

Should Ontology be Explanatory?¹

Amie L. Thomasson

Department of Philosophy
Dartmouth College
Hanover NH, USA
AmieLThomasson@gmail.com

Abstract: The central question of ontology has long been thought to be ‘What is there?’. The central way of answering it has been to consider which entities we *must posit* as part of a best total explanatory theory. This paper argues against this ‘explanatory’ conception of metaphysics, by showing that it relies on an unarticulated assumption that all the terms at issue in these metaphysical debates serve an explanatory function. Making use of work in systemic functional linguistics enables us to identify the many different functions played by terms of interest in metaphysics. And that makes it clear that ‘contribution of explanatory power’ should be rejected as an across-the-board criterion in ontology. This work in functional linguistics also enables us to see why it is *useful* to have a language that entitles us to use redundant inferences to introduce terms for properties, numbers, and the like, giving us new reason to accept ‘easy’ inferences that there are such things. As a result, we should give up thinking that ‘what is there?’ provides a deep and interesting question for a discipline called ‘ontology’ to answer, and give up thinking that the task for ontology is to determine which entities to ‘posit’ as part of a best total explanatory theory.

Key Words: Metaontology, Quine, Easy Ontology, Systemic Functional Linguistics

The central question of ontology, as Quine presented it, could be expressed in three words: “What is there?” (1953/2001, 1). But those who aim to answer this question generally do not take their goal as simply to generate a list or inventory, and tend to deny that answers are easy to come by. Instead, the project of ontology has been thought of as a matter of determining what entities we *should* or *must* ‘posit’ as part of a best *total explanatory theory*. As a result, it has become common to treat *whether or not the alleged entities add explanatory power* to our theories as a central criterion in determining whether to accept them; we ask whether Ps would have any explanatory power, as (at least partial) grounds for answering the ontological question about whether *there are* Ps. Sometimes metaphysical ‘posits’ are alleged to explain certain facts or observations (such as why the red barn and the red house have something in common, or why it is true that grass is green). In other places, they are taken to explain *norms*, such as why we *should* act in certain ways (because this is dictated by the moral facts), or why we *should* theorize using some concepts rather than others (because these concepts pick out the *natural* properties and relations).

Many have found this conception of the project of ontology tempting, since it seems to lend a sort of scientific respectability to metaphysics. The thought is supposed to be that, *just as*

¹ Portions of this material are taken from my book, *Rethinking Metaphysics* (forthcoming), especially chapter 2. Thanks go to Peter Lewis, Ted Locke, Michelle Moody Adams, Mark Povich, Katherine Ritchie, and audiences at Columbia and University of New Hampshire for very helpful input and comments on prior versions of some of this material.

*we can have reasons to accept (or reject) the existence of a new planet, fundamental particle, or microbe given the explanatory power it brings (or fails to bring) to a relevant scientific theory, so can we have reason to accept (or reject) metaphysical ‘posits’ such as universals or natural properties, based on the explanatory power they add (or fail to add).*² Of course, it is contested whether contributing explanatory power (or other theoretic virtues) licenses us in positing unobservable entities even in science. But let’s put that to the side and ask instead: Is addition of explanatory power a legitimate criterion for accepting or rejecting entities in ontology (*assuming it is in empirical matters*)? More broadly, is it legitimate to think of the project of ontology as an *explanatory* project in which we aim to say what there is by determining what we must ‘posit’ in our best explanatory theories?

I will argue that the answer is ‘No’. I will begin by going back to Quine, who insisted that meanings and properties would not add any explanatory power and should be rejected. I will argue that Quine is right that meanings and properties do not add explanatory power to his imaginary opponent McX’s view. I will argue, moreover, that this is typical *wherever* reference to certain entities can be introduced via trivial or ‘easy’ inferences from claims we already accept.³ This might lead one to ask whether we should accept that there are such legitimate ‘easy’ inferences.⁴ But the deeper point here is that, once we ask *why* language should have the built-in capacity for redundancy that makes ‘easy’ inferences possible, we can begin to see why the *demand* for additional explanatory power is entirely out of place. Thus, I will argue, if our goal is to answer the question ‘What is there?’, we should *not* pursue it via a project that aims to ‘posit’ only those entities that contribute ‘explanatory power’ to our ‘theories’.

This brings us back to the first question: Is ‘What is there?’ a deep and interesting question for ontology? I will argue that the answer again is ‘no’. I will close with some suggestions of where we might better focus our efforts.

1. When ‘positing’ entities does *not* add explanatory power

Quine’s official view is that we are ontologically committed to all and only those entities over which we must quantify in our best total explanatory theories. And this, of course, is not *quite* the same as requiring that any given putative (kind of) entity itself contribute explanatory power in order for us to posit it (or things of its kind). Perhaps if some total explanatory theory that quantified over Ks were sufficiently better—in terms of possessing theoretic virtues *other than* explanatory power—we would have reason to accept it over a rival that didn’t quantify over Ks. Nonetheless, it is not unreasonable to move from this to a *prima facie* case for accepting entities (and reducing parsimony) in a theory only if it enables us to increase other theoretic virtues (such as explanatory power) in a way that more than compensates for the loss in parsimony.

In any case, Quine himself certainly does make use of the idea that we should appeal to the *explanatory power* to be gained as a crucial criterion for whether or not we should accept that entities of a given kind exist. And his use of the idea initiated an extensive tradition in

² Along these lines, Stathis Psillos treats as “the only workable criterion of reality” that “something is real if its positing plays an indispensable role in the explanation of well-founded phenomena” (2005, 398).

³ Thus, it seems to apply to many debates about whether there (really) are properties, numbers, events, fictional characters, propositions, etc. Reference to a traditional God, or to libertarian free will, by contrast, plausibly cannot be introduced by trivial inferences from uncontroversial claims.

⁴ I make the case that we should accept that there are such easy inferences, and respond to various objections to that view, in my (2015).

metaphysics of arguing for/against certain entities on grounds of their allegedly contributing/failing to contribute explanatory power.

In ‘On What There Is’ Quine appeals to failures to contribute to explanatory power in his argument against his imagined opponent McX, who holds that “Ontological statements follow immediately from all manner of casual statements of commonplace fact,” so that, for example, “There is an attribute” follows from “There are red houses, red roses, red sunsets” (1953/2001, 10). In response, Quine writes:

“That the houses and roses and sunsets are all of them red may be taken as ultimate and irreducible, and it may be held that *McX is no better off, in point of real explanatory power, for all the occult entities which he posits under such names as ‘redness’*” (1953/2001, 10, italics mine).

Similarly, in denying that there are meanings, Quine writes that while we may speak of utterances as significant or insignificant, synonymous or heteronomous, “...*the explanatory value of special and irreducible intermediary entities called meanings is surely illusory*” (1953/2001, italics mine).

Since Quine, this form of argument has become ubiquitous in metaphysics—both among those who accept and those who reject the existence of contested entities. So, while Quine rejects ‘positing’ attributes or meanings on grounds that they would not add any explanatory power, others have argued that we must accept the existence of universal attributes to ‘explain’ why certain predicates apply to certain individuals (see Katz 1998). This form of argument has come to play a crucial role among those who reject (or accept) entities from properties and meanings, to moral properties, fictional characters and social groups.

Quine doesn’t say why he thinks that accepting attributes or meanings would not give McX any (additional) explanatory power. So why might one, with Quine, think that such entities don’t add explanatory power to McX’s ‘theory’?

Let us begin by considering ordinary empirical cases in which we think we *are* justified in positing an entity, or a kind of entity, because it would explain certain facts or observations. The paradigmatic cases for good inference to the best explanation arguments in ordinary empirical and scientific reasoning go like this: If we notice ragged holes in our packets of food and tiny feces on the floor, and hear scuffling in the night, we might make an inference to the best explanation and hypothesize that there are mice in the pantry, which are the cause of all of these observed effects. Doctors often hypothesize about the presence of microbes in the body, as the best explanation of various observed symptoms. And Leverrier and Adams (independently) famously inferred the existence of a planet beyond Uranus as the best explanation of unexpected deviations in Uranus’ orbit. These are classic arguments for positing the existence of something in order to explain a phenomenon. And I think Quine is right in suspecting that (even if they are good in their home empirical contexts) arguments of this form can’t be used to justify ‘positing’ attributes or meanings.

Since Quine doesn’t say why, let me give some reasons why we might doubt that accepting attributes or meanings (as McX does) adds explanatory power.⁵ McX says that statements, say, about the existence of properties “follow immediately from all manner of casual

⁵ See Scott Shalkowski (2010) and Juha Saatsi (2017) for other reasons to doubt the appropriateness of inference to the best explanation arguments in philosophical ontology, even if we allow them in everyday life and the empirical sciences.

statements of commonplace fact,” (1953/2001, p.10). That is, we can apparently derive nominative property talk or meaning talk, or come to quantify over properties and meanings (in the way that Quine treats as bringing ‘ontological commitment’ to ‘new entities’) by means of trivial inferences from sentences that don’t refer to them. So, from ‘the house is red and the barn is red’ we can apparently infer ‘the house has the property of redness and the barn has a property of redness’ and ‘there is a property of redness that the barn and house have in common’. Similarly: from ‘dog’ means the same as ‘Hund’, we can infer that ‘dog’ has the same meaning as “Hund”, and so that there is a meaning the words have in common.⁶

But if we are (as McX assumes) entitled to introduce new referring terms for these entities by hypostatizations from these simpler sentences, then the entities cannot be used to *explain* why the original sentences are true.⁷ To avoid getting caught up in debates about how to understand explanation in scientific cases, we can begin from a simple and uncontroversial observation: whatever an explanation is, dormitive virtue ‘explanations’ are *not* genuine explanations. As Moliere’s old joke has it, the doctor offers only the illusion of explanation if you ask, ‘Why do poppies make us sleepy?’ and the doctor replies, ‘Because they have the dormitive virtue’. You don’t need a full-blown theory of explanation to recognize that, whatever explanation may be, dormitive virtue ‘explanations’ are not explanatory. Yet the following all follow the pattern of dormitive virtue explanations: ‘The fact that barn is red is *explained by the fact that there is a property of redness possessed by the barn*’; ‘The fact that the predicate ‘red’ applies to the barn *is explained by the fact that there is a property of redness possessed by the barn*’; and “‘Dog’ means the same as ‘Hund’” *is explained by* ‘Dog’ and ‘Hund’ having the same meaning”. In each case, we introduce a fancier, nominative form of speech (referring to a property or meaning), but that cannot provide anything analogous to a scientific, causal explanation such as might justify us in positing the explainer.⁸ *Whatever* explanation may be, to restate a fact in fancier terms (or terms of new grammatical categories) is not to explain it.⁹

Quine’s McX thought that existence claims followed immediately—perhaps trivially—from all manner of commonplace fact. And I have aimed to flesh out Quine’s intuition here that

⁶ Of course, Quine would also reject the idea that there are such conceptual truths or analyticities underlying these inferences—but that is not the central point at issue here. Rather, it’s that, if we accept McX’s view, we do (as Quine insists) have reason to doubt that the relevant entities could add explanatory power.

⁷ At least not in the ordinary scientific or everyday empirical sense of ‘explanation’. ‘Explanation’ is a term used in many ways. Here I will be concerned only with a sense of ‘explanation’ relevant to our scientific explanations and empirical explanations in everyday life—since these are the contexts in which explanatory arguments for the existence of something seem relevant. This leaves open that there might be some other sense of ‘explanation’ in a *metaphysical* sense. Such purported ‘metaphysical explanations’, however, aren’t commonly thought to give us reason to think that the explainers exist. For example, while some will say that Socrates explains the existence of the singleton set {Socrates}, no one thinks that we have reason to posit the existence of Socrates because that would provide the best explanation of the existence of the set. Thus, I will leave questions about a distinctively ‘metaphysical’ form of explanation to one side here, to be treated separately.

⁸ As will become clear later, I don’t mean to suggest that the introduction of these nominal terms, and shift to a different grammatical form, serves *merely* to make the talk (or speaker) sound ‘fancier’. There may be a variety of functions served by the change in grammatical form.

⁹ Serious metaphysicians typically resist the idea that these are just restatements in alternative grammatical forms. But the fact that such trivial inferences are readily permitted in ordinary speech suggests that this *does* capture the introduction rules for the terms in ordinary English, and that those who think more is required are implicitly changing the rules. As I have argued elsewhere (2015), pressure may then be put on the serious metaphysician to clarify *what more is required* for the nominative statement to be true, beyond the truth of the uncontroversial statement with which we began.

if we are thinking along McX's lines, it seems that accepting the existence of properties and meanings will fail to add any real explanatory power to our theories.

The lesson to draw so far is that, if we do take claims that there *are* properties or meanings to follow by simple hypostatizations from claims like 'the house is red and the barn is red' or "'dog' can be translated as 'Hund'", then the relevant entities cannot aid in *explaining* the original facts. Thus far, Quine was right. Nonetheless, as we will see, a deeper understanding of 'easy' arguments will also give us reason to think that contribution of novel explanatory power was a misguided criterion to begin with.

2. Should we accept easy inferences to existence?

This lesson is of wider relevance today, given the recent interest in 'easy' ontological arguments for the existence of things of various kinds.¹⁰ For it reminds us of an important consequence of accepting easy arguments: if you think the existence of certain metaphysically controversial entities follows trivially from other (non-controversial) claims, you should not think of them as *explaining* the truth of those claims (or the facts stated in those claims).¹¹

But should we accept that there are such valid 'easy' arguments for the existence of properties, meanings, and other contested entities? This is of course a controversial matter.¹² Metaphysicians typically reject the trivial arguments. Instead, they think it is a highly *non-trivial* matter whether there are properties, and a non-trivial question whether the best 'total metaphysical theory' will 'posit' Platonic properties, Aristotelian *in rebus* universals, tropes, etc. Some (following Quine (1953/2001), or more recently Williamson (2007)) reject altogether the idea that there are analytic entailments.¹³ Others think it would take something like 'magic' for there to be analytic entailments that lead to conclusions that something (not mentioned in the original claim) *exists*.¹⁴ Still others appeal to long-standing 'bad company' objections, that allege that the supposedly analytic inferences can lead us to contradiction or conflict with known fact.¹⁵

There is no space here to review or respond to all of the reasons given against accepting easy arguments. For that I must refer readers to my (2015). Nonetheless, it may be worth reminding ourselves of the sorts of evidence we can appeal to *in favor* of the idea that there are such trivial entailments.

What sort of evidence can we give that a conditional of ordinary English is analytic? One way to go is to appeal to our common responses and linguistic intuitions: We can (with Schiffer (2003)) appeal to the felt redundancy of 'the house is red and the house has the property of redness'. We can appeal to our standard epistemic norms, which don't require further investigation to move from 'the house is red' to 'the house has the property of redness', whereas they do require further investigation to move from 'the house is red' to 'the house has a door'. We can (with Strawson and Grice (1956)) appeal to the different reactions we might have to someone who said "the house is red and it's not the case that the house has a door" (doubt or disbelief), versus to someone who said "the house is red and it's not the case that the house has the property of redness" (bewilderment, not knowing what to make of it). Or we can appeal to our different reactions to someone who says, "the house is red, *so* the house has the property of

¹⁰ See, e.g. Hale and Wright (2001), Schiffer (2003) and Thomasson (2015).

¹¹ Cf. my (2015, Chapter 3).

¹² For criticisms of easy arguments see, e.g., Yablo (2000), Hofweber (2005), Bennett (2009), Eklund (2009), and Evinne (2016).

¹³ I respond to these arguments in my (2007, Chapter 2) and (2015, Chapter 7) respectively.

¹⁴ For discussion and response, see my (2015, Chapter 6).

¹⁵ For discussion and response, see my (2015, Chapter 8).

redness”, versus “the house is red, *so* the house has a chimney” (acceptance versus puzzlement at the non-sequitur). We could appeal to the way in which we would correct those who violated or denied this rule—by aiming to understand what they might be trying to say, or attempting to teach them the rule, rather than questioning their evidence. We could appeal to the inappropriateness of inserting ‘probably’ in the ‘the house is red, so the house *probably* has the property of redness”, although such insertions are normally appropriate where mere evidential support is in question (“the house is in a snowy region, so it *probably* has a pitched roof”). And we can note that these are precisely the kinds of clue we use to identify analyticities in other cases—cases metaphysicians have not made contentious. In any case, we can each either do this in the first person, by reflection on our own linguistic intuitions (and on the assumption, for each of us, that we are reasonably typical competent speakers), or we can engage in empirical analysis or experimental work to see how broadly shared these reactions and intuitions are across the linguistic community.

Another point in favor of accepting that property terms, number terms, etc. are introduced by these sorts of rules—rules that take us from a non-referring use of a term (‘red’ in an attributive use or ‘two’ in a determiner use) to the relevant noun terms (‘the property of redness’, ‘the number two’)—is that it gives us the basis for a plausible account of how the latter terms could be introduced to language, and how the relevant entities could be known. That is, it makes clear how we can bridge up from the ‘core language’ we use in describing the perceptible world around us to acquire terms (on what Friederike Moltmann (2020) calls the ‘periphery’) for abstract entities, and to acquire knowledge of abstracta. If there are easy inferences licensed by introduction rules for the relevant terms, we can account for both of these by appealing to how terms for properties, numbers, and the like are introduced to the language by way of inferences from ‘core’ parts of language (speaking of two red apples) to introduce derivative terms via these rules. But if we deny that there are such rules, then it is hard to see how terms for such abstract entities could be *learned*—for clearly nothing like worldly observation, tracking, or ostension will work here. And how could we possibly come to know about such entities, if we can’t get there via an analytic route from other forms of talk, and other truths, to truths about numbers, properties, or other abstracta? It is clear how we can hope to get an account of how we could learn the relevant forms of language and come to know truths about the relevant kinds of objects, if there are trivial inferences. But those who deny that there are the relevant rules owe us an account of both how we could learn the terminology, and how we could come to know about the relevant entities.

As we will see, work in linguistics gives us some new considerations in favor of the view that the trivial inferences are perfectly acceptable. More importantly, it also helps show why it was misguided to ever expect things like properties, meanings, and the like to refer to ‘posits’ that could add ‘explanatory power’ to our theories, playing a similar role to posited mice, planets, or microbes.

3. Why would we want a language with such redundancies?

If we take seriously the linguistic evidence that it is redundant to say (for example) ‘The house is red *and* the house has the property of redness’, a new question arises: *Why would we have, or want, a language with such redundancies built in? What would introducing such vocabulary, via trivial inferences, do for us?*

To answer these questions, we should turn to work in linguistics that addresses *functional* questions about what expanding the language via such trivial inferences would enable us to do.

The work in linguistics that most directly and fully addresses these functional questions is work in systemic functional linguistics—an approach to the study of language begins from the idea that “language has evolved in the service of certain functions” (Halliday 1973, 14). As Simon Dik puts it:

Functional Grammar is based on a functional view of the nature of natural language. A view, that is, in which a natural language is first and foremost regarded as an instrument of social interaction by means of which human beings can communicate with each other... (1980, 1).

Systemic functional linguistics goes on to ask questions about ways in which the functions of language are reflected in “the internal organization of language as a system” (Halliday 1973, 15), and ways in which different features of language enable it to fulfill these functions.

Systemic functional linguistics may be unfamiliar to many philosophers who (if they have background in linguistics at all) have been more attuned to work in the Chomskian tradition. It has (so far) been far more influential in anthropology and education than in philosophy.¹⁶ And it focuses on a different range of questions about language than those addressed in the Chomskian tradition--questions about the functions language serves in human social life and culture, and about the ways language is structured that enable it to fulfill these functions: the very questions that we need to address here.¹⁷ As Halliday puts it, “I would defend the view that different coexisting models in linguistics may best be regarded as appropriate to different aims, rather than as competing contenders for the same goal” (1964, 13).

Systemic functional linguistics does share some common roots with philosophy. The tradition emerged from work in the Prague School of Linguistics from the 1920s and 30s, led by Vilem Mathesius, and with Roman Jakobson centrally involved. Jakobsen in turn was influenced by Husserl’s *systematic* approach to language in the third and fourth of the *Logical Investigations*¹⁸ where Husserl identified linguistic laws of ‘compounding’ that establish which meaning categories can be sensibly conjoined.¹⁹ Husserl also identified laws of ‘modification’ that govern the ways in which one meaning category can be transformed into another “as in the cases of nominalization, in which verbs, adjectives, adverbs, etc. are used as nouns” (Aurora 2015, 12). The emphasis on *function* developed under the influence of the anthropologist Bronislaw Malinowski, who insisted that we must study meaning by analyzing the functions of language in its culture, and whose work on language influenced both Wittgenstein and the English linguist J. R. Firth. Firth (who drew on work by Wittgenstein and Ryle (Bateman 2017, 14) and repeatedly refers to Wittgenstein’s focus on types of speech function (Firth 1962, 10)), was the teacher of Michael Halliday. Halliday in turn was influenced not only by Firth, but also by Sapir, Whorf, and work in the U.S. anthropological tradition. Halliday’s work played a central

¹⁶ In fact (except for some of my own recent work) I have seen zero references to it in philosophy. It has been particularly directed towards work in education—in helping to determine what children with language delays, or learning a second language, cannot do as effectively as their peers, with a goal of helping remedy these deficits.

¹⁷ Bavali and Sadighi (2008) argue that systemic functional linguistics is not a rival to familiar work in the Chomskian tradition. On the lack of rivalry, see also (Halliday 2014, 56). For further discussion of the place of systemic functional linguistics in the history of linguistics, see also Bateman (2017).

¹⁸ See Aurora (2015). Roman Jakobson “explicitly considers Husserl’s phenomenology [especially the third Logical Investigation] as one of the main sources underlying Prague structuralism and Russian formalism” (Aurora 2015, 14). Husserl also gave a lecture to the Prague Circle (invited by Jakobson) on November 18, 1935.

¹⁹ As I draw out elsewhere (2002), this Husserlian work also influenced Ryle’s later work on linguistic categories and category mistakes.

role in developing systemic functional linguistics in its modern form, which I shall draw on here.²⁰

Work in systemic functional linguistics helps make evident why we should have a language that introduces capacities for redundancies, such as those that are made explicit in trivial ‘easy’ inferences. Many of the terms that are introduced through trivial inferences are instances of what Michael Halliday calls ‘grammatical metaphors’.²¹ I should say immediately that I do *not* want to read—or think it’s appropriate to read—what Halliday and others in the systemic functional linguistics literature call ‘grammatical metaphors’ as involving *merely metaphorical speech*, to be interpreted in a fictionalist or simulating way. The term ‘metaphor’ derives from the Greek term *metaphora* for ‘transfer’ or ‘carrying over’. While familiar (lexical) metaphors carry over one *semantic* meaning, transferring it for another, *grammatical* metaphors carry over a term from one *grammatical category* to another, so that we shift grammatical category from adjective to noun (‘the barn is red’ to ‘the barn has the property of redness’), from noun to verb (‘he had a knife’ to ‘he knifed someone’), and so on. While grammatical metaphors aren’t limited to nominalizations, they include (nominalized) talk of properties, processes, events, causes, time, possibility, thought, etc.

Here is the idea. In early language—both developmentally and evolutionarily (and across languages)—we begin with ‘congruent’ meanings: nouns for things, verbs for processes... (Halliday 2009, 117).²² An early subject-verb-object utterance like ‘man clean car’ is congruent. Congruent terms are “evolutionarily and developmentally prior” (Halliday 2009, 117). In early language development, these more basic ‘congruent’ nouns such as ‘stick’, ‘car’, ‘hole’, and ‘ball’ are acquired in response to observations of interest in the environment: “In contexts of observation, recall, and prediction” (Halliday 1975, 27). Many of these congruent terms seem to be something like those Huw Price identifies as ‘e-representational’. That is, they are terms that have the job “to *co-vary* with something else—typically, some *external* factor or environmental condition” (Price 2011, 20). As Halliday puts it, these terms function to contribute to the child’s learning about the environment (1975, 27-28).

But by departing from congruent formulations, grammatical metaphors may be introduced, adding to the ‘meaning potential’ of language (Halliday 2009, 116).²³ “In grammatical metaphor, one grammatical class takes over from another” (Halliday 2009, 126)—so we no longer have congruence (Halliday 2009, 116). “A meaning that was originally construed by one kind of wording comes instead to be construed by another”, e.g. when we move from saying ‘wash the car’ to speaking of ‘a carwash’ or ‘taking the car in for a wash’ (Halliday 2009, 117). So, for example, we may shift from speaking of washing the car, to saying we ‘gave the car a good clean’, and go on to refer to ‘a carwash’, ‘the cleaning of the car’, and so on.²⁴

²⁰ For a brief historical overview, see Malmkjaer (1991), 141-6.

²¹ For further discussion of grammatical metaphors and the various forms they may take and effects they can have, see Taverniers (2017).

²² For a history of the use of the term ‘congruent’ and discussion of criteria for congruence, see Taverniers (2003) and (2017).

²³ This also seems like a way of making good on the idea that there are optional additional frameworks that can be added onto the ‘thing’ language, in Carnap’s (1950/1956) terms.

²⁴ P.F. Strawson seems to be onto a similar point. He discusses the way the subject/predicate form is applied in the basic case to spatio-temporal particulars, but by ‘imaginative extension’ this logical form is also carried over to ‘higher levels’, enabling us to make predications of colors, numbers, etc. The basic case is a model for the other cases, but “From this fact spring both the delusions of Platonism and the delusions of anti-Platonism. They are indeed, but two sides of the same delusion” (1974/2016, 30).

With these alternative formulations, speakers have options about how they express themselves, and what they can go on to do with those forms of expression. For once we have these nominalizations we can go on to say, ‘*The carwash* was a very messy enterprise’ or ‘*The cleaning of the car* took all afternoon’.

So, suppose we begin *not* by asking apparently metaphysical questions, such as: “Are there meanings, properties, causes, time, or kinds?” or “What *are* meanings, properties, causes, time, or kinds?” Instead, suppose we begin by asking: *What functions do these nominalized forms of language serve?* Though it may also sound more ‘learned’ and be used for rhetorical effect, the primary motivation for nominalization is, “a functional one: by nominalizing we are able to do things with the text that we cannot do in unnominalized text” (Eggins 2004, 95). Everyday grammatical metaphors (which, for example, enable us to shift from saying ‘she bathed’ to ‘she took a bath’, or from ‘they danced’ to ‘they did a dance’), enable us to quantify and qualify processes. That in turn enables us to say things such as: “she had a relaxing hot bath” or, “they did this well-known Romanian folk-dance”... (Halliday 2009, 135).²⁵ Consider how much more cumbersome it would be to express the former without grammatical metaphor (“She bathed with hot water and found it very relaxing”?). As Suzanne Eggins notes, with nouns, we can expand our clauses to do many things that we can’t do with other parts of a clause (like verbs), viz: count, specify, describe, classify, qualify (Eggins 2004, 96). We can, for example, go on to say that the act after intermission included three folk dances, that the bath was very relaxing or the car wash very messy, or that it was *the bath* (not the shower) which caused the downstairs flood. This enables greater lexical density: that is, a greater proportion of content carrying words (main verbs, nouns, adjectives, adverbs) in the text.²⁶ It also (as Thomas Hofweber (2005) noted, and Eggins (2004) argues) can help us organize texts rhetorically, putting the focus on different parts of information—as we might correct the prior speaker in insisting that *The number of folk dances after intermission was four*.

But while everyday language does contain some grammatical metaphors (such as those above), grammatical metaphors get much more numerous, sophisticated and layered when we enter scientific theory construction and bureaucracy. The ‘learned’ kind of grammatical metaphor functions similarly in Chinese and in English. And these “seem to represent tendencies common to the elaborated discourse of science and technology, government and bureaucracy, in all languages” (Halliday 2009, 138). There is historical evidence that use of grammatical metaphors explodes with the development of science, technology, and bureaucracy. As a result, *many such terms play an important role in explanation and theory construction—though not as posited explainers*. Indeed, Halliday emphasizes that the function of much grammatical metaphor is to enable theory construction:

It is no exaggeration to say that grammatical metaphor is at the foundation of all scientific thought. You cannot construct a theory—that is, a designed theory, as distinct from the evolved, commonsense theory incorporated in the grammar of everyday discourse—without exploiting the power of the grammar to create new, ‘virtual’ phenomena by using metaphoric strategies of this kind (Halliday 2009, 119).

²⁵ Halliday also shows that parallel nominalization constructions appear in Chinese to similar effect (Halliday 2009, 135).

²⁶ Non-content carrying words include prepositions, conjunctions, auxiliary verbs and pronouns.

“Such discourse evolved as the language of technology and science” and was “moulded by the demands of the physical sciences into its modern form” (Halliday 2009, 125). Such grammatical metaphors were also in ancient Greek and Chinese and Sanskrit—indeed Halliday claims that “this metaphoric shift from the clausal to the nominal construal of experience seems to be a characteristic of scientific discourse in every language” (Halliday 2009, 123n21). Grammatical metaphors provide “the cornerstones of a *theory*, which is a designed semantic subsystem for reorganizing experience in a technological environment” (Halliday 2009, 120); “with the emergence of experimental science, another major semiotic shift took place, further exploiting the stratal potential of language”, e.g. Newton’s *Opticks* is flush with novel nominalizations, as can be seen in passages such as:

“*For in all whites produced by Nature, there uses to be a mixture of all sorts of Rays, and by consequence a composition of all Colours*” (156-7, cited in Halliday 2009, 120).

These nominalizations (‘whites’, ‘mixture’, ‘composition’...) enable the flow of information and construction of theories (Halliday 2009, 120)—so that we can identify the *mixture* of rays of various sorts as the *cause* of *whites* produced in nature.²⁷

If this is the right account of how such terms enter our language and of the functions they serve, we can draw some important conclusions. The first is that we should *accept* that the inferences are trivial—that there is *nothing more it takes* to have a bath than to bathe; and nothing more it takes for these three cardiac patients to share the property of having high blood pressure than for Smith, Lopez, Simonov, etc. to each have their blood pressure elevated. The alleged reasons against thinking that the inferences are trivial tend to arise from biases inherited from a misguided approach to metaphysics—or from being misled by grammatical metaphors.

The second conclusion is that we can see that the requirement that entities such as properties must ‘add explanatory power’ is misguided. As long as we have these terms, introduced to our vocabulary by a process of grammatical metaphor, we should (indeed, to be coherent with the rules of use for the introduced terms, we must) accept that *there are* properties, processes, events, and the like. To think that, if we are to accept that there are such things, we must think of them as *contributing explanatory power* is to mistake the theoretic role of the relevant terms, as well as to misconstrue the rules by which they work. These terms may be serving fundamentally different functions than terms for mice or microbes, yet it could nonetheless be a complete mistake to try to eliminate them from our theories—as they play other essential roles in our ability to theorize in ways that require us to construct generalizations and explanations.

Given that we have these terms in our language, and that we are licensed to make inferences that introduce them given the built-in nominalizing potential of the language, further questions about whether there ‘really are’ such things as properties, meanings, etc. are otiose. Just as easy ontology has it.²⁸

Nonetheless we can also add something deeper than what easy ontology alone gave us. For we can add a functional story about *why it is useful* to have a language with capabilities for such redundancy, and what functions are served by adding such nominalized forms of speech.

²⁷ Grammatical metaphors also play crucial roles in enabling us to construct *bureaucracies* with generalized rules giving permissions and entitlements—but I will leave that to the side here.

²⁸ Nor is there any ‘algorithm for elimination’ (in the words of Stephen Schiffer) for them—we won’t turn up empirical counter-evidence against their existence or some ‘mistake’ in our chain of reference, since the successful introduction of such nominalizations does not (like that for congruent nouns introduced observationally) require new empirical evidence.

4. Other roles in explanation

The idea from systemic functional linguistics, that grammatical metaphors play an important role in explanation and theory construction, fits well with ideas that have occasionally been voiced in philosophy: that nominalizations of various kinds (including terms for numbers, properties, propositions and other abstracta) can figure in and aid in our scientific-explanatory theories—though in quite other ways than serving to name observed or unobserved ‘explainers’.

Cohesively with the results of systemic functional linguistics, several philosophers have raised interesting suggestions about useful roles such terms play in explanations, by lending expressive power to the language, simplifying our statements of theory or law, enabling us to give a unified articulation of what is to be explained, and so on. Laws are involved in explanations, and Stephen Yablo (2005) argues that reference to numbers enables us to finitely state physical laws that it would otherwise take an infinitely long series of infinitely long sentences to state—thereby simplifying and improving our explanations. Alejandro Perez Carballo (2016) argues that mathematics gives us a picture of logical space “rich in structural propositions”, which can help us get propositions with “high explanatory value” and systematize data and make predictions without getting lost in irrelevant detail (2016, 475). Augustin Rayo (2009, 25) argues that learning logical truths is useful because it increases one’s ability to distinguish between intelligible and unintelligible scenarios, and so to use old information in new ways. For example, one can move from the information that there are 7 apples and 12 pears to the information that there are 19 pieces of fruit (Rayo 2009, 26). This isn’t a matter of acquiring new information, but of using old information in new ways, enabling us to answer new questions like ‘how many pieces of fruit?’ (Rayo 2009, 27). And I have argued elsewhere (2015, 157 n. 23) that introducing new noun terms can pragmatically enhance our ability to formulate explanations. For example, medical researchers might do better to have a noun term ‘heart attack’ to enable them to demand, and perhaps ultimately give (in general terms), an explanation of why *heart attacks* occur, rather than having to put things in terms of explaining why Smith’s heart stopped beating and Lopez’s heart stopped beating and... Similarly, while we may not have much use for explaining what a red house and a red barn have in common, we might have far more use for property talk in explaining *what properties all of the relevant cardiac patients had in common* (was it the property of *having high cholesterol*, or of *having a certain genetic mutation*, or...?) Using property terms can help us in *formulating* explanations, as it may (for example) enable us to *express* certain kinds of question, such as, ‘what do those flowers favored by hummingbirds have in common?’, or ‘what is it that humans can see but rabbits cannot?’ Similarly, as Sally Haslanger’s (2012) work suggests, some important *questions* (why is it that African Americans have a higher rate of police stops?) cannot even be formulated as *demands for explanation* without having something like generalized race terms in our vocabulary.

In short, we need not deny that reference to entities like properties, events, numbers and propositions may *figure in* and play important roles in enabling, enhancing, and simplifying our explanations. But this emphatically does *not* mean that they are ‘posited’ as metaphysical ‘explainers’ in the way that mice or microbes are—they do not (to use terminology from Stathis Psillos (1999, 110)) ‘fuel’ the explanatory success of the theories they figure in.²⁹

How can we distinguish what entities fuel a theory’s explanatory success? Katherine Hawley offers a helpful criterion:

²⁹ For more on the distinction between figuring in and fueling explanations, and arguments that numbers, properties, and the like do not fuel relevant explanations, see my (2019).

“If a claim H is to be involved in generating a prediction in a way that entitles it to share in the confirmation which successful prediction brings [and so to ‘fuel’ its success]... H must satisfy two conditions with respect to the generation of the prediction. First, it must be the case that the theory minus H cannot generate the prediction alone. Second, it must also be the case that there is no available, sensible alternative to H which could have done the work just as well” (2006).

Along these lines, we might add, if an entity ‘posited’ is to ‘fuel’ a theory’s explanation, (rather than just to figure in it) it must be the case that the theory *minus the posit* could not generate the explanation alone. I have argued elsewhere (2019) that typical entities contested in metaphysical debates, such as meanings, properties, and numbers, do not pass this test, and so cannot be thought to ‘fuel’ a ‘theory’s’ success (even though the relevant terms may usefully *figure* in our theories).

Now you might say: if reference to such entities can do good explanatory work, introducing new *kinds of* explanation, improving or simplifying old, then whether or not they ‘fuel’ the explanation, shouldn’t we posit them? Well, (though I don’t favor the term ‘posit’ in such contexts) in a sense *yes*, we should accept that *there are* such things. But the argument will go via a different route, and won’t involve anything like ‘positing’ explainers.

We should begin by noticing an important point central to systemic functional linguistics, but often overlooked by serious metaphysicians. That is: our terms may serve different *functions*. It is an idea clearly visible in work by Wittgenstein, Ryle, and Sellars, and more recently in work in the neo-Pragmatist tradition by people like Huw Price, Michael Williams, and Robert Brandom.³⁰ Huw Price has introduced a distinction between terms that serve an *e[xternal]-representational* function (those whose job “is to co-vary with something else—typically, some *external* factor, or environmental condition” (2011, 20)) and those that are merely *i-representations* (in the sense that they have a certain inferential role). I suspect that we should subdivide these categories significantly further.³¹ In the case of e-representations: there are some terms, observationally introduced, whose job is to co-vary with some external factor or environmental condition (perhaps: ‘wolf’, ‘river’), and others that are *theoretically* introduced to refer to theoretic entities that would fuel an *explanation* of our observations (perhaps: ‘electron’, ‘black hole’).

The point here is that many of those terms that refer (if at all) to entities contested in specifically *ontological* debates appear to serve functions *other than* tracking entities or fueling explanations. Perhaps nominative number terms function to enable us to state laws in finite form that couldn’t be otherwise stated, and/or to systematize data in ways that enable us to see patterns and get explanations we wouldn’t otherwise get. Perhaps talk of properties or events enables us (*inter alia*) to state in briefer and more unified (generalized) terms what it is that needs explaining, or to unify and/or simplify the explanations given. I don’t want to commit to any specific functional theses here—I would prefer to leave that to the experts on these areas of discourse. I aim only to point to the idea that *alternative functional stories* come naturally once we see these terms not as aiming to refer to entities that would *fuel* our explanations, but as terms that nonetheless may *figure* usefully in them.

A common response at this stage is to insist that my opponents, too, could accept that certain terms may (also) play these roles in explanation or elsewhere. There, is, it is true, nothing

³⁰ The idea also plays a key role in expressivist work on moral discourse, for example by Simon Blackburn (1993) and Allan Gibbard (1990).

³¹ And I suspect that Price would be open to this further ramification of functional pluralism.

to stop them from accepting these assessments. But if (in acknowledging these other roles) they accept that the role of the relevant terms is *not* to refer to a posited entity to fuel explanations, then why should ability to contribute explanatory power be a relevant criterion for accepting the entities? Alternative functional stories should be equally plausible for everyone, but the point here is that they undermine the across-the-board use of an explanatory power criterion, and undermine the explanatory view of the project of metaphysics.

5. Roles outside of explanation

This brings us to a further question: *should metaphysics explain?* The project of ontology has often been presented as a matter of determining what entities we *should* or *must* ‘posit’ as part of a best *total explanatory theory*, and it has become common to treat *whether or not the alleged entities add explanatory power* to our theories as a central criterion in determining whether to accept them. But we can now see why it is misguided to take ability to fuel scientific explanations as a criterion for accepting whether or not certain entities (properties, meanings, etc.) exist.

As (I hope) has become clear from the discussion so far, those who take contribution of explanatory power to be an important criterion for whether or not we accept that there are entities of any kind fail to acknowledge an important point: *that there may be different functional roles* for different forms of language, and different activities we engage in when using language. When we are theorizing, we may sometimes introduce terms to track things we observe, and we may sometimes introduce terms for (as yet) unobserved entities that are ‘posited’ to explain various observed phenomena. But we are not *always* doing that.

If the proper role of terms for properties is not to fuel explanations, but rather to enable us (for example) to even simply ask *questions* about what various patients had in common, then the failure of ‘posited’ properties to ‘fuel our explanations’ should be no mark against introducing and making full use of property vocabulary. If the functional role of nominative number terms includes enabling us to simplify our statements of laws or to structure logical space in a way that is epistemically useful (Perez Carballo 2016, 461), then the fact that numbers don’t themselves serve as causal explainers is no knock against them. On this view, in short, where Quine (and many who followed him) went wrong is not in *denying* that such entities can be *explainers* (carefully stated, there is a sense in which this is exactly right). Instead, we go astray if we assume that contribution of explanatory power *is a relevant criterion* for accepting forms of language that enable us to quantify over meanings, properties or numbers.

In fact, we are not always *theorizing* at all. Consider debates about whether there are chances, moral properties, and natural properties. Such debates often have centered around questions about whether such entities could *explain* why our credences should be guided in certain ways, why we should act in certain ways, or why we should theorize in certain terms (see Sider 2011). But there are familiar problems with these attempted ‘explanations’. Shamik Dasgupta (2018) argues that, even if we accept such metaphysical posits, they can’t do the work of explaining *why* our credences, theorizing, or actions should be guided by them. As he puts it (roughly): posit whatever whatnots you like, they can’t explain why we *ought to* believe, theorize, or act this way. This is what Dasgupta calls the ‘problem of missing value’. As Dasgupta nicely emphasizes, the problem of missing value arises in questions about why moral facts should be action-guiding, why chances should be credence-guiding, and (the case he newly presses) why natural facts or properties *should* guide our theorizing. In each case, the metaphysician ‘posits’ some feature of the world that is supposed to have normative import: for

how we should act, believe, or theorize. In each case, the problem of missing value arises with the demand to *explain why* our action, beliefs, or theorizing should be guided by this worldly feature (Dasgupta 2018, 287). And in each case, there seems to be no adequate explanation available, since the entities ‘posited’ are (in Dasgupta’s words) ‘normatively inert’ (2018, 289). Here again, the metaphysical ‘posit’ seems unable to *explain why we should* act, believe, or theorize in certain ways. And here again, we might be able to see why we can’t get genuine explanation, if we take talk of moral facts and properties (say) to arise from hypostatizations that take us from, for example, from “You shouldn’t kill”³² to “Killing has the property of being morally wrong” and “It is a moral fact that killing is wrong”. The latter forms of expression may provide useful alternatives that enable us to do new things with moral language (enabling us to say, for example, that killing is *more wrong* than theft; or to generalize in saying that those who do *things that are morally wrong* should be punished). But such hypostatized claims cannot *explain* the original ‘should’ claims from which they are derived.

The problem of missing value gives us good reason to deny that ‘positing’ such things can ‘explain’ our norms. Does that mean that we shouldn’t accept that there are moral properties, chances, or natural properties? I think that’s a good conclusion to draw *only if* these terms have *the function of referring to entities to serve as explainers*.

But we have good reason to deny this. I have argued above that many nominalized terms, introduced as ‘grammatical metaphors’, may play useful roles in explanation and theory construction, *even when they do not function to refer to posited ‘explainers’*. But now we can go beyond this. For attending to moral discourse makes it clear that the functions of an area of discourse *needn’t have anything to do with explaining at all*. It is plausible that moral discourse fundamentally serves a *normative* function, acting as a kind of “chaperone for human behavior” (Warren 2015) or providing a kind of “social glue, bonding individuals together in a shared justificatory structure and providing a tool for solving many group coordination problems” (Joyce 2006, 117). Introducing *nominalized* moral terminology serves important functions towards these ends: enabling us to compare *moral requirements*, to demand and give *justifications* for various requirements, to *reason about* what the requirements are, to comparatively evaluate the strengths of *competing reasons for action*, and so on.³³

As a result, even though adding reference to moral properties doesn’t add explanatory power to a theory, the move to say ‘such entities wouldn’t add explanatory power, so we shouldn’t say there are such things’ is a complete wrong turn, a *non sequitur*. If I am right that such terms were never *supposed to function* to refer to explainers, then the failure of their (alleged) referents to explain is no knock against them. We should reject the explanatory power *criterion*, not the entities.³⁴

6. Rethinking the project—and the question—of ontology

³² Plausibly, this inference is only valid when the first (“You shouldn’t kill”) is uttered as a *categorical imperative*, not a mere recommendation or imperative of prudence.

³³ For further development of these ideas, see Warren and Thomasson (2023) and my (in progress).

³⁴ Another consequence that’s interesting and important from my point of view is this: I have argued elsewhere (2015, 156ff) that the ‘easy approach to ontology’ can get us a form of ‘simple realism’ about many kinds of entity—saying that there are numbers, properties, events, modal facts, etc. *in the only sense that has sense*. What it can’t get us is a form of *explanatory realism* that would take these entities to explain certain observations, or to ‘explain what makes the relevant discourse true’. But now we can see that this is no loss—that the tendency to think such explanations should be forthcoming is the product of an unreflective assumption that the function of all forms of discourse is explanatory.

We should also reject the view that the project of ontology is to determine which entities we must ‘posit’ as part of a best total *explanatory* theory. Whether we formulate this as a demand that, to accept the existence of Ps, Ps themselves must fuel our explanations, or as the looser demand that we should accept the existence of Ps only if we must quantify over them in our best explanatory theories, the demand can only look suitable if we are blind to the other roles terms serve in explanation (for the first formulation), and (for the second) if we are blind to the other things we do with language other than to formulate explanatory theories. Perhaps we should even give up using quasi-scientific terminology to characterize what we’re doing, and give up talking in terms of metaphysical ‘theories’ and ‘positives’.³⁵

Some readers may have noticed my use of scare quotes whenever I have used the term ‘posit’ above. I can now clarify the reason for distancing myself in this way.³⁶ To treat entities such as properties, meanings, numbers and other entities contested in metaphysics as ‘positives’ of certain metaphysical ‘theories’ is to suggest that the existence of these entities is something like an empirical hypothesis, to be confirmed or disconfirmed by further evidence. Some people are undoubtedly attracted to this—perhaps in hopes that it will lend metaphysics the respectability of the sciences. But it leads us down the wrong path, encourages us to ask the wrong questions (“what facts can these entities enable us to explain?”), and to employ inappropriate standards in answering questions about what there is (insisting that entities we accept should contribute explanatory power to our ‘theories’). The misleading terminology is symptomatic of the underlying problems that come from assuming that metaphysical views should be *like scientific theories*.

If the above is correct, we also need to rethink the idea that addressing the question ‘What is there?’ provides an interesting and deep question to be addressed by ontology.³⁷ For if there are easy inferences to existence claims, then ‘What is there?’ is not a deep or interesting question for ontology to pursue. Instead, many of the contested questions (“are there properties/ numbers/ facts...”) can be answered via trivial inferences from uncontroversial truths. Many of the remaining, more difficult existence questions (“are there weapons of mass destruction in Iraq?”, “are there cases of COVID-19 in my county?” “are there social octopuses?” “is there life on other planets?”) are best addressed by journalists, scientists, or detectives—not by ontologists.

If we persist in addressing questions about the existence of properties, numbers, and the like, then the route to answering them will take a different path: If adding such terms to our vocabulary enables us to fulfill an important function (whether in *figuring in* explanations, in communicating, justifying, or renegotiating norms, or...)³⁸—then we have at least *prima facie* reason to add (or retain) these terms in our vocabulary. And once we do so (if we just add these

³⁶ I have similar reservations about speaking of metaphysical ‘explanations’. Even if this is defined in some way that makes it quite different from the causal explanations of the sciences, the terminology may encourage an inappropriate model of what metaphysics can do.

³⁷ Of course, Quine himself would have denied that there is work here for a *distinctive discipline of ontology* to do—though it has been common for metaphysicians since Quine to think of their task in this way.

³⁸ Note that this shouldn’t be read as an ‘if and only if’. In line with the tenets of easy ontology, I accept that there may be valid trivial inferences that entail the existence of entities of certain kinds, *regardless of* whether or not the relevant terms serve an important function. It is simply that the case against eliminativism will be easier to make where we can see that the terms clearly *do* fulfill a useful function. (Thanks to Joshua Gert for suggesting that I note this point.)

terms as conservative extensions, and in ways that entitle us to make trivial inferences to the existence of the relevant entities), we can often get ‘easy’ arguments that such things exist.³⁹

The underlying trouble with thinking of metaphysics on a scientific model may be seen as the following. Metaphysics has traditionally been concerned with all kinds of topics in life—or, put in the linguistic mode—all kinds of areas of discourse. Accordingly, it enquires about moral properties, modal facts, objects and properties, persons and actions, causes and events, mental states, and (increasingly) interesting social phenomena such as genders, races, classes and the like.

In adopting a quasi-scientific epistemology, serious metaphysics subjects ontological questions to the standards suitable for scientific theories. This is a mistake. Science may aim to causally explain and predict. But not all discourse does. Subjecting all discourse to the same standards is a bit like evaluating all artifacts—from knives to salad bowls to hairdryers—according to how well they cut steak. Even *within* science, as we have seen, terms may fulfill many different functions—not just serving to quantify over entities that are meant to fuel our explanations. *Beyond* science we have a far wider range of functions our language serves to fulfill—whether in regulating our conduct, conveying rules of use for our terms, assigning praise and blame, guiding our methods of reasoning, etc.

If metaphysics is to address all these topics, then we need to begin by assessing how each of these forms of discourse works, and what functions it serves—not just forcing them all into a quasi-scientific mold, and assessing them by criteria suitable for a scientific explanatory theory.

Many of the entities at issue in metaphysical debates do not add (or aim to add) ‘explanatory power’ that ‘fuels’ a theory’s explanation. And appeals to metaphysical whatnots typically can’t explain why we *ought to* do anything either. But that is not a knock against the view that there are such things—it’s a knock against a scientific conception of ontology that thinks of the project of ontology as ‘positing’ entities to fuel our explanations.⁴⁰

If the above is correct, then we should give up thinking that ‘what is there?’ provides a deep and interesting question for a discipline called ‘ontology’ to answer. We can of course continue to ask and answer existence questions. Questions about what exists may be perfectly legitimate, but the *philosopher’s* existence questions can often be answered by trivial inferences, and *everyday* existence questions (about the existence of weapons of mass destruction, distant planets, or nearby microbes) are to be answered by journalists or scientists—not by practitioners of some distinctive discipline of ‘ontology’. Nor should we think that the *project* for ontology is to determine which entities to ‘posit’ as part of some explanatory enterprise. Perhaps we should, as Carnap suggested, even give up the word ‘ontology’, as likely to mislead us into thinking there is such a distinctive task, or worthwhile project.

But what of other candidate projects for ontology? Those will have to be evaluated separately. Nonetheless, the above considerations should lead to us to approach some prominent suggestions with caution. For example, a popular alternative is to think (with Armstrong (2004) or Cameron (2010)) that the project of ontology is not to determine what exists, but to determine what the ultimate *truthmakers* are for our claims, or to determine ‘what makes what true’. But

³⁹ For a formal statement of the conditions that must be met to ensure that trivial inferences are acceptable (in response to ‘bad company’ problems), see Chapter 8 of my (2015). This is of course not to say that introducing such vocabulary *brings the relevant things into existence*—see my (2015, Chapter 6).

⁴⁰ As Price puts it, “this kind of functional pluralism challenges a kind of mono-functional conception of language that seems implicit in Quine’s own view—for Quine, *the* significant task of the statement-making part of language is that of recording the conclusions of an activity that is ultimately continuous with natural science” (2011, 13).

this project again seems in danger of assuming that all indicative statements *have the function of tracking certain features of the world*⁴¹—such that the relevant statements are *defective* if there fails to be some feature of the world that can *explain what makes them true*. But as we have seen, many of the philosophically contested areas of vocabulary (concerning properties, numbers, propositions, modality, morality etc.) are precisely areas of vocabulary that apparently serve entirely different functions—in which case the insistence that we should ‘search for their truthmakers’ may lead us astray. This, of course, is just to provide preliminary words of caution. I give a fuller examination of various metaphysical projects in *Rethinking Metaphysics* (forthcoming).⁴²

What should we do instead? We should begin by aiming to understand how the various areas of human discourse we are interested in work. What are the functions of the relevant terms? How do they work—what rules do they follow? Should we keep them in our vocabulary, revise them, or replace them? As I have argued elsewhere (2015), once we accept the terms and understand how they work, we can typically answer the existence questions easily and affirmatively—often just by undertaking trivial inferences. As I have suggested, we can look to linguistics, especially systemic functional linguistics, for help in determining the functions of and rules followed by various forms of discourse—and this is broadly empirical work. But that is not to say that we should think of *ontology* as a quasi-scientific project of determining which entities to posit as part of the best total explanatory theory. On the contrary, the empirical work in linguistics may help reveal why such an ontological project is misguided, relying on misunderstandings about the functions of the relevant forms of discourse.

But that is not the only work that remains to be done. The most interesting work, in my view, which metaphysics often has done and certainly can do, lies not in answering existence questions, but rather in conceptual research and development. That may involve, in part, interesting descriptive linguistic/conceptual questions about the functions and workings of individual terms, or about their interrelations as part of a more total linguistic and conceptual scheme (I think of this as a form of conceptual reverse engineering). We might, for example, (continue to) ask: What are the relations between our concepts of freedom, responsibility, and morality? Are there a few basic conceptual categories we must have in order to cognize an ‘external world’ at all? And it may involve interesting and difficult *normative* conceptual and linguistic questions, about what terms and concepts we *should* employ for various purposes, and how we should employ them (should we make use of race concepts at all? Should we think of freedom in compatibilist terms, and if so which? How should we conceive of art and its relation to society, intention, and beauty?)... Our philosophical work may also involve still deeper questions about what purposes we should pursue. We cannot hope to answer such normative questions *merely* by empirical inquiry. These will be questions for conceptual engineering and

⁴¹ That is, roughly, that they serve as e-representations in something like Price’s sense (2011, 20) of representations that have as their job to “*co-vary* with something else—typically some *external* factor, or environmental condition” (2011, 20).

⁴² See my (2020) for further discussion of the truthmaker project. In *Rethinking Metaphysics* (forthcoming), I argue that the traditional conception, explanatory conception, structural conception, truthmaker conception, and many fundamentality and grounding conceptions of metaphysics all suffer from failing to note the functional pluralism of language.

conceptual ethics, not for ‘ontology’ as it is usually conceived. But they are questions we might far more fruitfully pursue than the traditional ‘question of ontology’.⁴³

Works Cited:

- Armstrong, David M. (2004). *Truth and Truthmakers*. Cambridge: Cambridge University Press.
- Bennett, Karen. (2009). “Composition, Colocation, and Metaontology”, in Chalmers et al.
- Blackburn, Simon (1993). *Essays in Quasi-Realism*. New York: Oxford University Press.
- Cameron, Ross (2010). “How to have a radically minimal ontology”, *Philosophical Studies* 151: 249-64.
- Dasgupta, Shamik (2018). “Realism and the Absence of Value”. *The Philosophical Review* 127.3 (2018): 279-322.
- Dik, Simon C. (1980). *Studies in Functional Grammar*. London: Academic Press.
- Eggs, Suzanne (2004). *An Introduction to Systemic Functional Linguistics*. Second edition. London: Continuum.
- Eklund, Matti (2009). “Bad Company and Neo-Fregean Philosophy”, *Synthese* 170 (3): 393-414.
- Evnine, Simon (2016). “Much Ado about Something from Nothing; or, Problems for Ontological Minimalism”. In Stephan Blatti and Sandra Lapointe, eds., *Ontology After Carnap*. Oxford: Oxford University Press: 145-64.
- Gibbard, Allan (1990). *Wise Choices, Apt Feelings*. Cambridge, Massachusetts: Harvard University Press.
- Hale, Bob and Crispin Wright (2001). *The Reason’s Proper Study: Essays towards a Neo-Fregean Philosophy of Mathematics*. Oxford: Clarendon.
- Halliday, Michael (1973). *Explorations in the Functions of Language*. New York: Elsevier.
- Halliday, Michael (1975). *Learning How to Mean: Explorations in the Development of Language*. New York: Elsevier.
- Halliday, Michael (2009). *The Essential Halliday*. Ed. By Jonathan J. Webster. London: Continuum.
- Haslanger, Sally (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford University Press.
- Hawley, Katherine (2006). “Science as a Guide to Metaphysics” in *Synthese* 149: 451-470.
- Joyce, Richard. (2006). *The Evolution of Morality*. Cambridge: Cambridge University Press.
- Katz, Jerrold J. (1998). *Realistic Rationalism*. Cambridge: MIT Press.
- Hofweber, Thomas (2005). “A Puzzle about Ontology”. *Nous* 39:2: 256-83.
- Kraut, Robert (2010). “Universals, Metaphysical Explanations, and Pragmatism”. *Journal of Philosophy* 107 (11): 590-609.
- Moltmann, Friederike (2020). “Natural Language Ontology”, in R. Bliss and J. Miller, eds. *Routledge Handbook of Metametaphysics*.
- Perez Carballo, Alejandro (2016). “Structuring Logical Space”, *Philosophy and Phenomenological Research* 92/2: 460-91.
- Price, Huw (2011). *Naturalism without Mirrors*. Oxford: Oxford University Press.
- Psillos, Stathis (1999). *Scientific Realism: How Science Tracks Truth*. London: Routledge.

⁴³ For further discussion of the idea that much work of metaphysics should be reconceived as work in conceptual engineering (including both backwards-looking reverse engineering, and forward-looking constructive engineering), see my (2017) and (2016) and, for fullest discussion, see my *Rethinking Metaphysics* (forthcoming).

- Psillos, Stathis (2005). "Scientific Realism and Metaphysics". *Ratio* 18: 385-404.
- Quine, W. V. O. (1953/2001). "On What there Is". In *From a Logical Point of View*. Cambridge, Massachusetts: Harvard University Press.
- Rayo, Augustin (2009). "Towards a trivialist account of mathematics", in Bueno O., Linnebo Ø. (eds) *New Waves in Philosophy of Mathematics*. New Waves in Philosophy. Palgrave Macmillan, London.
- Saatsi, Juha (2017). Explanation and explanationism in science and metaphysics *Metaphysics and the Philosophy of Science: New Essays*, M. Slater and Z. Yudell (eds.), Oxford University Press, 162–191, 2017.
- Schiffer, Stephen (2003). *The Things we Mean*. Oxford: Oxford University Press.
- Shalkowski, Scott (2010). "IBE, GMR and Metaphysical Projects" in *Modality: Metaphysics, Logic and Epistemology*, ed. Bob Hale and Aviv Hoffmann. Oxford: Oxford University Press.
- Sider, Theodore (2011). *Writing the Book of the World*. Oxford: Oxford University Press.
- Strawson, P. F. and H. P. Grice. (1956). "In Defense of a Dogma". *Philosophical Review* 65, 2: 141-58.
- Strawson, P. F. (1974/2016). *Subject and Predicate in Logic and Grammar*. New York: Barnes and Noble.
- Swoyer, Chris (1999). "How Ontology Might be Possible: Explanation and Inference in Metaphysics". *Midwest Studies in Philosophy* XXIII: 100-31.
- Taverniers, Miriam (2003). "Grammatical Metaphors in SFL: A Historiography of the Introduction and Initial Study of the Term". In *Grammatical Metaphor: Views from Systemic Functional Linguistics*. Ed. Anne-Marie Vandenberg, Miriam Taverniers, and Louise Ravelli. John Benjamins.
- Taverniers, Miriam (2017). "Grammatical Metaphor". In Tom Bartlett and Gerard O'Grady, eds. *The Routledge Handbook of Systemic Functional Linguistics*. London: Routledge.
- Thomasson, Amie L. (2007). *Ordinary Objects*. New York: Oxford University Press.
- Thomasson, Amie L. (2015). *Ontology Made Easy*. New York: Oxford University Press.
- Thomasson, Amie L. (2016). "Metaphysical Disputes and Metalinguistic Negotiation", in *Analytic Philosophy* (July 2016): 1-28.
- Thomasson, Amie L. (2017). "What can we do, when we do metaphysics?", in Giuseppina d'Oro and Soren Overgaard, eds. *The Cambridge Companion to Philosophical Methodology*, Cambridge: Cambridge University Press: 108-29.
- Thomasson, Amie L. (2019). "Replies to comments on *Ontology Made Easy*", *Philosophy and Phenomenological Research*. July 2019. <https://doi.org/10.1111/phpr.12619>
- Thomasson, Amie L. (2020). "Truthmakers and Easy Ontology", in *Oxford Studies in Metaphysics*, ed. Karen Bennett and Dean W. Zimmerman, Oxford: Oxford University Press: 3-34.
- Thomasson, Amie L. (forthcoming). *Rethinking Metaphysics*. New York: Oxford University Press.
- Warren, Mark (2015). "Moral Inferentialism and the Frege-Geach Problem". *Philosophical Studies* 172 (11): 2859-2855.
- Warren, Mark and Amie L. Thomasson (2023). "Prospects for a Quietest Moral Realism", in *Oxford Handbook of Moral Realism*, ed. David Copp and Paul Bloomfield, Oxford: Oxford University Press: 526-553.

- Williams, Michael (2011). "[Pragmatism, Minimalism, Expressivism](#)". *International Journal of Philosophical Studies* 18 (3):317-330.
- Williamson, Timothy (2007). *The Philosophy of Philosophy*. Oxford: Blackwell.
- Yablo, Stephen (2000). "A Priority and Existence", in Paul Boghossian and Christopher Peacocke, eds. *New Essays on the A Priori*. Oxford: Oxford University Press: 197-228.
- Yablo, Stephen (2005). "The Myth of the Seven". In Mark Eli Kalderon, ed. *Fictionalism in Metaphysics*. Oxford: Oxford University Press.