Ontological Nihilism*

Jason Turner

December 8, 2009

Ontological Nihilism is the radical-sounding thesis that there is nothing at all. Almost nobody believes it. 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 are some things.

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 (§1), I present (§2) and clarify (§3) a crucial challenge for it. I show what is wrong with two less plausible attempts to meet this challenge (§4), and then argue that the proposal thought by many to be much more promising (§5) succumbs to the same problems as the less promising attempts (§§6 and 7). It turns out that, in order to make Ontological Nihilism viable without smuggling an ontology in through the back door, we need a bloated ideology, indefinitely many 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 kinds 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,

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 think he endorses 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.

But it seems undeniable that the richness and structure of experience is somehow accounted for by structure in the world. And it seems reasonable that our ordinary beliefs, arising as they do from experience, 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 thinks the blob-less blob of reality has a rich structure which our ordinary beliefs (usually) track. But he insists that this structure does 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 describe and differentiate these things.

Such descriptions of the world 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 rubber bands. The pegs represent things, and the rubber bands represent ways these things are and are interrelated.

To say 'Bertrand thought about language', for instance, is to hang the *thought about language* rubber band on the peg labeled 'Bertrand'. And to say 'Some

---

[1]Ontological Nihilist, that is. For stylistic reasons, I'll often 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.

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) has an *admires* rubber band stretching from it to each of the pegs with the *didn't notice the inconsistency in Basic Law V* band on.[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 thinks instead that it has some other sort of structure which our experience latches on to.

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 Nihilist 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 the tapestry of experiences.

2   THE NEED FOR PARAPHRASE

2.1   The Challenge

Our Nihilist 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.

A similar challenge faces, e.g., the more conservative perforational nihilist, who claims that there are no holes. At first blush, the perforational nihilist's claim seems incredible, implying that we are radically deceived about the nature of the world. Suppose, for instance, that you just crossed a bridge like that in figure 1. If someone asks you why you crossed on the left, you will likely 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. Perforational nihilists insist that there are no holes; since (3) seems to entail

---

[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
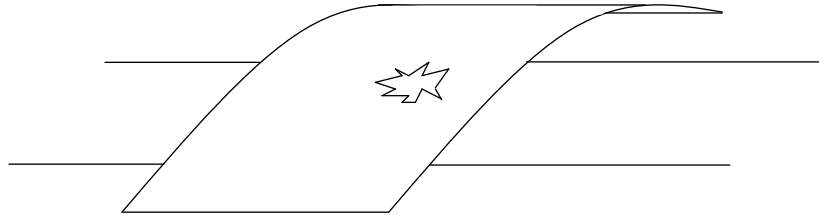
that there are, they should reject (3). So it seems they must say that you were radically mistaken about the nature of the bridge — and that crossing on the right-hand side would have been fine.

Of course, perforational nihilists want to say neither thing. They grant that some bridges are indeed shaped like the one in the diagram, and that you should not walk on the right-hand-sides of such bridges. 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, that you were getting at with (3) and which explains your reluctance to cross on the right. They object only to the idea that this important fact involves a special kind of entity called a 'hole'. On their view, crossing on the right-hand side of the bridge is not bad in virtue of its relation to one of reality's pegs with a 'hole' rubber-band on it. 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 good 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?

A 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 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 he *also* says that, despite (3)'s usefulness in this regard, it is nonetheless defective. It misrepresents the real metaphysics of the situation as involving a *hole*, and they don't. And we, in response to his denial, want to know what real, hole-free metaphysical facts of the matter make (3) useful despite this defectiveness.[5]

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

---

[4]Cf. Sider (2008: 132).

[5]I am being deliberately cagey about just what this 'metaphysical defectiveness' amounts to. Perhaps (3) is simply *false,* but can be used to convey true information in the neighborhood (cf. Merricks 2001: ch. 7), or perhaps (3) is true in ordinary contexts but false, in 'serious' philosophical contexts (cf. van Inwagen 1990: ch. 10). We need take no stand on that here.

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 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 talk of holes for talk of perforated objects. Then, whenever we would say something of the form

(H) There is a hole in _____,

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

(P) _____ is perforated.[6]

The perforational nihilist thus tells us what he thinks this hole-free world *is* like — he thinks it is filled with perforated things — and how apparent talk of holes is really getting at these perforation facts.[7]

The term 'paraphrase scheme', may call to mind a certain philosophy of language according to which (4) cannot be an analysis of, or synonymous with, ordinary uses of (3) (see, e.g., Quine 1960*b*: 250). And it may suggest that the proposed scheme must meet certain conditions: that it be finitely specifiable, for instance, or intelligible to anyone who understands the paraphrased claims.

But let's not bind the nihilist to the commitments of any particular philosophy of language. 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 he thinks it is getting at — regardless of whether it is finitely stateable, easily recognizable by anyone capable of talking about holes, etc.[8]

A question: Must similar hole-sentences receive similar paraphrases? The proposal above is fairly systematic, but how poorly should we view a nihilist who offers a more gerrymandered scheme, paraphrasing some sentences of the form (H) in one way, and paraphrasing others in another?

---

[6]We will need an even more general scheme to deal with e.g. 'There are six holes in this cracker'; see Lewis and Lewis (1970) for a fuller treatment and critical discussion.

[7]A paraphrase scheme might not be the only way a perforational nihilist can give us this account. Perhaps instead he could explain why hole talk is a useful fiction using a non-paraphrase methodology similar to that found in Field 1980. But I stick to paraphrase schemes in the text because I cannot see how to even begin a Fieldian treatment of Ontological Nihilism.

[8]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'.

'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 look askance at strategies which paraphrase them differently. Nobody thought the facts gotten at by 'There is a hole in the bridge' and 'There is a hole in the argument' were similar in the first place. But we do think a bridge's having a hole is somehow similar to 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 the nihilist must do to convince us of our error.[9]

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 also get at important facts.[10] 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 also needs a paraphrase scheme: a way to trade in

---

[9]There are subtleties. Perhaps the perforational nihilist uses one recipe to paraphrase 'There is a hole in the bridge' as ϕ and uses a very different recipe to paraprhase 'There is a hole in the door' as ψ. Nonetheless, if ϕ and ψ are themselves clearly very similar facts, then the differences in the recipes used to get to them from the original hole-sentences do not mean the nihilist is denying any intuitive similarity.

[10]Of course, as a Nihilist, he denies that there are any facts at all. But we should think of his 'fact'-talk as merely a useful turn of phrase for explaining 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.

ontological, pegboard-and-rubber-band presupposing claims for 'ontologically innocent' facts — facts which do not entail that there is anything.

We can simplify the Nihilist's task by pretending the *target language* to be paraphrased is that of first-order logic without names (but with identity). This language is generally thought sufficient for talking about ontological structure: its existential quantifier, '∃', 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 also ask the Nihilist to paraphrase only part of our ontology-involving talk. In particular, we make him paraphrase only claims from well-established scientific theory (or, at least, first-order consequences of well-established scientific theory).

We thus make the Nihilist's task both easier and harder. Easier, because we take from his shoulders the burden of deciding which sentences deserve paraphrase. A Nihilist ought not paraphrase *everything* we say: some of what we say just isn't getting at any important fact. (Nihilists need not paraphrase e.g. 'Phlostigon is emitted during combustion'.) 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 focusing on these sentences gives us the right to demand that 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 sentences get at similar facts, then they should 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 Nihilist's paraphrase is to do what he needs it to, it must generate sentences that are 'ontologically innocent'. But what does that mean? And what makes an expression ontologically innocent?

3.1 Ontological Commitment

Some (interpreted) sentences have a feature philosophers are pleased to call 'ontological commitment', and a sentence is 'ontologically innocent' if and only if it carries no ontological commitments. I do not intend to spill more ink over the proper 'criterion of ontological commitment',[11] but I do want to be clear about what I take commitment to consist in.

The core idea is that we somehow manage to linguistically convey that the world has a certain ontological structure — that there are some things of a certain kind *K*, 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.[12] But the sentence has to be the right one — I cannot commit myself to unicorns just with any old sentence. 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.[13]

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.

3.2 Variable Binding and Quantification Proper

The Nihilist needs to find an ontologically innocent language with which to paraphrase the ontologically guilty target. If we can see what linguistic resources give a language its guilt we will know what resources the Nihilist must avoid.

We all learned at Quine's knee that, in first-order languages, the existential

---

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

[12]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.

[13]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, 2008) suggests instead that 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.

quantifier '∃' makes for ontological guilt. So let's get to know this expression a little better. We can learn about it by considering the two roles it plays: both as a *variable binder* and as a *quantifier proper*.

What does it mean to say that '∃' is a 'variable binder'? In first-order languages, we can take a sentence open in a variable 'x' and prefix it with '∃x' or '∀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. Variable binding is this this turning of open expressions into closed (or less open) ones.

Quantifiers 'quantify proper' by saying *how many* pegs are arranged the way its postfixed formula says. '∃x...', for instance, says that least one peg is ...; '∀x...' says that every peg is .... If we have more sophisticated quantifiers than these, we can also say, for instance, that infinitely many pegs are ..., or that most pegs that are ... are also _ _ _.

Is first-order logic ontologically guilty because '∃' binds variables, or because it quantifies proper? To sharpen the question, we divide '∃'s dual burdens between two different expressions. This is what *lambda-abstraction languages* do.[14] 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 '∃x(... x ...)' as a sort of quasi-English expression, meaning

There is something that is an x such that ... x ...

Likewise, '∀x(... x ...)' can be translated as

Everything is an x such that ... x ... [15]

But we could do the same work with separate expressions: one which means 'there is something', one which means 'everything', and a third which means 'is an x such that ... x ...'. The third expression would be a variable-binder; the first two, quantifiers proper.

In a lambda-abstraction language, the work of that third expression is done by a predicate-forming operator, 'λ', that combines with a variable and an open expression to make a predicate: where ϕ is an open expression, $\ulcorner \lambda x \phi \urcorner$ means

---

[14]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.

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

⌜is an x such that ϕ⌝. Such languages also have quantifiers proper, '∃$_p$' and '∀$_p$', which mean 'there is something that' and 'everything', respectively.

Lambda-abstraction languages are no less guilty than 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 to be exonerated.

That the variable-binder is innocent: 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: '∃$_p$' means *there is something*. We commit ourselves ontologically when we say that there is something which is some way or another, and '∃$_p$' is the expression we use to say this.

We can better 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. We might think of (8) as picking out a single rubber-band structure — one involving three rubber-bands glued together, as in figure 2. But 'λ' does not fasten that rubber band to any

Figure 2: A Rubber Band for (8)
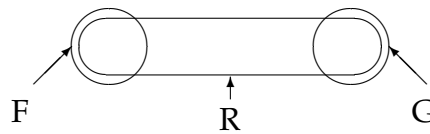
pegs. If you want to say that an *F* thing *R*s a *G* thing — if you want to stretch this complex rubber band between a pair of pegs — you've got to find a way to plunk down some pegs to stretch it between.

The 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:

> The target language is easily paraphrased. 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. 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.

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 is ontologically innocent and can be uniformly replaced for 'there are' to turn falsehoods into truths.

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

4.1.1 A Warm-up Exercise

Imagine meeting a man — Eustace — who, to your surprise, tells you nothing is blue. Incredulous, you point at something you had always thought of as blue and ask: 'What color is that?' Eustace responds, 'Eulb'.

Bewildered, you ask further: 'Is eulb a color?' He says, 'Yes'. He tells you that eulb is a cool color, the color of the sky, lying on the spectrum between red and green. You ask what color complements eulb, and 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 to assert using the word 'blue', but happily assents to the sentence that results from it by a systematic replacement of 'blue' for 'eulb'.

You will soon think that when Eustace says 'eulb', he means *blue* — 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' for it, 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 certain shades. You certainly aren't disagreeing about anything's *color*.

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

> Eustace 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 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'.[16]

Unfortunately, both the 'green'- and the 'grue'-speaking communities are annihilated by an asteroid strike before the future date. 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'. But their linguistic dispositions don't guarantee that they ever in fact say such things; 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 across a wide range of possible circumstances; the 'green'-speakers are disposed to call green things 'green' in circumstances

---

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

after the envisaged future date, and the 'grue'-speakers are not.[17]

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. Imagine two communities which differ only in one's using '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.[18]

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 Eustace insisted vehemently that 'eulb' did *not* mean the same as 'blue' in our mouths. It is then as though he stipulates the following:

---

[17]One caveat: we should *not* be concerned with the speakers' dispositions to utter sentences containing both disputed words. 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 these assertions themselves. Cf. §4.1.3 below.

[18]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 isomorphic contexts (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, context-insensitive class of languages, so we can ignore these details in what follows. Thanks here to Ted Sider.

(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 constraints are jointly satisfiable. If there is a property B that applies to exactly those things that ordinary people would call 'blue' under ordinary circumstances, an ideal interpreter will be pressured to interpret 'blue' as meaning B. When Eustace 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 Eustace that sort of control over the interpretation of everyone else's 'blue' — any reasonable principle of charity will have her make Eustace, rather than the rest of us, speak falsely. So, insofar as she ensures (S1)'s satisfaction, she will have good reason to leave (S2) unsatisfied.

What goes for Eustace goes for the Quiet Nihilist: having insisted that sentences such as (10) are true in exactly those situations where we assert (9), he cannot also insist that 'schm$\exists$' has a different meaning than '$\exists$'. Insofar as we grant his first insistence, we have good reason to think that 'schm$\exists$' means (is intensionally equivalent to) 'there is' after all.

4.1.4 (∗) and Charity Arguments

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 (∗) argument above is simply that the charity arguments 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

16

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*. 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 A and B use the words α and β in the same way, against a background of other, shared words with the same interpretations, then α and β must be interpreted the same way, too. No naturalness constraint or other word-level constraint can conflict with (∗): the shared words have the same interpretations and hence 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's language 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:

---

[19]This might very well suggest that the 'grue'-speakers from §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.

[20]The situation is even better: if the rest of the respective languages' semantic backgrounds are the same, a lot of *other* interpretative issues between both parties must 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 that could pressure us to assign α a meaning intensionally inequivalent to β.

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 '∼', '&', etc. A sentence such as

(11) There is one electron in a hydrogen atom,

will be paraphrased into an atomic sentence — 'P', say — 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.

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 *primitive* — that is, meaningful and undefined. But no matter how many 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 that he shouldn't. 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

$$R$$

$$\vdots$$

Since each of these atomic propositions constitutes a primitive bit of ideology, the Nihilist needs indefinitely many primitives.

4.2.2 Lack of Systematicity

Second, the view is *inferentially unsystematic* — unable to explain or systematize large swaths of inferences. Consider:

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

(14) There are some electrons orbiting a proton.

Suppose the Propositional Nihilist paraphrases these as 'A' and 'B', respectively. Since (13) entails (14), 'A' ought to entail 'B'.

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, the Nihilist cannot explain why the inference from 'A' to 'B' is valid although the inference from 'A' to, say, 'C', the paraphrase of

(15) Some neutron is in region R,

is not. 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.[21]

---

[21]Andrew Cortens and Dean Zimmerman have suggested a variant on Propositional Nihilism which first makes use of the Tractarian reduction of quantifiers to (possibly infinite) conjunctions and disjunctions (Wittgenstein 1921: §§5.52–5.5262) and then trades in the residual atomic predications for 'P', 'Q', 'R', etc. Here the quantificational inferences reduce to truth-functional ones, and so the inferential unsystematicity charge will not stick. If the proposal can be made to work — which will involve making Nihilistically acceptable sense of the Tractarian insistence that every object have exactly one name — I suspect it is the best Propositional Nihilistic view on the market (although it, too, is ideologically extravagant). But I lack the space to consider the view further here.

4.2.3 Holism

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

   We ontologically-minded folk think, more or less, that there are only a limited number of ways things can be and can be related, and the way reality is *in toto* is fixed by how each thing is and relates 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
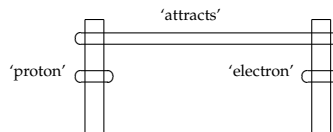
Figure 3: The Rubber Band Structure of (16)
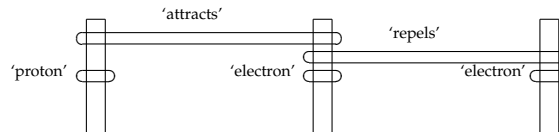
it is easy to see how this complex pegboard-and-rubber band structure is built up out of two simpler structures, one that involves the leftmost and center pegs, and one that involves the center and rightmost pegs. In a certain way, the fact expressed by (16) is built up out of 'smaller' facts — in particular, the facts expressed by

 (17) A proton attracts an electron (figure 4),
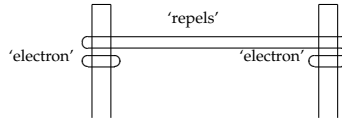
Figure 4: The Rubber Band Structure of (17)

and

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

Figure 5: The Rubber Band Structure of (18)

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

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, a simple structureless I-know-not-what. But no propositional paraphrase of (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, 'So what? I've bitten bullets enough in my time not to mind three 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. 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 that he has not really incurred all these costs. I can see no way for him to keep holism and ideological bloat off the bill. But he can at least try to quit the charges of inferential unsystematicity. Let's look at two ways he might do this.

Syntactic Unsystematicity is No Big Deal

Here is his first attempt:

> 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.

But he here waves a red herring. §4.2.2's complaint wasn't that the inferences between the atomic sentences were *syntactically indiscernible* — not worn on their syntactic faces, as it were. The complaint 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 language. Hawthorne and Cortens (1995: 151) point out that while the inference

(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 note on the same page) this does not make the inferences 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. (19)'s validity is neither syntactically discernible nor 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

---

[22]Hawthorne and Cortens's original invalid example, 'He ran halfway up the hill; therefore, he ran up the hill' seems to have a different syntactic form than 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'.

atomic claim 'A', and that it is what we were getting at all along with (13), the Propositional Nihilist has nothing left to say. In particular, he has no story about what 'A' *means* that would let him explain why it entails 'B' but not 'C'.

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 has two components: a *content tag* and an *inference tag*. A content tag is a syntactically simple expression, such as a capital letter (with or without subscripts). An *inference tag* is syntactically complex, made up out of various *pseudo-expressions:* pseudo-variables ('x', 'y', 'z', ...), pseudo-predicates ('=', 'P', 'Q', 'R', ...), and pseudo-quantifiers ('∀', '∃'). 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 target-language sentence. Sentences of the Propositional language have the form $\ulcorner P^\phi \urcorner$, where P is a content tag and ϕ an inference tag. But not *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 $\{\phi_1, \phi_2, \ldots\} \vdash \psi$. In this case, since we always paraphrase a target-language sentence into one with an isomorphic inference tag, our inference rule will tell us that an inference of the tag language is valid exactly when it is the paraphrase of a valid first-order instance.

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 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

---

[23]They are thus like the semantically empty 'it' of weather reports; see §5 below.

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?

So if the Propositional Nihilist endorses all the inferences we expect him to, his system will be unsystematic. But he now tells us he doesn't endorse all the inferences we expect him to, saying:

> 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 get at with ϕ cannot be true if the important fact we get at with ψ is false. Anyone who paraphrases ϕ so that it does not entail the paraphrase of ψ effectively says that we are wrong about this relationship: the important fact we get at with ϕ *could* be true even if the important fact we get at with ψ is false.

It is certainly no metaphysical desiderata that every inference ordinary folk make turn out valid. But it's one thing to say that ordinary folk tend to be wrong about certain troublesome inferences, and another thing to say that ordinary folk tend to be wrong about almost every inference they're inclined to make. We make paradigmatically quantificational inferences — 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.

Let's spell this out more carefully. Ontological nihilists of any stripe want to 'save the appearances' — to explain why it is sometimes useful to talk as though there are certain kinds of things even though there aren't. But part of

---

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

saving the appearances is 'saving the practices' — explaining why certain natural transitions involving these kinds of things are so useful.

Consider again the bridge in figure 1. We noted that someone who doesn't believe in holes needs to tell us why you can point to

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

even in the absence of holes, to explain why you crossed on the left instead of the right.

In one sense, to explain an action is to 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 challenge to the perforational nihilist was not to give a descriptive explanation of your bridge-crossing behavior. That's easy: he needs only to point to your beliefs in:

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

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

and your desire not to plummet into the river to give such an explanation.

The challenge was instead for him to give a *justifying* explanation — an explanation that shows why your so crossing was a *good idea*. Consider Hal, who tends to hallucinate 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 1. The descriptive explanation of Hal's crossing the figure 6 bridge on the left points to his (21)- and (22)-like beliefs. If there are holes, we can explain why your bridge-crossing behavior was better motivated than Hal's by pointing out that there *was* a hole in your bridge, so that your (21)- and (22)-like beliefs were true and Hal's false. But, since the perforational nihilist thinks there are no holes, he cannot tell this story about the difference between your behavior and Hal's. His challenge is to tell some other story instead.

The story he in fact tells points not to (3) but to

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

---

[25]See Davidson (1963).

Figure 6: A Perfectly Good Bridge
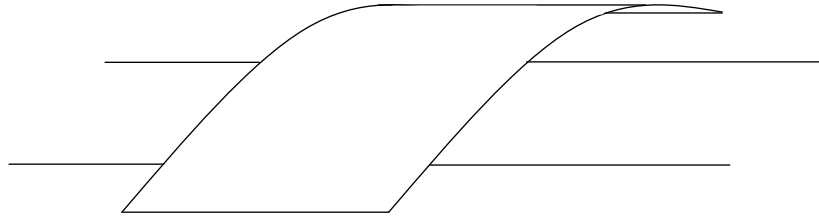
Your (21)- and (22)-like beliefs were true, and your behavior well-motivated, because your bridge was perforated; Hal's counterpart beliefs were false, and behavior ill-motivated, because his bridge was not.

But this move works only if (4) (plus some background assumptions) entails (22). If you could walk over perforated parts of bridges unharmed, the truth of (4) wouldn't be relevant to how you crossed. So the paraphrase scheme 'explains the practices' only if it preserves the validity of certain inferences.

The Propositional Nihilist needs a justifying explanation for our bridge-crossing behavior just as much as the perforational nihilist — more so, since he thinks there are no bridges, either. But he paraphrases (3) and (22) as atomic propositions. So in order to offer the needed explanation, he needs some (presumably brute) inferences between his atomic propositions.

And this is no local phenomenon; quantificational inferences — inferences between sentences the Propositional Nihilist will paraphrase as atomics — are ubiquitous in scientific reasoning as well as everyday life. It is difficult, in fact, to imagine successfully navigating our environment without them. And it is hard to see what, short of the inferences being valid under paraphrase or a cosmic coincidence, could explain why these inferences are useful. The charge of inferential unsystematicity thus stands: if he is to meet the challenge of §2.1, the Propositional Nihilist needs 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 rain, snow, or coldness are manifest — without saying that any particular object is manifesting them. 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 throwing rubber bands on peg-free areas of the board.[26]

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 about particular things. For instance, instead of saying

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

we could say

(24) It is catting.

Just as 'it is raining' says that rain is going on without saying that there is any thing which is raining, (24) says that catting is going on without saying that any particular thing is a cat.

Following Strawson, we call sentences such as (24) *feature-placing sentences*, and if a language only allows sentences (and truth-functional compounds of sentences) of this sort, we call it a *feature-placing language*. The idea is that the Nihilist can paraphrase every apparently quantificational sentence of the target language into one of a feature-placing language and thereby account for all

---

[26]This may not be quite right. The semantics of 'is raining' may make it a predicate of places. Even though the 'it' must be semantically empty (see Seppänen 2002: 445–453), '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). Out of charity towards the Nihilist, though, we ignore these complications here; cf. note 27 below.

the undeniable facts without appealing to any pegboard-like structure. (See Hawthorne and Cortens 1995)[27]

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

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

as

(26) 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:

(27) $\exists x$(x is positively charged & x is negatively charged).

(28) $\exists x$(x is positively charged) & $\exists x$(x is negatively charged).

(28) 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 (27) (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 (27) and (28) respectively as:

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

(30) 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

---

[27]If the predicate-of-places account of the 'it' in 'It is raining' described in note 26 is right, then a Nihilistically acceptable reading of (24) won't be strictly parallel to 'It is raining': (24) will predicate cattingness of *places*, and thus invoke pegboard structure at that level. I think, however, that we can still make sense of the Nihilist's intended place-free use of (24); and even if we can't, we can learn much from pretending we can and seeing how far the Nihilist can push his proposal. So I make no hay over these otherwise problematic linguistic considerations here.

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 'llueve'.) 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.[28] (29) and (30) are equivalent, so they can't respectively paraphrase both (27) and (28).

We do better if we make some logically complex predicates out of the simple predicates *first*, before turning them into 'feature' expressions to be placed. We could then construct a predicate 'is positively charged and negatively charged', and turn *that* into a single feature expression 'is (positive-charge and negative-charge)-ing', distinct from the conjunction of 'is positive-charging' and 'is negative-charging'. Then the Nihilist could paraphrase (27) as

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

equivalent to neither (29) nor (30).[29]

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 '∼': if *P* is a predicate, ⌜∼*P*⌝ is its negation.

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 give the feature-placing language a huge stock of primitive expressions relative to the target language.

---

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

[29]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 (27) 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.

And making all of these expressions primitive obliterates logical relations we might well want to keep. (31), for instance, ought to entail

(32) It is positive-charging.

But if 'is (positive-charge and negative-charge)-ing' and 'is positive-charging' are disparate, semantically simple items, it's hard to see how it could without a brute, necessary connection.

If he wants to avoid bloating his ideology and de-systematizing his inferences, the Nihilist can do better by helping 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*. 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. We can make the semantics and the syntax march in step by letting the predicate-functor combine with predicates to make *sentences*. Suppose we write this functor 'Δ': then for any predicate A, whether simple or complex, ⌜$\Delta(A)$⌝ is the Nihilist's sentence meaning ⌜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

(33) ⌜$\exists x(\ldots x \ldots)$⌝,

where '…x …' is some truth-functional compound of atomic predications ⌜Ax⌝.[30] Say that such truth-functional compounds are in *existential normal form*.

Second, every truth-functional compound of atomic predications of the form ⌜Ax⌝ 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. (⌜Ax & Bx⌝ becomes ⌜$(A\&B)x$⌝, ⌜$\sim Ax$⌝ becomes ⌜$(\sim A)x$⌝, and so on.) Call this the *functor reduction* of the original truth-functional compound.

So we paraphrase a first-order sentence ϕ by first putting ϕ in existential normal form, and then replacing each subsentence of the form (33) with

(34) ⌜$\Delta(P)$⌝,

---

[30]This follows from the fact that a sentence of the form (33) is equivalent to one using only one variable; cf. Boolos et al. (2002: 274–275).

where P is the functor reduction of '...x...'. This gives us feature-placing replacements for each sentence of the target language without any of the costs of Quiet or Propositional Nihilism. 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 don't yet know how to deal with relational predicates. Our best science will endorse relational claims such as

∃x∃y(x repels y),
∃x∃y(x orbits y & x attracts y),
etc.

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; can we extend it to deal with relational predicates? We will need 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.

But in just what sense? What happens to a many-placed predicate when it gets prefixed with 'Δ'? The Nihilist has only two viable options: either say that the prefixing 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. Then 'Δ(repels)', for instance, would be a one-placed predicate — the Nihilist's 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 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, attaching 'Δ' to, say, the predicate 'orbits', makes a new predicate, 'Δ(orbits)'. Since 'orbits' has two places, this new complex

predicate has just one. And, although it is hard to say in any straightforward way 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, we can attach another 'Δ' to it: 'something orbits something' is paraphrased as '$\Delta\Delta$(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 existential 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 trade in quantifiers for 'Δ'-functors to get:

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

(keeping negations interspersed between the 'Δ's).

This will work so long as we can always come up with a complex predicate for each quantifier-free open sentence. The Nihilist already has many resources needed 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, $\ulcorner A \& B \urcorner$ is an *i*-placed predicate, where *i* is the greater of *n* and *m*, so that $\ulcorner(A \& B)x_1, \ldots, x_i\urcorner$ is equivalent to $\ulcorner Ax_1, \ldots, x_n \& Bx_1, \ldots, x_m \urcorner$.[31] Then, for instance, he can turn the open sentence

x is a proton & x orbits *y*

---

[31]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.

into the predicate

> (is a proton & orbits)

and paraphrase

> ∃x∃y(x is a proton & x orbits y)

as

> ΔΔ(is a proton & orbits).

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

(35) ∀x∃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 (35), the natural candidates available for that paraphrase are:

(36) ΔΔ(orbits)
      ∼Δ∼Δ(orbits)
      Δ∼Δ(∼orbits)
      ∼ΔΔ(∼orbits)

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

(37) ∃x∃y(x orbits y)
      ∀x∃y(x orbits y)
      ∃x∀y(x orbits y)
      ∀x∀y(x orbits y)

Since (35) is not equivalent to any of the sentences in (37), it should not be paraphrased as anything in (36).

The problem is perfectly general. We are starting from

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

(perhaps with negations sprinkled through the block of quantifiers) and trying to 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 the variables in this equivalent one-predicate open sentence occur in the same order that they do 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 (35) does, such switching affects meaning. We get problems in precisely these cases.

In (35), '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'. No truth-functional compound of 'orbits' will do the trick. We need something else.

If we had 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 we could paraphrase (35) as

(38) $\sim\!\Delta\!\sim\!\Delta(\text{orbited by})$

Where will we find this predicate? We might add it to our stock of primitives. But we have seen that we do better, avoiding ideological bloat and inferential brutality, if we find a way to build it up from 'orbits'. And indeed we can, by introducing another predicate functor: the *inversion functor*, 'INV'. Where R is any two placed predicate, $\ulcorner \text{INV}(R) \urcorner$ is a predicate that means $\ulcorner \text{is } R\text{-ed by} \urcorner$. That is, for any two-placed predicate R, the open sentence $\ulcorner R(x,y) \urcorner$ is equivalent to $\ulcorner \text{INV}(R)(y,x) \urcorner$. With INV in hand, (35) gets paraphrased as

(39) $\sim\!\Delta\!\sim\!\Delta(\text{INV}(\text{orbits}))$.

'INV' is well-defined only for binary predicates.[32] It tells us to 'swap' predicate positions when applied to a predicate. But if R has more than two positions, we could swap these positions in different ways.

Two particular ways of swapping positions turns out to be very useful. One way is to simply invert the last two of a predicate's positions; call this *minor inversion*. Or we could move a predicate's last position to the front and bump the rest down a notch; call this *major inversion*. If 'Inv' is a minor inversion functor, then (for quaternary R, say) $\ulcorner \text{Inv}(R)(w,x,z,y) \urcorner$ is equivalent to $\ulcorner R(w,x,y,z) \urcorner$. And if '$\mathbb{I}\text{nv}$' is a major inversion functor, $\ulcorner \mathbb{I}\text{nv}(R)(z,w,x,y) \urcorner$ is equivalent to $\ulcorner R(w,x,y,z) \urcorner$. It turns out that these two functors, wisely deployed, can generate *any* rearrangement of predicates' positions we might like.

One final issue needs resolving before we will be ready to paraphrase everything in the target language. Consider the open sentence

---

[32] And perhaps unary ones, if we take $\ulcorner \text{INV}(P) \urcorner$ to be equivalent to P for unary P.

(40) x attracts y & y attracts z.

In order to paraphrase sentences involving (40), we need a predicate P where $\ulcorner P(\text{x},\text{y},\text{z})\urcorner$ is equivalent to (40). 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 '$\mathbb{Inv}$' 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.

To deal with (40), we give ourselves a *padding* functor, which adds a 'dummy' place to a predicate. That is, for any predicate P and variable y, $\ulcorner P(x_1, \ldots, x_n)\urcorner$ is equivalent to $\ulcorner \text{Pad}(P)(y, x_1, \ldots, x_n)\urcorner$. (The new variable, y, is a dummy in that it does no work — as we ontologically minded folk would say, whether or not some objects satisfy $\ulcorner \text{Pad}(P)\urcorner$ has nothing to do with what object gets assigned to y.)

How does this help with (40)? First note that '$\text{Pad}(\text{attracts})(\text{x}, \text{y}, \text{z})$' is equivalent to 'y attracts z', because the object assigned to the place added by 'Pad' — in this case, 'x' — doesn't make a difference to the predicate's satisfaction-conditions. So (40) is eqivalent to

(41) x attracts y & $\text{Pad}(\text{attracts})(\text{x},\text{y},\text{z})$

because we replace (40)'s second conjunct with something equivalent. But when we conjoin a two-placed predicate and a three-placed one, we get a new three-placed predicate that is satisfied by some things only if the first two of them satisfy the two-placed predicate and all three satisfy the three-placed one. So (41) will be equivalent to

(42) (attracts & $\text{Pad}(\text{attracts})$)(x,y,z),

making '(attracts & $\text{Pad}(\text{attracts})$)' just the predicate we're looking for.

6.2 *A Nihilist's Paradise?*

In fact, with these six functors — 'Δ', '$\sim$', '&', 'Inv', '$\mathbb{Inv}$', and 'Pad'[33] — 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 (1960a, 1971) *Predicate Functor Language*, or *Functorese*. It has not only the expressive resources needed to translate anything we say in a first-order language, but its own attendant logic, besides. (Cf. Kuhn 1983, Bacon 1985)

---

[33]And an identity predicate, shared (I assume) by the target and paraphrase languages.

Call the Nihilist who paraphrases the target language into functorese the *Functorese Nihilist*. The Functorese Nihilist avoids the costs of Propositional Nihilism: He avoids ideological bloat by limiting his new primitive expressions to six. And he avoids inferential unsystematicity by appealing to logical relations between complex predicates — relations encoded in functorese's own attendant logic. This logic mirrors the predicate calculus: if ϕ entails ψ in the predicate calculus, then the functorese paraphrase of ϕ entails the functorese paraphrase of ψ in predicate functor logic.

Some (e.g., Jonathan Schaffer (2009: 368–370) and David Chalmers (2009: 118); see also Burgess and Rosen 1997: 185–188), perhaps led by reasons such as these, treat functorese as the Nihilistic language of choice. 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*. 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 Functorese's six predicate functors, 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 §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 §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'. 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 §3.2.

Step Two: Introduce another 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 '$\Delta\lambda$'.[34]

Step Three: We appeal to ($*$) from §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.

6.4 An Objection

It is tempting to think that the ($*$) Argument proves too much and so can't be right. The main idea runs something like this:

> We all agreed back in §5 that when we attach 'Δ' to a one-placed predicate A we got an expression that meant ⌜It is A-ing⌝. And

---

[34]More precisely, (open or closed) sentences of the form ⌜$\delta x\phi$⌝ are to be interpreted as ⌜$\Delta\lambda x\phi$⌝.

we all agreed that $\ulcorner$It is A-ing$\urcorner$ did not mean, and did not entail, $\ulcorner \exists x A x \urcorner$. 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. But 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 that there is an in-principle possible, ontologically innocent, one-placed-predicate-only language like the one discussed in §5.[35] If $(*)$ entails 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 had the same one-placed predicates as T, 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 '$\Delta_{FP}\lambda$'. Then we 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 §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.

These counterfactual circumstances will include ones in which the tribe's language is enriched with all of the multiple-placed predicates that we have in T. So we must ask how the tribe is disposed to extend their language to one with many-placed predicates. They *might* be disposed to to extend 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 §7. Or they may have some other dispositions entirely, or have no such dispositions at all. And if they are disposed to extend their

---

[35]At least, we ought to agree insofar as we are not troubled by, or are setting aside, worries that we can only make sense of the feature-placing languages on the model of weather sentences, which in turn must be thought of as covert predicates of places as discussed in note 26.

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 §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 §5, and that its '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* when we extend feature-placing languages 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 (Evans 1975, Quine 1992), or how features placed yesterday relate to features placed today (Evans 1975, Strawson 1954, 1963), etc. So 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 or not they're on the same peg.

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. We say 'It is electroning', and throw the 'electron' rubber band on the board; 'It is protoning', and throw the proton rubber band on the board; etc.

The other predicate functors let us make complex rubber bands out of simpler ones. But we have 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 the new complex predicate 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 objects in a model of reality is that they mark out certain locations on the board as special, as 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.) Since pegs represent objects, this is to say that *objects* are special landing-sites for features. Once it matters where one part of a rubber band has been stuck, we've smuggled in ontology. So, even if the Functorese 'Δ' is the 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

These last observations suggest that the Ontological Nihilist got into trouble when he decided 'Δ' should turn n-placed predicates into $n-1$-placed predicates. So let's go back to that point and try something else.

We could let Δ turn many-placed predicates into *sentences*. Just as 'It is raining' says that rain is going on, and 'Δ(proton)' says 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 orbits something else. Thus for any n-placed predicate R, $\ulcorner \Delta(R) \urcorner$ will be the Nihilist paraphrase of $\ulcorner \exists x_1 \dots \exists x_n (R(x_1, \dots, x_n)) \urcorner$.

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

(43) ∃x∃y(x is an electron & y is a proton & x orbits y),

(44) ∃x∃y∃z(x attracts y & y repels z),

and so on. We can make considerable headway on this by helping ourselves to the predicate functors '$\sim$', '&', 'Inv', $\mathbb{Inv}$', and 'Pad' from §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 (43)–(44) as

(45) Δ(electron & PAD(Proton) & orbits),

(46) Δ(attracts & PAD(repels)),

respectively.[36]

So long as we stick to target-language sentences beginning with blocks of existential quantifiers, this will do fine. But how will we paraphrase, for instance, 'Something orbits nothing'? It begins with a quantificational block like this:

(47) ∃x∼∃y. . .

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 'Δ∼Δ(orbits)'. But down that path lies predicate functorese and, as we saw, disaster. 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 (47).

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

(48) ∃x∼∃y∃z(x attracts y & x repels z)

(49) ∃x∃y∼∃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, using just 'Δ' and 'Σ', give these two different paraphrases.

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.

---

[36]A different option involves complicating the 'Δ'-functor, giving it extra 'slots' for more predicates and paraphrasing (43) 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.

Inferential unsystematicity: Note that these 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 taken as semantic black boxes — the Nihilist 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 with no further story as to what these expressions mean, he cannot explain the inferences they participate in.

Holism: In §4.2.3, we noted that we ontologically minded folk can think of the more global fact expressed by (44) 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 is the same y as the one doing the repelling.

But a Nihilist who paraphrases (44) as (46) 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

Figure 7: A Rubber Band for (44)
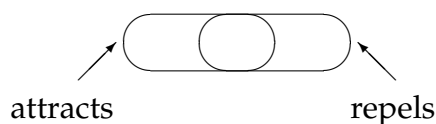
attracts   repels

feature by gluing together the 'attracts' 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. A deployment of an 'attracts' rubber band corresponds to an x attracting a y, and the deployment of a 'repels' rubber band corresponds to a y repelling a z. But to 'build up' the right complex fact from these deployments, the Nihilist also needs a fact corresponding to 'the y being attracted in the first deployment is the y doing the repelling in the second'. And there is no Nihilistically acceptable, object-free way to make sense of that claim.[37] 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*, which a Nihilist cannot do.

---

[37]If throwing two rubber bands on the board so that they look like figure 7 did the trick, it would be by making 'places' on the board important and thereby smuggling ontology back into the picture, as discussed at the end of §6.4.

8 Conclusion

We have not, of course, canvassed every way an Ontological Nihilist might try to paraphrase our target language. But 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 holistic picture of reality, a bloated ideology and indefinitely many brute entailments.

This gives us some reason to reject Ontological Nihilism. But this is hardly headline news. After all, only a few metaphysicians would ever have suspected Ontological Nihilism of truth 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. It is through ontology that we can think of reality's global structure as built up out of more local structures, for it turns out that a set of 'pegs', of *things*, is crucial for this sort of bottom-up picture. 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 with the structure 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 (2009). "Ontological Anti-Realism." In Chalmers et al. (2009), 77–129.

Chalmers, David, David Manley and Ryan Wasserman (eds.) (2009). *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.

— (2008). "On Specifying Truth-Conditions." *The Philosophical Review* 117(3): 385–443. Unpublished manuscript.

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

Schaffer, Jonathan (2009). "On What Grounds What." In Chalmers et al. (2009), 347–383.

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

Sider, Theodore (2008). "Monism and Statespace Structure." In Robin LePoidevin (ed.), *Being: Developments in Contemporary Metaphysics*, 129–150. Cambridge University Press.

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

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.

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.

Wittgenstein, Ludwig (1921). *Tractatus Logico-Philosophicus*. London: Routledge & Kegan Paul. Translated 1961 by D. F. Pears and B. F. McGuinness.