

Conceptual Ethics and the Work of Metaphysics

Version for Chapel Hill Colloquium
(Draft as of 10/7/2016)

Amie L. Thomasson

I have argued elsewhere (2015) that ontological questions can be answered easily: that questions about whether things of a given kind exist can often be answered (in the affirmative) by trivial inferences from uncontested premises. Metaphysical modal questions, too, I have argued (2007, 2013 and in progress), can typically be addressed by a combination of empirical work and conceptual analysis.

But if such large swaths of the territory of traditional metaphysics can be addressed so easily, this naturally leads to the question of what metaphysics can do, and what value we can see in the traditional and contemporary work of metaphysics. I have argued elsewhere (2016 and forthcoming) that we should see metaphysics as fundamentally involved in conceptual work: not merely the kind of descriptive conceptual work found in the work of such figures as Gilbert Ryle, Peter Strawson, or Frank Jackson, but also—and often more interestingly—*normative* conceptual work: work in what is often called ‘conceptual engineering’ or ‘conceptual ethics’.¹ This echoes suggestions by others—classically, Carnap; more recently, Simon Blackburn, who writes that “just as the engineer studies the structure of material things, so the philosopher studies the structure of thought” (1999, 2). More recently still, similar views have been defended by David Plunkett (2015) and Matti Eklund, who suggests more broadly that “Philosophy should... be thought of as *conceptual engineering*” (forthcoming, 2 in draft).²

Now, there are two ways of reading this claim. One is as claiming that a great deal of historical and contemporary work in metaphysics can be aptly interpreted as (overtly or covertly) engaging in conceptual work. So, on this model, traditional debates about personal identity, for example, can at least in some cases be interpreted as debates about what concept of *person* we have, or what concept we *should* adopt for some purpose or other; debates about whether numbers exist can be seen as debates about whether we *should* adopt what Carnap called the ‘number framework’—including noun terms for numbers; debates about the persistence conditions of works of art may be seen as debates about when we should count works of art as preserved or destroyed. The other way of reading the claim is that (putting aside any issues about the proper interpretation of past debates) that we *should* come to understand the legitimate work

¹ There does not seem to as of yet be any consensus about how to use the terms ‘conceptual engineering’ and ‘conceptual ethics’. I will follow Burgess and Plunkett in mainly using the latter, since ‘conceptual engineering’ may suggest a construction of new concepts than (also) issues of choosing among, repairing and restructuring old concepts—all of which I have in view as the normative conceptual work here at issue.

² Eklund (forthcoming, 14-15 in draft) argues that philosophy could be seen as conceptual analysis, but that it would be better to take it as conceptual engineering in the sense of “a study of what concept best plays the theoretical role of our concept of [e.g.] truth and what features this concept has...” (15)). He also cites Blackburn in his 1999 introduction to philosophy, and Brandom (2001) as characterizing philosophy as conceptual engineering—as (he says) David Chalmers (2012) does, without using the label.

that can be done by those in metaphysics as involved in (descriptive and normative) conceptual work. Whether we call it ‘metaphysics’ or not, here the idea is that determining how our actual conceptual scheme works, and what conceptual scheme we *should* adopt, are legitimate and important projects that remain, even if traditional metaphysical versions of the relevant questions are abandoned. I have argued elsewhere for both the descriptive and normative versions of this claim, but I won’t argue for them here.³

For me, the idea that metaphysics can be understood as (to be) engaged in conceptual work—the most interesting part of which is normative conceptual work—is particularly appealing, since it provides a way of adding strength and depth to the deflationary metametaphysical position. For this enables even deflationists to re-interpret and still see some value in much of the work of traditional metaphysics, to the extent that we can see it as implicitly engaged in conceptual analysis and/or normative conceptual work, and leaves much work to be done in determining how central concepts such as *person*, *freedom*, *death* or *art* should be employed.

But how should we go about engaging in this kind of normative conceptual work? Here, a threat is likely to arise again for the deflationist. For many of those attracted to serious metaphysics will accept that issues in conceptual choice are important—even central to the metaphysical or, more broadly, philosophical enterprise. But, they will insist, the proper way to approach issues of conceptual choice is *via metaphysics*. Call these ‘metaphysics-first’ approaches to conceptual choice. For (some might say) we should, for example, only adopt concepts that turn out to refer to the things that *exist*—but then we must first settle the metaphysical existence question before we can settle the question in conceptual ethics. (Call this the ‘existence first’ approach). Others might take a different route (call it a ‘structure first’ approach): Ted Sider argues that the central questions of metaphysics involve ideological questions about which conceptual scheme we ought to adopt, where that, however, is seen as a matter of asking “which notions carve perfectly at the joints”, revealing the structure of reality (2011, 5).

If the right approach to conceptual choice is a metaphysics-first approach, then of course the metaontological deflationist is in trouble. For the deflationist’s idea was to replace the traditional conception of metaphysics as engaged in discovering deep metaphysical facts with a more modest conception of metaphysics as engaged in conceptual work—the most interesting and contentious side of which is normative conceptual work. But if normative conceptual work, properly done, relies on metaphysical work and is answerable to metaphysical facts of the world (say, about existence or structure), then the deflationary project doesn’t get off the ground.

Of course, even the critics of deflationism acknowledge that there is an alternative pragmatic approach to conceptual choice. On this view, we do not first answer metaphysical questions about, say, the existence of numbers to figure out whether we should adopt nominative number terms into our conceptual scheme. Instead, we may have pragmatic reasons for introducing noun terms for numbers (e.g. to simplify our statements of scientific laws). Given the

³ David Plunkett (2015) argues that many classic debates in metaphysics can be understood as—at least *implicitly*—involved in negotiating what terms or concepts we should use, and I have made a similar case elsewhere (2016 and forthcoming). In my view, ordinary language philosophy and phenomenology can easily be understood as engaged in conceptual work, but so, as I have argued elsewhere, can classic historical work including Hume on necessity, Locke on personal identity, Berkeley on material objects, as well as more recent work such as Baker on persons, Bell or Davies on art, Gert, Culver and Clouser (2006) on death, Haslanger on race and gender, Millikan (1984) on function, etc

way number terms are introduced, once we have them in our scheme we are entitled to make inferences that tell us that there are numbers—apparently answering the metaphysical question about the existence of numbers, in the only sense it has. In brief, on the pragmatic view our conceptual scheme may be chosen pragmatically, and its choice will (often combined with empirical investigations) *entitle* us to make claims about existence or structure.

But pragmatic approaches to conceptual choice are widely distrusted, and rejected on grounds that they cannot sufficiently account for our intuitions about worldly structure, or worse still that they make our conceptual choices merely arbitrary, subjective, or power-driven. If that is the best we can do, then a ‘metaphysics’ (or replacement for metaphysics) that can only proceed on that basis might seem less than desirable.

I will argue, however, that metaphysics-first approaches to conceptual choice face more difficulties than their defenders often acknowledge—particularly epistemological difficulties. Then I will move on to sketch a pragmatics-first approach. I will argue that the pragmatic approach can overcome the perceived shortcomings better than is commonly thought, and that it can preserve the virtues of its competitor while avoiding its difficulties. The goal here is not to show that the metaphysics-first approach is unworkable (I’m better at defense than offense). Instead, the goal is to show that the pragmatic approach to conceptual choice is more plausible than you might have thought—perhaps even so plausible as to make the difficulties that come with the metaphysics-first approach come to seem avoidable and unnecessary. And if the pragmatic approach to conceptual choice is a viable (and maybe even preferable) alternative to metaphysics-first approaches, then the deflationist can indeed appeal to conceptual choice as a useful model for what metaphysics can and should do.

Metaphysics-first Models

Someone who takes the ‘metaphysics-first’ approach to conceptual ethics, in broad strokes, is someone who accepts that many of the debates in metaphysics may be seen as involved in something like negotiating which terms, concepts, or conceptual scheme we ought to use, but who thinks that it *is the metaphysical facts of the world* that determine which conceptual scheme we ought to use. Many metaphysicians who are tempted to think that conceptual ethics should play a role in metaphysics might be inclined towards a metaphysics-first view over a pragmatic view of conceptual choice. For this approach seems better suited to capture the metaphysician’s self-conception that metaphysics is about the world, that it is not just a matter of investigating our concepts or terms or how we use them (Sider 2011, 6). It is also thought to ensure that) it does not make our conceptual choices arbitrary, subjective, or power-driven: for there is an objective worldly standard to which such choices are answerable.

The existence-first view

Metaphysics-first views may take various forms.⁴ The existence-first version of the approach is relatively familiar. Peter van Inwagen (2016, Synthese) suggests something along these lines, when he considers Carnap’s reading of a nominalist (‘Norma’) as suggesting that we switch to a nominalist language, and responds:

⁴ There may be other versions available, too, such as one that would say that we should adopt a conceptual scheme that properly characterizes the *essences* of things in the world... There may also be room for versions that are neither metaphysical nor pragmatic, but transcendental. For brevity, I will leave those options to one side here.

...if Norma is engaged in a project in logic or semantics, *why* is she engaged in it? Why, obviously, because she does not think that there *are* any things but concrete objects, because she thinks that there *are* no numbers (2016, 17).

In general, one metaphysics-first approach to conceptual choice is to think that one ought first to do the metaphysical work, of determining what things (really) exist, and then adopt only those terms or concepts that succeed at referring. So one can take the nominalist or organicist to be recommending revisions in our conceptual scheme alright (that we cease using nominative number language or terms for composite inanimate objects—at least in positive atomic assertions). But, as such metaphysicians will insist, they are doing so *on metaphysical grounds: that we should make these conceptual changes because there are no numbers or composite inanimate material objects to refer to*. In this way, metaphysics (here, in particular, ontology) remains primary, as it is the standard by which we are to engage in and assess conceptual choices (or at least many conceptual choices: those that involve the choice of nominative terms and concepts, which are *supposed to refer*).

Carnap rejects the existence-first approach directly:

Many philosophers regard a question of this kind [a philosophical question about the existence or reality of a system of entities] as an ontological question which must be raised and answered *before* the introduction of the new language forms. The latter introduction, they believe, is legitimate only if it can be justified by an ontological insight supplying an affirmative answer to the question of reality. (1950/1956, 214).

An existence-first approach to conceptual choice, of course, faces certain familiar problems. One such problem arises from the increasing resistance to thinking that the primary job of metaphysics is to answer existence questions, or even to thinking that these are deep questions for metaphysics. I have attempted to fuel that resistance elsewhere (2015), arguing that those existence questions that can be answered can be addressed by straightforward empirical and/or conceptual means—often via trivial inferences from uncontested premises. I won't try to make that case again here, but clearly taking an existence-first approach to conceptual choice requires thinking that there are deep metaphysical facts about what 'really' exists to be discovered.

This leaves such an approach with obvious liabilities, the most pressing of which are epistemological problems: how are we supposed to determine the answers to these ontological existence questions, which we can then put to work in engaging in conceptual choice? In certain cases, one can defend a negative existence claim by alleging a contradiction in the concept, but such cases are rare and generally contestable—as the elements that are supposed to figure in the conceptual explication can often be rejected or reinterpreted. More often, those who defend ontological views do so by appealing to general metaphysical principles to justify their choice over alternatives (e.g. 'no co-location', 'no causal redundancy', 'nothing that is not in space and time', 'nothing without causal powers')—principles that push the epistemological question back a notch.

An increasingly popular approach is to take the epistemology of metaphysics to parallel the epistemology of science. That is: we find our best metaphysical theories, by appeal to the theoretic virtues such as empirical adequacy, simplicity, explanatory power, unity, and so on, and take these virtues to give us reason to think that our theories are true. This is an approach popularized by David Lewis' arguments for possible worlds on grounds of inference to the best explanation, and recently defended by Laurie Paul who argues that "We use theoretical

desiderata as guides to truth in metaphysics just as we use such desiderata as guides to truth in science” (2012, 21).

A series of worries might be raised about the analogy, however. The first is that competing metaphysical theories are, as their defenders acknowledge, typically empirically equivalent. By contrast, in science ‘even approximate empirical equivalence is very rare’ (Paul 2012, 12). If it is only theoretic virtues other than empirical adequacy at issue in debates between distinctively metaphysical rival theories, it is unclear why we should think that these differences—in the simplicity of the theory or its ontology, in its explanatory power, etc.—are really apt to track truth, rather than just marking the usefulness of the theory for limited creatures like ourselves (see Bricker forthcoming). A related worry, as Karen Bennett (2009) and Uriah Kriegel (2013) have argued at length, is that quite typically competing metaphysical theories involve simply trading off one theoretic virtue for another, leaving us at sea in aiming to determine which metaphysical theory to choose. Paul accepts that there is a difference of degree here, but insists that this difference in degree doesn’t undermine the truth-conduciveness of appeals to the theoretic virtues in metaphysics, “If such theoretical desiderata are truth conducive in science, they are also truth conducive in metaphysics” (2012, 21).⁵

But even if one is prepared to accept that the theoretic virtues (other than empirical adequacy) are truth-conducive in the sciences, and one is willing to live with relative indeterminateness and uncertainty for metaphysics, deeper worries can be raised. For there are grounds for thinking there may be a difference not just in degree, but in kind, between metaphysical theories and scientific theories—differences that prevent us from thinking that any truth-conduciveness the theoretic virtues (beyond empirical adequacy) have in the scientific case carries over to metaphysical theories. Huemer (2009) works through four different accounts of the evidential value of parsimony in empirical theorizing, and argues that none applies to the philosophical cases, suggesting that in typical philosophical contexts ontological simplicity has no evidential value. Shalkowski (2010) argues that inference to the best explanation can be empirically shown to be a reliable mode of inference where it concerns observables (so that there is the possibility of independent access to confirm its results), but not where its conclusions concern unobservable facts (2010, 177). Shalkowski concludes that, while inference to the best explanation may be perfectly good in the ordinary empirical cases that motivate it, “there is little hope to be found in the use of [inference to the best explanation] to settle metaphysical questions” (2010, 184).⁶ Saatsi (forthcoming) argues that there are differences in kind for the use of inference to the best explanation in science (whether one deals with claims about observable or unobservable entities) versus in metaphysics. These differences, he argues, suggest that its reliability in science doesn’t carry over to metaphysics, in part since it is hard “to conceive of a naturalistically acceptable account of the truth-conduciveness of explanationism in metaphysics” given the absence of empirical feedback to guide our explanatory practices in metaphysics—

⁵ Like Paul, Sider accepts that there are differences in degree between using the theoretic virtues in science and metaphysics, “We employ many of the same criteria—whatever those are—for theory choice within metaphysics that we employ outside of metaphysics. Admittedly, those criteria give less clear guidance in metaphysics than elsewhere; but there’s no harm in following this argument where it leads: metaphysical inquiry is by its nature comparatively speculative and uncertain.” (2011, 12).

⁶ This argument, it seems, would also potentially undermine the idea that we can use inference to the best explanation to argue from the success of our best ‘total’ theory to the joint-carvingness of its ideology. For we similarly lack any kind of independent access to confirm whether or not this sort of inference is good as a way of settling questions about structure.

though we may have such an account available for its use in inferences about both observables and unobservables in science.

Unfortunately, there is not space to try to settle this issue here. But it is worth noting that relying on parallels between metaphysical and scientific appeals to theoretic virtues relies on much wider issues about whether those parallels hold—and whether, even if we assume the theoretic virtues other than empirical adequacy are truth-conducive (not merely pragmatic) in the sciences, we have reasons to think that carries over in metaphysics. That is an investigation I shall have to leave for another occasion, noting only that major questions remain about how and whether one can solve the epistemological difficulties facing those who adopt an existence-first approach to conceptual choice.

The structure-first view

But a different metaphysics-first approach to conceptual choice is available. Rather than appealing to a metaphysical inventory of what really exists to guide conceptual choice, one could instead look to facts about structure to serve as the metaphysical measure of what conceptual scheme we ought to adopt. Sider develops the approach, arguing that metaphysics is not ‘about what there is’ but rather about ‘structure’ (2011, 1)—and the central questions of metaphysics involve asking “which notions carve *perfectly* at the joints” (2011, 5). On this model, one can see the work of metaphysics as involving conceptual choice, where this is a matter of seeking a best *ideology*—best in the sense of best mapping the real joints of the world. As Sider writes:

“A central task of metaphysics has always been to discern the ultimate or fundamental reality underlying the appearances. I think of this task as the investigation of reality’s structure” (2011, vii).

Investigating reality’s structure, in Sider’s view, is a matter of investigating what is the ‘privileged description’ of reality—one that uses the right concepts. The view is motivated strongly by intuitions about structure—that it is just objectively better to employ a concept like *blue* than *bleen*, or like our chemical concept of lithium rather than a concept that applies to lithium on earth but not to the same kind if found on Mars (Sider 2011, 7).

And so at least part of the task of metaphysics on this view does involve conceptual choice:

“For a representation to be fully successful, truth is not enough; the representation must also use the right concepts, so that its conceptual structure matches reality’s structure” (2011, vii)

This is a view, then (as Plunkett 2015, and Burgess and Plunkett (2013, 1093) also note) that coheres with the idea that metaphysics should centrally involve issues of conceptual choice. But it is clearly not a deflationary position that takes conceptual choice to be a merely pragmatic matter: instead, the goal of conceptual choice (at least when doing fundamental metaphysics) is to choose the conceptual scheme that best matches reality’s structure.

This approach differs from the existence-first approach in at least two ways. First, it is insulated from arguments that metaphysical existence questions are misguided, for we do not try to first answer existence questions, on this model, to guide our conceptual choice. Secondly, it is potentially of much broader relevance to issues in conceptual or linguistic choice. For existence questions can only be asked using noun terms, or perhaps nominalized predicates (e.g. ‘does the property of being tall exist?’). By contrast, on Sider’s view, questions of structure can be asked regarding terms of any grammatical category, “we can ... ask [not only which predicate terms

but also] which function symbols, predicate modifiers, sentence operators, variable binders, and so on, get at the world's structure" (2011, 85).

But the structuralist approach also retains the heavyweight metaphysical spirit, in thinking of our conceptual choices as answerable to metaphysical facts. For we ought to choose those concepts which best carve at the joints of reality, or which map onto bits of structure in the world (viii). "The goal of metaphysics is to give a fundamental description of the world, and doing so requires more than merely saying what there is" (2011, viii).

To think that our conceptual choices should be answerable to facts about structure, one must of course accept that there *are* such facts about structure. These must be distinctly *metaphysical* facts, not ordinary facts about *empirical* similarities and differences. The talk of 'joints of nature' might obscure this point: it is an empirically discoverable matter what the joints are in a side of beef, so one might be tempted to think that it is a straightforward empirical matter what the joints of reality are, and so what concepts we should have—at least provided we adopt this one normative principle: that we should adopt a conceptual structure that carves at the joints. But this, of course, misses the point of Sider's generalized notion of structure. Answering questions like whether quantificational notions like 'there exists', modal concepts, tensed concepts, and logical concepts 'carve at the joints of reality' (Sider 2011, 7-8) clearly not a matter of simply tracking empirically detectable similarities and differences. Instead, facts about structure in this sense are fundamental *metaphysical* facts—structure is a metaphysical 'posit'. This leaves us with a 'placement problem' for these structural facts, analogous to placement problems we are familiar with for moral and modal facts: how do we 'place' these facts about (metaphysical) structure in the natural/empirical world? How do they relate to other sorts of facts?

Once we see the breadth of the notion of structure, the epistemological difficulties become clear, too: for discovering structure is not a straightforward empirical matter analogous to discovering the joints of a side of beef by wiggling one's knife. Sider gives an answer to how we can come to know these facts about structure: we can begin from our best *scientific* or *total* theories, and take those theories to give us reason to believe that their ideology carves at the joints, revealing structure:

We should believe generally what good theories say; so if a good theory makes an ontological claim, we should believe it. The ontological claim took part in a theoretical success, and therefore inherits a borrowed luster; it merits our belief. This is all familiar; but a believer in structure can say more. A good theory isn't merely likely to be *true*. Its ideology is also likely to carve at the joints. For the conceptual decisions made in adopting that theory—and not just the theory's ontology—were vindicated; those conceptual decisions also took part in a theoretical success, and also inherit a borrowed luster. (2011, 12).

So, "We have defeasible reason to believe that the conceptual decisions of successful theories correspond to something real: reality's structure" (2011, 21).

While this isn't straightforwardly empirical, it *is*, Sider claims, analogous to work done in the deep theoretical corners of science. This then gives us "an epistemology for structure":

Claims about structure can be supported by evidence, and so are not inferentially isolated, and so are not in danger of unintelligibility... What is needed for progress in these issues is a more sophisticated and detailed understanding of the epistemology of our more theoretical endeavors, such as mathematics and theoretical physics, including their foundations (2011, 15).

So understood, the approach is to engage in scientific theory choice according to the usual methods, and then “regard as joint-carving the ideology that is indispensable in your best theory” (2011, 14).

This is a somewhat different appeal to the theoretical virtues than we saw above. The idea here isn’t that competing *metaphysical* theories can be evaluated by their possession of theoretic virtues, but rather that we weigh up our *total* theories (prominently including our best scientific theories⁷)—and that the confirmation of the best theory gives the epistemic backbone to confirm the ideology: “... the confirmation of a theory confirms its ideological choices and hence supports beliefs about structure. A theory’s ideology is as much a part of its worldly content as its ontology” (2011, 13).⁸

The best grounds for accepting our scientific theories is generally thought of as coming from their empirical successes. To make Sider’s suggestion work requires an extraordinarily broad view of what is confirmed with the success of a theory: that success confirms not just the empirical content, but also the ontology of a theory and, more crucially for present purposes, also confirms the ideology—giving us reason to think of it as reflecting the structure of the world. But, as scientific realists are painfully aware, if we think of empirical success as confirming our scientific theories *in toto*, we run into trouble with the pessimistic induction. For our best-confirmed past theories have all been discarded as science progresses. This, then, threatens to undermine the idea that success (even of our best current theories) is a sign of their total truth.

Commonly, scientific realists respond to the pessimistic induction by aiming to narrow down what they claim is confirmed by a theory’s predictive success, distinguishing which parts, claims or aspects of a theory are, and which are not, confirmed by it. If we can identify the aspects of a theory that are confirmed, the idea is, we can retain grounds for realism at least about those aspects of the theory (and see them as what is retained in subsequent theories), even in the face of theoretic changes that might involve rejecting certain of its posits or replacing portions of its conceptual scheme. There are different ways to go here: some, with John Worrall (1989), have suggested that one can resist the pull of the pessimistic induction by taking the success of scientific theories to confer confirmation not on their full theoretical content or ontological posits, but only on the ‘structure’ they reveal—where ‘structure’ is used in quite a different sense from Sider (it is the structure captured in, for example, a series of equations about the behavior of optical phenomena—not the ideology or ontology of a theory of light) (1989, 117). Katherine Hawley (2006) (following Stathis Psillos 1999) suggests a reasonable criterion for distinguishing what *claims* are and are not confirmed with a theory’s predictive success: a claim’s ‘involvement’ in a successful scientific theory can be thought of as giving some reason

⁷ Perhaps, as Sider tentatively suggests, not including the special sciences (2011, 14).

⁸ Sider acknowledges that sometimes our evidence doesn’t support a unique pair of ideology and theory, and acknowledges moreover that sometimes “it is hard to see what evidence *could* be mustered in favor of one pair rather than another” (2011, 13). But he is relatively sanguine about this problem. His relaxed attitude about the epistemological problems for metaphysics ultimately relies on a companions in guilt argument: that metametaphysical deflationists don’t get the epistemic high ground: ‘since metametaphysical critiques are just more metaphysics, they raise all the old epistemic questions’ (2011, 83). But, as I have argued elsewhere (2015, Chapter 10) that there is a way for the deflationist to retain the epistemic high ground, and reduce the epistemological mysteries that beset both traditional and Siderian heavyweight metaphysics. (Roughly, one can reject the idea that the quantifier ‘carves at the logical joints of the metaphysics-first approach to conceptual choice are genuine disadvantages as compared with a deflationary view. world’ not by taking a metaphysical stand on what structure the world has and lacks, but rather by arguing that this misinterprets the *function* of logical terms like the quantifier (see my 2015, Chapter 10.) The epistemic mysteries that come with this version of a structure-first view are a genuine disadvantage as compared with a deflationary approach.

to think it's true when it is involved in generating a prediction in a way that entitles it to share in the confirmation (2006, 462), when it (as Psillos puts it) "fuels" the theory's predictive success (1999, 110).⁹ Philip Kitcher distinguishes the essential 'working posits' of a theory from its idle 'presuppositional posits' (1993, 149). There is no space to enter into the details here. The crucial point is that taking as broad a view as Sider does, in which all aspects of a successful theory are confirmed by its success, may unwittingly undermine attempts to defend scientific realism against the pessimistic induction. We need a more fine-grained view of what does and does not inherit the luster of predictive success—and whether ideology is to be on that list remains an open question.

There are some reasons for doubting that it should. Stephen Yablo (2009, 519) has argued that presuppositions about the existence of abstracta are 'fail-safe' in the sense that, even if we think that certain sentences (e.g. 'the number of planets is eight') presuppose the existence of numbers, that presupposition could fail without mattering to the truth of the assertive content of the sentence.¹⁰ That is, even if there 'turn out to be' no numbers, the assertive content of the sentence—that there are eight planets—does not change in truth-value. Such fail-safe presuppositions of our theoretical claims make them unlike other presuppositions, e.g. if saying 'Vulcan orbits the sun five times in an Earth year' presupposes that there is a planet Vulcan, then failure of *that* presupposition, by contrast, (in Yablo's words) 'wrecks' the whole assertive enterprise. I have argued (2014) that if we accept Yablo's approach, his conclusions can and should be generalized—that in general *ontological* assumptions fit his model as fail-safe presuppositions.¹¹ That is to say, one might suppose that (positive, atomic) use of terms such as 'table', 'mereological sum', or 'property' presuppose that there are tables, mereological sums, or properties. Or, alternatively, using Sider's approach, one might think that using these terms in our best scientific theories presupposes that they are *joint-carving* terms. But even if we think that our use of such terms has such ontological presuppositions—either regarding the existence of things referred to, or regarding the joint-carvingness of the terms—those presuppositions are fail-safe, in the sense that they could fail without altering the truth-value of the assertive content of the claims of the theory. For if I say 'the mereological sum of the particles in solution weighs 29 grams', the assertive content of this (that the particles in solution jointly weigh 29 grams) retains its truth-value regardless of whether or not there 'really are' mereological sums, or whether the terminology of mereological sums is joint-carving. This gives one natural place to divide what is and is not confirmed with a theory: for it is natural to think that it is the assertive

⁹ A claim H is properly involved in generating a prediction when it satisfies two conditions: It "must be the case that the theory-minus-H cannot generate the prediction alone" and "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, 462). But while Sider is clearly right that our ideological choices can aid theoretical success, it seems, almost a category mistake to ask whether our ideological choices can 'fuel' predictive success in the manner at issue here. For a conceptual scheme is used in stating a theory, including its predictive claims and those from which we infer them; it is not itself a claim of a theory with respect to which we can ask whether or not it is needed to generate the prediction. I will have to leave further development of this idea for another occasion.

¹⁰ The assertive content of the sentence is its analytic implications that remain when we 'subtract' the presupposition of the existence of numbers.

¹¹ Yablo goes on to argue that where such presuppositions are fail-safe, ontological questions regarding them are 'moot' in the sense that there is nothing to settle them. I disagree with Yablo about this metaontological conclusion (see my 2014 and his 2014), but nonetheless think that the mechanism of 'subtraction' and identifying the assertive content of statements may be useful for a variety of purposes—perhaps among them distinguishing what aspects of a theory are and are not confirmed by its predictive success.

content of a theory that is confirmed by a theory's predictive success. Where structural or existential ontological presuppositions are fail-safe, making no difference to the truth of the assertive content of the claims of the theory, we may not be entitled to think that they inherit confirmation from its success. At any rate, this is one avenue to explore in examining more deeply whether one can hold—as the epistemology of the structure-first view requires—that the success of a theory gives us reason to think that its ideology is joint-carving.

Generality worries

Finally, even if one puts aside the ontological and epistemological worries, there are reasons to think the structure-first model cannot provide a *general* guide to conceptual choice. To be fair, it is not clear that it was ever intended to do so. Certainly Sider does not present it as a general approach to conceptual choice, though he does think that we should to accept as *fundamental* those notions that carve *perfectly* at the joints (2011, 5); these, Sider thinks of as including certain concepts of physics, logic and math.

But we still need a method for choosing among such social concepts as race and gender, economic and political concepts, moral concepts, aesthetic concepts, etc. Though they are not at all plausibly thought to carve perfectly at the joints, as Sider himself notes, such concepts (as 'person' and 'cause') may be deeply "embedded in our conceptual scheme. These concepts play a role in many central aspects of our thought: thought about ourselves, moral responsibility, deliberation, control, prediction, explanation, and myriad others" (2011, 73).

Those who favor a structure-first approach have several options. One could suggest that we should then jettison such concepts—at least for purposes of metaphysics. Or we might replace each such expression with an improved expression 'that we stipulate is to stand for the joint-carving meaning in the vicinity', and then ask a new question using the new term that is '*superior* to the original question, for it concerns reality's fundamental structure' (2011, 74).

But given the important role such concepts play in our lives, one might well think that, even if they turn out to not be joint-carving, it would be a mistake to abandon or replace them.¹² It also won't do to simply say, in each case, that we should adopt that concept that carves at the most fundamental joints possible. For, as Sider notes, they might not be the right sort of joints: political concepts, for example, might not be properly interpreted as carving at a natural land-based boundary: even if there is a "physically special line through the Ural mountains, that line isn't thereby the determinate boundary of Europe" (2011, 32).

A deeper reason it would seem to be a mistake to replace our non-joint-carving concepts with the joint-carving concepts 'in the vicinity' (even if there are some) is that many of these terms seem not to be *aiming at carving the world at its joints at all*. Sider accepts that a term "like 'amulet'... is not 'trying' to stand for a joint-carving meaning..." (2011, 74; see also 32), so that it would be a mistake to interpret it as carving at some joint, as "the correct interpretation must maximize the combination of charity and eligibility; thus a highly joint-carving

¹² In his response to Elizabeth Barnes, Sider acknowledges that he should have made clear that he was speaking only of a 'certain portion of metaphysics'; which 'we may call 'ultimate metaphysics'', and that he does not mean to deny that metaphysics includes "many questions other than those about fundamental reality" (forthcoming, 1). These presumably include questions we might ask with race or gender concepts, or perhaps with other concepts of ordinary life, whether or not they are, or even aim to be, joint-tracking. And this is a good move. For even if we limit it to eliminating those concepts *for purposes of metaphysics*, that would dramatically constrain what metaphysics can do, and its relevance to human life.

interpretation might nevertheless be incorrect if it is too uncharitable. Not all words carve at the joints after all” (2011, 74). 32). As we might naturally say, it all depends what the term was *supposed* to do for us, what its *function* is.

This observation—that terms may have a plurality of different functions, not all of which involve an aim to track worldly structure—gives grounds for more general doubts about the appropriateness of taking a generalized structure-first approach to criteria for concept choice. Although the structural approach seems natural for natural kind terms (where Lewis proposed it), for many terms it is at the least a philosophical open question whether the terms are *supposed to be* joint-carving, or to supposed to be *tracking worldly features* at all. Consider moral terms. On an expressivist view, the function of having these terms in our vocabulary is not to track moral features of reality, but (roughly) to enable us to express and coordinate our plans or attitudes in specific ways that enable us to better live together. Suppose expressivists are on the right track about the function of moral terms—if so, then having moral terms in our ideology is very important, and it would be a mistake to jettison such terms or concepts on grounds of their failure to track joints of reality. And there are many similar views that ascribe some function other than joint-tracking to certain philosophically central terms: Consider Yablo’s view that mathematical terms serve the function of enabling us to express scientific laws more succinctly (2005), Ryle’s (1949) view that dispositional terms and mental predicates serve to license inferences, or Brandom’s view that modal talk serves to make explicit certain norms governing the use of our non-modal terms (*Between Saying and Doing*, Lecture 4). Even logical terms, as I have argued elsewhere (2015, Chapter 10, following McFarlane) arguably serve a very different function than tracking joints (logical or otherwise) in reality—namely, enabling us to make assertions with and reason using *other* terms.

In each case, a non joint-tracking view of the function of the relevant discourse is at least a live option—but if it turns out to be right, it would seem entirely misguided to rule out the relevant concepts on grounds of their failing to carve at joints or track features of the world they never even aimed to track. It would be equally misguided to replace them with some substitute concept that did track at some joints (indeed where there is such a huge difference in function, it is hard to see by what rights we could consider it a replacement for the original concept). In short, then, we can’t take the structure-first proposal seems as a *general* guide to conceptual choice without presupposing a kind of functional monism about our terms and concepts: that all do (or should) aim at tracking joints in reality or matching themselves to the kinds and categories of the world.¹³

In sum, even if one accepts that there is ‘structure’ or ‘joints of reality’, and is not concerned with the epistemology, it is not at all clear that *all* terms have a joint-tracking function—and so, it is not clear that the structure-first approach can provide a general guide in conceptual choice. To get a more global approach, we need to begin by assessing the function the term or concept is to serve. Now it is of course open to the friend of structure to propose a hybrid view: taking a structure-first approach to choosing some of our concepts, and a pragmatic approach to others. But, as I will now turn to argue, once we properly bring in the role of function, we have the route open for a comprehensive pragmatic approach that can avoid the

¹³ Incidentally, if one thinks of at least some aspects of a theory’s ideology (such as the mathematical or logical systems employed) as not having a joint-tracking function, that is further reason for thinking that, even if we take the total theory’s predictive success to be evidence that we have made good ideological choices, these need not be thought of as ‘good’ in the sense of tracking the structure of the world.

problems long thought to plague it, and do a respectable job even with the cases that motivate the structure-first view—perhaps rendering a hybrid approach and its associated difficulties unnecessary.

The Pragmatics-First Model

As I mentioned above, my particular interest in the idea that much of the central work in metaphysics can be understood as normative conceptual work is as a way of showing how a deflationary metametaphysical approach can nonetheless capture the difficulty and importance of much work in metaphysics. But this won't work if the only or best answer to the question: "how should we engage in this normative conceptual choice?" is that we should choose those concepts that best track certain metaphysical facts or features of reality. Those who take a pragmatic approach to conceptual choice instead, with Carnap, insist that we don't require an *ontological justification* for introducing a new form of speech.

Those with heavyweight metametaphysical inclinations are prone to see a threat here: if the metaphysical facts of the world cannot determine which of our concepts or terms we ought to use, are our choices then bound to be subjective, just a reflection of something 'about us', not worldly? One concern is that a pragmatic approach leaves us unable to account for Siderian intuitions that the world has structure, and that accordingly some conceptual choices (e.g. lithium rather than *lithium on earth*) are just objectively better than others (2011, 18-19). Sider puts the concern most strongly "If there is no sense in which the physical truths are objectively better than the scrambled truths, beyond the fact that they are propositions that we happened to have expressed, then the postmodernist forces of darkness have won" (2011, 65).

Moreover, if there are not metaphysical features to serve as standards for conceptual choice, one might worry that the pragmatic approach "...suggests that normative issues about what concepts we should use can be settled by voluntary choices that we ourselves make", leaving such choices subjective (Plunkett 2015, 860-61). And if they are then subjective, not constrained by metaphysical facts, one might worry that these conceptual choices are doomed to being arbitrary, or determined only by power relations or accidental facts of nature or history. If so, then the conceptual ethics view, taken in a deflationary direction, might threaten to make metaphysics inherently uninteresting, even suspect—with each practitioner simply attempting to wield his or her power in imposing her own perhaps idiosyncratic conceptual choices. And the deflationist might not be able to make sense of the plausible idea that there are *standards* governing conceptual choice, which enable us to justifiably engage in critique of some (but perhaps not other) conceptual schemes.

But these fears are misplaced. There are far better options available for the deflationist to take in sketching a methodology for normative conceptual work—options that can preserve the idea that conceptual choice is often, in an important sense, worldly, and that can respect our *ordinary* intuitions (though perhaps not all metaphysical intuitions) about structure. A pragmatist approach can also, I will argue, give constraints for conceptual choice that ensure that it is not merely arbitrary or subjective, and needn't be merely power-driven, leaving room for critique.

We can begin from that we found was needed above to get a plausible global approach to conceptual choice: the idea that our terms or concepts may serve many different functions. If we aim to engage in conceptual engineering, it is natural here, as with other engineering problems, to begin by determining the function that is to be served by the relevant term, concept, or conceptual scheme. This is cohesive with Haslanger's approach, as she introduces her

‘analytical’ approach by suggesting that we begin by considering “What is the point of having these concepts? What cognitive or practical task do they (or should they) enable us to accomplish? Are they effective tools to accomplish our (legitimate) purposes; if not, what concepts would serve these purposes better?” (2012, 224)

In asking these functional questions, we need not take a stand on whether the concepts and terms we are investigating are natural products of evolution, explicit artifactual creations, or something in between (as with the nests of birds or cities of ants). In any of these cases, we can profitably ask questions about the function of the heart, the Nimbus 3000, or the nest. We also needn’t take a stand on what view of function to adopt here, though one might appeal to either a Millikan style view of proper function, or a Cummins style view of system function in addressing questions about why it is useful to have the relevant term in our vocabulary (or concept in our repertoire): what role it plays (perhaps along with allied terms and concepts) in our overall conceptual system, what we would be missing if we lacked such a term or concept, or what having that concept to hand did for societies that enabled them to carry on and reproduce their conceptual system, including use of the concept at issue.¹⁴

The analyses of the functional pluralists fit quite naturally in these molds—consider, e.g. Stephen Yablo’s (2005) analysis of what we can do using mathematical discourse that we couldn’t otherwise (simplify our statements of laws), or Simon Blackburn’s analysis of the point of moral discourse in enabling us to coordinate our attitudes in certain sophisticated ways that we couldn’t do without having this normative vocabulary to hand, or Paul Horwich’s view of the role of a truth predicate serves as a device of generalization... In each case, these analyses purport to identify something that this range of vocabulary *does* or (better) enables us to do, that we couldn’t do (or couldn’t do as effectively or efficiently) without it—these are something like system-functional analyses, which of course can serve as clues to proper function analyses: to why it would have been useful to have terms like this, why they might have been preserved and perpetuated in our culture.

Supposing our concepts or terms do serve functions, a natural first place to turn to in evaluating them—in engaging in the normative work of determining what terms and concepts we *should* have, which we should retain, jettison, or revise (and how), is to begin by assessing what function these terms or concepts serve. It is no good engineering a boat, or deciding which boats should be kept, modified, or replaced without a clear idea of whether the boat is to function in providing a fast and nimble escape from police boats, in transporting masses of heavy cargo across the wide oceans, or in undertaking exploration in the icy arctic. Some features would require repair for any boat—failure to float for example. So similarly some features of a concept would require repair regardless of purpose—such as its propensity to ensnare us in contradictions. But no detailed evaluation can be made without an assessment of a function or functions to be served, or more broadly (as it tends to be expressed in engineering primers) the problem to be solved. Once a purpose (or multiple purposes) is/are identified, we can go on to use that in evaluating whether the term or concept in question should be retained, rejected, or revised, and what sort of rules or constraints would best (or better) enable it to fulfill its function(s).

Plausibly, some words do serve *something like* a joint-tracking function. The prime candidates for these are the terms Lewis and Armstrong began with in developing the joint-carving idea: predicates that aim to pick out ‘natural’ properties and relations. These are those

¹⁴ Thus, importantly, this is not a view of function on which a thing’s function is tied to its creator’s intentions about its function.

that will figure in laws (Armstrong) and in our natural-scientific theories. Since these terms have the function of serving in explanatory and predictive scientific theories, which in turn aim to predict and explain, there are worldly constraints on what concepts we ought to adopt.

Given the relevant function and the constraints that come with it, it is easy to see why we do better to have in our chemical theory the current chemical concept of ‘lithium’ rather than a concept that would apply to lithium on earth, but not to the same chemical kind if found on Mars (cf. Sider 2011, 7). Geographic constraints in themselves are not helpful to chemical explanations and predictions, the limited ‘lithium’ concept would not be as useful in a chemical theory. Similarly, it is easy to see, on these grounds, why the concepts of ‘fish’ and ‘mammal’ would be more useful than ‘sea creature’ and ‘land creature’ for explanations and predictions in biology: more predictions of behavior, internal construction, disease susceptibility, reproduction, and so on will be facilitated by use of the former concepts than the latter concepts.

I have respect for natural kind structure. I have a kid with a nut allergy. It is a matter of life and death (‘death in seven minutes’, her allergist tells us) whether something is biologically a tree nut or is called a ‘nut’. It is a matter of life and death because it enables us to *predict* whether ingesting something will cause a life-threatening allergic reaction. It is not just a subjective matter whether ‘tree nut’ is a better concept than one that includes all and only things *called* ‘nut’ (including hazlenuts, peanuts, coconuts and nutmeg, but not cashews, pistachios, and almonds). That one concept but not the other is usefully, and efficiently *predictive* in this way, which has life-or-death consequences (as well as consequences for other aspects of biological theory, farming, etc.), is all I need to be fully convinced that one set of concepts is objectively better—it does not require additional metaphysical vindication.

The pragmatic approach to conceptual choice can and does take into account all of these perfectly objective, worldly, empirically driven reasons for choosing one set of concepts over another, *where these concepts are designed to figure in our explanatory and predictive theories*. And the advantages are not just in the simplicity of stating our biological theories, but also in the living of our lives, communication with others, safety, formulation of laws of state (as well as expression of laws of nature), etc. We can take into account our ordinary intuitions that some concepts are ‘objectively better’ than others, and that conceptual choice in such cases must be world-responsive. For when concepts are (as natural kind concepts are) designed to be useful in our empirical explanatory and predictive theories, we are thereby committed to being deferential to the world—letting experimental evidence help determine which *do* best serve in our predictions and explanations. Those concepts that function well in these explanatory and predictive roles will tend to be those we think of as ‘carving at the joints’ in the ordinary sense of marking those similarities and differences that are most relevant to our overall body of predictions and explanations.

Sider insists that “Joint-carving thought does not have merely instrumental value” (2011, 61) but rather is a constitutive aim of inquiry. But the instrumental value of employing concepts in our scientific theories that are particularly useful in predicting and explaining does seem sufficient to account for the ordinary intuitions about structure used to motivate the theory: that we ought to employ concepts like the purely chemical concept of water, the biological concept of mammal, and even the color concept of blue (rather than bleen)—without positing extra metaphysical facts about ‘structure’, without invoking epistemic difficulties about how we could know such facts. Once we can account for the worldliness and objectivity of criteria for conceptual choice in these cases, it’s worth asking how very powerful are any remaining intuitions about ‘real structure’ that aren’t accommodated in this way, and whether we need any

further metaphysical vindication of our intuitions that some concepts are objectively better than others.

In any case, given functional pluralism, the criteria for evaluating, retaining, rejecting or rejigging extant concepts will vary according to the function such concepts serve. It won't *always* be an apt criticism of a concept to say that doesn't 'track the joints of reality' or serve in our best scientific theories, anymore than it is always an apt criticism of a boat to say that it couldn't carry more than 1000 tons of cargo. For social and institutional terms like 'married', 'citizen', 'person' or 'voluntary', or philosophically interesting and contested terms like terms for the moral, modal, or mathematical, the proper criteria for evaluation might not be whether the terms or concepts serve well in building explanatory and predictive theories, but whether they properly serve their own functions—say, endowing certain close human relationships with legal protections, enabling us to assign legal and moral rights and responsibilities, enabling us to coordinate our plans and attitudes, or enabling us to simplify our expressions of laws.

Even where the function of a concept is not predictive/explanatory, however, the pragmatic approach can nonetheless allow that way we choose or develop a concept to fulfill its function is not a merely arbitrary or subjective matter. For such conceptual choices are typically constrained by both worldly factors and by 'site constraints' involving the rest of our conceptual scheme and surrounding practices. Consider as an example the concept of death, as examined by Bernard Gert, Charles Culver, and Danner Clouser. They argue that the function of the concept of death is to enable us to determine when medical care should cease, funeral preparations should begin, survivors' benefits put into effect, and so on (and there is no precise joint in nature marked by it, but rather a continuum of changes). Choosing, precisifying or engineering such a concept must be responsive to worldly matters. Certain empirical discoveries might place new pressure on our old vague concept of death: first, the old way of treating cessation of spontaneous breathing and circulation as a criterion for death comes into question with the ready availability of artificial ventilation, and puts pressure on finding new ways of identifying criteria for death. At the same time, the use of new and increasingly expensive medical technologies, and the critical need for organ transplantation to be done quickly, puts new pressure on determining the time of death more precisely than heretofore, so that expensive treatments can be stopped and organs harvested with greater chance of success for the recipient. These are empirical factors that put pressure on the old concept, and give reasons for at least precisifying the concept and altering the criteria typically used in applying it, so that it may continue to serve its function. Where the function of a concept like 'death' involves, at least in part, enabling us to make decisions about when medical care should cease, and that medical care becomes increasingly costly or scarce, we may have a need to precisify the concept of death beyond the vaguer concept that served us well a hundred years ago (Gert et al 2006, 284). The pragmatic approach to normative conceptual work is certainly worldly in that, to do it well, one must be responsive to worldly constraints and new empirical situations. In conceptual engineering no less than civil engineering, the question of which design (of concept or bridge) will best fulfill the relevant function, given the requirements, does not leave room for a merely 'arbitrary' or power-driven answer.

Normative conceptual work, on this model, isn't just constrained by worldly facts, however, nor is it just function that plays—or rather, should play—a role in our decisions in conceptual engineering or in civil engineering. Instead, civil engineering projects must be resolved while taking into account constraints of the site: what the relevant land and geography are like, what the constraints are on surrounding extant structures and geographic features, etc.

Similarly, when we engage in conceptual engineering, we must be mindful of the multiple inferential connections our concepts bear to other concepts and practices.¹⁵ One can see this again in Gert et al (2006) on the concept of death. The term ‘death’ is, they emphasize, conceptually tied to a range of important concepts and practices: If someone is dead, it is appropriate to terminate all medical care, to initiate funeral proceedings, to read the will, to appoint a successor if she is president; other changes in legal status follow: insurance policies, social security benefits. There are also other, more internal conceptual connections between ‘dead’ and other terms: If someone is dead they cannot at the same time be alive; it marks the boundary between the process of dying and the process of disintegration (2006, 286). Gert et al appeal to these conceptual connections as a way of criticizing the conceptual revisionism of certain medical doctors, who aimed to (re-)define death in such a way that they would be permitted to harvest organs sooner, when they would have a greater chance of success with transplantation. Such physicians, he argues, make the mistake of noticing only the connections between ‘death’ and other *medical* terms and practices, not the wider system of concepts and social practices in which ‘death’ plays a central role.¹⁶ Thus in addition to the worldly constraints, these ‘site constraints’ provide additional constraints on conceptual development and revision, again reducing any sense that—without metaphysical guidance to appeal to—the pragmatist is left treating such choices as arbitrary.

¹⁵ I suspect that this is related to the point Eklund makes as he argues that one cannot ‘selectively’ engineer the quantifier (or, presumably, other concepts) (forthcoming, 19).

¹⁶ Gert, et. al use this example as part of a generalized argument for conservatism in conceptual change: while new circumstances (such as new medical technologies in keeping patients alive using artificial respiration, and in enabling organ transplantation) may require new precisifications of terms like ‘death’, they argue:

When a term plays an important part in social and legal practices, as ‘death’ does, then the greater the change in the meaning of the term, the greater the likelihood that there will be significant social and legal problems. (2006, 284)

“It is almost impossible to describe a situation in which it is appropriate to redefine a term with widespread ordinary use in order to change any particular medical (or even social or legal) practice, in which that term plays a significant role” (2006, 285)

I think this is an important and under-appreciated point: that conceptual engineering, no less than civil, does not take place in a vacuum, and that it is extremely important to note and be responsive to the inferential connections between the term in question (whether one we are considering revising or eliminating) and our other terms and broader practices. Nonetheless, while site constraints must be taken into account, I would not take this as grounds for an overriding principle of conservativeness in conceptual engineering. In civil engineering it may be a good—but defeasible—principle to interfere as little as possible in building your new bridge or building with the layout and design of surrounding roads and structures. But when problems get bad enough, or there are overriding social, medical or moral issues at stake, there are times for a more complete clearing. So similarly, though ‘marriage’ is connected to a wide range of social practices, changing the legal definition to not precisify but rather expand the applicability of the term to same-sex partnerships was the right thing to do. As was bulldozing the whole network of race concepts such as ‘octaroon’, ‘quadroon’, ‘mulatto’ that played an influential social and legal role in slaveholding societies. (As Burgess and Plunkett note, one question in conceptual ethics is “whether we ought to be using a given concept *at all*” (2013, 1095)).

This brings to light another, deeper side of normative conceptual work: conceptual engineering may assess what functions our terms *do* serve and how they *ought to be* structured in order to serve those functions well, but it also may go beyond that to assess what functions they *should* serve—an area where not just practical but also distinctly moral considerations come into play. Sally Haslanger’s work on race and gender concepts undertakes this work explicitly. While the old concepts may have served to give artificial pseudo-scientific legitimacy to discriminatory practices, Haslanger argues that concepts in this vicinity *should* be retained to serve a different function: serving as “effective tools in the fight against injustice” (2012, 226). Similarly, one might argue for rejecting pejorative terms from our vocabulary on grounds of their serving insidious functions of insult and exclusion.

We are now in a position to reply to the common criticism that the pragmatics-first view makes conceptual choices too arbitrary or left to be determined only by power relations or accidental facts of nature or history. One common fear about taking a pragmatic approach is that it leads to an arbitrariness whereby, if my practical purpose is to enrich my own pockets, I (perhaps I am on the committee for revising the DSM) should simply accept any bribe to do what I can to shift the meaning of, say, ‘disease’, in a way that will benefit the bribers (suppose they are representatives of major insurance companies who hope to pay out for treatments of fewer conditions). But the thesis here isn’t that in engaging in conceptual ethics should proceed by attending to the *goals* (whatever they might happen to be) of whoever might happen to be involved in the conceptual negotiation, but rather that the *function* of the terms in question should be taken into account in evaluating various proposed changes in (or preservation of) its use. Taking a bribe in such a case could be criticized on parallel grounds to a public official who takes a bribe (with the goal of lining her own pockets) to permit an engineering company to build a substandard and unsafe (but cheaper) bridge. So taking function into account in evaluating proposals in conceptual ethics does not entail a crass ‘pragmatism’ that encourages us to think that anything goes as long as it meets our own individual goals.

Conclusion

Where have we come? I have aimed to show that a pragmatic approach to normative conceptual work is a viable, and perhaps even preferable, alternative to metaphysics-first approaches. And I have aimed to lay out an initial blueprint for how normative conceptual work *should* be done, on a pragmatics-first model. No doubt it requires a great deal of expansion and revision. But already here we can note some important features of the proposal.

The first is that there is every prospect of giving an account of normative conceptual work that does not involve an appeal to *metaphysical facts* for guidance or as a standard against which success should be judged. This is not to say that I embrace a metaphysical view according to which *there are no such* metaphysical facts. As Sider has made clear, that would be a metaphysical view like any other. Instead, it is a model on which we have no *need for appeal to* such metaphysical facts to give standards for conceptual choice, no need to discover them to be doing conceptual ethics well, no need to think of us, in making our conceptual choices, as aiming to track such metaphysical facts. Such deeper metaphysical facts are here not ruled out or denied, but rendered *unnecessary*. If they are unnecessary to do the work of providing a good answer to how conceptual choice should go—in ways that match our intuitions and practices—and positing them is subject to ontological and epistemological costs, then we may be better off without them, and accepting a pragmatic model for conceptual choice.

And at any rate, I hope to have shown that such a pragmatic approach to conceptual choice is a viable alternative for metaontological deflationists who still hope to make some sense of the difficulty, depth, and value of work that has often gone under the heading ‘metaphysics’. On this model, all that is required is descriptive and normative conceptual work, alongside empirical work. To do normative conceptual work explicitly, we should figure out, empirically, what function the terms *do* serve, do conceptual analytic work in figuring out how they do work, how they are related to other concepts and practices, do normative conceptual work, combined with empirical work, in determining whether certain modifications or precisifications would better enable the term to fulfill the relevant functions, and perhaps ultimately reach back to deeper questions about what functions we want our concepts to serve, or what functions they ought to serve. By appealing to nothing more mysterious than straightforward empirical work

and conceptual work, this approach to conceptual choice remains cohesive with a deflationary metaontological approach, and available as a way to add strength and depth to the deflationary view.

Like civil engineering, conceptual engineering is not a matter for discovery but for invention. But also like engineering, that does not mean that the choices we make are arbitrary, unconstrained, merely subject to our will, or ‘subjective’. Which boat, or development of a concept, will work best *given our needs and situation* may often be an objective matter, once all constraints are in. Of course this is not to suggest that there will always be a *uniquely best* solution to a problem in civil or conceptual engineering.

The crucial point here is that, once we understand the approach better, we can easily see that the problems thought to plague the pragmatics-first approach are avoided. If we can properly develop and understand an approach to conceptual ethics in this regard, it does leave difficult, important, and principled work for us to do, when we do metaphysics.

Works Cited

- Baker, Lynne Rudder (2000). *Persons and Bodies: A Constitution View*. Cambridge: Cambridge University Press.
- Bennett, Karen (2009). ‘Composition, Colocation and Metaontology’, in David J. Chalmers, David Manley, and Ryan Wasserman, eds. *Metametaphysics*, Oxford: Oxford University Press.
- Bennett, Karen (2015). ‘There is no special problem with metaphysics’. *Philosophical Studies* DOI 10.1007/s11098-014-0439-0.
- Blackburn, Simon (1993). *Essays in Quasi-Realism*. New York: Oxford University Press.
- Brandom, Robert (2008). *Between Saying and Doing: Towards an Analytic Pragmatism*, Oxford: Oxford University Press.
- Bricker, Phillip (forthcoming). ‘Realism without Parochialism’. *Modal Matters: Essays in Metaphysics*. Oxford: Oxford University Press.
- Burgess, Alexis and David Plunkett (2013a). ‘Conceptual Ethics I’. *Philosophy Compass* 8 (12): 1091-1101.
- Burgess, Alexis and David Plunkett (2013b). ‘Conceptual Ethics II’. *Philosophy Compass* 8 (12): 1102-1110.
- Carnap, Rudolph. (1950/1956). “Empiricism, Semantics, and Ontology”. Reprinted in *Meaning and Necessity*, Second edition. Chicago: University of Chicago Press. (second edition 1956)
- Davies, David (2004). *Art as Performance*. Oxford: Blackwell.
- Eklund, Matti (forthcoming) ‘Intuitions, Conceptual Engineering, and Conceptual Fixed Points’, in Christopher Daly, ed. *Palgrave Handbook of Philosophical Methods*.
- Foucault, Michel (1965). *Madness and Civilization*. Translated by Richard Howard. New York: Vintage Books.
- Gert, Bernard, Charles M. Culver, and K. Danner Clouser, *Bioethics: A Systematic Approach*, second edition (OUP 2006)
- Haslanger, Sally (2012). *Resisting Reality*. Oxford: Oxford University Press.
- Hawley, Katherine (2006). “Science as a Guide to Metaphysics” in *Synthese* 149: 451-470.
- Horwich, Paul (1998). *Truth*. Second edition. Oxford: Oxford University Press.

- Huemer, Michael (2009). "When is Parsimony a Virtue?", *Philosophical Quarterly* 59 (235): 216-236.
- Jackson, Frank (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Kitcher, P. (1993). *The Advancement of Science*. Oxford: Oxford University Press.
- Kriegel, Uriah (2013). 'The Epistemological Challenge of Revisionary Metaphysics'. *Philosopher's Imprint* 13:12, pp. 1-30.
- Locke, John (1690). *An Essay Concerning Human Understanding*. (Reprinted by William Collins and Sons, 1964).
- Millikan, Ruth Garrett (1984). *Language, Thought and Other Biological Categories*. Cambridge, Massachusetts: MIT Press.
- Paul, L. A. (2012). "Metaphysics as Modeling: The Handmaiden's Tale" *Philosophical Studies* 160: 1-29.
- Plunkett, David and Tim Sundell (2013). 'Disagreement and the Semantics of Normative and Evaluative Terms'. *Philosopher's Imprint* 13 (23): 1-37.
- Plunkett, David. (2015). "Which Concepts Should We Use?: Metalinguistic Negotiations and The Methodology of Philosophy", *Inquiry*. Vol. 58, Nos. 7-8, 828-874. (2015)
- Price, Huw (2011). *Naturalism without Mirrors*. Oxford: Oxford University Press.
- Psillos, Stathis (1999). *Scientific Realism: How Science Tracks Truth*. London: Routledge.
- Ryle, Gilbert (1949). *The Concept of Mind*. Chicago: University of Chicago Press.
- Ryle, Gilbert (1970). "Autobiographical", in Oscar P. Wood and George Pitcher, eds. *Ryle*. New York: Doubleday and Co.
- Saatsi, Juha (2016). "Explanation and explanationism in science and metaphysics", in *Metaphysics and the Philosophy of Science: New Essays*. 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, Peter (1992). *Analysis and Metaphysics: An Introduction to Philosophy*. Oxford: Oxford University Press.
- Thomasson, Amie L. (2007a). *Ordinary Objects*. New York: Oxford University Press.
- Thomasson, Amie L. (2007b). 'Modal Normativism and the Methods of Metaphysics'. *Philosophical Topics* 35/1&2: 135-60.
- Thomasson, Amie (2012). 'Experimental Philosophy and the Methods of Ontology'. *The Monist* Vol. 95, no. 2: 175-200.
- Thomasson, Amie L. (2013). 'The Nancy D. Simco Lecture: Norms and Necessity'. *Southern Journal of Philosophy* 51, No. 2: 143-60.
- Thomasson, Amie L. (2014). "Quizzical Ontology and Easy Ontology", in *Journal of Philosophy* CXI, No. 9/10: 502-528.
- Thomasson, Amie L. (2015). *Ontology Made Easy*. New York: Oxford University Press.
- Thomasson, Amie L. (forthcoming). 'Metaphysical Disputes and Metalinguistic Negotiation', in *Analytic Philosophy*.
- Van Inwagen, Peter (2016). "The Neo-Carnapians", *Synthese*. DOI 10.1007/s11229-016-1110-4.

- Yablo, Stephen (2005). "The Myth of the Seven". In Mark Eli Kalderon, ed. *Fictionalism in Metaphysics*. Oxford: Oxford University Press.
- Yablo, Stephen (2009). "Must Existence-Questions have Answers?", in David Chalmers, David Manley and Ryan Wasserman, eds. *Metametaphysics*. Oxford: Oxford University Press.
- Yablo, Stephen (2014) "Carnap's Paradox and Easy Ontology", in *Journal of Philosophy* CXI, No. 9/10: 502-528.