* role. If the nominalist can only account for the descriptive role of mathematics, then we have reason to reject nominalism in favour of platonism. To see what Colyvan has in mind when he discusses the explanatory role of mathematics, consider one of his examples. He writes:¹⁴¹

The Kirkwood gaps are localized regions in the main asteroid belt between Mars and Jupiter where there are relatively few asteroids. The explanation for the existence and location of these gaps is mathematical and involves the eigenvalues of the local region of the solar system (including Jupiter). The basic idea is that the system has certain resonances and as a consequence some orbits are unstable. Any object initially heading into such an orbit, as a result of regular close encounters with other bodies (most notably Jupiter), will be dragged off to an orbit on either side of its initial orbit. An eigenanalysis delivers a mathematical explanation of both the existence and location of these unstable orbits.

Here, mathematics is taken not to merely represent features of the Kirkwood gaps, but to *explain* them. If the fictionalist strategy is to succeed, then it must also be able to account for the explanatory role of mathematics.¹⁴²

It is important not to get the idea that mathematical entities play a role in the *causal* story about why the Kirkwood gaps exist – mathematical entities cannot play any such causal role. But, nevertheless, Colyvan thinks that mathematics is playing a genuinely explanatory role. As he notes, the (mathematics-free) causal story about why

contend with is mathematical discourse, precisely because mathematical discourse seems to be ineliminable from scientific practice.

Easy road nominalism, on the other hand, refers more specifically to the attempt to defend nominalism by conceding that mathematical discourse is ineliminable from scientific practice but arguing that the use of such discourse does not commit one to the existence of abstract mathematical discourse. As fictionalism both focuses on abstract objects and moves away from any particular analogy to fiction, and as easy road nominalism generalizes to non-mathematical discourse, the difference between them begins to blur.

¹⁴¹ Colyvan, 2010, p. 302.

¹⁴² This challenge by Colyvan has generated a number of responses on behalf of easy road nominalism. See Bueno (2012), Leng (2012), and Yablo (2012). See Colyvan (2012) for replies to all three.

the Kirkwood gaps exist will be needlessly complicated: each asteroid will have its own complicated, contingent, causal history. Even if the causal story could be completely reconstructed, there are a couple of obvious advantages to the mathematical explanation.¹⁴³ For one, the causal explanation is at the wrong level of generality and doesn't help us to understand what is going on – focusing on the details of the causal histories of particular asteroids obscures the big picture. The mathematical explanation, on the other hand, allows us to ignore irrelevant causal details. Second, the causal explanation doesn't explain why *no* asteroid can maintain a stable orbit in the Kirkwood gap; there are *modal* facts (about the possibility of asteroids in a stable orbit in the Kirkwood gap) that are not captured by giving the causal history of every actual asteroid. As a result, Colyvan concludes, the mathematical explanation is to be preferred. Given that the mathematical explanation involves mathematical entities (the eigenvalues), we have reason to posit the existence of those entities.

Colyvan is surely right about the advantages of the mathematical explanation over an explanation given in terms of the causal history of all relevant asteroids. But I contend that he is mistaken in thinking that the advantages of the mathematical explanation give us any reason to posit the existence of mathematical entities. As with the utility of mathematical discourse as a descriptive aid, the utility of mathematical discourse as an explanatory aid does not depend on the existence of mathematical entities.

To see why not, let's begin by looking at the fact that the mathematical explanation Colyvan gives is able to capture a certain modal fact about the world: no asteroid could have a stable orbit in any of the Kirkwood gaps. It is worth asking how,

¹⁴³ Colyvan (2010), pp. 302-3 and Colyvan (2012), pp. 1034-6.

exactly, mathematical objects are supposed to be helpful here. The idea cannot be that the reason why no asteroid can have a stable orbit in any of the Kirkwood gaps is *because* the Kirkwood gaps can be mathematically modeled in a certain way. Abstract entities, it bears repeating, are causally inert: they *cannot* influence the behavior of physical objects. Nomological possibilities do not hold (even in part) in virtue of how things stand with abstract entities.

Why, then, is it the case that no asteroid could have a stable orbit in any of the Kirkwood gaps? Colyvan himself gestures at the answer in the quote given above: “The basic idea is that the system has certain resonances and as a consequence some orbits are unstable. Any object initially heading into such an orbit, as a result of regular close encounters with other bodies (most notably Jupiter), will be dragged off to an orbit on either side of its initial orbit.” Surely such facts do not hold in virtue of the existence of eigenvalues, but rather in virtue of the laws of nature and the arrangement and properties of the large celestial bodies in our solar system. Mathematical discourse may help us model modal facts about the Kirkwood gaps, but this does not require the existence of mathematical entities.

What of the other supposed benefit of mathematical explanation – that it offers an explanation at the right level of generality? Colyvan focused on two explanatory options: mathematical explanation (at the right level of generality) and causal explanation focusing on the complicated and contingent causal histories of a large number of celestial bodies. But nominalistically acceptable features of the world are not so constrained: not only are the causal histories of particular asteroids nominalistically acceptable, but so are higher-level regularities and structural features of the world. The Kirkwood gaps are not

best explained by the causal histories of actual asteroids, but by a nominalistically acceptable structural feature of the world: the instability of particular orbits around Jupiter. Mathematical discourse may help us represent these structural features, and at the right level of generality, but this is no reason to posit the existence of abstract mathematical entities.

An illuminating question to ask here is how mathematics, on the platonist supposition that there are abstract mathematical entities, is supposed to be useful as an explanatory aid. It won't do to simply appeal to the truth of mathematics. There are many true claims that are of no scientific relevance. It is true that I like chocolate, but this is of no use in explaining the Kirkwood gaps. What makes math different? The applicability of mathematics to science shouldn't simply be left as a miracle. In fact I think the platonist has a good answer to give, but it is the same answer that Mary Leng provides on behalf of the nominalist: mathematical explanations "explain a phenomenon by showing it to result from the structural features of the empirical set up."¹⁴⁴ The idea is that a mathematical explanation succeeds by attributing to the world a certain structure. If this is what makes for successful mathematical explanation, it is hard to see how mathematical entities could be required. Mathematical discourse is helpful for explaining the physical world insofar as there is some appropriate similarity between the mathematical structure being described and the structure of the physical system being talked about. In this explanatory aim, it is irrelevant whether the mathematical structure being described really exists in

¹⁴⁴ Leng (2012), p. 988.

some platonic realm, or is a mere fiction concocted for its utility. Elsewhere, I call this account of the explanatory utility of mathematics the *Structural Similarity Account*.¹⁴⁵

The structural similarity account gives us a ready explanation for why the nominalist cares about the fiction of mathematics. The nominalist needn't puzzle over sentences that tell us nothing apart from what's true according to the fiction of mathematics. Such sentences are useful precisely because they tell us what's true according to the fiction of mathematics. Mathematics enables us to describe and analyze structural properties abstractly, and then to apply that knowledge when we encounter a physical system with the appropriate structure.

This explanation of the utility of mathematical discourse naturally extends to a number of other areas of discourse about abstract objects. We can think of the role of pure mathematics as modeling abstract structures. It can do so regardless of whether or not such structures are known to be present in the physical world, or are even thought to be possible structures for the physical world. But abstract models needn't be purely mathematical. Abstract models of objects such as atoms, or airplanes, or the solar system, are likewise a central component of scientific practice. And their utility can be explained in much the same way: studying abstract models is useful insofar as they bear some appropriate similarity to the physical system being modeled.¹⁴⁶ This explanation can then be given the standard fictionalist treatment; the utility of talking about abstract models is independent of the actual existence of such abstract models. The mere fiction suffices.

¹⁴⁵ See my paper "The Problem of Insufficient Concrete Ontology".

¹⁴⁶ Spelling out what this "appropriate similarity" amounts to is not a straightforward question. Abstract models are often greatly simplified compared to the complex physical systems they model. Often the simplified elements of a model are physically impossible – frictionless surfaces, mass points, or so on. For more on scientific models, a good entry point to the literature is Frigg (2012).

This fictionalist approach to abstract modeling has a theoretical flexibility that is unmatched except by the most plenitudinous platonist theories. Let us consider again the fiction of possible worlds. Talking about the structure of the space of possible worlds is a good way of describing and explaining patterns in the modal facts, provided that the fiction of possible worlds we are working with adequately reflects the modal facts. But whereas the platonist is constrained by which possible worlds exist, the fictionalist has no such constraint. For the fictionalist, the metaphysically impossible world needn't be treated any differently than the possible world. A variety of models can be explored: those with or without trans-world individuals, those with or without mereological composition, and so on. Metaphysicians are not shy about claiming that the world is necessarily some way or another. Yet it would be nice if we could reason about both options – we often don't know which theory is the correct one. And in practice, philosophers regularly do just that.

The platonist has a number of options for dealing with these apparent metaphysical impossibilities. One is to deny that we can fruitfully consider these metaphysical impossibilities. This strategy does not strike me as especially promising, since it seems to conflict with actual philosophical practice: philosophers regularly disagree about which metaphysical theories are correct and are able to theorize about mutually incompatible theories. Another option is to opt for a plenitudinous theory according to which there are 'impossible' worlds corresponding to *all* these metaphysical impossibilities. But a third option is for the platonist to endorse the fictionalist approach on which we can reason about metaphysical impossibilities without thinking that there

are any corresponding worlds. Fictionalism thus promises to be useful not just to the nominalist but to the platonist as well.¹⁴⁷

3.7 Fictionalism and Content

We've now seen two roles that reference to abstract objects can play for the nominalist. First, it can function as a descriptive aid. And second, it can function as an explanatory aid. The utility of talk about abstract objects as an explanatory aid gives us an explanation for why sentences that seem to tell us nothing about nominalistically acceptable reality except for what is true about the relevant fiction are of use to the nominalist. The overall abstract model being developed can play a useful role even if the individual statements we make when describing that model carry no useful content about nominalistically acceptable reality. It is completely acceptable, and expected, for sentences about abstract objects to tell us nothing apart from what is supposed to be true in the relevant fiction.

In light of this, let us reconsider relative reflexive fictionalism. Relative reflexive fictionalism draws a distinction between sentences in which abstract objects function solely as representational aids and sentences in which abstract objects function as things-represented. In the former case, the sentences really aren't about the abstract objects and the utility of the sentence is best explained by appealing to the content it carries about the nominalistically acceptable world. In the latter case, the sentences are, at least in part,

¹⁴⁷ Arguably, the fictionalist approach can even be used to reason about *logically* impossible worlds. By engaging in the fiction of some alternative logical framework, worlds can be considered that violate the logical constraints that govern the actual world.

about those abstract objects. Here a more holistic approach to the utility of the sentences is more appropriate. We can understand these sentences as telling us what is true according to the platonic fiction. And this fiction can in turn be used as an explanatory aid. If the explanatory utility of mathematics and abstract models is best explained by the structural similarity account – and I believe that this is the best account available to both the platonist and the nominalist – then this use of abstract objects as an explanatory aid is nominalistically acceptable.

Unfortunately for relative reflexive fictionalism there are other sentences that prove more difficult to handle. Relative reflexive fictionalism relies on drawing a distinction between sentences that tell us something about the fiction-independent world, and those that tell us about the fiction. It treats the former sentences much as object fictionalism does – as carrying content about the way the world needs to be in order to make the literal content of those utterances true according to the relevant fiction. And it treats the latter sentences much the way meta-fictionalism does – as carrying content about the fiction. But not all sentences are so neatly categorized. Consider again the modal case, and the sentence “There is a possible world in which pigs fly”. In this sentence, the possible world is a thing-represented: the sentence is telling us something about what is true in the fiction of possible worlds, namely that there exists a world in which such-and-so happens. Yet this sentence *also* clearly tells us something about the bare modal facts: that it is possible that pigs fly. The sentence appears to be doing double-duty: telling us both about the fiction and the fiction-independent world. And this is not an unusual sentence. Indeed, this is precisely the sort of sentence that modal meta-fictionalism was designed to account for.

Which of the two categories of sentences should relative reflexive fictionalism slot this sentence into? Is the content about the fiction, or about the fiction-independent world? Either way, something important is left out of the content. Or perhaps the sentence should be thought to carry both contents. The difficulty here comes from the fact that we understood relative reflexive fictionalism as a form of content fictionalism, and content fictionalism generally aims to capture what the sentence is able to tell us in the ‘content’ of that sentence.

There are a couple of important points to make here. The first is that any fully developed version of content fictionalism is going to have to take a stand on the nature of content. And our theory of content should not be overly influenced by metaphysical considerations: the content fictionalist shouldn’t adopt a particular theory of content just because it fits well with fictionalism. Content fictionalists should not be overly confident that the best theory of content will be one that fits all of the claims they make.

Consider again a sentence we discussed earlier: “(B) For all positive real numbers M and R , the escape velocity from a sphere of mass M and diameter $2R$ is the square root of $2GM/R$, where G is the gravitational constant”. This sentence was used as an example of a mathematical sentence that allows a physicist to convey the physical facts in a concise way, when no non-mathematical sentence could do the same. Now consider how Yablo phrases the point: “... (B) succeeds in gathering together into *a single content* infinitely many facts...”.¹⁴⁸ Yablo is making a stronger claim than the more general point that (B) is useful in conveying the physical facts. That (B) is useful in this more general way is clear from scientific practice – science depends on the ability of scientists to be

¹⁴⁸ Yablo (2005), p. 95. Emphasis added.

able to determine what sentences like (B) entail about the physical world. But this does not require that what (B) can be used to tell us about the physical world is the *content* of (B). And we might be suspicious of this stronger claim. After all, Yablo admits that there is no way to specify this content without referring to mathematical entities. Is it plausible for a sentence to have a content that cannot be fully specified? Could a language user really have such content in mind? My objective here isn't to argue that the answer to these questions is "No". But the answer is not obviously "Yes" – bold claims about the nature of content are being made.

The second point to make here is that nothing about the utility of referring to abstract objects as a descriptive aid relies on the idea that the facts of interest to the nominalist that can be gleaned from the sentence are part of the *content* of the sentence. This is a point I tried to emphasize in my discussion of modal meta-fictionalism. I noted that while the content of the sentences can be understood to be about the fiction of possible worlds, the facts of primary interest to the nominalist are the bare modal facts. This is not a lesson that content fictionalists often take to heart; much of the literature on content fictionalism is focused on how to pack all of the facts of interest into the content of a sentence while excluding any content about the undesirable entities. But there is nothing in the general fictionalist approach that requires the content fictionalist to do so. And if we give up on the idea that we must do so, then sentences like "There is a possible world in which pigs fly" that carry information both about the fiction and fiction-independent reality pose no real problem for relative reflexive fictionalism: nothing crucial to the fictionalist is lost if some of this information is not included in the content of the sentence.

Indeed, nothing crucial to the fictionalist is lost if *none* of this information is included in the content of the sentence. Attitude fictionalism takes advantage of this point and simply takes the simplest possible view about the content of non-literal utterances: non-literal and literal utterances of the same sentence have the same content. If we want to know the content of the sentence “The number of even prime numbers = 1”, it doesn’t matter if we are considering the literal utterances of the platonist or the non-literal utterances of the nominalist. According to attitude fictionalism, the nominalists needn’t explicitly talk about the fiction at all; it suffices to engage in the fiction and to talk directly about abstract objects. In this way, the nominalist just like the platonist learns about the supposed properties of abstract objects and can employ that knowledge as a descriptive and explanatory aid.

To drive home the adequacy of attitude fictionalism, it is fruitful to go even further and consider what the nominalist ought to say about the literal utterances made by committed platonists. Thus far we’ve been focused on the non-literal utterances made by nominalists. After all, we wanted to know what justified nominalists in uttering sentences that refer to abstract objects given that they don’t believe that any such objects exist. But suppose that the nominalist is right, and that there are no abstract objects. Then we don’t just need an explanation for why nominalists are able to nevertheless use sentences that refer to abstract objects. We also need an explanation for why everyone else, including platonists, have been able to fruitfully use such sentences. Platonists have not been making non-literal utterances about abstract objects. They have been making literal assertions about those objects, and have made considerable scientific progress while doing so.

Fictionalists have, for the most part, been too narrowly focused on how to account for non-literal utterances. But the nominalist must also account for the undeniable scientific utility of literal utterances about abstract objects. For surely, though many scientists may have no interest in ontological questions altogether, there are those scientists who believe in the existence of numbers and intend for their mathematical utterances to be taken as literally as possible.

Fortunately for the fictionalist, the descriptive and explanatory utility of referring to abstract objects does not require any particular attitude towards those entities on the part of the speaker. Nor does it require that the speaker is making non-literal utterances. A committed platonist speaking literally *still* tells us something about the physical world when she utters a sentence like “The number of Martian moons is 2” or “The average star has 2.4 planets”. Indeed, even for the platonist, the primary aim of uttering such sentences is to communicate something about the physical world. And the committed platonist still tells us something about the fiction of numbers when she utters “2 is the only even prime number”, even if she strongly believes that numbers really exist.

This being the case, we can see that any debate over the content of non-literal utterances is tangential to the fictionalist explanation of the utility of referring to abstract objects. If there is reason to prefer content fictionalism over attitude fictionalism, or vice versa, it will not come from general nominalistic considerations. The core of the fictionalist approach can be made to work either way.

If that's so, then why be a content fictionalist? What advantage is there to adopting a more complicated account of the content of non-literal utterances?¹⁴⁹ There

¹⁴⁹ It is worth noting that content fictionalism, with its more complicated account of content, is subject to difficulties that attitude fictionalism is not. In particular content fictionalism may make a mess of inference,

are a couple of potential motivations. One is that content fictionalism allows the nominalist who is making non-literal utterances to be making *true* claims. According to attitude fictionalism, the nominalist who utters “2 is the only even prime number” is uttering, but not asserting, something false about numbers (but which conveys something true about the fiction of numbers). According to (some versions of) content fictionalism, on the other hand, the nominalist who utters “2 is the only even prime number” is asserting something true about the fiction of numbers. On both views, the literal content of the sentence is false. And on both views the nominalist manages to convey something true. But only according to content fictionalism does the nominalist manage to assert something true. Is this a difference worth caring about? I confess that I don’t see why it would be. But I expect that there will be those who view this as an advantage of content fictionalism.

A second, but related, potential motivation for content fictionalism emerges if we focus in particular on descriptive content fictionalism. Considerable debate has taken place on the extent to which the mathematical claims of ordinary speakers, scientists, and mathematicians ought to be understood as non-literal utterances. For example, as I noted earlier, Yablo advances relative reflexive fictionalism in a descriptive spirit: he holds that the mathematical utterances of ordinary speakers are much more akin to figurative language than to literal assertions. An assessment of this descriptive claim is beyond the scope of this paper,¹⁵⁰ but suppose that Yablo is correct about the ubiquity of non-literal

since what is asserted differs from the literal content of the sentences uttered. Inferences that are valid given the literal content of the sentences may not be valid given the content they have according to content fictionalism. For more on this issue, see Eklund (2015) and Joyce (2005), pp. 292-3.

¹⁵⁰ The analogy between ordinary mathematical language and figurative language is far from perfect. For one, even though metaphor use often doesn’t seem to require an *explicit* intention to engage in non-literal language, on some level metaphor does seem to require an intention to speak non-literally. A speaker who

utterances in ordinary discourse about abstract objects. Then, according to (some variants of) descriptive content fictionalism, these speakers making non-literal utterances are nevertheless making true assertions. According to attitude fictionalism, on the other hand, the content of the non-literal utterances being made is strictly speaking false. The content fictionalist may well argue, however, that any good theory of the content of non-literal utterances ought to attribute true content to those utterances.

However, I hope the discussion in this paper has made it salient just how difficult it is to construct a tenable version of descriptive content fictionalism. The problem stems in particular from sentences such as “The number of even prime numbers = 1” that don’t seem to carry any content about the nominalistically acceptable world apart from the fact that such and so is true according to the fiction. There simply seems to be no suitable account of the nominalistic content of such a sentence apart from the meta-fictionalist approach of taking it to carry content about what is true in the fiction. But as we’ve seen, the meta-fictionalist approach does not seem plausible as a descriptive account of the content of fictional utterances. Consequently, it seems to me that that descriptive content fictionalism ought to be rejected.

If descriptive content fictionalism is untenable, and there is no real advantage to prescriptive content fictionalism over attitude fictionalism, then one might think that

utters the sentence “They put a lot of hurdles in your path” metaphorically clearly intends to be interpreted quite differently from someone who utters that sentence literally, even if this intention never rises to the level of conscious awareness. This is in contrast to mathematics. Normally, a speaker who utters the sentence “2 is the only even prime number” doesn’t intend to engage in any sort of figurative language. For more on this, see Burgess (2004), p. 26.

A second disanalogy between mathematical language and figurative language is that, as Jason Stanley points out (Stanley 2001, p. 13), our mathematical utterances do not even seem non-literal *on reflection*. Introspect as hard as one might, the mathematical claims that we make simply do not seem akin to figurative utterances (apart, of course, from those fictionalists who do intend to utter mathematical sentences figuratively).

That said, Yablo does point to a variety of ways in which ordinary mathematical utterances differ from ordinary literal utterances. See, in particular, Yablo (2000) and Yablo (2001).

some version of attitude fictionalism ought to be adopted. But that is not necessarily the case. The matter could be more complicated still. There is no guarantee that all non-literal utterances ought to be treated alike. And indeed, we've seen analogies to starkly different types of non-literal utterances in this paper. Take an instance of figurative language, such as the sentence "I had butterflies in my stomach". It does not strike me as unreasonable in this case to claim that the speaker asserted something true about a sensation they were feeling in their stomach when nervous. So if we are convinced by the descriptive fictionalist that ordinary utterances that involve abstract objects solely as representation aids are akin to figurative utterances, then it would be reasonable to retain a limited version of content fictionalism restricted to such sentences.

Now consider instead an instance of fiction, such as the sentence "Santa Claus lives at the North Pole." If someone engaged in the fiction utters this sentence, it doesn't seem as natural to claim that the speaker has really asserted something about the Santa Claus fiction. This is precisely the sort of conflation between engaging in a story and talking about that story that the meta-fictionalist was accused of making. In this case, attitude fictionalism seems like a more fitting approach: non-literal utterances of this sort of sentence simply retain their literal content. So, perhaps non-literal figurative utterances and non-literal fictional utterances require different treatment. And if that is so, then a hybrid of content and attitude fictionalism may be the best option available to the fictionalist. On such a view, non-literal utterances of the figurative variety would have content other than their literal content, but non-literal utterances of the fictional would retain their literal content.

It is also possible that there is considerable indeterminacy in the content of many non-literal utterances. Nothing about the fictionalist approach excludes vagueness or variability in the contents of non-literal utterances. For suppose that some variant of content fictionalism is correct about the content of mathematical utterances made by nominalists. A nominalist and a platonist could still use mathematical language to successfully communicate about the physical world, despite the difference in the contents of their utterances (in particular, with regards to abstract mathematical entities). Considerable variability in content can be tolerated while allowing for successful communication about some topic of common interest, provided that the variability in content is not contextually salient to the discussion at hand.

This flexibility allows for some rather interesting proposals concerning possible indeterminacy in the content of utterances. Matti Eklund has defended a view which he calls *Indifferentism*. According to the indifferentist, “with respect to much that we say or imply we do not commit ourselves either to its literal truth or to its truth in any fiction; we are, simply, non-committed.”¹⁵¹ So, for example, a speaker who asserts “my car is parked in the driveway” is simply not committed any way on mereological matters, even if the literal truth of the sentence requires the existence of a car rather than mereological simples arranged-car-wise. Likewise, even though speakers accept the sentence “ $2 + 2 = 4$ ” as true, they do not thereby commit themselves to the existence of numbers, nor to the claim that the sentence is literally true, given that the literal truth of the sentence requires the existence of numbers. On this view, utterances about abstract objects are not

¹⁵¹ Eklund (2005), p. 558.

especially non-literal. Rather, they are about as non-literal as typical non-figurative assertions.

Indifferentism is primarily a view about the attitudes of speakers – and is advanced in a descriptive spirit as a competitor to descriptive content fictionalism. But Eklund suggests that this approach can be paired with a disjunctive view of content, on which the content of an utterance can be identified with the disjunction of all the propositions that the speaker is neutral between.¹⁵² Given the indifference with which most speakers treat ontological questions, on this account typical mathematical utterances will have highly disjunctive contents.

Yablo briefly considers a related position on several occasions. He suggests that sentences such as “ $2 + 3 = 5$ ” could be thought of as ‘maybe-metaphors’: to be interpreted literally if so interpreted they are true, and otherwise metaphorically.¹⁵³ He compares this to sentences such as “Nixon has a stunted superego” and “I have tingles in my legs”, and writes that speakers “are far more certain that [the sentence] S is getting at something right than that the thing it is getting at is the proposition that S, as some literalist might construe it.”¹⁵⁴

My intention here is not to defend one of these accounts in particular, but rather to point to the great deal of flexibility available to the fictionalist. While a detailed account of the content of non-literal utterances is of considerable interest, and has been a central topic in the literature on fictionalism, the core fictionalist strategy is independent of how these details are worked out. Whatever the case may be, it is clear enough that speakers

¹⁵² Eklund (2005), p. 272.

¹⁵³ Yablo (2005), p. 111.

¹⁵⁴ Yablo (2000), p. 300.

do use sentences that refer to abstract objects in order to convey information about the physical world. And that is enough to guarantee the utility of reference to abstract objects even if such objects do not exist.

3.8 Fictionalism and Truth

Let us assume that the general fictionalist strategy that I've been considering in this paper can be made to work, and that the nominalist has an explanation of the utility of speaking as if there are abstract objects. Then the widespread use (and utility) of language that refers to abstract objects does not give us reason to posit the existence of abstract objects: we would expect to find the same linguistic behavior regardless of whether or not abstract objects exist.

Is there any further reason for preferring platonism over nominalism? Here is one that many platonists find compelling:¹⁵⁵ the existence of abstract objects is required for our scientific theories to be *true*. After all, the thought goes, we should accept our best scientific theories as literally true. And if those theories include sentences that carry an ontological commitment to abstract objects, then we should accept that commitment. From this perspective, the fictionalist strategy simply misses the point: it doesn't matter for what reason we refer to abstract objects, it just matters that we do.

In reply, the fictionalist must *deny* that we should accept our best scientific theories as literally true – those portions of the theories that carry unwanted ontological

¹⁵⁵ I do not mean to suggest that this is the only other reason a platonist might reject nominalism. I cannot hope to cover every possible argument that the platonist might bring to bear. But this objection is a common refrain among platonists, so it seems especially important to consider it here.

commitments must be taken to be literally false.¹⁵⁶ This feature of fictionalism has given rise to the charge that fictionalism, at least with regards to entities (such as numbers) that play an important role in scientific theories, is *comically immodest*.¹⁵⁷

This objection has its genesis in David Lewis' remark in *Parts of Classes* that he would be "moved to laughter at the thought of how presumptuous it would be to reject mathematics for philosophical reasons", even if we reject it 'gently' by "explaining how it can be a most useful fiction, 'good without being true'".¹⁵⁸ Burgess elaborates on this remark by suggesting that once we focus on the comparative historical track records of philosophers and scientists, it is clear that we ought to defer to the scientists; science has a proven record of progress and success, whereas philosophy has not fared quite so well.¹⁵⁹

In one way, Burgess undersells the strength of his position. As Chris Daly notes, science does not need to have a strong track record of making true mathematical claims to pose a problem for the fictionalist about numbers. It suffices for science to have established *one* true claim about numbers. That alone would be enough to commit us to the existence of numbers.¹⁶⁰ The question is whether or not we have reason to believe that our best scientific theories include one true claim that carries a commitment to numbers.

¹⁵⁶ See, for example, Yablo (2005), p. 98, and Balaguer (2009), p. 134.

¹⁵⁷ It's often assumed that descriptive content fictionalism is immune to this objection. I fail to see why. After all, the descriptive content fictionalist does hold that our best scientific theories are literally false. The distinctive claim of the descriptive content fictionalist is that scientists often utter their theories non-literally, and that the content of those non-literal utterances are true. However, I fail to see how it is any less 'comically immodest' to say that scientific theories are literally false but that scientists are speaking non-literally (and truly), than it is to say that scientists are speaking literally but falsely.

¹⁵⁸ Lewis (1991), pp. 58-59.

¹⁵⁹ Burgess (2004), p. 30.

¹⁶⁰ Daly (2006), pp. 210-11.

Burgess insists that we do – that science has a track record of success – but this is simply to beg the question against the fictionalist. The fictionalist will agree that science has a proven track record of *practical* success, and that scientists and mathematicians are experts at determining what follows from certain postulates or axioms. But nothing about this success is evidence in favour of thinking that sentences that carry a commitment to mathematical entities are true.

It is true that philosophers should defer to mathematicians about what is true according to mathematics. But the relevant question, Daly points out, is “whether mathematicians can reliably establish that $2+2=4$ rather than (crudely) *that it’s not the case that $2+2=4$ but science works as if $2+2=4$.*”¹⁶¹ Nothing in the track record of mathematics or science gives us any reason to think that mathematicians are experts with regards to that question.

Indeed, the ontological status of numbers is entirely mathematically unimportant and uninteresting. Accordingly, fictionalism is an issue that most mathematicians do not think or care about, and hence have no real expertise in. Moreover, there is nothing about mathematical or scientific practice that provides any evidence for thinking that mathematical objects exist, or that mathematical theories are literally true as opposed to true in the mathematical fiction.¹⁶²

There is, in fact, no particular reason for the scientist or mathematician to care if their theories are literally true. The physicist studying escape velocity is interested in constructing a theory that paints an accurate picture of the physical world. Although she uses mathematical discourse in doing so, the existence or non-existence of numbers is

¹⁶¹ Daly (2006), p. 213.

¹⁶² For more on this point, see Leng (2005), pp. 291-2, and Balaguer (2009), p. 155.

entirely irrelevant for her purposes. Her sole concern in using mathematical discourse is to correctly convey information about the behavior of physical objects in motion. So, if the literal truth of her mathematical claims requires the existence of numbers – as the platonist and the fictionalist agree they do – then the literal truth of her mathematical claims are likewise entirely irrelevant for her purposes.

This is so even though mathematical claims form an indispensable part of her theory of escape velocity. She would be just as well served by taking a fictionalist stance towards her mathematical claims as she would by taking those claims to be literally true. From a scientific viewpoint, the difference simply does not matter. If the truth or falsity of the mathematical claims that are part of our scientific theories make no difference to the empirical success of those theories, then we simply have no grounds on which to infer that those theories are true. If the ontological status of numbers is as irrelevant to the scientific enterprise as it seems to be, then denying the literal truth of the mathematical discourse that is part of our best scientific theories is scientifically respectable. It would seem to me, then, that a fixation on literal truth is misguided, and is not supported by scientific practice.

Nevertheless, I suspect that many philosophers are inclined to think that the attribution of error, especially of widespread error, is theoretically costly – a theory that avoids an attribution of error is superior, all else being equal. So let me say a little more about why I don't believe that the 'comical immodesty' of fictionalism is grounds for rejecting nominalism.

The concept of truth plays (at least) two important roles. First, it tracks whether what is said corresponds to what is in fact the case.¹⁶³ And second, it plays a normative role – distinguishing the ‘good’ sentences that we ought to utter from the ‘bad’ sentences that we ought to avoid. But if the fictionalist is correct, then we have good reason to utter sentences that do not correspond fully to how things are – even when our goal is to describe the world as truthfully as possible. In particular, for the nominalist, we have reason to utter sentences that get things right about the physical world but that get things wrong about abstract objects. So these two roles of our concept of truth can come apart. It is unfortunate for the fictionalist that the normative force of the concept of truth is so hard to shake.

And even if it is difficult to shake, it nevertheless does not give us good reason to believe in the existence of abstract objects. I’ve taken it as common ground in the debate between the fictionalist and the platonist that ontological questions are not ‘easy’ – that the ontological status of abstract objects is a substantive issue. But if that’s the case, it is inappropriate to posit the existence of abstract objects for purely semantic reasons. The ontological status of mind-independent entities is not determined by what people believe or by what they utter. It would be bizarre indeed to argue for the existence of God on the basis that the majority of people are religious and we are thus able to maximize the number of utterances about God that are true. This style of argument strikes me as no

¹⁶³ Here I’m assuming a correspondence theory of truth. Deflationists about truth will not be happy with this characterization, but I cannot fully engage in that debate here. I will just say the following. Insofar as one of the aims of language is to allow us to communicate about what the world is like, we must have some notion that tracks the connection between what is said and what the world is like. I take this to be one of the roles of the concept of ‘truth’. Those who think otherwise are free to reframe my arguments here.

more plausible when it comes to abstract objects.¹⁶⁴ We simply cannot deduce the existence of entities from the fact that doing so leads to more utterances being true.

If ontological questions are not easy, and abstract objects don't exist, then sentences that carry an ontological commitment to abstract objects are false. The fictionalist accepts that sentences that refer to abstract objects uttered literally carry an ontological commitment to abstract objects, and hence must take those sentences, understood literally, to be false. Can the nominalist nevertheless avoid an attribution of widespread error to ordinary speakers who refer to abstract objects? Yes, in one of two ways. First, the fictionalist could take the irrelevance of the ontological status of abstract objects to most speakers as evidence that those speakers are not speaking literally. Or second, the nominalist could hold that the fictionalist made a mistake in thinking that the literal content of such utterances carries an ontological commitment to abstract objects; perhaps it is the platonist, who means to take on an ontological commitment to abstract objects when making an assertion, that departs from the literal content of the sentence.

I don't want to put too much stress on these points, except to note than once again there is a great deal of flexibility available to the nominalist. The utility of sentences that refer to abstract objects can be explained in much the same way regardless of how these issues are settled. Moreover, as was the case in the debate between the attitude fictionalist and the content fictionalist, there is nothing here preventing the notion of content from being vague or indeterminate. I'm fairly sympathetic to this last possibility, but I won't

¹⁶⁴ Matters would be different if we had good reason to suppose that language users were reliable at tracking the existence of abstract objects. In that case, there would be a significant cost to the claim that people were uttering falsehoods. But it wouldn't merely be because accepting the existence of the entities made more utterances true. It would be because if speakers reliably track whether a type of entity exists, and claim that that type of entity exists, then we have evidence for the existence of that type of entity. However, humans cannot reliably track whether abstract objects exist. Indeed, the vast majority of speakers are altogether insensitive to, or unconcerned with, the ontological status of abstract objects. And so no similar argument is available.

argue that point here. I merely wish to emphasize that it is not a difficulty for the nominalist if that turns out to be the case.

BIBLIOGRAPHY

- Armour-Garb, B. and Woodbridge, J. (2012). "The Story About Propositions", *Nous*, vol. 46:4, pp. 635-674.
- Armstrong, D. (1978). *Universals and Scientific Realism, Volume II: A Theory of Universals*, Cambridge University Press.
- Armstrong, D. (1988). "Are Quantities Relations? A Reply to Bigelow and Pargetter", *Philosophical Studies* 54, pp. 305-316.
- Armstrong, D. (1989). *Universals: An Opinionated Introduction*, Colorado: Westview Press.
- Armstrong, D. (2010). *Sketch for a Systematic Metaphysics*, Oxford University Press.
- Azzouni, J. (2004). *Deflating Existential Consequence: A Case for Nominalism*, Oxford: Oxford University Press.
- Balaguer, M. (1998). "Attitudes without Propositions", *Philosophy and Phenomenological Research*, vol. 58:4, pp. 805-826.
- Balaguer, M. (2009). "Fictionalism, Theft, and the Story of Mathematics", *Philosophia Mathematica III*, vol. 17, pp. 131-162.
- Balaguer, M. (2010). "Fictionalism, Mathematical Facts and Logical/Modal Facts", in J. Woods (ed.) *Fictions and Models*, Philosophia.
- Balaguer, M. (2011). "Fictionalism in the Philosophy of Math", *Stanford Encyclopedia of Philosophy*. First Published 2008.
- Bigelow, J. and Pargetter, R. (1988). "Quantities", *Philosophical Studies*, vol. 54, pp. 287-304.
- Bricker, P. (2015). "Composition as a *Kind* of Identity", *Inquiry*, vol. 59.3, pp. 264-294.
- Bricker, P. (2017). "Is There a Humean Account of Quantities?", *Philosophical Issues*, vol. 27, pp. 26-51.
- Brock, S. (1993). "Modal Fictionalism: A Response to Rosen", *Mind*, vol. 102, pp. 147-150.
- Bueno, O. (2012). "An Easy Road to Nominalism", *Mind*, vol. 121:484, pp. 967-982.

- Burgess, J. (2004). "Mathematics and Bleak House", *Philosophia Mathematica* III, vol. 12, pp. 18-36.
- Colyvan, M. (2010). "There is No Easy Road to Nominalism", *Mind*, vol. 119:474, pp. 285-306.
- Colyvan, M. (2012). "Road Work Ahead: Heavy Machinery on the Easy Road", *Mind*, vol. 121:484, pp. 1031-1046.
- Crimmins, M. (1998). "Hesperus and Phosphorus: Sense, Pretense, and Reference", *The Philosophical Review*, vol. 107:1, pp. 1-47.
- Daly, C. (2006). "Mathematical Fictionalism: No Comedy of Errors", *Analysis*, vol. 66:3, pp. 208-216.
- Eddon, M. (2007) "Armstrong on Quantities and Resemblance", *Philosophical Studies*, vol. 137, pp. 385-404.
- Eddon, M. (2013). "Fundamental Properties of Fundamental Properties", *Oxford Studies in Metaphysics*, vol. 8, pp. 78-104.
- Eklund, M. (2005). "Fiction, Indifference, and Ontology", *Philosophy and Phenomenological Research*, vol. 71:3, pp. 557-579.
- Eklund, M. (2015). "Fictionalism", *Stanford Encyclopedia of Philosophy*. First published 2007.
- Ellis, B. (1966). *Basic Concepts of Measurement*, Cambridge University Press.
- Frigg, R. (2012) "Models in Science", *Stanford Encyclopedia of Philosophy*. First published 2006.
- Field, H. (1980). *Science Without Numbers*, Princeton, NJ: Princeton University Press.
- Joyce, R. (2005). "Moral Fictionalism", in M. Kalderon (ed.) *Fictionalism in Metaphysics*, Oxford: Clarendon Press.
- Leng, M. (2005). "Revolutionary Fictionalism: A Call to Arms", *Philosophia Mathematica* III, vol. 13, pp. 277-293.
- Leng, M. (2010). *Mathematics and Reality*, Oxford: Oxford University Press.
- Leng, M. (2012). "Taking it Easy: A Response to Colyvan", *Mind*, vol. 121:484, pp. 983-995.

- Lewis, D. (1975) "Languages and Language", *Minnesota Studies in the Philosophy of Science*, University of Minnesota Press, pp. 3-35.
- Lewis, D. (1986a). "Against Structural Universals", *Australasian Journal of Philosophy*, 64, pp. 25-46
- Lewis, D. (1986b). *On The Plurality of Worlds*, Oxford: Blackwell.
- Lewis, D. (1991). *Parts of Classes*, Oxford: Basil Blackwell.
- Liggins, D. (2012). "Weaseling and the Content of Science", *Mind*, vol. 121:484, pp. 997-1005.
- Melia, J. (1995). "On What There's Not", *Analysis*, vol. 55:4, pp. 223-229.
- Melia, J. (1998). "Field's Programme: Some Interference", *Analysis* vol. 58:2, pp. 63-71.
- Melia, J. (2000). "Weaseling Away the Indispensability Argument", *Mind*, vol. 109:435, pp. 455-479.
- Monton, B. and Mohler, C. (2017). "Constructive Empiricism", *Stanford Encyclopedia of Philosophy*. First published 2008.
- Mundy, B. (1987). "The Metaphysics of Quantity", *Philosophical Studies*, vol. 51:1, pp. 29-54.
- Mundy, B. (1989). "Elementary Categorical Logic, Predicates of Variable Degree, and Theory of Quantity", *Journal of Philosophical Logic*, vol. 18:2, pp. 115-140.
- Nolan, D. (2016). "Modal Fictionalism", *Stanford Encyclopedia of Philosophy*. First published 2002.
- Nolan, D. and O'Leary-Hawthorne J. (1996). "Reflexive Fictionalisms", *Analysis*, vol. 56:1, pp. 23-32.
- Pérez Carballo, A. (2014). "Structuring Logical Space", *Philosophy and Phenomenological Research*, vol. 17:2, pp. 460-491.
- Putnam, H. (1975). *Mathematics, Matter and Method*, Cambridge: Cambridge University Press.
- Quine, W.V. (1948). "On What There Is", *The Review of Metaphysics*, vol. 2:1, pp. 21-38.
- Rosen, G. (1990). "Modal Fictionalism", *Mind*, vol. 99:395, pp. 327-354.

- Rosen, G. (1993). "A Problem for Fictionalism About Possible Worlds", *Analysis*, vol. 53:2, pp. 71-81.
- Rosen, G. (2017). "Abstract Objects", *Stanford Encyclopedia of Philosophy*. First published 2001.
- Stanley, J. (2001). "Hermeneutic Fictionalism", in P. French and H. Wettstein (eds.) *Midwest Studies in Philosophy Volume XXV: Figurative Language*, Oxford: Blackwell.
- Sidelle, A. (1989). *Necessity, Essence, and Individuation: A Defense of Conventionalism*, Ithaca: Cornell University Press.
- Van Fraassen, B. (1980). *The Scientific Image*, Oxford: Oxford University Press.
- Yablo, S. (1998). "Does Ontology Rest on a Mistake?" *Aristotelian Society, Supplementary volume*, vol. 72, pp 229-61.
- Yablo, S. (2000). "A Paradox of Existence", in A. Everett and T. Hofweber (eds.) *Empty Names, Fiction and the Puzzle of Non-Existence*, Stanford: CSLI Publications.
- Yablo, S. (2001). "Go Figure: A Path through Fictionalism", in P. French and H. Wettstein (eds.) *Midwest Studies in Philosophy Volume XXV: Figurative Language*, Oxford: Blackwell.
- Yablo, S. (2005). "The Myth of the Seven", in M. Kalderon (ed.) *Fictionalism in Metaphysics*, Oxford: Clarendon Press.
- Yablo, S. (2012). "Explanation, Extrapolation, and Existence", *Mind*, vol. 121:484, pp. 1007-1029.