

How to have a radically minimal ontology¹

Ross P Cameron
University of Leeds

I am attracted to a radically minimal ontology. Many of the entities we quantify over in everyday speech do not, I hold, really exist. Complex objects are one such case: there is no mereology in reality – our ontology is one of entities lacking proper parts. However, I do not want to embrace an error-theory of talk about tables, chairs, etc: it is, even speaking strictly and literally, true to say such things exist. Rather, I suggest, we should view the (strict and literal) truth of such claims as not bringing an ontological commitment to tables, chairs, etc. It is true to say that there *are* such things; but *that* it is true does not commit us to admitting such things into our ontology. The purpose of this paper is to further elucidate this metaontological view and to defend it from some recent objections.

In previous work I have defended the claim that if you want to hold that ‘there are Xs’ is strictly and literally true whilst resisting ontological commitment to the Xs, you should show that one can provide grounds for the truth of such claims without appealing to the Xs: that is, you should show that an ontology lacking in Xs can nonetheless *make true* sentences proclaiming the existence of, or attributing features to, the Xs.² The thought is that the ontological commitments of a sentence are those entities that are needed as truthmakers for the sentence: those entities that must number amongst the ontology of the world if the world is to provide an adequate *grounding* for the truth the sentence. So, for example, to resist commitment to tables I do not need to reject the truth of table-talk but rather show that table-talk can be made true by something other than tables.

Is this the familiar project of *analysing* or *paraphrasing* sentences quantifying over objectionable entities into sentences that do not do so? In my opinion the terms ‘analysis’ and ‘paraphrase’ mean such different things in different philosophers’ mouths as to be too misleading to be useful, so I won’t describe my position using these terms. Certainly, on my view, for every sentence quantifying over the offending entities there is a sentence that in some sense corresponds to it that does not so quantify, but rather only quantifies over the truthmakers for the former sentence. But I make no claim that the two sentences mean the same thing or play the same theoretical role. Does that disqualify me from providing an analysis or paraphrase? Well, I’ll let you decide how to use the words. But certainly, as will emerge below, my strategy has metaphysical baggage that a paradigmatic paraphraser like Quine would not accept.

So the ontological commitments of a sentence S are, on my view, exactly those entities needed to ground the truth of S. What entities are needed to ground the truth of ‘there is a table’? Not a table, say I, for then we would be forced either to reject its truth or to admit tables into our ontology. What then? The natural thought, if we want to avoid commitment to complex objects, is to look to the simple parts of the

¹ Thanks to an audience at the Centre for Metaphysics and Mind, University of Leeds, for helpful discussion, especially Elizabeth Barnes, Jason Turner, and Robbie Williams.

² Cameron (2008a).

table. But can we say that it is just these simples that are the ontological commitments of the sentence or – since those simples could exist and it not be true that there is a table (because of the simples being arranged differently) – must we admit into our ontology states of affairs of these simples being arranged a certain way?

The answer to this depends on how serious we want to be about truthmaking. If you think that *all* the demands a sentence makes on the world can be captured in terms of what demands it imposes on *ontology* (this is the characteristic claim of the truthmaker maximalist), then you should hold that the existence of the truthmakers for S settles the truth of S. In that case, if you want to resist commitment to tables, you should hold that the truthmaker for – and hence the ontological commitment of – ‘there is a table’ is the state of affairs of the table’s simple parts being arranged as they are (i.e. table shaped). But this is not the only option. If one rejects truthmaker maximalism one will hold that there are sentences whose truth makes demands on the world other than ontological demands. The truthmakers for a sentence S, if one goes this route, may not then settle the truth of S, for the truth of S demands of the world more than that its truthmakers exist. On this view, the ontological commitments of a sentence are given by *certain* of the demands the sentence makes on the world: namely, those demands that concern ontology.³ Someone sympathetic to this approach could hold that the truth of ‘there is a table’ demands of the world not that there be a table, and not that there be a state of affairs of simple entities arranged table-wise, but simply that there be simples arranged table-wise. Here, the ontological commitments of ‘there is a table’ are just the table’s simple parts. That’s all that needs to be amongst our ontology for the world to ground the truth of that sentence. The table’s simple parts are all the ontology needed for the truth of the sentence: they are its sole truthmakers and hence its sole ontological commitment. Of course, having such entities in our ontology does not *suffice* for the truth of that sentence, but that is not because the ontological commitments of the sentence are not exhausted by these entities but rather because the truth of the sentence demands more of the world than that its ontological commitments exist: it also demands that they are arranged a certain way; but since such a demand does not concern ontology, it does not indicate a further ontological commitment of the sentence.

I do not wish to take a stand here on whether all the demands of a sentence can be captured in terms of the demands it makes on ontology – that is, on whether or not truthmaker maximalism is true. My church is an ecumenical one, and I wish to welcome both maximalists and non-maximalists to the fold. The general principle I wish to defend can be accepted by both: that the ontological commitments of a sentence are not what the sentence quantifies over but rather what entities must be included in our ontology to ground the truth of the sentence – what entities must exist to make the sentence true. Whether or not *all* that needs to happen for the sentence to be true is that such entities exist – which depends entirely on whether or not the sentence imposes demands on the world other than ontological demands – is another question, and one that need not be decided today.

Let me be clear about what the view is *not*. Brogaard identifies my position as follows⁴: S is committed to Fs iff, Necessarily, if S has a truthmaker, there are Fs.

³ Cf. Rayo (2007).

⁴ Brogaard (2008, p28-29).

This would be an absurd position to hold. For starters, it renders any sentence that necessarily lacks a truthmaker – the possibility of which I want to stay neutral on, given my above neutrality on maximalism – committed to anything you like. It is also incapable of securing the benefits I seek. ‘There are elephants’ is committed to elephants, on this criterion, because it’s necessary that if that sentence has a truthmaker then there are elephants (just because it’s necessary that if it’s true then there are elephants, and a sentence can’t have a truthmaker without being true); but I do *not* think this sentence is committed to elephants precisely because I don’t think that elephants are involved in what *makes* it true.

My view is not that the ontological commitments of S are what must exist if S is made true but rather *what must do the truthmaking*. Brogaard’s objections against what she calls the Truth-Maker criterion of ontological commitment, being levelled against a straw man, do not touch this view. But some objections have emerged that are aimed at what is actually my view, and it is the purpose of this paper to rebut them.

It is not in general true that a sentence has a unique truthmaker. Given this, Jonathan Schaffer offers me a dilemma.⁵ Is a sentence committed to *all* of the things that make it true, or just some of them? If we answer the former, we get over-commitment; if we answer the latter, we get under-commitment, at least on the most natural way of saying *which* of a sentence’s truthmakers are its ontological commitments.

The first horn goes as follows. Take the sentence ‘A is F’. Suppose for the sake of illustration that that sentence is made true by the state of affairs of A’s being F. And suppose ‘B is G’ is made true by the state of affairs of B’s being G. And ‘A is F and B is G’, suppose, is made true by the *conjunctive* state of affairs of A’s being F and B’s being G. But this last conjunctive state of affairs is also a perfectly adequate truthmaker, seemingly, for our original truth ‘A is F’. So ‘A is F’ is ontologically committing to the conjunctive state of affairs of A’s being F and B’s being G. But that is implausible! This truth has nothing to do with B being G, and shouldn’t bring commitment to a state of affairs that has B’s being G as a constituent.

Suppose on the other hand that we say that not *every* truthmaker is an ontological commitment. Which ones *are* then? The only plausible answer appears to be: the minimal truthmaker(s). The conjunctive state of affairs above is not a *minimal* truthmaker for ‘A is F’, only the state of affairs of A’s being F is, so only the latter is an ontological commitment. But now, what to say of sentences that have *no* minimal truthmaker(s), like ‘there are denumerably many electrons’? For any collection of electrons you point to that make this true, I can point to a subplurality that make it true, so there is no *minimal* plurality that make it true. So is this sentence just not ontologically committing? But that’s absurd! Of course this sentence makes demands on our ontology: it makes a *huge* demand – that there be denumerably many electrons!

I find it hard to be worried by this dilemma. We should make a distinction, as indeed anyone must, between a sentence bringing an ontological commitment to some particular thing(s), and it ontologically committing you to some things or other. I say that a sentence S commits you to some *particular* thing A when A *has* to make S true

⁵ Schaffer (2008, p14-16).

if it is true; that S commits you to there being some things or other amongst the Xs when it *has* to be the case that some things or other amongst the Xs make S true if it is true; and that the ontology of the world is the entirety of the entities that make true the sentence that truly and completely describes the world.

‘A is F’, we might think, thus brings a particular commitment to the state of affairs of A’s being F, because it *has* to be made true by that state of affairs if it is true. It *certainly* doesn’t bring a particular commitment to the conjunctive state of affairs of A’s being F and B’s being G since B might not be G and A still be F, in which case the conjunctive state of affairs wouldn’t exist, and so wouldn’t make ‘A is F’ true. Some sentences, however, don’t bring any *particular* ontological commitment. Consider the sentence ‘A is F or B is G’. It can be made true by the state of affairs of A’s being F, or by the state of affairs of B’s being G, or by the conjunctive state of A’s being F and B’s being G, but none of these *has* to make that sentence true if it is true, and so it makes no demand on ontology that it contain some *particular* thing(s), and hence brings no *particular* ontological commitment. But nevertheless, it is still ontologically committing: it is ontologically committing to there being *one or other* of the state of affairs of A’s being F or of B’s being G, since one or other of them must make that sentence true if it is true. And so if this sentence is a conjunction of the sentence that truly and completely describes the world, one of these states of affairs must be in our ontology – and so while the sentence doesn’t bring a *particular* ontological commitment, it is ontologically committing.

And likewise with ‘there are denumerably many electrons’. Let’s suppose ‘there is an electron’ is made true by any electron. ‘There are denumerably many electrons’, then, is made true by any denumerable plurality of electrons. But there is no *particular* plurality of electrons that *must* make this true if it is true, so this sentence brings no *particular* ontological commitment. But it nonetheless is ontologically committing, since *some denumerable plurality or other* of the (possible) electrons must make this sentence true if it is true. And so if our complete and true theory of the world contains this sentence, our ontology had better contain denumerably many electrons.

Indeed, I find it hard to see why one might think there’s any kind of worry here. The Quinean has to say similar things. The Quinean says that the ontological commitments are what *have* to be included amongst the values of the variables in a statement of the theory. But what *has* to be amongst the values of the variables if our theory says that there are denumerably many electrons? Well, there had better be denumerably many electrons amongst the values of the variables, but there are no *particular* electrons that have to be amongst the values of the variables, since for any plurality of electrons you point to I can point to a subplurality that would do the job just as well; so there’s no *particular* ontological commitment this sentence brings on the Quinean criterion, but it nevertheless commits you to there being denumerably many electrons. Compare that to the truthmaker story: there’s no particular truthmaker(s) whose existence this claim commits you to, but it does commit you to there being *some* truthmaker(s) for that claim. So I just can’t see why there’s a particular problem for the truthmaker criterion here. This is exactly what Quine says with respect to the difference between, e.g., ‘there is an elephant’ and ‘there is the unique Socratiser’. The latter is ontologically committing to Socrates, since Socrates *has* to be amongst the values of the variables if ‘there is the unique Socratiser’ is true; but there’s no particular thing that *has* to be amongst the values of the variables if

‘there is an elephant’ is true, since *any* elephant will do; so there is no *particular* ontological commitment of ‘there is an elephant’, but it is nevertheless ontologically committing to elephants, and if our true theory says that sentence is true, our ontology (if Quine is right) had better contain elephants. So I’m in exactly the same position as the Quinean in making a distinction between a sentence’s bringing a particular ontological commitments and its bringing an ontological commitment to there being some thing(s) or other, and so I think that there is no objection here.

Schaffer levels another objection against my approach. I say that ‘there are Fs’ might be true without bringing commitment to Fs. In that case, existential quantification cannot in general be ontologically committing to what is quantified over. But I also say that (in many cases at least) when ‘there are Fs’ is true ‘there is a truthmaker for ‘there are Fs’’ is true, and in this case I *do* think that there is commitment to the truthmaker. But why treat the quantifiers differently from one case to the next? Either quantification in general brings ontological commitment, in which case ‘there are Fs’ brings commitment to Fs, or it doesn’t, in which case there’s no explanation for why ‘there is a truthmaker for S’ brings commitment to the truthmaker. This is a very good challenge⁶, and answering it will, I hope, help to clarify my position.

I see it as part of the benefit of my approach to ontological commitment that one can pull apart what there is from what things must be elements of one’s ontology. Such a view might initially sound like nonsense: what can it mean for there to be an F in one’s ontology other than that there is an F?

Compare what I just said with two other claims: (i) there are modal truths, but modality is no part of reality; (ii) there are tensed truths, but tense is no part of reality. Those claims might initially sound incoherent as well, but nonetheless, such claims are made. The way of making sense of them is to appeal to the notion of naturalness. The first of these claims is made by the modal reductionist, who thinks that there are in fact modal truths – which is just to say that sentences in English using modal vocabulary can be used to say true things – but that the modal notions do not carve the world at its joints: a language that did not invoke modal notions could do just as good a job of describing the structure of reality as our own. And likewise, *mutatis mutandis*, for tense for the holder of the second claim.

This is the notion of naturalness familiar from Lewis’ discussion of gruesome predicates⁷, but extended to other parts of the language in the manner suggested by Sider.⁸ Consider a tribe who classify things using the terms ‘grue’ and ‘bleen’ but not ‘green’ and ‘blue’. They can do so successfully (the terms are not inconsistent), and can speak truly when they say that things are grue and bleen; but, according to Lewis, they do not do as good a job of describing the structure of the world as we do when we (just as truly – no more, no less) say that things are green and blue. The green things and the blue things belong together in a way that the grue things and the bleen things don’t. ‘Green’ and ‘blue’ are *more natural* predicates than ‘grue’ and ‘bleen’: classifying things using the former terms is better than classifying things using the latter terms, because we will not only make true predications but reflect the objective groupings of things; green things are *in reality* of a kind, grue things are not.

⁶ Ibid. p16.

⁷ Lewis (1983).

⁸ See Sider (2009).

If you're happy with speaking about the naturalness or otherwise of predicates, there's no reason not to be happy to speak about the naturalness or otherwise of other parts of language, in particular the quantifier.⁹ Suppose one tribe consistently utters the sentence 'there is a table' in circumstances where another tribe would utter only 'there are some simples arranged table-wise'. It's plausible that both tribes speak truly but simply speak a different language: that they just mean something different by 'there is'.¹⁰ And so if there's really a genuine ontological question to be asked regarding whether or not the world's ontology is one of merely simples or rather one of complex objects then we can't ask it simply by asking whether 'there are tables' (or some such sentence) is true – since that just depends on whether what we mean by 'there is' is what the first tribe or the second tribe means by it – rather, we have to ask which of the two tribes' languages (if either) gets the ontology *right*. Both tribes are saying true things, but whose language carves the world at its joints? If we were speaking a language that used the *perfectly natural* quantifier, would we truly be able to describe the world as one where 'there are' complex objects? *That* is the ontological question regarding complex objects.¹¹

Neo-Carnapians hold that there is no most natural quantifier, and hence that there are simply different ways of carving up the world into objects, none better than any other at describing the quantificational structure of the world. For the neo-Carnapian, there simply *is* no ontological question beyond what we in fact mean by 'there is' in English. Ontological realists, on the other hand, think that there is a unique most natural quantifier: a language with this quantifier does a better job of describing the quantificational structure of the world than a language with a less natural quantifier.¹²

I am an ontological realist. But note that it is no part of the doctrine thus stated that the English 'there is . . .' is the most natural quantifier. As an ontological realist I am committed to saying that there is a possible language such that the true existence claims in that language correspond to the quantificational structure of the world – this is a language whose existence claims carve reality at its joints: one where you can truly and completely list your ontology by saying what there is. But I do not think English is such a language.¹³ The ontology of the world is what can be said to exist using the most natural quantifier; but if this is not the English 'there is . . .' then it might be the case that there are some things that are not amongst our ontology.

The terms in our language do not (all) carve the world at the joints; nevertheless, we can use them to say true things, and the question arises as to what is the relationship between the truths expressed by the sentences of English and the truths expressible in the fundamental language – call it 'Ontologese' – whose terms all carve perfectly at the world's joints. Perhaps the relation is identity: that is, that the sentence in English and the sentence in Ontologese express the same proposition; or perhaps the sentences express different propositions, and there is a primitive notion of *grounding* such that the proposition expressed by the sentence of Ontologese *grounds* the proposition expressed by the sentence of English. I don't want to take a stand on this here,

⁹ *ibid.*

¹⁰ See Hirsch (2009).

¹¹ Cf. Cameron (forthcoming).

¹² See Sider (2009).

¹³ See Cameron (forthcoming) for further discussion.

although we'll mention below areas where the difference might play out. All I commit to here is that the sentence in Ontologese describes reality in a metaphysically perspicuous way, and that reality's being that way explains *why* the sentence in English is true.

I now stipulatively define 'truthmaker' as follows: the truthmaker(s) for S are those things, if any, said to exist by the sentence in Ontologese whose truth explains the truth of S.

We are ontologically committed to all the truthmakers of the true sentences of English because those are, by definition of 'truthmaking', the entities whose existence is carved from the quantificational structure of the world by the most natural quantifier, and whose existence is at least part of the explanation of the truths expressible in English. We are not ontologically committed to anything *else* because, ex hypothesi, no other entities are demanded to be amongst our ontology by the truths we accept, *even if* the truths we accept quantify over other entities, or involve singular terms that refer to other entities.

Were there nothing we could say about the difference between the quantifier in 'there is a truthmaker for S' and 'there is an elephant' I would agree with Schaffer that it is a mystery why the former sentence should bring commitment to a truthmaker but the latter not bring commitment to an elephant. But there *is* something we can say about why they are different: the former but not the latter sentence would still express a truth if by 'there is . . .' we meant the most natural quantifier rather than the unnatural one we in fact mean when we quantify in English. And I know this last claim is true because it's simply analytic of 'truthmaker': A wouldn't *be* a truthmaker for S if a speaker of the language that carves at the joints wouldn't agree that 'A exists' expresses a truth.

So Schaffer's objections both fail. Another objection to my general metaontological view recently emerged as a criticism to my particular application of that view to the case of the ontology of musical works¹⁴, and while the criticism if successful would be damaging beyond this application of it, it will prove useful to consider it via this case.

The central ontological puzzle concerning works of music is this. It is overwhelmingly intuitive – to the point, I think, of being a datum that any successful theory must account for – that an utterance of 'Beethoven's ninth exists' says something true if uttered at a time after that at which Beethoven completed composing his ninth, and that an utterance of the same sentence type says something false if uttered at a time before Beethoven had any thoughts about his ninth or had started any of the work that led to its composition. It seems to follow *either* that some entity – Beethoven's ninth – exists now but didn't exist before Beethoven was born, or that there is some entity that is now Beethoven's ninth but that wasn't always Beethoven's ninth. The latter sounds odd: it seems that Beethoven created a musical work, not that he took something that wasn't a musical work and made it into one. But Beethoven's ninth doesn't appear to be identical with any of the ordinary concrete objects of everyday acquaintance or empirical scientific discovery, and so it seems it

¹⁴ Cameron (2008b).

must be an abstract entity. But abstract entities do not appear to be the kind of entities that can exist at some times and not at others, and so we have a puzzle.

Solutions to the puzzle tend to bite one of two bullets: they either deny the common-sense intuition that Beethoven's ninth exists after the composition but not before¹⁵, or they complicate our ontology with abstract entities which can be brought into existence by the actions of humans.¹⁶ I suggest that the puzzle is best solved by denying that musical works are needed in our ontology to account for the truth of sentences proclaiming their existence and/or attributing features to them. In that case we can agree that 'Beethoven's ninth exists' says something true at time *t* and false at time *t** but deny that the truth of 'Beethoven's ninth exists' brings an ontological commitment to some thing that is Beethoven's ninth, thus allowing us to resist the troublesome conclusion that the ontology of the world includes an abstract entity that exists at *t* but not at *t**.

Just as the reductionist about modality, or tense, or laws, believes that there are true sentences in English about how things might have been, or how things were and will be, or what generalisations are lawlike, but doesn't believe that reality is modal, or tensed, or law-governed (which is to say, a language using perfectly natural vocabulary wouldn't contain modal or tense operators, or law vocabulary, and could give a complete and true description of reality), so do I hold that there are true sentences in English regarding the existence of musical works, but I don't hold that ontology contains musical works (which is to say, a language using a perfectly natural quantifier wouldn't quantify over musical works). Just as modal truths are sensitive to non-modal features of reality (such as what's going on at spatio-temporally isolated worlds, perhaps¹⁷) and tensed truths sensitive to tenseless features of reality (such as what's going on simultaneously with *this* thought, etc) so, say I, are truths proclaiming the existence of musical works sensitive not to whether our ontology contains musical works but rather to features of reality not involving musical works. Which features? Well we'll come back to this, but let's for the moment say the following: that the truth of 'Beethoven's ninth exists' is sensitive not to ontology containing an entity that is Beethoven's ninth but rather to our ontology containing (or having contained) the event of Beethoven indicating a certain abstract sound structure: that type such that a token of it is heard when we hear a performance of Beethoven's ninth. 'Beethoven's ninth exists' is true *because* Beethoven indicated this abstract sound structure, and not because our ontology contains some abstract entity that is Beethoven's ninth.

This proposal has been criticised by Stefano Predelli.¹⁸ Predelli makes two accusations against my view. Firstly, my approach isn't original. This is the familiar game of *analysing* a problematic sentence – 'Beethoven's ninth exists' – in terms of an unproblematic one – 'Beethoven indicated the abstract sound structure E'. Second, the approach doesn't work, as there are possible worlds in which one of these sentences is true and the other false. Why? Because a world where Beethoven indicated the very sound structure E for performance but in which the musical history of the world prior to that event was very different from how it actually was prior to

¹⁵ Dodd (2007).

¹⁶ Levinson (1980).

¹⁷ As in Lewis (1986).

¹⁸ Predelli (forthcoming).

his writing his ninth might not be a world in which Beethoven's ninth exists, and that's because the identity conditions of a musical work are informed by the context in which it was composed: change the context and the same sound structure can be indicated but a different musical work be composed.¹⁹

This claim – that the identity of musical works is informed by the context in which it is composed – is one I do not want to be forced into denying. Fortunately, I don't have to. As I said above, I don't want to get into what it takes to be providing an 'analysis', but I think Predelli's second criticism here is only at all worrying if you mean something strong by it – and if you *do* mean something that strong by 'analysis', I am certainly not intending to give one. Predelli's rush to see me as engaging in the familiar project of providing an analysis has, I suspect, blinded him to the possibility that it is the very differences in my approach that allow me to resist exactly the kind of counter-example he presses against me.

How exactly to respond to Predelli depends in part on an issue I refused to take sides on above: whether the relationship between the proposition expressed by a sentence of English and the relevant sentence of Ontologese is one of identity or of grounding. If it is one of grounding, then we can resist Predelli's counter-example by denying the constraint that the sentence of Ontologese whose truth explains the truth of the English sentence must be true in all and only the same counterfactual circumstances as the latter. If the relationship is identity we can't take this route, since presumably two propositions can only differ with respect to what counterfactual circumstances they would be true in if they are indeed *two* propositions and not one; but if the relationship is identity we should respond to Predelli's counter-example simply by invoking a more complex sentence of Ontologese that *is* true, necessarily, if and only if the English sentence is true.

Suppose we go the grounding route. My claim then is simply that the truth of 'Beethoven's ninth exists' is *actually* grounded in Beethoven's having indicated a certain abstract sound structure, and that this is compatible with the fact it *might* not have been so grounded.

As I said above, I am not aiming to give an *analysis* of 'Beethoven's ninth exists', in the sense of saying that what we *mean* when we say that Beethoven's ninth exists is that Beethoven indicated a certain sound structure for performance: the man on the street (if they've had a good enough musical education) might not have any conception of an abstract sound structure, but will be perfectly happy to proclaim the existence of Beethoven's ninth. And once the synonymy claim is abandoned, and only the grounding claim accepted, we make no commitment to the two sentences truly describing exactly the same counterfactual circumstances.

Consider again the case of complex objects. I claim that there aren't really any complex objects in our ontology, but that we speak truly when we utter English sentences proclaiming their existence, because the English 'there is' is not perfectly natural. 'The table exists' is true, then, but what grounds this is not the existence of any table but rather, say, the existence of, and spatio-temporal distribution of, a certain collection of simples, call them the Xs. But the table might have been made of

¹⁹ See Levinson (1980).

different matter: a different collection of simples might have composed it. Had that been the case, ‘the table exists’ would still have been true but the Xs might not have existed. No problem, say I: I wasn’t providing an *analysis* of ‘the table exists’ but merely pointing to its actual grounds. It is perfectly compatible with ‘the table exists’ being actually true due to the existence and arrangement of the Xs that it could have been true due to the existence and arrangement of the Ys instead. Likewise, the Xs being arranged as they actually are is not sufficient for the truth of ‘the table exists’, because if there’s enough stuff in the Xs surroundings we wouldn’t have a table but a block of wood. No problem say I: ‘the table exists’ being actually true due to the existence and arrangement of the Xs is perfectly compatible with it being possible that the Xs are so arranged without grounding the truth of ‘the table exists’.

Perhaps it will be objected that in that case ‘the table exists’ isn’t true because of the existence and arrangement of the Xs after all, but that it is true because *either* (1) the Xs are so arranged *and* aren’t surrounded by anything else *and* there are no other defeaters *or* (2) the Ys exist and are so arranged *and* there are no defeaters *or* . . .

If you are attracted to that thought I think you should hold that the sentences of English and Ontologese express the same proposition, but that when expressed in Ontologese the proposition is an infinite disjunction. (The disjunction would, I presume, have to be infinite, since it’s implausible that there are finitely many ways the world might be with regard to the existence and distribution of simples in order for ‘the table exists’ to be true.)

I think there’s a perfectly good sense in which ‘the table exists’ is actually grounded simply in the existence and arrangement of the Xs, but if you think this infinite disjunction is the *real* explanation or even just a *better* explanation, let me simply concede that. That concession would be damaging were I in the business of giving an *analysis*, in the sense of providing the same meaning, since it’s implausible that we *mean* by ‘the table exists’ some infinite disjunction with very complex disjuncts, or that this infinite disjunction plays the same theoretical role as the English sentence ‘the table exists’. But again, I’m not in that business. I’m in the business of pointing to the features of reality that ground the truth of our English sentences proclaiming the existence of complex objects. Those features are, I claim, not the existence of complex objects but the existence and distribution of simples. The fact that someone speaking a perfectly natural language could only provide a sentence necessarily equivalent with our English sentences if they had the resources to provide us with infinite disjunctions is neither here nor there.

As with complex objects, so with musical works. I say that the truth of ‘Beethoven’s ninth exists’ is *actually* grounded in Beethoven’s having indicated a certain abstract sound structure (E) for performance. But there are counterfactual circumstances in which Beethoven indicates E for performance but where his ninth doesn’t exist: where he indicates that E is to be performed as the interlude during a 24-hour techno epic, perhaps. And perhaps, if musical works can survive slight variation in the notes that make them up, just as tables can survive slight variation in the matter that composes them, there are counterfactual circumstances where ‘Beethoven’s ninth exists’ is true but where Beethoven doesn’t indicate E for performance but rather the slightly different sounding E*. I am unworried: the truth of ‘Beethoven’s ninth exists’ being

actually grounded in Beethoven indicating E for performance is compatible with the possibility of it being grounded in another manner.

And again, perhaps one would wish to conclude in that case that the ‘real’ explanation of the truth of ‘Beethoven’s ninth exists’ is that one of a number of conditions obtain: either (1) Beethoven indicated E for performance *and* didn’t indicate that it was to be performed as the interlude of a techno epic *and* the musical history of the world is such-and-such *and* there are no other defeaters, or (2) he indicated E* for performance *and* . . . etc. Provided this infinite disjunction written using only perfectly natural terms never talks about the existence of musical works – and it needn’t, it only needs to talk about the possible patterns of indication of sound structures by composers – then we have an adequate grounding of the facts concerning the existence of musical works in facts which don’t concern such entities. The fact that one can only state the exact way reality must be in order for the musical work to exist if one can state an infinite disjunction is, again, neither here nor there: it would be worrying were I attempting to analyse, in the sense of providing a synonymous sentence, the English sentence in terms of this infinite disjunction whose terms carve the world at the joints. But again, that is not my project. I am merely trying to indicate the *grounds* of the English sentence: the features of the structure of the world that account for the fact that the musical work exists. These features, I claim, are the facts concerning what sound structures have been indicated for performance and in what context. The fact that I cannot, in a finitary language, give necessary and sufficient conditions for the truth of sentences proclaiming the existence of musical works in terms of these fundamental features of reality does not count against the proposal that it is these features that explain the facts concerning the existence of musical works: it counts only against the claim that our English sentences mean the same thing as the sentences in a language whose terms carve at the joints and whose truth explains the truth of the sentences of English. But that’s a claim I never made.

Now, what *would* refute my account is if the global facts concerning who indicated what sound structures for performance and in what context could be the same but the facts concerning what musical works exist be different. Just as the reductionist about laws demands that if you fix the facts concerning what regularities hold you thereby fix the facts concerning what laws hold; just as one who holds that complex objects are no addition of being over and above their simple parts demands that if you fix the facts concerning what simples exist and the properties they have and relations that hold between them you thereby fix the facts concerning what complex objects exist; so do I, in grounding facts concerning the existence of musical works in facts concerning the indication of sound structures, demand that if you fix the facts concerning what sound structures have been indicated when, and by whom, so you thereby fix what musical works exist. But this seems overwhelmingly plausible to me. And certainly nothing Predelli says rules out this supervenience claim: the examples where a composer indicates the same sound structure but doesn’t compose the same work depend on there being some change *elsewhere* in the facts concerning who indicated what and when. And so the reductionist program is yet to be threatened.

A final issue. Predelli complains that I haven’t proven that musical works do not number amongst the ontology of the world. For all I have said, it’s a necessary truth that if the event of Beethoven indicating E is an element of ontology so is a musical work that is Beethoven’s ninth. That is to say, for all I have said, the sentence using

the perfectly natural quantifier ‘There is an event of Beethoven indicating E’ strictly implies the sentence, also stated using the perfectly natural quantifier, ‘Beethoven’s ninth exists’. In that case, I haven’t proven that the English sentence ‘Beethoven’s ninth exists’ (here with our non-natural quantifier) can be true without Beethoven’s ninth being an element of ontology. True; nor have I said anything that rules out the claim that, necessarily, if the English sentence ‘Beethoven’s ninth exists’ is true then the ontology of the world contains pink elephants. But ontologists are not in general in the business of *proving* that our ontology contains such-and-such and not thus-and-so; our aim is usually the more modest one of providing reasons to believe such claims. Here is my reason for not believing ontology contains either Beethoven’s ninth or pink elephants: it is Ockham’s razor. Don’t admit something to your ontology unless you have a good reason *to* admit it! One reason to admit Beethoven’s ninth to your ontology would be if it were needed to ground the truth of the obviously true sentence of English ‘Beethoven’s ninth exists’, but this I have argued against. There are other potential reasons to admit Beethoven’s ninth into our ontology: perhaps it’s needed to ground the truth of something *else* we think is true, or perhaps it’s a necessary truth that if our ontology contain something we *do* think is needed that our ontology contain Beethoven’s ninth as well. I don’t see much prospect for justifying either of these claims, but I concede my position is vulnerable to future arguments that might convince me of their truth. The mere suggestion that there could be such a necessary truth, however, is not a justification for accepting that there is one.

Now, in order for Ockham’s razor to do the work I intend it to do here in establishing musical nihilism, the correct formulation of Ockham’s razor must tell us to minimise what entities are needed to do *grounding* work rather than simply what entities *exist*. And in general if both compositional and musical nihilism are to be advantageous over their rivals then the principle of ontological parsimony must tell us to judge theories by what entities they admit to their fundamental ontology – that is, what entities do truthmaking work; what entities are ranged over by the perfectly natural quantifier – rather than what entities they claim exist. After all, I admit that Beethoven’s ninth and tables *exist*, so if that was enough to incur an ontological cost, I wouldn’t have any advantage over my rivals.

But that’s exactly how we should understand Ockham’s razor. What entities are ranged over by the perfectly natural quantifier is *all that there really is*! That’s why I speak of those items, and only those, as being amongst the ontology of the world. Speaking a language with an unnatural quantifier, we quantify over things that *don’t really exist*. We say truly that there are complex objects, musical works, etc, but it is merely a matter of convention that we can truly describe reality thus: such claims don’t limn the quantificational structure of reality. That’s why the initial puzzle concerning musical works is resolved on my view. It might still turn out to be a truth that musical works are abstract entities that are created. But the thought is that the truth of this claim doesn’t conflict with the intuition, properly understood, that abstract entities don’t come into or go out of existence. The *ontological* scruples we have against created abstracta should mandate only that there are no created abstracta in our ontology; if the intuition is correct, the perfectly natural quantifier will not quantify over temporarily existing abstracta; but this is compatible with the conventions governing our usage of the unnatural quantifier in English resulting in the truth of the English sentence ‘there is an abstract musical work that didn’t exist

previously'. Compare the following. You might have an intuition against vague existence. But if our ontology is one of determinately existing simples and it is a vague matter whether their existence and distribution leads to the truth of some claim in English proclaiming the existence of a complex object, this does not seem like the allegedly objectionable kind of ontic vagueness that people were objecting to concerning vague existence.²⁰

Since it is merely a matter of convention that we can truly describe the world as one containing musical works, many questions concerning the 'metaphysics of musical works' are not really metaphysical questions but rather conceptual/linguistic ones. For example, the question as to whether someone other than Beethoven could have composed Beethoven's ninth, or whether Beethoven could have composed it slightly differently, or whether two composers could, in the same world, compose the same work of music, etc. These questions might *sound* like they're concerned with the essence of some entity; but only real entities have real essences²¹ – in fact, these questions are sensitive only to how we think/talk about musical works. The question is simply whether our usage of 'Beethoven's ninth' is flexible enough that we would count a world where Beethoven wrote a similar sounding work as a world containing Beethoven's ninth, etc. And some such questions might simply be indeterminate, of course, since our usage of the term and our concept of the work might not settle the issue. But while such indeterminacy might be worrisome were it posited concerning the real essence of some entity, such indeterminacy arising from linguistic or conceptual unsettledness is utterly unproblematic.

It's because musical works and complex objects don't have any real ontological status that it is no ontological cost to admit the existence of such things: ontological cost is incurred not by truly describing the world using a sentence of the form 'there is an F' but rather by admitting entities to *ground* the truths of our language – the entities that *really exist* and would be quantified over even if we used the perfectly natural quantifier that carves the world at its quantificational joints. That's why my position *really is* compositional and musical nihilism. Yes, complex objects and musical works exist; but they are no part of ontology, and thus the truth of the sentences 'tables exist' or 'musical works exist' carries no ontological commitment to tables or to musical works respectively; and hence my view is ontologically parsimonious over rivals who admit such entities into their ontology. That is why you should accept compositional and musical nihilism, and the metaontology that underwrites them.

References

- Brogaard, Berit, (2008), 'Inscrutability and Ontological Commitment', *Philosophical Studies* Vol. 141, No. 1, p21-42
- Cameron, Ross, (2008a), 'Truthmakers and Ontological commitment: or, how to deal with complex objects and mathematical ontology without getting into trouble', *Philosophical Studies* Vol. 140, No. 1, p1-18
- Cameron, Ross, (2008b), 'There are no things that are musical works', *The British Journal of Aesthetics*, Vol. 48, p295-314

²⁰ See Cameron (2008a, p15-16).

²¹ I owe the phrase 'real entities have real essences' to Kris McDaniel.

- Cameron, Ross, (forthcoming), 'Quantification, Naturalness and Ontology', in *New Waves in Metaphysics*, edited by Allan Hazlett, Palgrave-Macmillan
- Dodd, Julian, (2007), *Works of Music: An Essay in Ontology*, Oxford: Oxford University Press
- Hirsch, Eli, (2009), 'Ontology and Alternative Languages', in Chalmers, Manley and Wasserman (eds) *Metametaphysics*, Oxford: Oxford University Press
- Levinson, Jerrold, (1980), 'What a Musical Work Is', *Journal of Philosophy* Vol. 77, p5-28
- Lewis, David, (1983), 'New Work For a Theory of Universals', *The Australasian Journal of Philosophy* 61, p343-377
- Lewis, David, (1986), *On the Plurality of Worlds*, Oxford: Blackwell
- Predelli, Stefano, (forthcoming), 'Ontology and Musical Nihilism: A Reply to Cameron', *British Journal of Aesthetics*
- Rayo, Agustín, (2007), 'Ontological Commitment', *Philosophy Compass* 2/3, p428-444
- Schaffer, Jonathan, (2008), 'Truthmaker Commitments', *Philosophical Studies* Vol. 141, No. 1, p7-19
- Sider, Ted, (2009), 'Ontological Realism', in Chalmers, Manley and Wasserman (eds) *Metametaphysics*, Oxford: Oxford University Press, p384-423