Ontological Nihilism*

Jason Turner

March 17, 2009

Ontological Nihilism is the radical-sounding thesis that there is nothing at all. Almost everyone believes that it is false. But this does not make it philosophically uninteresting: we can come to better understand a proposition by studying its opposite. By better understanding what Ontological Nihilism is — and what problems beset it — we can better understand just what we say when we say that there is something in the world.

This paper explores Ontological Nihilism, and for just this reason. After discussing what the thesis would amount to if it were to have any plausibility whatsoever (section 1), I present (section 2) and clarify (section 3) a crucial challenge for it. I show what is wrong with two less plausible attempts to meet this challenge (section 4), and then argue that the proposal thought by many to be much more promising (section 5) succumbs to the same problems as the less promising attempts (sections 6 and 7). It turns out that, in order to make Ontological Nihilism viable without smuggling in an illicit appeal to things, we need a bloated ideology, indefinitely brute, necessary connections, and a deep-seated holism about the structure of reality. Once we understand why these are costs of the view, we see just what we gain by thinking of the world as built up out of things.

1 Ontological Nihilism

Ontology, Quine tells us, asks what there is; and while this ontological question can be answered in a word — 'everything' — there is still room for disagreement about cases. (1948: 1) When we encounter this case-by-case disagreement, we occasionally come across views that can best be described as versions of ontological nihilism. *Compositional nihilists*, for instance, hold that there are no composite objects: nothing has parts. So-called nominalists (of the good, old-fashioned 'nothing is abstract' type) could just as well be called *abstractional nihilists*: they claim that there are no abstract objects. *Perforational nihilists* are those who, like the Lewis' (1970) Argle, say that there are no holes. And so on.

These run-of-the-mill ontological nihilists do something that every good metaphysician wants to do at one time or another — deny that there is anything of such-and-such a kind. But another kind of ontological nihilist goes further,

*Thanks to Ross Cameron, Cian Dorr, John Hawthorne, Hud Hudson, Starr Turner, Ted Sider, Robbie Williams and Dean Zimmerman for helpful comments and discussion.

denying that there is *anything at all*. He answers Quine's ontological question not with 'everything', but with 'nothing'. He is not just an ontological nihilist, but an *Ontological Nihilist*, complete with capital letters.

When an Ontological Nihilist says that there isn't anything at all, we might naturally take him as committing himself to the following claims:

(1) Our ordinary beliefs — such as that some electrons are attracted to some protons or that there are buildings in Portugal — are radically mistaken.

(2) Reality is a blank void — an unstructured and undifferentiated blob, but without the blob.

Each of these is incredibly hard to believe. It seems undeniable that our experiences are richly structured and differentiated, and that the structure of our experiences will somehow be accounted for by structure in the world. And it seems reasonable that our ordinary beliefs, formed as they are on the basis of our richly-structured experiences, will thus track this worldly structure. If the Nihilist[1] endorses (1), he rejects the reasonable. And if he endorses (2), he denies the undeniable.

The Nihilist need not be quite as crazy as all that, though — he can hold that there isn't anything at all without endorsing either of (1) or (2). He can agree that our experience exhibits structure, and that the organization of reality accounts for this structure. And he can think that this structure connects up in important ways with our ordinary beliefs, since these beliefs are formed in large part by our interactions with this structure. What he insists is that this structure will not involve any *things*, any *entities* — any *ontology*.

Let me explain. At the simplest level, we describe the world by combining two different types of expressions.[2] We take noun phrases — paradigmatically, proper names such as 'Bertrand' and 'Gottlob', and quantifier phrases such as 'every philosopher' or 'some logicist' — and combine them with predicates — such as 'thought about language' or 'didn't notice the inconsistency in Basic Law V'. The noun phrases latch on to some things, and predicate phrases then describe these things and differentiate them from one another.

By describing the world in this way, we implicitly suppose that it has a certain sort of structure — an *ontological* structure. Ontological structure is the sort of structure we could adequately represent with a pegboard and some rubber

[1] Ontological Nihilist, that is. For stylistic reasons, I'll sometimes drop the 'Ontological', letting the capital 'N' do the disambiguating work.

[2] At least, we do when using the languages with which I have any familiarity; perhaps some languages do not at bottom operate this way. If so, it would be interesting to see what kind of metaphysics native speakers of these languages produce.

bands. The pegs would represent the things, and the rubber bands would represent the ways these things are and are interrelated.

To say 'Bertrand thought about language', for instance, is essentially to hang the *thought about language* rubber band on the peg labeled 'Bertrand'. And to say 'Some logicist admired every philosopher who didn't notice the inconsistency in Basic Law V' is to say that, somewhere on the pegboard, there is a peg which (a) has a *logicist* rubber band hanging on it, and (b) if you take any peg which has the *didn't notice the inconsistency in Basic Law V* band on, there's will be an *admires* rubber band stretching between those two pegs.[3]

Pegboard-and-rubber-band structure — *ontological* structure — is one kind of structure. The Ontological Nihilist denies that reality exhibits *this* kind of structure. But he needn't thereby claim that reality is completely unstructured. He grants that the rich texture of our experiences is grounded in some external facts about how the world is and isn't. He simply denies that these facts have to do with any *ontology*. According to him, however the universe does its thing, it *doesn't* do it by hanging rubber bands on pegs.

The pegboard model — the ontological model — of structure is fairly natural and well-understood. We know what reality would be like if it were structured that way. On the other hand, we don't come pre-equipped with any other way of thinking; simply saying that reality isn't like a pegboard leaves us with no clue of how it might be instead. So the Ontological Nihlist owes us a story: a story about the kind of structure reality *does* have, and how this structure manages to account for the richness and variety of our experiences.

2 The Need for Paraphrase

2.1 *The Challenge*

The Nihilist, we imagine, denies each of (1) and (2). So, in lieu of (2), he needs to tell us what structure reality *does* have, if not pegboard-and-rubber-band-like. And in lieu of (1), he must tell us how this structure hooks up to our ordinary beliefs and practices.

Let's compare this challenge to a similar challenge for a more conservative sort of nihilist: the perforational nihilist, who claims that there are no holes.

[3]This pegboard-and-rubber-band image is helpful, but imperfect. In particular, it leaves little room for non-symmetric predicates (such as 'loves') or predicates with a fixed adicy: rubber bands do not have a direction, and can be hung on as many or as few pegs as its elasticity will allow. Nonetheless, the image has its uses, and for our purposes here we can manage this model without these technicalities getting in the way.

Figure 1: A Defective Bridge

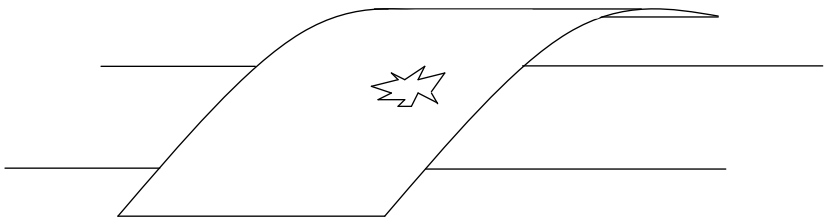

At first blush, the perforational nihilist's claim may seem incredible, implying that we are radically deceived about the nature of the world. Suppose, for instance, that you just crossed a bridge like the one in figure 1. If someone asks you why you crossed on the left, you will probably say

(3) There is a hole in the right-hand side of the bridge,

and point out that you were not keen on dropping through to the river below.

The perforational nihilist insists that there are no holes; since (3) seems to entail that there *are* holes, perforational nihilists should reject (3). So it seems they must say that you were radically mistaken about the nature of the bridge — and that crossing on the right-hand side would have been fine.

Of course, perforational nihilists want to say neither thing. They are happy to grant that there are indeed bridges shaped like the one in the diagram, and that walking on the right-hand-sides of such bridges is a bad idea. And they will say that there is *something* right about your utterance of (3): even though there are no holes, there is some important fact, relevant to bridge-crossing activities, that you were getting at with (3). This fact adequately explains your reluctance to cross on the right. The perforational nihilist's complaint is only with the idea that this important fact involves a special kind of entity called a 'hole'. They do not think that crossing on the right-hand side of the bridge is bad because among reality's pegs there is one with a 'hole' rubber-band on it, and the 'in' rubber-band was stretched between it and the side of the bridge. Whatever fact you were getting at with (3), it didn't involve a special class of hole-ey entities in this way.

Perforational nihilists can convince us, by saying all of this, that they do not

think we are radically mistaken about the nature of certain precarious bridges and the like. But they will have told us nothing about how the world *is* in virtue of which (3) is a good thing to say in the circumstances. If crossing on the right-hand side of the bridge isn't a bad idea thanks to its being related to some separate entity, some hole, in its right-hand side, then why *is* it a bad idea?

The perforational nihilist could refuse to answer this question. If he did, he would endorse a certain *negative* metaphysical thesis: the appropriateness of saying (3) in the circumstances isn't thanks to an entity rightly called a 'hole'. But he then would provide us with no *positive* metaphysical thesis about how the world is structured, perforation-wise; he would say nothing about how to fill the gap that we would otherwise fill with holes.

There are two reasons perforational nihilists should go further. First: doubters may worry that if there were no entities deserving to be called 'holes', the world just wouldn't *have* enough structure to guarantee that (3) is a good thing to say in the envisaged circumstances. Perforational nihilists can assuage these doubts by giving a positive account of the world's perforation-relevant structure that provides this guarantee.[4]

Second, and to my mind, more important: if we stop with a negative thesis, we do only half the job of metaphysical inquiry. Metaphysics asks what the fundamental structure of the world is and how this structure accounts for the richness and variety of experience. To simply tell us what the world *isn't* like is not yet to tell us what the world *is* like. A complete metaphysical picture will tell us what the world is like, and if it is indeed not a blatant error to appeal to (3) when explaining how we cross bridges like the one in the diagram, a complete metaphysical picture will tell us why.

2.2 *How to Respond To the Challenge*

The perforational nihilist thinks we get at some important fact about the world when we assert (3) in the presence of bridges like the one in figure 1. But the perforational nihilist *also* says that, despite (3)'s usefulness in this regard, there is nonetheless something defective about it. It misrepresents the real metaphysical facts of the matter as involving a *hole*, and they don't. And we, in response to their denial, want to know what the real, hole-free metaphysical facts of the matter *are* that make (3) useful but nonetheless defective.[5]

[4]Cf. Sider (forthcoming: §2).

[5]I am being deliberately cagey about just what this 'metaphysical defectiveness' amounts to. It may be that (3) is simply *false*, but can be used to convey true information in the neighborhood; see Merricks (2001) for this sort of view about table-and-chair talk. Or perhaps (3) says something true in ordinary contexts but says something else, which is false, in 'serious' philosophical

The perforational nihilist answers our question in the simplest way by telling us what useful but hole-free fact (3) is getting at. For instance, the perforational nihilist might think that, although there are no holes, certain physical objects have a special shape property, that of being *perforated*. Furthermore, he claims,

(4) The right-hand side of the bridge is perforated,

is true. And he will say that (4) is the true and metaphysically perspicuous fact we have been getting at with (3) all along.

If (3) is the only useful hole-involving sentence we ever say, this will be enough. But it is not; we communicate many other important facts by talking about holes. So we need more than just this particular, one-off explanation — we need an account of how hole-talk communicates important facts generally.

A perforational nihilist can give us this account is by providing a *paraphrase scheme*: a systematic recipe for taking claims about holes and specifying the important hole-free facts we communicate with those claims. For instance, they may decide to trade in apparent talk of holes and the objects they are in for talk about which objects are perforated. Then, whenever we would say

There is a hole in _____,

the perforational nihilist will tell us the important fact we are communicating is

_____ is perforated.[6]

The perforational nihilist thus tells us what he thinks this hole-free world *is* like — he thinks it is filled with things with certain perforated shapes — and how apparent talk of holes is really getting at these perforation facts.

The term 'paraphrase scheme', may call to mind a certain philosophy of language according to which (4) cannot in any sense count as an analysis of, or be synonymous with, ordinary uses of (3) (see, e.g., Quine 1960*b*: 250). And we may thereby implicitly suggest that the proposed scheme must meet certain conditions: that the proposed paraphrases must be finitely specifiable, for instance,

---

contexts; see van Inwagen (1990) for this sort of view about table-and-chair talk. The differences between these views can be set aside for our purposes.

[6] More generally, whenever we would say

There are *n* holes in _____,

the perforational nihilist can say

_____ is *n*-perforated.

See Lewis and Lewis (1970) for a fuller treatment of this sort of paraphrase scheme, and for some of the troubles it encounters.

or that anyone who understands the claim to be paraphrased away must be able to understand the claim it is paraphrased into.

But let's not foist any particular philosophy of language on the nihilist or bind him to its peculiar commitments. We demand merely that the perforational nihilist tell us, for any claim involving holes that he takes to be getting at some important fact, exactly what important, hole-free fact the claim he thinks it is getting at. We do not insist that this hole-free fact be finite, easily recognizable by anyone capable of talking about holes, etc.[7]

Given our liberality about paraphrase schemes, what should we say about a proposal's *systematicity*? Must similar hole-sentences receive similar paraphrases? The proposal above is relatively systematic, but how poorly should we view a nihilist who offers a more gerrymandered scheme, paraphrasing some sentences of the form

There is a hole in _____,

in one way, and paraphrasing others in another?

We should not automatically dismiss a gerrymandered paraphrase scheme. 'There is a hole in the bridge' says something very different about the bridge than 'There is a hole in the argument' does about the argument, so we shouldn't demand that the perforational nihilist paraphrase these in the same way. But, insofar as the nihilist thinks that various claims about holes *are* getting at similar facts, they ought to paraphrase them in similar ways. And insofar as *we* think that various claims about holes are getting at similar facts, we ought to take any paraphrase strategy that paraphrases them in different ways as accusing us of some sort of mistake. Nobody should worry if perforational nihilists paraphrase 'There is a hole in the bridge' and 'There is a hole in the argument' differently, since nobody thought these facts were similar in the first place. But we do think that there is a certain sort of similarity between a bridge's having a hole and a door's having a hole; if a perforational nihilist paraphrases 'There is a hole in the bridge' and 'There is a hole in the door' in radically different ways, he thereby denies that these claims are getting at similar facts after all. And, the more convinced we are of these facts' similarity, the more work the nihilist must do to convict us of error in this.[8]

[7]Likewise, we need not insist that 'paraphrases' avoid semantic ascent; anyone who thinks there are no ultimately egocentric facts — no facts that must be stated using terms like 'I' or 'you', for instance — may fairly take Kaplan's (1989) semantics for indexicals as providing a 'paraphrase', in our sense, of tokens of sentences of the form ⌜I am F⌝ even though, as Kaplan argues, there is no way to provide the account as a translation from sentences to sentences all in the 'material mode'.

[8]There are subtle issues to be sensitive to, though. Perhaps the perforational nihilist uses

2.3 *Paraphrase and Ontological Nihilism*

Just as the perforational nihilist does not want to deny that (3) gets at some important fact, the Ontological Nihilist does not want to deny that claims such as

(5) There are buildings in Portugal,
Some people have several shirts,
There are more marshmallows in my hot chocolate than in yours,

and so on are also getting at important facts.[9] But since the Ontological Nihilist denies that there is anything at all — and hence denies that there are buildings, people, shirts, or marshmallows — he must think the sentences in (5) are somehow misleading. He needs to tell us what this building-, people-, shirt-, and marshmallow-free world is like, and why its being this way makes the sentences in (5) worth saying. So he, like the perforational nihilist, needs a paraphrase scheme: a method for taking ontological, pegboard-and-rubber-band presupposing claims and trading them in for 'ontologically innocent' facts — facts which do not entail that there is anything.

This will be a complex and difficult business for the Nihilist. We can simplify it a bit by pretending the *target language* — the language he is going to be 'paraphrasing away' — is the relatively simple language of first-order logic without names (but with identity). This language is generally thought sufficient for talking about ontological structure: its existential quantifier, '$\exists$', means *there is*, and it can form all sorts of sentences that talk about what there is, what there isn't, and how things are interrelated.

We can also help the Nihilist by making him paraphrase only a portion of our ontology-involving talk. In particular, we make him paraphrase only claims from well-established scientific theory (or, at least, simple first-order consequences of well-established scientific theory). By making him do this, we makes his task both easier and harder.

Easier, because it takes from the Nihilist's shoulders the burden of deciding which sentences deserve paraphrase. A Nihilist ought not paraphrase *everything*

one recipe to paraphrase 'There is a hole in the bridge' as P and uses a very different recipe to paraprhase 'There is a hole in the door' as Q. Nonetheless, if P and Q are themselves clearly very similar facts, then the differences in the formulas used to get to them from the original hole-sentences do not mean the nihilist is denying any intuitive similarity.

[9]Of course, insofar as he is a Nihilist, he wants to deny that there are any facts at all. But the 'fact'-talk he gives us should be thought of as merely a useful turn of phrase for trying to explain his view to us doubters. He will talk about facts only while trying to get us into the spirit of his view; once we are fully converted to Nihilism, he promises to show us how to understand what he was saying without any 'fact'-talk at all. Similarly for his talk about 'the world', 'structure', 'sentences', and so on.

we say: some of what we say just isn't getting at any important fact. (Nihilists need not give us a paraphrase for 'Phlogiston is emitted during combustion', for instance.) But Nihilism is plausible only if it can recover at least the claims of our (incredibly fruitful) best science — surely if any claims ever get at important facts, these do.

Harder, because by focusing on these sentences gives us the right to demand the Nihilist paraphrase systematically. Even if

(6) An electron orbits a proton, and

(7) Two electrons orbit a proton,

are metaphysically misleading, they clearly get at very similar facts. But if similar sentences are getting at similar facts, then we should expect them to be paraphrased in similar ways.

With this in mind, we demand the following of our Ontological Nihilist: give us a systematic recipe for taking any sentence of a first-order language (with predicates assumed to be predicates of our best science) and cooking up the ontologically innocent claim it was supposed to be getting at all along.

3 Ontological Guilt: An Aside

If the proposed paraphrase scheme is to be acceptable, it must be 'ontologically innocent'. But just what does that mean? And what is it about an expression that makes it ontologically innocent?

3.1 Ontological Commitment

Some (interpreted) sentences have a feature philosophers are pleased to call 'ontological commitment'. A sentence is 'ontologically innocent' if and only if it carries no ontological commitments. Unfortunately, though, this term tends to get used more often than it gets defined, and I fear as a result it tends to be heard more often than understood. I do not intend to spill any more ink over the proper 'criterion of ontological commitment',[10] but I do want to be clear about just what 'ontological committment' is supposed to be.

The core idea is that we somehow manage to convey, semantically, by our linguistic activity, that the world has a certain ontological structure. In particular, we convey that there are some things of a certain kind K — that there are some

[10] Although see Cartwright (1954), Rayo (2007) Richard (1998), and van Inwagen (1998: Thesis V) for discussion.

pegs with the '*K*' rubber-band hanging from them. When someone performs the right sort of linguistic activity, we say that the individual is *ontologically committed to Ks*.

The 'right sort' of linguistic activity is sincere assertion of the right sentences, properly understood.[11] But the sentence has to be the right one — I cannot commit myself to unicorns just with any old sentence. I have to say 'There are unicorns' or something like that. That is, I have to use a sentence that *says* that there are unicorns. So I am ontologically committed to unicorns if and only if I understand and sincerely assert a sentence that says that there are unicorns; and in general I am ontologically committed to *K*s if and only if I understand and sincerely assert a sentence that says that there are *K*s.[12]

From this, we can extract a derivative notion of sentential commitment: a sentence *carries ontological commitment to Ks* if and only if anyone who understands and sincerely asserts it would thereby be ontologically committed to *K*s. So we can identify languages that are *ontologically guilty*: they allow us to form sentences that carry ontological commitments to some kind *K* or another. And a language will be *ontologically innocent* if and only if it isn't ontologically guilty.

The Nihilist needs to find an ontologically innocent language with which to paraphrase the ontologically guilty target. But we can get a better picture of what this innocent language must be like by getting a better understanding of why guilty languages are guilty.

3.2 *Variable Binding and Quantification Proper*

We all learned at Quine's knee that, in first-order languages, the existential quantifier '$\exists$' makes for ontological guilt. But in these languages it does two jobs: it manages variable-binding, and it says something about how many values of its bound variable satisfy the postfixed formula. For which of these tasks do we find it ontologically guilty?

Let's clarify the natures of these two tasks. In first-order languages, we can take a sentence open in a variable 'x' and prefix it with '$\exists$x' or '$\forall$x' to get a new sentence. If the original sentence was open in other variables, the new sentence is open in those variables, too. Otherwise, the sentence is *closed* and can be evaluated for truth.

[11] If Joe mistakenly thinks that 'unicorn' means *zebra*, he doesn't ontologically commit himself to unicorns when he says 'There are unicorns'. Thanks here to Ted Sider.

[12] This, more or less, is how Mark Richard (1998) seems to understand the notion; and Peter van Inwagen (1998: Thesis 5) is perhaps best interpreted this way, too. Agustín Rayo (2007, MS) suggests a different formulation, according to which I'm ontologically committed to *K*s iff I understand and sincerely assert a sentence with truth-conditions which demand that there are *K*s. Little hangs on this distinction in what follows.

The turning of open sentences into closed (or at least less open) ones is what we call *variable binding*. Variable binding is what lets us make complex predications about a single peg. We can use

$\exists x(Fx \,\&\, Gx)$

to say that it's one and the same thing which satisfies both 'F' and 'G'. We do this by binding two instances of the same variable. Semantically — from the point of view of the pegboard — variable binding is what lets us hang *two* rubber bands from one peg.

In addition to variable-binding, quantifiers also *quantify proper*: they say *how many* pegs are arranged the way the postfixed formula says. '$\exists x \ldots$', for instance, says that least one peg is $\ldots$; '$\forall x \ldots$' says that every peg is $\ldots$. If we have some more sophisticated quantifiers than first-order languages allow, we can also say, for instance, that infinitely many pegs are a certain way, or that most pegs that are one way are also some other way.

We want to know whether the guilt of first-order languages comes from their variable-binding or their quantification proper. We can sharpen the question by dividing the dual burdens of the first-order quantifiers between two different expressions of another language. This is what *lambda-abstraction languages* do.[13] These languages have the predicates and truth-functional constants of first-order languages. But instead of the first-order quantifiers, they have two separate symbols: a variable-binder and a proper quantifier.

Here's the idea. Introductory logic texts often tell us that we can read '$\exists x(\ldots x \ldots)$' as a sort of quasi-English expression, meaning

> There is something that is an x such that $\ldots$ x $\ldots$

Likewise, '$\forall x(\ldots x \ldots)$' can be translated as

> Everything is an x such that $\ldots$ x $\ldots$ [14]

But we could do the same work with separate expressions: one which means 'something', one which means 'everything', and a third which means 'is an x such that $\ldots$ x $\ldots$'.

This is what lambda-abstraction languages do. They have a predicate-forming operator, 'λ' that combines with a variable and an open expression to make a

[13] There are, in fact, quite a few languages that go by the name of 'lambda-abstraction'. We are here concerned with the first-order fragment of typed lambda-abstraction languages with categoramical quantifiers (Gamut 1991: 102–115; see Hindley and Seldin 1986: 266-286 for the untyped counterpart). Since the languages have categorical quantifiers, only 'λ' can bind variables; since they are typed and first-order, 'λ'-abstraction can only form first-order predicates.

[14] Quine (1960*b*: 162); cf. van Inwagen (1998: 18–22).

predicate: where ϕ is an open expression, $\ulcorner \lambda x \phi \urcorner$ means $\ulcorner$is an x such that $\phi \urcorner$. They also have expressions '$\exists_p$' and '$\forall_p$' that mean 'there is something that' and 'everything', respectively.

These languages are just as ontologically guilty as first-order ones. But we can meaningfully ask which term — the variable binders or the quantifiers proper — give rise to their guilt. And I think the answer is straightforward: the quantifiers proper are to blame, and the variable-binders should be exonerated.

That the variable-binder is not to blame: suppose that we had a language with 'λ' and only one sentence-making operator, 'B', which means 'It is possible for there to be someone who believes that something…'. No ontologically committal sentence could be formed in that language. We could only use it to talk about what possible believers could or couldn't believe. But we can talk about *that* all day without ever saying anything about what there *is*. The variable binder doesn't suffice for ontological guilt.

That the quantifiers proper are to blame: '$\exists_p$' means *there is something*. We commit ourselves ontologically when we say that there is something which is some way or another, and '$\exists_p$' is the expression we use to say this.

We can see why quantifiers proper make a language ontologically committal by thinking about why variable-binders don't. Consider a complex predicate such as

(8) $\lambda x \lambda y$(F(x) & G(y) & R(x,y)).

A pair of pegs will satisfy this predicate exactly when one of them has the 'F' rubber band hanging from it, the other has the 'G' rubber band hanging from it, and the 'R' rubber band is stretched between them. If we wanted to identify a single rubber-band structure picked out by (8), we might think of it as three rubber-bands glued together as in figure 2. But in this case, 'λ' creates a complex

Figure 2: A Rubber Band for (8)

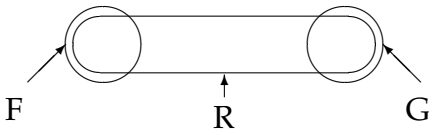

rubber band. It does not fasten that rubber band to any pegs. If you want to say that an *F* thing *R*s a *G* thing — if you want to stretch this complex rubber band between a pair of pegs — you've got to find a way to plunk down some pegs to stretch it between.

The 'λ'-operator essentially gives us the power to create complex rubber bands. This is not enough to impart ontological structure: making rubber bands doesn't give us somewhere to hang them. We have to use a quantifier proper to drop a peg down into the board, giving the rubber band something to hook on to.

We now better understand the relationship between quantification and ontological structure. The existential quantifier proper is ontologically committing because it, and no non-quantificational expression, has the job of plunking pegs down on the board. It is the existential quantifier, not the variable-binder or any other semantic gizmo, that both requires and semantically communicates that the ontological structure of reality *includes pegs* — pegs of a certain type, pegs with rubber bands corresponding to the expressions prefixed by the quantifier. And this is why quantifiers proper — especially *existential* quantifiers proper — make a language guilty.

4 Two Less Plausible Strategies

Let's return to our search for an ontologically innocent way to paraphrase our ontologically guilty target language. We will begin by considering a couple of clearly unattractive proposals. When we see the problems that beset these strategies, we will better know which pitfalls a more nuanced strategy must avoid.

4.1 Quiet Nihilism

Consider first a Nihilist who says:

> I don't see what all the fuss is about. It's easy to paraphrase the target language. We simply introduce an ontologically innocent expression, 'there schmare' (and its obvious cognates), which we can use to capture all the truths we might ever need to capture. For example, when an ordinary speaker utters
>
> (9) There are two electrons in every helium atom,
>
> she speaks falsely but manages to communicate the true
>
> (10) There schmare two electrons in schmevery helium atom,
>
> where the only difference between (9) and (10) lies in the meanings of 'there are' and 'there schmare' (and cognates, like 'every' and

> 'schmevery'), respectively. And 'there schmare' doesn't invoke peg-board structure in any way whatsoever.

When we press him on the meaning of 'there schmare', this Nihilist refuses to say anything informative. He merely insists over and over again that it can be uniformly replaced for 'there are' to turn falsehoods into truths and that it is ontologically innocent.

Call this fellow a *Quiet Nihilist*. He seems to be cheating — surely it can't be *that* easy to get by without ontology. But just what, exactly, is wrong with his strategy?

4.1.1 A Warm-up Exercise

Imagine meeting a man — Eustance — who, to your surprise, tells you nothing is blue. 'What?' you cry in amazement. Pointing at something you had always thought of as blue, you ask: 'What color is that, then?' And Eustance responds, 'Eulb'.

Trying to figure out what's going on, you ask further: 'Is eulb a color?' He says, 'Yes'. He tells you that eulb is a cool color, the color of the sky, and that it lies on the spectrum between red and green. When you ask what color complements eulb, he replies, 'yellow'. He even insists, 'Contrary to what most people think, purple is not a combination of red and blue. It's a combination of red and eulb.' He denies any sentence that you are willing to assert using the word 'blue', but will happily assert the sentence that results from it by a systematic replacement of 'blue' for 'eulb'.

It won't be long before you start thinking that when Eustance says 'eulb', he means *blue* — he is talking about the color you have known and loved all along, the color of the sky and of bluebirds, the color you have always called 'blue'. And so, even though he won't use the word 'blue' to describe those things, you will suspect that, insofar as the two of you have any real disagreement at all, it is only disagreement about which word to use for the color blue. You certainly aren't disagreeing about anything's *color*.

Suppose we think of you and Eustance as speaking subtly different languages: the 'blue'-language and the 'eulb'-language, respectively. Then your understandable attitude towards the shallowness of the 'eulb'-speaker's claims seems underwritten by the following line of thought:

> Eustance and I seem to mean the same thing by all of our terms *other* than 'blue' and 'eulb', and he uses 'elub' in exactly the same way that I use 'blue'. But, since our words get to mean what they mean thanks to the way we use them, 'blue' in my mouth and 'eulb' in his should

have the same meaning. Since 'blue' in my mouth means *blue*, 'eulb' in his mouth must mean that, too.

There is a general lesson here. Suppose L_1 and L_2 are languages that are exactly alike except that, where L_1 has an expression α, L_2 has a different expression, β. If ϕ is a sentence in L_1 that uses α, we write it as ϕ_α, and ϕ_β will be the expression of L_2 that is just like ϕ_α except that β is replaced everywhere for α. The line of thought just sketched relies on the following principle:

($*$) If every term (other than α and β) is interpreted the same way in L_1 as it is in L_2, and if the speakers of L_1 utter ϕ_α in all and only the circumstances in which speakers of L_2 utter ϕ_β, then α and β have the same interpretation also.

In the above case, of course, the 'blue'-language was L_1, the 'eulb'-language L_2, 'blue' was α and 'eulb' was β. Since you and your interlocutor meant the same thing by your *other* expressions, ($*$) licenses the conclusion that 'blue' and 'eulb' mean the same thing in your respective mouths.

4.1.2 The Status of ($*$)

Let's clear up a few points about ($*$) before going on. First, it talks about circumstances in which speakers of L_1 utter ϕ_α and in which speakers of L_2 utter ϕ_β. This talk ought to be understood dispositionally: to say that you and I utter ϕ in just the same circumstances is to say that our dispositions are such that, for any circumstance C, I am disposed to utter ϕ in C iff you are disposed to utter ϕ in C.

If we don't understand ($*$) in this way, it will prove too much. Imagine two communities that differ linguistically only in that one uses 'green' and one uses 'grue'. The green speakers are just like us, except they have never read Goodman 1979/1983 and never entertained the predicates 'grue' and 'bleen'. The 'grue'-speakers are just like the 'green'-ones, except (a), they have never entertained a predicate that works the way 'green' does, and (b) although this community calls things 'grue' exactly when the 'green'-speaking community calls them 'green', they have different linguistic intentions. The 'grue'-speakers fully intend, when they encounter green-looking things for the first time after the set future date, to not call them 'grue' anymore. And they fully intend to call blue-looking things encountered for the first time after this date 'grue'.[15]

[15]Since they are otherwise just like their 'green'-speaking counterparts, they intend that after this future date they will be able to say things like 'this sapphire is both blue and grue'. But they have not yet introduced a term to use for green things observed after this date.

Unfortunately, both the 'green'- and the 'grue'-speaking communities are annihilated by an asteroid strike before the future date is reached. So, when it comes to actual tokens of 'green' and 'grue' uttered, the two communities agree entirely. (This isn't guaranteed: the 'grue'-speakers might say things like 'emeralds observed now are grue, but emeralds observed after the special future date won't be'. Nonetheless, these referential dispositions don't guarantee that they will ever in fact say anything like this; let's suppose they never do.) If $(*)$ is understood just about what speakers *in fact* say, it will tell us that 'grue' and 'green' in these communities' respective mouths have the same interpretation. This looks implausible. Fortunately, though, $(*)$ will not license this result if it is understood as talking about the way speakers are disposed to use the expression in any possible circumstance; in any circumstances involving green things after the future date, the 'green'-speakers are disposed to call them 'green' and the 'grue'-speakers are not.[16]

Here is a second observation: $(*)$ will only seem plausible if 'interpretation' in the consequent is understood in a coarse-grained way, so that intensionally equivalent interpretations have the same interpretation. We can easily imagine two communities which differ only in that one uses the term 'triangular' whenever the other would use 'trilateral'. We should expect these communities to together satisfy the antecedent of $(*)$, but it is at least contestable that, in some sense, we don't want to say that 'triangular' means the same thing as 'trilateral'. However, we *do* want to say that these two expressions are at least intensionally equivalent — that they at least apply to the same things in the same possible circumstances. We ought to understand $(*)$ so that it says nothing more than this.[17]

[16]One caveat: we should *not* be concerned with the speakers' dispositions to utter sentences containing both of the disputed words in question. For instance, we shouldn't demand that $(*)$'s antecedent not be satisfied in the above 'blue'/'eulb' case simply because the 'eulb'-speaker is disposed to assert 'eulb things are not blue' and you, at least after serious reflection, are not disposed to assert 'eulb things are not blue'. The question is whether, *setting aside the way the speakers think these terms interact*, we should interpret them the same way; $(*)$ is supposed to give us a guide for determining whether speakers' assertions of this sort are plausible, and as such it should not be overly sensitive to the mere fact that they make these assertions. Cf. section 4.1.3 below.

[17]A third observation: 'circumstances' and 'interpretation' will both have to be understood in a fairly specific way if we are to make room for context-sensitive expressions. In particular, two speakers 'being in the same circumstances' should be understood as entailing their being in the same context (in as narrow a sense as possible, so that if John truly says 'I am tired', Bill can only count as being in the same circumstance if it is one in which Bill is tired). And two expressions 'having the same interpretation' should be understood as their having the same *character*, as opposed to the same content (in Kaplan's (1989) terms). But our focus here is on a narrower class of languages — a class that is context-insensitive — so we can ignore these details in what follows. Thanks here to Ted Sider.

4.1.3 (∗) and Quiet Nihilism

(∗), of course, makes trouble for Quiet Nihilism. Consider the first-order Quiet language the Nihilist will use to paraphrase the first-order target language. It has all the same predicates and truth-functional connectives as our first-order language, but whereas we use the existential quantifier '∃', which means 'there is something that...', he uses his 'schmexistential' quantifier, 'schm∃', which he says means 'there schmis something that...'. But he grants that his predicates and truth-functional connectives mean what ours do, and recommends using 'schm∃' in all and only the circumstances in which we are disposed to use '∃'. So (∗) tells us that 'schm∃' in his mouth means what '∃' does in ours.

Could the Quiet Nihilist defuse the appeal to (∗) by his mere insistence that 'schm∃' *doesn't* mean the same thing as '∃' does? I doubt it. Suppose Eustance insisted vehemently that 'eulb' did *not* mean the same as 'blue' in our mouths. He then places the following stipulations on the meaning of 'eulb':

(S1) 'Eulb' applies to exactly those things ordinary people would call 'blue' under ordinary conditions.

(S2) 'Eulb' is not interpreted the same way as (is not intensionally equivalent to) 'blue'.

It is not at all clear that these stipulations are jointly satisfiable. If there is a property B that applies to exactly those things that ordinary people would call 'blue' under ordinary considerations, an ideal interpreter will be under pressure to interpret 'blue' as meaning B. When Eustance comes by and makes stipulation (S1), the ideal interpreter will have no choice but to interpret 'eulb' as B. But then she will have no way to satisfy (S2) without *re*-interpreting 'blue' as something other than B. No ideal interpreter would give Eustance that sort of control over the interpretation of everyone else's 'blue' — any reasonable principle of charity will have her make Eustance, rather than the rest of us, speak falsely. So, insofar as she makes sure (S1) is satisfied, she will have good reason to leave (S2) unsatisfied.

What goes for Eustance goes for the Quiet Nihilist: he can insist all he wants that 'schm∃' does not mean '∃', but this gives us no reason to think both that it does not and that sentences such as (10) are true in exactly those situations where we think we ought to assert (9). If we grant that his 'schm∃'-using sentences are true in the circumstances he says they are, we will have good reason to think that 'schm∃' means 'there is' after all.

4.1.4 $(*)$ and Charity Arguments

One final comment is in order. My argument against Quiet Nihilism bears some superficial similarities to some other interpretative arguments that philosophers (e.g., Eli Hirsch (2002, 2005, 2007)) have run in other cases of metaphysical dispute. These 'charity' arguments run more-or-less as follows: party A insists that every one of party B's sentences ϕ is false, but can be translated into a true sentence $t(\phi)$ of party A's preferred idiom. But party A will assert $t(\phi)$ in exactly the situations where party B asserts ϕ, so (the argument goes) if $t(\phi)$ really is true in the circumstances where A would utter it, a charitable interpreter will interpret ϕ as synonymous with $t(\phi)$ and therefore as true in those circumstances as well. Since our sentences mean whatever ideal interpreters say they mean, A should think that ϕ in B's mouth has the same meaning, and hence the same truth-value, of $t(\phi)$ after all.

The crucial difference between these charity arguments and my above argument relying on $(*)$ is simply that they take place at the level of sentences whereas mine takes place at the level of words. There are thus ways to resist the charity arguments that do not likewise affect the $(*)$ argument. To take one well-discussed example, there might, as Lewis (1983*a*: 45–55, 1984) argued, be a so-called *naturalness* constraint on interpretation: try, *inter alia*, to give each word as natural and un-gerrymandered a meaning as possible. This constraint will of course be balanced against other interpretative constraints like charity. But it is crucially a constraint about the interpretation of *words* rather than the interpretation of *sentences*.[18] It may very well be that every interpretation of parties A and B that makes all of B's sentences ϕ synonymous with A's $t(\phi)$ does so by assigning overly gerrymandered meanings to the individual words of A's or B's language. And so the naturalness constraint may, as a result, require some of B's ϕs to have different meanings than A's $t(\phi)$s after all.[19]

But the $(*)$ argument relies on an interpretative principle about the meanings of *words*: when parties A and B uses the words α and β in the same way, against a background of other, shared words all understood as having the same

[18]This, at least, is how the constraint is interpreted by many; see e.g. Sider (2001*a*,*b*: xxi–xxiv), Stalnaker (2004), and Weatherson (2003). Wolfgang Schwarz (MS) argues persuasively that Lewis' considered view had little concern for the interpretation of individual terms and cared primarily about the assigning of contents to individuals' mental states. In this case, the constraint is not about *words*; but neither is it about sentences, and so it too may override pressure from charity to interpret ϕ and $t(\phi)$ as synonymous — for such an assignment may, for example, require assigning overly gerrymandered mental states to either A or B.

[19]This might very well suggest that the 'grue'-speakers from section 4.1.2 really meant 'green' by 'grue'. But this is perhaps to make the constraint too strong: an ideal interpreter's injunction to give words an ungerrymandered interpretation ought not outweigh a community's *explicit intention* to use an expression in a gerrymandered way.

interpretation in A's and B's mouths, then α and β must be interpreted the same way, too. No naturalness constraint or other word-level constraint is likely to conflict with $(*)$: the shared words have the same interpretations and hence all the same semantic properties in both languages, and if we give α and β the same interpretation, they will have the same semantic properties in both languages, too. If β gets a highly natural interpretation, for instance, then nothing about naturalness can keep α from getting the same interpretation. And if β has a gerrymandered interpretation, the fact that this interpretation snuck in against the semantic background of B means that there can be no naturalness-inspired bar to giving α the same interpretation against the same semantic background.[20]

The considerations underwriting the $(*)$ argument are much more fine-grained than those supposed to underwrite standard charity arguments, and the two kinds of argument ought not be confused. Even those who suspect interpretative charity arguments in general ought to find $(*)$ plausible and thus reject Quiet Nihilism.

4.2 Propositional Nihilism

A second Nihilist says instead:

> I am going to paraphrase the target language into the language of propositional logic. It has 'atomic' sentences 'P', 'Q', 'R', ..., and truth-functional connectives '$\sim$', '&', etc. A sentence such as
>
> (11) There is one electron in a hydrogen atom,
>
> will be paraphrased into an atomic sentence — 'P', for instance — and
>
> (12) There are two electrons in a helium atom,
>
> will be paraphrased into another atomic sentence, say 'Q'. But these atomic sentences don't invoke any pegboard structure. They just say that thus-and-so is the case, where thus-and-so is some ontology-free state of reality.

[20]The situation is really slightly better than this; by insisting that the rest of the respective languages' semantic backgrounds are the same, we essentially ensure that a lot of *other* interpretative issues between both parties have been settled in favor of similar meanings for their languages. Given that this much has been fixed, and given that the only remaining expression is used the same way by all parties, it becomes very hard to think of *any* interpretative constraint, whether at the level of sentences, words, or mental states (see note 18 above), that could pressure us to interpret α differently than β.

When we press the Propositional Nihilist to tell us more about what these sentences mean, he also refuses to say anything helpful.

I doubt that Propositional Nihilism is untenable in the way that Quiet Nihilism is. But it has several defects that we should not pass over.

4.2.1 Exploded Ideology

First, the view is *ideologically extravagant*. A theory's ideology consists of the expressions the theory takes as meaningful and undefined — the expressions, as it is often put, that the theory takes as *primitive*. But no matter how many (or how few) primitive expressions the target language has, the Nihilist's propositional paraphrasing language needs many, many more. With just a few predicates and standard first-order resources we can construct indefinitely many logically distinct sentences, e.g.:

> There is one electron in region R.
> There are two electrons in region R.
> There are three electrons in region R.
> ⋮

Since these sentences are not truth-functional compounds, they must each be paraphrased as some atomic proposition. And each of these is logically distinct, so if the Nihilist translates two of these as the same sentence, he will collapse distinctions we can make. So, insofar as he wants his paraphrases to preserve our ability to make these sorts of distinctions, he will need to paraphrase each of these by a *different* atomic proposition:

> P
> Q
> R
> ⋮

Since each of these atomic propositions constitutes a primitive bit of ideology, the Nihilist has an enormous number of primitives.

4.2.2 Lack of Systematicity

Second, the view is *inferentially unsystematic*: it endorses tremendously many inferences the validity of which it cannot explain. Consider, for instance, the sentences:

(13) There are exactly two electrons orbiting a proton,

(14) There are some electrons orbiting a proton.

The Propositional Nihilist paraphrases these as atomic sentences, say, 'A' and 'B'. And presumably, as (13) entails (14), 'A' will entail 'B', too.

The inference from (13) to (14) is underwritten by a nice, systematic theory — the first-order predicate calculus. But nothing underwrites the Nihilist's counterpart inference from 'A' to 'B'. In particular, there is nothing the Nihilist can say to explain why the inference from 'A' to 'B' is valid although the inference from 'A' to, say, 'C', where 'C' is the paraphrase of

(15) Some neutron is in region *R*,

is not. The fact that 'A' entails 'B' but not 'C' is, according to the Nihilist, a brute fact: it admits of no more basic explanation. And, although everybody has some brute facts somewhere or another, the Propositional Nihilist has more than his share: presumably there will be indefinitely many valid (and indefinitely many invalid) inferences between atomic propositions, and the validity (or invalidity) of *each one* will be a further brute fact.

4.2.3 Holism

Finally, the view is *holistic*: it cannot make sense of the thought that reality's global structure is somehow 'built up' out of its various local structures.

We ontologically-minded folk think something like the following: there are a limited number of ways things could be, and a limited number of ways things could be interrelated. And the way reality is *in toto* is determined by the way each thing is and is interrelated to its fellows.

For instance, when I say

(16) An electron attracts a proton and repels another electron,

I say that there are three pegs, arranged with rubber bands as in figure 3. And

Figure 3: The Rubber Band Structure of (16)

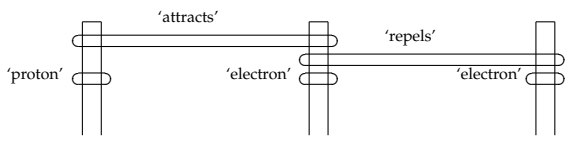

it is easy to see how this complex pegboard-and-rubber band structure is built

up out of two simpler structures, one that involves the leftmost and center pegs, and one that involves the center and rightmost pegs. In a certain way, the fact expressed by (16) is built up out of 'smaller' facts — in particular, the facts expressed by

(17) A proton attracts an electron (figure 4),

Figure 4: The Rubber Band Structure of (17)

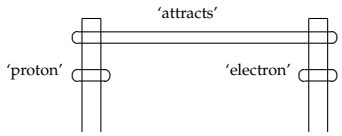

and

(18) An electron repels another electron (figure 5),

Figure 5: The Rubber Band Structure of (18)

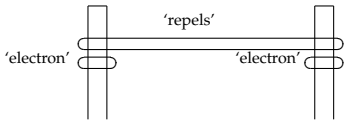

along with the fact that one of the electrons doing the repelling in (18) is also doing the attracting in (17).

Clearly the Nihilist cannot use this flagrantly ontological explanation of how the structure described by (his paraphrase of) (16) is 'built up' out of more basic structures. But what other explanation could he give? Whenever we start to talk about what looks like a distinctly ontological fact, he produces a new 'atomic' fact. Presumably, the fact is *atomic* because it encodes no further structure — it is, rather, simple, a structureless I-know-not-what. But no such paraphrase of a claim like (16) admits of an explanation of its structure in terms of more local structures — because any such paraphrase won't encode any structure to be explained.

4.2.4 Should the Propositional Nihilist Be Worried?

The Propositional Nihilist might shrug his shoulders and say, 'So what? I've bitten bullets in my time — what's a few more?' For my part, I think that the combination of inferential unsystematicity, ideological bloat, and rampant holism are troubling enough to prompt us to look elsewhere. I also think that most metaphysicians would — and should — agree, but I am not going to argue about it here. If the Propositional Nihilist is comfortable paying these prices for his Nihilism, so be it; but let it be known that he must indeed pay them.

A Propositional Nihilist might instead complain, though, that he has not really incurred one cost or another. As far as I can tell, both the explosion of ideology and the holism are straightforward consequences of the Propositional proposal[21] — I can see no hope for keeping them off the bill. But there are a few ways a Propositional Nihilist might try to claim that there is no cost from inferential unsystematicity. Let's look at these in turn.

Syntactic Unsystematicity is No Big Deal

The Propositional Nihilist's first appeal insists that inferential unsystematicity is no big deal:

> So my Propositional language has no good syntactic recipe for determining which inferences are valid. So what? Lots of perfectly good languages have this feature. Incompleteness results, for instance, tell us that higher-order languages cannot provide sound, finite inferential systems that license every valid inference. And even in natural language, many valid inferences are syntactically indistinguishable from invalid ones. So insofar as my language is unsystematic, it is no worse off than higher-order or natural languages.

This appeal is a red herring. Section 4.2.2's observation wasn't that the inferences between the atomic sentences were not *syntactically discernible*; inferential systematicity doesn't demand that sentences' inferential relations be worn on their syntactic faces. The observation was that the inferences have to be *semantically brute*: there is no explanation *whatsoever*, syntactic or otherwise, for their validity.

Let's look at the appeal to natural languages. Hawthorne and Cortens (1995: 151) point out that while the inference

[21] At least as formulated; I am aware of some other proposals that manage to push the bump in the carpet around, making the view more inferentially systematic at the cost of making it even more ideologically extravagant or holistic, for instance.

(19) He happily robbed the bank.
Therefore, he robbed the bank.

is clearly valid, the inference

(20) He allegedly robbed the bank.
Therefore, he robbed the bank.

is clearly not. And these two inferences are *syntactically* indistinguishable; the validity of (19) and invalidity of (20) are not worn on their syntactic faces.[22]

But (as Hawthorne and Cortens point out on the same page) this does not mean the inferences are brute: there is a simple *semantic* explanation for the difference in (19)'s and (20)'s validity. 'Happily' is an adverb which, when attached to a verb that picks out an action V, creates another verb which is still a kind of V-ing. But 'allegedly' is an adverb which, when attached to a verb that picks out an action V, does not create a new verb that picks out a kind of V-ing. The inference is not syntactically discernible, but that doesn't make it brute.

A similar point holds for higher-order languages. Even though they have no complete axiomatization, they do make room for semantic explanations of validity. The explanations come from the model theory for those languages, which makes then semantic, not syntactic, explanations.

Propositional Nihilism is not like either of these cases. It lacks not just a syntactic account of the inferences' validity, but it lacks a semantic one, too. Its atomic propositions, recall, do not have semantic values that encode any more detailed structure. They are propositional blobs — they can be true or false, but that's all we can say about them. After he has told us that there is a true atomic claim 'P', and that it is what we were getting at all along when we said that there was an electron orbiting a proton, the Propositional Nihilist has nothing left to say. In particular, he has no story about what 'P' *means* that would let him explain why it entails, say, 'Q' but not 'R'.

In fact, whether we have a *syntactic* way of systematizing the inferences is irrelevant. Suppose we supplement Propositional Nihilism with the following syntactic theory. Every sentence is composed of two syntactic components: a *content tag* and an *inference tag*. A content tag is simply a syntactically simple expression, such as a capital letter (perhaps with numbered subscripts, so that we can have more than 26 of them). An *inference tag* is syntactically complex, made up out of various *pseudo=expressions:* pseudo-variables ('x', 'y', 'z', ...),

[22] Hawthorne and Cortens's original invalid example, 'He ran halfway up the hill; therefore, he ran up the hill' is not clearly of the same syntactic form as their valid example, 'He ran quickly up the hill; therefore, he ran up the hill'. In the first case, 'quickly' modifies 'ran'; in the second, 'halfway' modifies 'up the hill'.

pseudo-predicates ('=', 'P', 'Q', 'R', ...), and pseudo-quantifiers ('$\forall$', '$\exists$'). There is one simple pseudo-expression in the language for every simple expression in the target language, and formulation rules for inference tags mirror those for sentences of the target language: ϕ is a pseudo-tag if and only if it is isomorphic to a sentence of the target language. Sentences of the Propositional language have the form $\ulcorner P^{\phi} \urcorner$, where P is a content tag and ϕ an inference tag. But we deny that *every* string of this sort is well-formed: each inference tag can be joined to only *one* content tag. That is, the syntax of the language specifies a function f from inference tags to content tags, and $\ulcorner P^{\phi} \urcorner$ is well-formed iff P is a content tag, ϕ an inference tag, and P is the value of ϕ for f.

Call this the *tag-language*. It has a fully specified syntax. It remains to give it a semantics. The semantics we give it is quite simple: every content tag is interpreted so as to encode one of the Propositional Nihilist's atomic facts. And inference tags, and all of their parts, are semantically empty.[23]

Propositional Nihilists can easily create a syntactic inference system that will mirror the inferential structure of the target language: paraphrase any sentence ϕ of the target language as $\ulcorner P^{\phi} \urcorner$, for some content tag P. Then say that $\{P_1^{\phi_1}, P_2^{\phi_2}, \ldots\} \vdash Q^{\psi}$ iff $\{\phi_1, \phi_2, \ldots\} \vdash \psi$.

This certainly gets the inferences right. For instance, the first-order renderings of (13) and (14) ,

(21) $\exists x \exists y \exists z (\mathrm{E}x \,\&\, \mathrm{E}y \,\&\, \mathrm{P}z \,\&\, x\mathrm{O}z \,\&\, y\mathrm{O}z \,\&\, x \neq y \,\&\, \forall w (w\mathrm{O}z \supset w = x \vee w = y))$ and

(22) $\exists x \exists y (\mathrm{E}x \,\&\, \mathrm{P}y \,\&\, x\mathrm{O}y)$,

(with '$\mathrm{E}x$' abbreviating 'x is an electron', '$\mathrm{P}x$' abbreviating 'x is a proton', and '$x\mathrm{O}y$' abbreviating 'x orbits y'), get paraphrased as

(23) $\mathrm{A}^{\exists x \exists y \exists z (\mathrm{E}x \,\&\, \mathrm{E}y \,\&\, \mathrm{P}z \,\&\, x\mathrm{O}z \,\&\, y\mathrm{O}z \,\&\, x \neq y \,\&\, \forall w (w\mathrm{O}z \supset w = x \vee w = y))}$ and

(24) $\mathrm{B}^{\exists x \exists y (\mathrm{E}x \,\&\, \mathrm{P}y \,\&\, x\mathrm{O}y)}$,

respectively. But since (22) is deducible from (21) in first-order logic, our Nihilistic inference rules let us deduce (24) from (23). This language can, in this manner, provide a full syntactic recipe for determining which inferences are valid.

But so what? This syntactic inference-encoding *has nothing to do with what the sentences mean*: the only part of the sentence that does any semantic work is also the only part of the sentence that is irrelevant to the syntactic validity-checking

[23] They are thus like the semantically empty 'it' of weather-sentences; see section 5 below.

procedure. The scheme tells us *which* inferences are valid, but does nothing to explain *why* those inferences deserve to be valid.

Syntactic systematizations of inferences are useful and informative when and insofar as variation in syntax corresponds to similar variation in semantics. The demand for 'inferential systematicity' is a demand for a semantic story about what underwrites the inferences — not merely a syntactic recipe for figuring out which inferences are the valid ones. Our ability to tell such a story depends ultimately on the structures encoded by the semantic values of the sentences involved. But the Propositional Nihilist denies that his atomic sentences encode any interesting structure; as a result, he denies his paraphrase languages the resources needed for inferential systematicity.

Why Should the Inferences Be Valid?

Our response to the last appeal tells us that, if the Propositional Nihilist endorses all the inferences we expect him to, his system will be unsystematic. But our Propositional Nihilist will now tell us he doesn't endorse all the inferences we expect him to. He says:

> It's no constraint on a paraphrase scheme that every inference supposed to be valid in the target language will remain valid under paraphrase. In fact, proponents of various paraphrase schemes often like them *because* they invalidate certain troublesome inferences.[24] The mere fact that ϕ entails ψ doesn't mean my paraphrase of ϕ must entail my paraphrase of ψ. So why can't I say that, in my language, atomic sentences typically *don't* entail other atomic sentences? Now my language is inferentially systematic again; it just does not license all the inferences you thought it would.

If we think that ϕ entails ψ, we think that whatever important fact we're getting at with ϕ cannot be true if the important fact we're getting at with ψ is false. If someone comes along with a paraphrase scheme according to which the paraphrase of ϕ does not entail the paraphrase of ψ, he says that we are wrong about this relationship: whatever important fact we're getting at with ϕ *could* be true even if the important fact we were getting at with ψ is false.

It is certainly no desiderata on a metaphysics that every inference ordinary folk are inclined to make turn out valid. But it's one thing to say that ordinary folk tend to be wrong about the validity of certain troublesome inferences, and another thing to say that ordinary folk tend to be wrong about the validity of

[24]See Hawthorne and Cortens (1995: 151–152); cf. van Inwagen (1990: 128).

almost every inference they're inclined to make. We tend to make paradigmatically quantificational inferences — inferences of the sort that cannot be captured in a purely Propositional language — all the time. If the Nihilist invalidates all of these, he comes dangerously close to saying that we are radically mistaken about the world and affirming thesis (1) after all. Since he wants to avoid (1), he ought to be careful about how many apparently valid inferences he wants to ultimately rule invalid.

The point can be put another way. Ontological nihilists of any stripe want to 'save the appearances' — that is, they want to explain why talking as though there are certain kinds of things is often useful even though there are no things of that kind. But along with this, they also need to 'save the practices', explaining why certain natural transitions involving talk of these kinds of things are so useful.

Consider again the bridge in figure 1. We noted there that someone who doesn't believe in holes needs to explain why, if there are no holes, you can point to

(3) There is a hole in the right-hand side of the bridge,

to explain why you crossed on the left instead of on the right. But this explanation will only count as a good explanation if certain inferences are valid.

In one sense, to explain an action it so explain why someone did it. Explanations of this sort usually cite some beliefs and desires: roughly, I can explain my *A*-ing by pointing out that I desired that *C* be the case and I believed that if I *A*-ed, *C*.[25] And I can explain why I *A*-ed instead of *B*-ed by pointing out that I desired that *C*, believed that if I *A*-ed, *C*, and believed that if I *B*-ed, not-C. Call this sort of explanation a *descriptive explanation* — it describes why somebody acted in a certain way.

The most natural way to give a descriptive explanation of your bridge-crossing behavior is to point out that you desired that you cross to the other side without falling through and that you believed both of:

(25) If I walk on the left-hand side of this bridge, I will cross without falling through.

(26) If I walk on the right-hand side of this bridge, I will fall through.

But while this explains why you crossed the bridge as you did, it does not explain why your so doing was a *good idea*. Consider Hal, who tends to hallucinate that there are holes where there aren't any. That is, he tends to hallucinate that bridges shaped like the one in figure 6 are instead shaped like the one in

Figure 6: A Perfectly Good Bridge

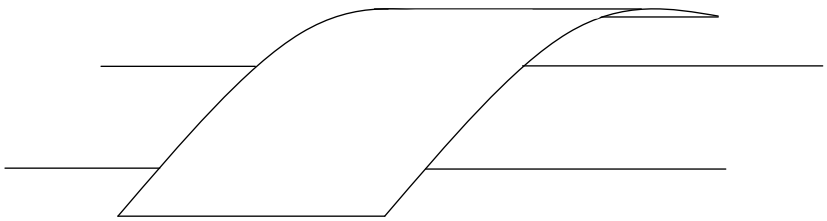

figure 1. Hal comes across a perfectly good bridge, and crosses it on the left. We can explain both your behavior and Hal's by pointing to your respective desire to cross without falling through and your (25)- and (26)-like beliefs. But we tend to think that you, unlike Hal, had a *good* reason for crossing your bridge: there *was* a hole in your bridge, and there wasn't one in Hal's.

We point to the truth of (3) to explain both why you crossed the bridge as you did and why your bridge-crossing behavior, unlike Hal's, made sense. That is, we point to the truth of (3) not only to give a descriptive explanation, but also to give a *justifying* explanation — an explanation as to why your actions were, under the circumstances, smart. Your recognition of the truth of (3) caused you to believe (25) and (26), which explains why you crossed as you did. And because your belief was based on the *truth* of (3), your bridge-crossing behavior was, unlike Hal's, reasonable.

The perforational nihilist could perhaps point to your belief in (3) to explain why you crossed as you did. But he cannot point to its *truth* to explain why your behavior was reasonable — for he does not think that (3) is true. This is where the paraphrase strategy comes in: he points instead to the truth of something in the neighborhood of (3) — namely,

(4) The right-hand side of the bridge is perforated,

that makes crossing on the right a bad idea. He believes that (4) is true of your bridge, and not of Hal's; that it is your recognition of the fact expressed by (4) that caused your behavior and something very different that caused Hal's; and that as a result your bridge-crossing behavior is well-motivated and Hal's is not.

[25]See Davidson (1963).

That is, he uses the paraphrases of hole-talk to provide the needed justifying explanations.

But this paraphrasis-explanation works only if it really is a bad idea to cross on the right-hand side of right-hand-side perforated bridges. More precisely, it works only if (4) (plus some very reasonable background assumptions) entails (26). If perforated sides of bridges were the sorts of things you could walk over unharmed, the truth of (4) just wouldn't be relevant to your bridge-crossing behavior. In other words, the perforational nihilist's paraphrase scheme meets the challenge only given the validity of certain transitions we are prone to make — so prone to make, in fact, that we often don't notice them until they are pointed out, as has been done here.

The Propositional Nihilist needs to give a justifying explanation for our bridge-crossing behavior just as much as the perforational nihilist — more so, since he thinks not only that there are no holes, but also that there are no bridges. But he will paraphrase (3) and (26) as atomic propositions. And the inference that takes us from the former to the latter is a paradigmatically quantificational inference.[26] So if the Propositional Nihilist wants to be able to offer the needed explanation, he will have to say that at least some inferences between his atomic propositions are valid.

This is no local phenomenon, either. Almost every instance of scientific discovery involves paradigmatically quantificational inferences — inferences of the sort the Propositional Nihilist must paraphrase as inferences between atomic claims. This sort of reasoning is built deeply into our natures — so deeply that it is hard to imagine successfully navigating our environment if it was not typically valid. One way to explain our success is to say that the facts we get at with the premises of our inferences typically entail the facts we get at with the conclusions of our inferences. But it is hard to see what else the Propositional Nihilist could point to, short of a cosmic coincidence, that could explain our success.

The charge of inferential unsystematicity thus stands: if the Propositional Nihilist's paraphrase strategy is to do what paraphrase strategies are supposed

[26]Presumably, the inference in question runs:

(i) There is a hole (of such-and-such a size) in the right-hand side of this bridge.

(ii) Anyone who walks on a side of a bridge with a hole (of such-and-such a size) falls through.

(iii) If I walk on the right-hand side of this bridge, I walk on a side of a bridge with a hole (of such-and-such a size). (from i)

(iv) Therefore, if I walk on the right-hand side of this bridge, I will fall through. (from ii, iii)

This inference is valid in quantificational logic. But the corresponding inference, with (ii) traded in for an atomic proposition, is not.

to do, it will need indefinitely many brute, inexplicable entailments between atomic propositions.

5 A Better Proposal: Feature-Placing Languages

5.1 Introducing Feature-Placing Languages

The Quiet proposal is untenable, and the Propositional proposal is unattractive. Perhaps a Nihilist can do better.

Consider first the sorts of sentences we use to report the weather:

> It is raining,
> It is snowing,
> It is cold,

and so on. Notice that, despite the 'it' in each sentence, none of these say that any *thing* is raining, snowing, or cold. These sentences simply 'place' certain meteorological features — simply say that raining or snowing is going on, or that coldness is manifest — without saying that any particular object is doing the raining or snowing or being cold. Unlike most English sentences, these are not talking about the arrangements of rubber-bands on pegs. If they are doing anything even in the neighborhood of that, they are simply throwing a rubber band onto the board between the pegs.[27]

P. F. Strawson (1954, 1963) noticed that we could, in principle, use sentences like this to make ontologically innocent (i.e., peg-free) claims in the neighborhood of claims we generally make about particular things. For instance, instead of saying

(27) $\exists x(x \text{ is a cat})$,

we could say

(28) It is catting.

[27]This may not be quite right. The semantics of 'is raining' may make it a predicate of places. The 'it' that syntax demands is definitely semantically empty (see Seppänen 2002: 445–453 for powerful arguments that it must be), but 'is raining' may nonetheless include a location 'slot' at the semantic level, filled in by context in a bare assertion of 'It is raining' but explicitly filled in constructions such as 'It is raining in Austin' or bound as in constructions such as 'Wherever Joe went, it rained' (cf., eg. Stanley 2002: 416–418 on binding). Out of charity towards the Nihilist, though, we will ignore these complications here; cf. note 28 below.

Just as 'it is raining' says that rain is going on without saying that there is any thing which is raining, (28) should be understood as saying that catting is going on without saying that any particular thing is a cat.

Following Strawson, we will call sentences such as (28) *feature-placing sentences*, and if a language only allows sentences (and truth-functional compounds of sentences) of this sort, we will call it a *feature-placing language*. The idea is that the Nihilist can paraphrase every apparently quantificational sentence we would ordinarily want to count as true into some sentence of a feature-placing language and thereby account for all the undeniable facts without appealing to any pegboard-like structure. (See Hawthorne and Cortens 1995)[28]

5.2 *The Proposal and Predicate Functors*

How do we turn this suggestion into a concrete paraphrase scheme? We begin by replacing every one-placed predicate (we will deal with relational predicates later) A with a feature-placing predicate ⌜is A-ing⌝. Then we paraphrase every sentence of the target language

(29) $\exists x A(x)$,

as

(30) It is A-ing.

We can now paraphrase very simple sentences. How do we deal with more complex ones? We need to tread carefully around them. Consider, for instance, the distinction between the following:

(31) ∃ x(x is positively charged & x is negatively charged)

(32) ∃ x(x is positively charged) & ∃ x(x is negatively charged).

(32) says that some things are positively charged and some things are negatively charged; this is the sort of sentence the Nihilist should paraphrase into something he takes to be true. But our best science rules out (31) (or so I am told), and so the Nihilist ought to paraphrase it as something he takes to be false.

It is initially tempting to paraphrase (31) and (32) respectively as:

[28] If the predicate-of-places account of the 'it' in 'It is raining' described in note 27 is right, then a Nihilistically acceptable reading of (28) won't be strictly parallel to 'It is raining': (28) will predicate cattingness of *places*, and thus invoke pegboard structure at that level. I think, however, that we can still make sense of what the Nihilist wants to say with (28); and even if we can't, we can learn much from pretending we can and seeing how far the Nihilist can run with his proposal. So I do not intend to make much hay over these otherwise problematic linguistic considerations here.

(33) It is positive-charging and negative-charging,

(34) It is positive-charging and it is negative-charging.

But the temptation should be resisted, for these sentences say the same thing. The semantics of feature-placing sentences treat the 'it' as empty and the predicate ⌜A-ing⌝ as expressing a proposition. The 'it' is needed simply to fill a syntactic requirement, but isn't doing any semantic work. (Some languages do not have this syntactic requirement, and their corresponding feature-placing sentences are simply verbs. The Spanish counterpart of 'it is raining' for instance, is the conjugated verb 'lluevre'.) But if 'is positive-charging' and 'is negative-charging' express propositions all by themselves, then any 'and' between them simply conjoins those propositions, regardless of where the 'it' shows up.[29] (33) and (34) are equivalent, so we can't use them to respectively paraphrase both (31) and (32).

We do better if we make some logically complex predicates out of the simple predicates *first*, before turning them into 'feature' expressions. We could then construct a predicate 'is positively charged and negatively charged', and turn *that* into a single feature-placing expression 'is (positive-charge and negative-charge)-ing' which is not to be understood as the conjunction of 'is positive-charging' and 'is negative-charging'. Then the Nihilist could paraphrase (31) as

(35) It is (positive-charge and negative-charge)-ing,

which is equivalent to neither (33) nor (34).[30]

Let's make this proposal more precise. Suppose we begin with a stock of simple predicates A, B, ... Then we help ourselves to some *predicate functors*, expressions that combine with predicates to make new predicates. For instance, we help ourselves to a *predicate conjunction* functor '&', which combines with any two predicates to create a third. If P and Q are any predicates, then ⌜P & Q⌝ is their conjunction. Likewise, we help ourselves to a *predicate negation* functor '$\sim$': if P is a predicate, ⌜$\sim P$⌝ is its negation.

[29]Compare, for instance, 'It is raining and it is cold' with 'It is raining and cold'. This transparency of the 'it' to truth-functional operators is one reason semanticists think 'it' is semantically empty; cf. Seppänen (2002: 448).

[30]Hawthorne and Cortens (1995: 148–149) suggest using adverbs to solve the problem, rendering a sentence such as 'There is a red cat' as 'It is catting redly'. While my suggestion here is similar in spirit, by not distinguishing between feature-placing verbs and adverbs, it is somewhat more streamlined: for instance, it can paraphrase (31) without deciding (as Hawthorne and Cortens's proposal would have to) which of 'is positively charged' and 'is negatively charged' to turn into a verb and which into an adverb.

We can build up any truth-functionally complex predicate we want with these two functors. But how will we turn these complex predicates into the sorts of expressions that the feature-placing language uses?

We might simply help ourselves to a large stock of primitive expressions: for every predicate A of the language to be paraphrased away, regardless of whether it is simple or complex, we introduce a primitive expression ⌜is A-ing⌝ of the feature-placing language. But that would be unlovely, incurring some of the costs of Propositional Nihilism. For instance, it would force the feature-placing language to have a huge stock of primitive expressions relative to the quantificational language it paraphrases. And making all of these expressions primitive in this way would obliterate logical relations that we might well want to keep around. Intuitively, since the existence of something both positively and negatively charged ought to entail the existence of something positively charged, for instance, the feature-placing (35) ought to entail

(36) It is positive-charging.

But if the expressions 'is (positive-charge and negative-charge)-ing' and 'is positive-charging' are simply two separate, semantically simple items, then either this entailment won't hold or the Nihilist will have to write it in by hand — along with a huge number of other entailments — as a brute necessary connection.

If he wants to avoid both bloating his ideology and de-systematizing his inferences, the Nihilist can do better: rather than take the many ⌜is A-ing⌝ expressions as primitive, help himself to a third predicate functor, 'is ...-ing', which combines with predicates (whether simple or complex) to produce the feature-placing predicates he needs for his paraphrases.

Actually, at this point we might as well drop the syntactic pretense that the feature-placing language's expressions ⌜is A-ing⌝ are *predicates*. As we have already noticed, from the perspective of the *semantics*, these things are sentence-like — they are truth-evaluable all on their own, and only demand a (semantically empty) 'it' to satisfy a quaint demand of English syntax. To make the semantics and the syntax march in step, we can let the predicate-functor combine with predicates to make *sentences*. Suppose we write this functor 'Δ': then for any predicate A, whether simple or complex, ⌜$\Delta(A)$⌝ will be the Nihilist's symbolic representation of the sentence ⌜It is A-ing⌝.

If our feature-placing language has the simple predicates of the target language and the three predicate functors '&', '∼', and 'Δ', we can paraphrase the target language into it simply and smoothly.

The paraphrase strategy relies on two facts. First, every sentence in a first-order language with only one-placed predicates is equivalent to a truth-functional compound of sentences of the form

(37) $\ulcorner \exists x(\ldots x \ldots) \urcorner$,

where '$\ldots$x$\ldots$' is some truth-functional compound of atomic predications of the form $\ulcorner Ax \urcorner$.[31] Say that sentences of this type are in *existential normal form*.

Second, every truth-functional compound of atomic predications of the form $\ulcorner Ax \urcorner$ can be turned into a predication of a single complex predicate made up from simple predicates and the truth-functional functors in a fairly obvious way. ($\ulcorner Ax \,\&\, Bx \urcorner$ becomes $\ulcorner (A\&B)x \urcorner$, $\ulcorner \sim Ax \urcorner$ becomes $\ulcorner (\sim A)x \urcorner$, and so on.) Call this the *functor reduction* of the original truth-functional compound.

In this case, to paraphrase a sentence ϕ of first-order logic in the feature-placing language, first put ϕ in existential normal form, and then for each subsentence of the form (37), replace it with

(38) $\ulcorner \Delta(P) \urcorner$,

where P is the functor reduction of '$\ldots x \ldots$'. Now we have feature-placing replacements for each sentence of the target language without any of the costs of Quiet or Propositional Nihilism. So the feature-placing option, it seems, gives us Nihilism on the cheap.

5.3 What About Relations?

But not so fast. We're not *entirely* done, because we have not yet said anything about how to deal with relational predicates. Our best science will endorse claims such as

> $\exists$x$\exists$y(x repels y),
> $\exists$x$\exists$y(x orbits y & x attracts y),
> etc.

which use relational predicates. But if science won't limit itself to a vocabulary of one-placed predicates, the Nihilist's language shouldn't either.

The paraphrase scheme already in place is nice; the Nihilist ought to try to extend it to deal with relational predicates. How would such an extension go? He will have to say that, just as we can 'place' the feature associated with a one-placed predicate A by prefixing it with a 'Δ', we can also somehow 'place' the relational feature associated with a many-placed predicate R by prefixing it with a 'Δ', too. Just as $\ulcorner \Delta(A) \urcorner$ says that it is A-ing, $\ulcorner \Delta(R) \urcorner$ will, in some sense or another, say that it is R-ing.

[31] This follows from the fact that a sentence of the form (37) is equivalent to one using only one variable; cf. Boolos et al. (2002: 274–275).

But in just what sense? What happens to a many-placed predicate when it gets prefixed with 'Δ'? The Nihilist really has only two useful options here: either say that prefixing a many-placed predicate with 'Δ' creates a new predicate, or say instead that it creates a sentence. On the first option, if R is an n-placed predicate, $\ulcorner\Delta(R)\urcorner$ is an $n-1$-placed predicate. On this proposal, 'Δ(repels)', for instance, is a one-placed predicate — the Nihilist's predicate paraphrase of our complex predicate 'repels something'. On the other option, attaching 'Δ' to a predicate always creates a sentence, no matter how many places the predicate had to begin with. On this proposal, 'Δ(repels)' is the Nihilist's sentential paraphrase of our sentence 'Something repels something'.

Let's examine each of these in turn.

6 Predicate Functorese

6.1 The Combinatorial Functors

According to the first proposal, when I attach 'Δ' to, say, the predicate 'orbits', I get a new predicate, 'Δ(orbits)'. Since 'orbits' has two places, this new complex predicate has just one. And, although it is difficult to find any predicate of natural language (or even of a hybrid natural language akin to Strawson's 'it is catting' language) that concisely communicates what this predicate *means*, the idea is straightforward: 'Δ(orbits)' is the Nihilist's feature-placing paraphrase of our one-placed predicate 'orbits something'. Then, to make a sentence out of this predicate, I can attach another 'Δ' to it: 'something orbits something' is paraphrased as 'ΔΔ(orbits).'

This proposal suggests a natural paraphrase strategy. Every first-order sentence is equivalent to one in *prenex normal form*: one which begins with a block of quantifiers followed by a quantifier-free open sentence. But any block of quantifiers can be converted to a block of existential quantifiers sprinkled with negations; say that a sentence that begins with quantifiers and negations which are then followed by a quantifier-free open sentence is in *prenex existential form*. Now, *if* we can find some n-placed predicate equivalent to any quantifier-free sentence open in n variables, we have a straightforward way to paraphrase any first-order sentence ϕ: first, convert ϕ to prenex existential form

$$\exists x_1 \cdots \exists x_i \cdots \exists x_n (\ldots x_1 \ldots x_i \ldots x_n \ldots)$$

(with negations interspersed between the various existential quantifiers if needed), convert the open sentence '$\ldots x_1 \ldots x_2 \ldots x_n \ldots$' to the equivalent n-ary predicate P to get

$$\exists x_1 \cdots \exists x_i \cdots \exists x_n (P(x_1, \ldots, x_n))$$

and paraphrase the quantifiers as 'Δ'-functors in the natural way to get:

$$\Delta \ldots \Delta \ldots \Delta(P)$$

(with negations interspersed between the 'Δ's in the obvious way). This will always work, so long as we can turn every quantifier-free open sentence into a complex predicate.

The real work is coming up with a complex predicate for each quantifier-free open sentence. The Nihilist already has many of the resources he needs for this job. For instance, he can turn any sentence open in only one variable into a complex predicate using just the functors '$\sim$' and '&'. And, via a natural extension of '$\sim$' and '&' to multi-placed predicates, he can trade in some other sentences, too. We extend '&' so that, if A is an n-placed predicate and B an m-placed predicate, $\ulcorner A \text{ \& } B \urcorner$ is an i-placed predicate, where i is the greater of n and m, so that $\ulcorner (A \text{ \& } B)x_1, \ldots, x_i \urcorner$ is equivalent to $\ulcorner Ax_1, \ldots, x_n \text{ \& } Bx_1, \ldots, x_m \urcorner$.[32] Then, for instance, he can turn the open sentence

x is a proton & x orbits y

into the predicate

(is a proton & orbits)

and paraphrase

$\exists$x$\exists$y(x is a proton & x orbits y)

as

$\Delta\Delta$(is a proton & orbits).

But some problematic first-order sentences remain. Begin with:

(39) $\forall$x$\exists$y(y orbits x).

Our current paraphrasing resources include the predicates of the target language, the 'Δ'-functor, (predicate and sentential) conjunction, and (predicate and sentential) negation. Assuming 'orbits' is the only predicate we use in paraphrasing (39), the natural candidates available for that paraphrase are:

[32] To say that one open sentence P is equivalent to another, Q, is to say that P can everywhere be replaced for Q *salva veritate* (at least in languages without opaque contexts). 'Equivalence', in this sense, is as dependent upon where variables are placed as it is upon where predicates are placed.

(40) $\Delta\Delta$(orbits)
$\sim\Delta\sim\Delta$(orbits)
$\Delta\sim\Delta(\sim$orbits)
$\sim\Delta\Delta(\sim$orbits)

But each of these are already tagged as respective paraphrases for:

(41) $\exists$x$\exists$y(x orbits y)
$\forall$x$\exists$y(x orbits y)
$\exists$x$\forall$y(x orbits y)
$\forall$x$\forall$y(x orbits y)

Since (39) is not equivalent to any of the sentences in (41), unless the Nihilist wants to run together claims that ought to be distinct, he will look for a paraphrase not found in (40).

Here is another way to see the problem. Our initial paraphrase strategy tells us to take sentences of the form

$$\exists x_1 \cdots \exists x_i \cdots \exists x_n (\ldots x_1 \ldots x_i \ldots, x_n \ldots)$$

(perhaps with negations sprinkled through the block of quantifiers) and then find a complex predicate P so that

$$P(x_1, \ldots, x_i, \ldots, x_n)$$

is equivalent to the open sentence

$$\ldots x_1 \ldots x_i \ldots x_n \ldots .$$

But it's crucial that, in this equivalent one-predicate open sentence, the variables occur in the same order that they occur in the original. If x_1 is the first variable bound in the block of quantifiers, it needs to be the first of P's arguments, if x_2 is bound second, it needs to be the second of P's arguments, and so on.

If the quantifiers in the sentence to be paraphrased are all existential, or all universal, then we can switch the order in which they bind variables without affecting the meaning of the sentence. But when the block has a mixture of existential and universal quantifiers, as (39) does, such switching affects meaning. We get problems in precisely these cases.

In (39), 'x' is bound first and 'y' is bound second. So we need to find a predicate P where $\ulcorner P$(x,y)$\urcorner$ is equivalent to the open sentence 'y orbits x'. Clearly, 'orbits' is not such a predicate: 'orbits(x,y)' is not equivalent to 'y orbits x'. And almost as clearly, no truth-functional compound of 'orbits' will do the trick either. We need something else.

If we had, in addition to the predicate 'orbits', the predicate 'is orbited by', our troubles would be over: 'y orbits x' is clearly equivalent to 'x is orbited by y' (or, in other notation, 'orbited by(x,y)'). Then 'orbited by' would be just the predicate we have been looking for, and we could paraphrase (39) as

(42) $\sim\Delta\sim\Delta(\text{orbited by})$

But where will we find this oh-so-useful predicate? We might just add a new primitive predicate, 'orbited by', to our stock. But as we have seen time and again, we do better, avoiding ideological bloat and inferential brutality, if we find a way to build up 'orbited by' from 'orbits'. And indeed we can, by introducing another predicate functor: the *inversion functor*, 'INV'. Where R is any two placed predicate, $\ulcorner\textsc{inv}(R)\urcorner$ is a predicate that means 'is R-ed by. With INV in hand, (39) is easy to paraphrase: it becomes

(43) $\sim\Delta\sim\Delta(\textsc{inv}(\text{orbits}))$.

For any two-placed predicate R, the open sentence $\ulcorner R(x,y)\urcorner$ is equivalent to $\ulcorner\textsc{inv}(R)(y,x)\urcorner$. 'INV' essentially tells the predicate's two positions to trade places. As a result, 'INV' is well-defined only for binary predicates.[33] Suppose that we have a predicate Q with, say, four places. We know that 'INV' tells predicates' positions to move around. But where will it tell them to move to? There isn't just a single rearrangement of positions that counts as 'trading places'; there are many.

It will, in fact, be useful to focus on two particular ways to trade predicates' positions' places. To see what they are, imagine our target language includes the primitive four-placed *congruence* predicate, 'Cong', which means '_ _ _ is as far from - - - as ... is from $\cdots$' (cf. Tarski 1959, Field 1980). Suppose we want to tell the 'Cong' predicate to move its positions around. One thing we might to is just tell it to swap the last two positions and leave the rest alone, so that $\ulcorner\text{Cong}(w,x,y,z)\urcorner$ will be equivalent to $\ulcorner\textsc{inv}(\text{Cong}(w,x,z,y))\urcorner$. But we might want instead for it to move the last position up to the front and bump everything back a notch, so that $\ulcorner\text{Cong}(w,x,y,z)\urcorner$ will be equivalent to $\ulcorner\textsc{inv}(\text{Cong}(z,w,x,y))\urcorner$.

Let's give ourselves predicate functors that will do each of these: *minor inversion*, 'INV', will swap a predicate's last two positions, and *major inversion*, '$\mathbb{Inv}$', will move a predicate's last position to the front. It turns out that these two functors, wisely deployed, can generate *any* rearrangement of predicates' positions we might like.

[33] And perhaps for unary ones: we might take $\ulcorner\textsc{inv}(P)\urcorner$ to be equivalent to P when P has only one place.

We are *almost* ready to paraphrase everything in the target language. But there is one final issue that needs to be resolved. Consider the open sentence

(44) x attracts y & y attracts z.

In order to paraphrase sentences involving (44), we need a predicate P where $\ulcorner P$(x,y,z)$\urcorner$ is equivalent to (44). But we have no way to build one out of 'attracts'. It is a two-placed predicate, and none of our functors let us get predicates with more places out of predicates with fewer. '$\sim$', 'Inv' and '$\mathbb{I}\mathrm{nv}$' leave the number of places alone, 'Δ' takes a place away, and even '&' only produces a predicate with as many places as its biggest argument. So anything we care to make from 'attracts' with our current resources will have no more than two places.

We can solve the problem by adding a *padding* functor: a functor that adds a 'dummy' place to the beginning of a predicate. That is, for any predicate P and variable y, $\ulcorner P(x_1, \ldots, x_n)\urcorner$ will be equivalent to $\ulcorner \textsc{Pad}(P)(y, x_1, \ldots, x_n)\urcorner$. (The new variable, y, is a dummy because it simply does no work — as we ontologically minded folk would say, whether or not some objects satisfy $\ulcorner \textsc{Pad}(P)\urcorner$ has nothing to do with what object gets assigned to y, but only which objects get assigned to the various x_i's.)

Now we can handle (44). First, note that 'Pad(attracts)(x, y, z)' is equivalent to 'y attracts z', so (44) is eqivalent to

(45) x attracts y & Pad(attracts)(x,y,z).

But (45) will be equivalent to

(46) (attracts & Pad(attracts))(x,y,z),

so '(attracts & Pad(attracts))' is just the predicate we're looking for.

6.2 *A Nihilist's Paradise?*

In fact, with these six functors — 'Δ', '$\sim$', '&', 'Inv', '$\mathbb{I}\mathrm{nv}$', and 'Pad'[34] — we can paraphrase absolutely any first-order sentence science might throw at us. And it gets better than that, for we have stumbled across Quine's (1960*a*, 1971) *Predicate Functor Language*, or *Functorese*. It not only has the expressive resources needed to translate anything we say in a first-order language, but it has its own attendant logic, besides. (Cf. Kuhn 1983, Bacon 1985)

Call the Nihilist who uses predicate functorese as his feature-placing language, best suited to show what reality's structure is *really* like and fitted to

[34] And an identity predicate; I assume our paraphrasing language has one of those to work with.

paraphrase science's first-order claims, the *Functorese Nihilist*. The Functorese Nihilist avoids the costs of Propositional Nihilism. For instance, he avoids the rampant ideological bloat that beset the Propositional Nihilist: where the Propositional Nihilist had to introduce indefinitely many new primitive expressions into his language to handle all of the first-order consequences of science, the Functorese Nihilist must introduce only six. And functorese has its own attendant, sound and complete logic, which mirrors the predicate calculus in the following way: if ϕ entails ψ in the predicate calculus, then the functorese paraphrase of ϕ will entail the functorese paraphrase of ψ in predicate functor logic. And this entailment is, intuitively, reflective of the interrelations of the meanings of the various complex predicates. So the Functorese Nihilist has no problems of inferential unsystematicity, either.

It is perhaps reasons such as these that have led some (e.g., Jonathan Schaffer (2008) and David Chalmers (2008); see also Burgess and Rosen 1997: 185–188) to suggest that functorese is *the* feature-placing language of choice for the Nihilist. But I think the Nihilist's hopes are misplaced if they are placed in functorese, for — despite its other laudable features — I doubt that functorese has the primary qualification for the Nihilist's paraphrasing job: that of being *ontologically innocent*. For, even though it avoids the ills that beset Propositional Nihilism, it falls straight into the ills of Quiet Nihilism.

6.3 *The Argument*

The main thrust of the argument is that 'Δ' means 'there is' and therefore that Functorese is not ontologically innocent. The idea is that, of the six predicate functors Functorese uses, only 'Δ' does any of the (alleged) ontology-avoiding work. The other functors — 'Inv', 'Pad', and the like — just give us a fancy way to handle variable-binding-like jobs in a variable-free way. But how we handle variable binding has nothing to do with ontological guilt, as we saw in section 3.2. So all the ontology-avoiding work must be done by 'Δ'. Unfortunately for the Functorese Nihilist, he will use 'Δ' exactly when we will say 'there is something', and he does so in a way that lets us conclude, by appeal to principle ($*$) from section 4.1, that 'Δ' means 'there is something' after all.

Let's make this argument more precise.

6.3.1 The ($*$) Argument

Let F be the functorese language that the Nihilist wants to paraphrase the target language, T, into. The argument proceeds in three steps:

Step One: Begin with F, and introduce a new language $F\lambda$ as follows: it has all the same primitive predicates and sentential connectives as F, and it retains the feature-placing functor 'Δ'. And these expressions are to be interpreted in the same way as they are in F. But $F\lambda$ does not have the other four predicate functors; instead, it has variables and the abstraction operator 'λ' from section 3.2.

Step Two: Define a new language, $F\delta$. $F\delta$ is just like $F\lambda$ except that, instead of having the 'Δ' functor and 'λ', it has one sentential variable-binding operator 'δ'. All of the expressions that $F\lambda$ and $F\delta$ share are to be interpreted the same way, and 'δ' is to be interpreted as '$\Delta\lambda$'.[35]

Step Three: We appeal to $(*)$ from section 4.1. If L_1 and L_2 are languages that differ only in that L_1 has a term α where L_2 has a term β, this principle says:

> $(*)$ If every term (other than α and β) is interpreted the same way in L_1 as it is in L_2, and if the speakers of L_1 utter ϕ_α in all and only the circumstances in which speakers of L_2 utter ϕ_β, then α and β have the same interpretation also.

Now consider the target language, T, that the Functorese Nihilist wants to paraphrase. It has all the same predicates as $F\delta$: F uses for simple predicates the predicates of T, and $F\delta$ inherits its simple predicates from F. Furthermore, these predicates are to be interpreted in the same way in T and $F\delta$, for the same reasons. Also, T and $F\delta$ share the same truth-functional connectives, which are also to be interpreted in the same way. The only expressions that T and $F\delta$ differ about are 'δ' and '$\exists$', and the Nihilist will say that ϕ_δ is true in exactly the cases where we say that $\phi_\exists$ is true. So, by $(*)$, 'δ' in $F\delta$ is interpreted the same way as '$\exists$' is in T.

We finish the argument with the following observations. We know that '$\exists$x' in T is interpreted as 'there is something that is an x such that…'. So the appeal to $(*)$ in Step Three tells us that 'δx' in $F\delta$ must also be interpreted as 'there is something that is an x such that …'. But, by the construction of Step Two, we know that 'δx' is interpreted in $F\delta$ as '$\Delta\lambda$x' from $F\lambda$. And we also know that 'λx' in $F\lambda$ is interpreted as 'is an x such that…'. So 'Δ' in $F\lambda$ must be interpreted as 'there is something that…'. But by the construction of Step One, 'Δ' in F has the same interpretation as 'Δ' in $F\lambda$; thus, 'Δ' in F is interpreted as 'there is something that…'. Hence, F is not ontologically innocent after all; its supposedly innocent expression 'Δ' is a quantifier proper in disguise.

[35] More precisely, sentences of the form $\ulcorner \delta x \phi \urcorner$ are to be interpreted as $\ulcorner \Delta \lambda x \phi \urcorner$.

6.4 *An Objection*

It is tempting to think that the $(*)$ Argument can't be right simply because, if it were, it would prove too much. The main idea runs something like this:

> We all agreed back in section 5 that when we attach 'Δ' to a one-placed predicate A we got an expression that meant $\ulcorner$It is A-ing$\urcorner$. And we all agreed that $\ulcorner$It is A-ing$\urcorner$ did not mean, and did not entail, $\ulcorner\exists x A x\urcorner$. But if the $(*)$ Argument were right, it would show that we were mistaken to even agree to this much — it would show that even the feature-placing language using only one-placed predicates was really quantificational all along. And this can't be right: surely we could use sentences like 'It is catting' and 'It is treeing' without thereby saying that there is a cat or that there is a tree!

We ought to agree with the objection that it is at least in principle possible for there to be an ontologically innocent, one-placed-predicate-only language like the one discussed in section 5.[36] If an appeal to $(*)$ shows that such a language is impossible, then we ought to reject $(*)$ and the anti-functorese argument given above.

But $(*)$ doesn't entail this impossibility. Suppose we came across a tribe of 'feature-placers' that spoke just such a language. That is, they had all of the same one-placed predicates that our target language T does, predicate-functors '&' and '$\sim$', and an expression 'Δ_{FP}' that attaches to predicates to create sentences.

By mimicking the steps gone through above, we can transform their language into a similar one that has the same predicates, no predicate functors, and a variable-binding operator 'δ_{FP}' that means '$\Delta_{FP}\lambda$'. And then we can compare this new language to the fragment T_1 of our first-order target language that uses only one-placed predicates to see whether or not the two satisfy the antecedent of $(*)$.

In order for both languages to satisfy this antecedent, the tribe must use 'δ_{FP}' in just the same way we use '$\exists$'. But recall from section 4.1.2 that 'use the same way' must be understood dispositionally: it's not enough that they *in fact* use 'δ_{FP}' whenever we use '$\exists$'. For any counterfactual situation C, they must be disposed to apply 'δ_{FP}' in C exactly when we are disposed to apply '$\exists$' in C.

We can ask, in particular, how the tribe is linguistically disposed to react to circumstances in which their language is enriched with all of the multiple-placed predicates that we have in our target language. That is, we can ask how the

[36] At least, we ought to agree insofar as we are not troubled by, or are setting aside, worries that (a) we can only make sense of the feature-placing languages on the model of weather sentences and (b) weather sentences are covert predicates of places as discussed in note 27.

tribe is disposed to extend their language to one with many-placed predicates. And they might be disposed to do it in a number of ways. The tribe *might* be disposed to extend their language in the predicate-functorese way, letting 'Δ_{FP}' turn n-placed predicates into $n-1$-placed predicates. If this is how the tribe is disposed, then $(*)$ does indeed say that 'Δ_{FP}' in their mouths means 'there is'.

But the tribe might not be thus disposed. They may, for instance, be disposed to extend their feature-placing language to many-placed predicates in the way to be suggested below, in section 7. Or they may have some other dispositions entirely, or have no such dispositions at all. And if they are disposed to extend their language in a way that doesn't turn it into functorese, then $(*)$ gives us no reason to think that 'Δ_{FP}' in their mouths means 'there is'.

Since it is in principle possible for there to be tribes like the one imagined with these sorts of dispositions, it is in principle possible for there to be an ontologically innocent feature-placing language like the one described in section 5. The argument does not prove too much after all.

But I anticipate a residual feeling of unease, along the lines of:

> Isn't it just clear that Predicate Functorese is the natural extension of the innocent feature-placing language from section 5, and that its expression 'Δ' is the ontologically innocent multi-placed extension of 'It is . . . -ing'?

In reply: no, it isn't clear at all. First, it is not clear that the functorese extension is the *natural* way to extend 'It is . . . -ing'; perhaps the extension to be discussed below is more natural. But even if it were the natural extension, this need not make it ontologically innocent. A number of philosophers have thought that ontologically guilty expressions *naturally emerge* out of an attempt to extend a feature-placing language to deal with troublesome cases. We start out saying things like 'it is catting over here' and 'it is dogging over there', but then run into various kinds of troubles expressing everything we want to express. For instance, we get into trouble deciding whether placed features 'go together' or not (Quine 1992, Evans 1975), or how features placed yesterday relate to features placed today (Evans 1975, Strawson 1954, 1963), etc. Then we extend our feature-placing language by adding some pegs to place these features on; we then know whether or not features go together, either right now or over time, based on whether they're on the same peg or not.

There is particular reason to think that something like this happens when we extend the one-placed version of the feature-placing language to full Functorese. We can think of 'placing features' as throwing rubber bands onto a peg-free board. If we say 'It is electroning', we throw the 'electron' rubber band on the

board; if we say 'It is protoning', we throw the proton rubber band on the board; and so on.

The other predicate functors let us make complex rubber bands out of simpler ones. But there is real conceptual difficulty understanding the Functorese Nihilist's preferred extension of 'Δ' to relational predicates. What have we done when we say 'Δ(orbits)'? We have somehow thrown part of the 'orbits' rubber band down on the board while keeping the other part up. But what are we going to do with the part that we've kept off the board? Suppose we prefix it with '$\sim\Delta\sim$'. Intuitively, this tells us that, for any other place where we could throw a rubber band down, we must to make sure the other half of *this* (kind of) rubber band gets thrown there also. But now it no longer matters simply *that* thus-and-so a rubber band has been thrown on the board — it also matters *where* it's been thrown, and where it could be thrown, too. In other words, certain locations of the board now matter. Certain locations on the board have to count as possible parking places for parts of polyadic rubber-bands, and it becomes significant when parts of two different rubber bands land on just one of these special locations.

Once we've gone this far, we've all but introduced pegs. The *point* of using pegs to represent things in a pegboard-and-rubber-band model of reality is that they make certain locations of the board special. They make those locations potential landing sites for parts of rubber bands. (We use *pegs* to mark these locations for practical reasons: they keep the rubber bands from sliding around.) But since pegs represent objects, this is just so say that *objects* are special landing-sites for parts of features. Once it starts mattering where one part of a rubber band has been stuck, we've smuggled in an ontology. So, even if the Functorese 'Δ' is in some sense a natural extension of the one-placed feature-placing language, there is good reason to think it is an extension that introduces ontology — and so good reason to think that the $(*)$ Argument was right all along.

7 Putting the Relations Inside the Functor

Let's recap. The above observations suggest that the Ontological Nihilist got into trouble by understanding 'Δ' as a functor that turns n-placed predicates into $n-1$-placed predicates. So let's go back to that point and try something else.

Instead of turning many-placed predicates into fewer-placed predicates, we could let it turn many-placed predicates into *sentences*. Just as 'It is raining' is understood as saying that rain is going on, and 'Δ(proton)' is understood as saying that protoning is going on, we can understand 'Δ(orbits)' as saying that

orbiting is going on and 'Δ(repels)' as saying that repelling is going on.

Saying that orbiting is going on will be the Nihilist's way of paraphrasing our claim that something is orbiting something else. Thus for any n-placed predicate R, $\ulcorner\Delta(R)\urcorner$ will be the Nihilist paraphrase of $\ulcorner\exists x_1 \ldots \exists x_n(R(x_1, \ldots, x_n))\urcorner$.

As before, we need to deal with more complex expressions, such as

(47) $\exists$x$\exists$y(x is an electron & y is a proton & x orbits y),

(48) $\exists$x$\exists$y$\exists$z(x orbits y & y orbits z),

(49) $\exists$x$\exists$y$\exists$z(x orbits y & x orbits z),

and so on. We can make considerable headway on this by helping ourselves to the predicate functors '$\sim$', '&', 'Inv', '$\mathbb{I}\mathrm{nv}$', and 'Pad' from section 6.1. (After all, it was the interpretation of 'Δ', rather than these five functors, that gave the Nihilist troubles in the previous section; with 'Δ' re-interpreted, the Nihilist may now return to these faithful friends.) Then we can find complex predicates equivalent to the embedded open sentences, and paraphrase (47)–(49) as

(50) Δ(electron & Pad(Proton) & orbits),

(51) Δ(orbits & Pad(orbits)),

(52) $\Delta(\mathbb{I}\mathrm{nv}(\textsc{Pad}(\text{orbits}) \;\&\; \textsc{Inv}(\text{orbits})))$,

respectively.[37]

So long as the only sentences the Nihilist wants to paraphrase are of the form, or equivalent to sentences of the form,

(53) $\exists x_1 \ldots \exists x_n(A(x_1, \ldots, x_n))$,

this will do fine. But how will he paraphrase, for instance, 'Something orbits nothing'? He can use 'Δ', plus the predicate functors, to paraphrase any sentence beginning with a block of existential quantifiers. But 'Something orbits nothing' isn't this kind of sentence. It begins with a quantificational block like this:

(54) $\exists$x$\sim\exists$y...

[37] A different option involves complicating the 'Δ'-functor, giving it extra 'slots' for more predicates and paraphrasing (47) as Δ(electron, proton | orbits). The idea here is that the predicates on the left side of the '|' indicate unary features to be placed, and those on the right side indicate many-placed features to be placed 'in between' the unary features, as it were. But it is not clear how to extend this to more complex cases; see Sider and Hawthorne (2003) for a version of this proposal and a discussion of some of the difficulties involved.

And the current proposal has nothing to say about sentences of this sort.

If we could prefix 'Δ' to n-placed predicates to get new predicates of a smaller adicy, we could paraphrase 'something orbits nothing' as '$\Delta{\sim}\Delta$(orbits)'. But down that path lies predicate functorese and, as we saw, Quiet Nihilism. So that path must be avoided. And no other path presents itself; there is nothing left for it but to introduce a new expression, say 'Σ', that the Nihilist will use whenever we ontologically-minded folk would begin a sentence with a block of quantifiers of the form (54).

The Nihilist won't be able to stop at 'Σ', either. Consider the following two sentences:

(55) $\exists x{\sim}\exists y\exists z$(x attracts y & x repels z)

(56) $\exists x\exists y{\sim}\exists z$(x attracts y & x repels z)

The first of these says that something neither attracts nor repels anything else; the second says that something attracts at least one thing but repels nothing. The Nihilist ought to be able to distinguish cases in which it is good to say one of these but not the other. But he cannot paraphrase either of these sentences with 'Δ' or with 'Σ'.[38]

We can mix negations into a block of quantifiers in indefinitely many ways, so the Nihilist will need an indefinitely large stock of primitive expressions in order to paraphrase away all of these sentences. So this Nihilist paraphrase strategy is already committed to one of the costs of Propositional Nihilism noted above: an exploded ideology.

This proposal is also susceptible to Propositional Nihilism's other difficulties: inferential unsystematicity and radical holism.

On inferential unsystematicity: note that his indefinitely many expressions will each be associated with inferences of their own type. And these inferences will resist any explanation, for the expressions 'Δ', 'Σ', and so on are for him a semantic black box — he has nothing to say about them except that, when attached to predicates of a certain sort, they produce sentences fit for certain sorts of paraphrases. But devoid of any further explanation of what these expressions mean, he has no resources for explaining the inferences they participate in.

On holism: consider the sentence

(57) $\exists x\exists y\exists z$(x attracts y & y repels z).

[38]He might perhaps decide that, when prefixed to three-placed predicates, Σ will act in a way so as to paraphrase one of these two; but he then still needs a new expression to attach to the predicate in order to paraphrase the other.

We noted in section 4.2.3 that we ontologically minded folk can think of the more global fact expressed by (57) as being somehow 'built up' out of the fact that an x attracts a y, the fact that a y repels a z, and the fact that the y being attracted in the first instance is the same y as the one doing the repelling in the second.

But a Nihilist who paraphrases (57) as

(58) $\Delta(\text{attracts} \;\&\; \textsc{Pad}(\text{repels}))$

thinks of this fact as essentially 'placing' a complex feature in reality — of deploying, in a peg-free way, a complex rubber-band of the shape in figure 7. But, although we make this complex feature by gluing together the 'attracts'

Figure 7: A Rubber Band for (58)

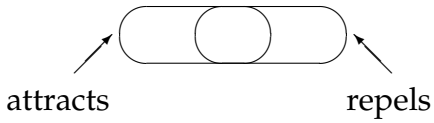

and 'repels' rubber bands, we cannot think of the deployment of this complex rubber-band structure as being somehow 'built up' out of the deployment of the 'attracts' and 'repels' rubber bands. The mere fact that these two rubber bands have been thrown on the board isn't enough to guarantee that they overlap in the required way. The fact that the 'attracts' rubber band has been deployed corresponds to our observation that an x attracts a y, and the fact that the 'repels' rubber band has been deployed corresponds to the fact that a y repels a z. But in order to 'build up' the right complex fact, the Nihilist will also need a fact that corresponds to 'the y being attracted in the first instance is the y doing the repelling in the second'. But there is no Nihilistically acceptable, object-free way to make sense of that claim.[39] That is, there is no way to *identify* the different parts of the 'attracts' and 'repels' rubber bands to say that they hook together in the right way — unless we plunk a peg down onto the board and say that the two rubber bands are each attached *to the same peg*. But this is precisely what a Nihilist cannot say.

[39]Simply throwing two rubber bands on the board so that they look like figure 7 cannot do the trick — if we said it did, we would make 'places' on the board important and thereby smuggle ontology back into the picture, as discussed at the end of the last section.

8 Conclusion

We have not, of course, canvassed every way an Ontological Nihilist might try to paraphrase away our target language. But, on reflection, it looks as though the considerations adduced here will extend to any Nihilist proposal. And if this is right, Ontological Nihilism faces a dilemma: if it is to be viable, avoiding the ills of Quiet Nihilism, it must embrace a particularly holistic picture of reality with an attendant bloated ideology and brute entailments.

This gives us some reason to reject Ontological Nihilism. But this should come as little surprise. After all, it is only metaphysicians, and a very small number of them at that, who would have ever suspected Ontological Nihilism of being true in the first place.

When we see the evils we must embrace in order to make Ontological Nihilism work, though, we gain a better appreciation for our ontology. Because we have one of those, we can think of the global structure of reality as being built up out of more local structures. It turns out that having an ontology — a set of 'pegs', of *things* — is crucial for this sort of bottom-up picture. For it is by *identifying* things across different local structures that we can build up more global structures. By picking out which things in *this* local structure are identical to which things in *that* one, we have a way to link those two structures together to come up with a more global one. And it is by thinking of the world ontologically that we can understand the validity of certain inferences: they are valid because the pegboard structure described by one claim fits or doesn't fit in fairly straightforward ways with the structures described by the other. Thinking of the world in an ontological way provides us with the resources to offer powerful systematic explanations of a wide variety of pervasive facts. That, perhaps, is part of why ontology has been important to philosophy all along.

References

Bacon, John (1985). "The Completeness of a Predicate-Functor Logic." *The Journal of Philosophical Logic* 2: 323–339.

Boolos, George, John P. Burgess and Richard C. Jeffrey (2002). *Computability and Logic*. 4th edition. Cambridge University Press.

Burgess, John P. and Gideon Rosen (1997). *A Subject With No Object*. Oxford: Oxford University Press.

Cartwright, Richard (1954). "Ontology and the Theory of Meaning." *Philosophy of Science* 21(4): 316–325. Reprinted in Cartwright (1987: 1–12).

— (1987). *Philosophical Essays*. Cambridge, Mass.: MIT Press.

Chalmers, David (2008). "Ontological Anti-Realism." In Chalmers et al. (2008).

Chalmers, David, David Manley and Ryan Wasserman (eds.) (2008). *Metametaphysics*. Oxford: Oxford University Press.

Davidson, Donald (1963). "Actions, Reasons, and Causes." *The Journal of Philosophy* 60(23): 685–700.

Evans, Gareth (1975). "Identity and Predication." *The Journal of Philosophy* 72(13): 343–363.

Field, Hartry (1980). *Science Without Numbers*. Malden, Mass.: Blackwell.

Gamut, L. T. F. (1991). *Logic, Language, and Meaning*, volume 2, Intensional Logic and Logical Grammar. Chicago: University of Chicago Press.

Goodman, Nelson (1979/1983). *Fact, Fiction, and Forecast*. 4th edition. Cambridge, Mass.: Harvard University Press.

Hawthorne, John (2006). *Metaphysical Essays*. Oxford: Oxford University Press.

Hawthorne, John (O'Leary-) and Andrew Cortens (1995). "Towards Ontological Nihilism." *Philosophical Studies* 79(2): 143–165.

Hindley, J. Roger and Jonathan P. Seldin (1986). *Introduction to Combinators and Lambda-Calculus*. Cambridge: Cambridge University Press.

Hirsch, Eli (2002). "Quantifier Variance and Realism." *Philosophical Issues* 12, *Realism and Relativism*: 51–73.

— (2005). "Physical-Object Ontology, Verbal Disputes, and Commonsense." *Philosophy and Phenomenological Research* 70(1): 67–97.

— (2007). "Ontological Arguments: Interpretive Charity and Quantifier Variance." In Dean Zimmerman, John Hawthorne and Theodore Sider (eds.), *Contemporary Debates in Metaphysics*, 367–381. Malden, Mass.: Blackwell.

Kaplan, David (1989). "Demonstratives." In Joseph Almog, John Perry and Howard Wettstein (eds.), *Themes from Kaplan*, 481–563. New York: Oxford University Press.

Kuhn, Steven T. (1983). "An Axiomatization of Predicate Functor Logic." *Notre Dame Journal of Formal Logic* 24: 233–241.

Lewis, David (1983*a*). "New Work for a Theory of Universals." *The Australasian Journal of Philosophy* 61: 343–377. Reprinted in Lewis (1999): 8–55.

— (1983*b*). *Philosophical Papers*, volume 1. Oxford: Oxford University Press.

— (1984). "Putnam's Paradox." *The Australasian Journal of Philosophy* 62: 221–236. Reprinted in Lewis (1999): 56–60.

— (1999). *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.

Lewis, David and Stephanie Lewis (1970). "Holes." *The Australasian Journal of Philosophy* 48: 206–212. Reprinted in Lewis (1983*b*): 3–9.

Merricks, Trenton (2001). *Objects and Persons*. Oxford: Clarendon Press.

Quine, Willard van Orman (1948). "On What There Is." *The Review of Metaphysics* 2: 21–38. Reprinted in Quine (1980): 1–19.

— (1960*a*). "Variables Explained Away." *Proceedings of the American Philosophical Society* 104: 343–347. Reprinted in Quine (1966): 227–235.

— (1960*b*). *Word and Object*. MIT Press.

— (1966). *Selected Logic Papers*. Cambridge, Mass.: Harvard University Press.

— (1971). "Algebraic Logic and Predicate Functors." In Richard Rudner and Israel Scheffler (eds.), *Logic and Art: Essays in Honor of Nelson Goodman*, 214–238. New York: Bobbs-Merrill. Revised and reprinted in Quine (1976): 238–307.

— (1976). *The Ways of Paradox and Other Essays*. 2nd edition. Cambridge, Mass.: Harvard University Press.

— (1980). *From a Logical Point of View*. Cambridge, Mass.: Harvard University Press.

— (1992). *Pursuit of Truth*. 2nd edition. Cambridge, Mass.: Harvard University Press.

Rayo, Agustín (2007). "Ontological Commitment." *Philosophy Compass* 2(3): 428–444.

— (MS). "On Specifying Truth-Conditions." Unpublished manuscript.

Richard, Mark (1998). "Commitment." *Philosophical Perspectives* 12: 225–281.

Schaffer, Jonathan (2008). "On What Grounds What." In Chalmers et al. (2008).

Schwarz, Wolfgang (MS). "Lewisian Meaning without Naturalness." Unpublished manuscript. Available online at <www.umsu.de/words/magnetism.pdf>.

Seppänen, Aimo (2002). "On Analysing the Pronoun IT." *English Studies* 83(5): 442–462.

Sider, Theodore (2001*a*). "Criteria of Personal Identity and the Limits of Conceptual Analysis." *Philosophical Perspectives* 15: 189–209.

— (2001*b*). *Four Dimensionalism*. Oxford University Press.

— (forthcoming). "Monism and Statespace Structure." In Robin LePoidevin (ed.), *Being: Developments in Contemporary Metaphysics*. Cambridge University Press.

Sider, Theodore and John Hawthorne (2003). "Locations." *Philosophical Topics* 53–76. Reprinted in Hawthorne 2006: 31–52.

Stalnaker, Robert (2004). "Lewis on Intentionality." *Australasian Journal of Philosophy* 82: 199–212.

Stanley, Jason (2002). "Context and Logical Form." *Linguistics and Philosophy* 23(4): 391–434.

Strawson, P. F. (1954). "Particular and General." *Proceedings of the Aristotelian Society* 54: 233–260.

— (1963). *Individuals*. London: Methuen.

Tarski, Alfred (1959). "What is Elementary Geometry?" In Leon Henkin, Patrick Suppes and Alfred Tarski (eds.), *The Axiomatic Method, with Special Reference to Geometry and Physics*. Amsterdam: North-Holland.

van Inwagen, Peter (1990). *Material Beings*. Ithaca, N.Y.: Cornell University Press.

— (1998). "Meta-Ontology." *Erkenntnis* 38: 223–250. Reprinted in van Inwagen (2001: 13–31).

— (2001). *Ontololgy, Identity, and Modality*. Cambridge: Cambridge University Press.

Weatherson, Brian (2003). "What Good are Counterexamples?" *Philosophical Studies* 115: 1–31.