

Eighteenth Excursus: Conceptual Analysis and Ordinary Language Philosophy

David J. Chalmers

This book has yielded a mixed verdict on the prospects for conceptual analysis. On one hand, the framework earlier in the book suggests that conceptual analysis is possible. On the other hand, the arguments regarding verbal disputes in chapter 9 suggest that its philosophical importance is limited.

We can start by stressing the positive. In the twentieth century, many argued that conceptual analysis is impossible: partly by arguing against the definitional model, and in partly arguing against an analytic/synthetic and apriori/aposteriori distinction. Here I have argued that a sort of conceptual analysis can survive these critiques. In particular, I have argued that a Carnapian intensional conception of conceptual analysis is unthreatened by Quine's critique of analyticity and by the failure of the definitional program.

On the intensional model of conceptual analysis, concepts are represented in effect by verdicts about cases. More specifically, they are represented by functions from scenarios to extensions. Nothing in the critique of definitions or Quine's critique compromises our ability to make armchair judgments about scenarios, at least given ideal reflection. With enough work and sufficient rational reflection, we can chart the intension for any concept that we possess and can map out its general structure.

All this may be useful for various purposes in the philosophy of language and mind. It allows us to theorize about the meaning of expressions and the contents of concepts in a general way, and it also allows us to investigate the meaning of many specific expressions and the content of many specific concepts.

Still, the project of conceptual analysis has often been associated with stronger claims than this. Conceptual analysis of specific expressions has sometimes be seen as the very heart of philosophy, and as the method that drives us to the most important philosophical conclusions. My own view is very different from this. I think that the conceptual analysis of specific expressions can

lead us to linguistic conclusions, and can thereby settle linguistic disputes and verbal disputes. But its power to yield substantive philosophical conclusions, those which settle substantive debates, is highly limited.

To see this, we can start by noting that much of traditional conceptual analysis centers on questions of the form “What is *X*?” or “What is it to be *X*?”. Much of contemporary naturalistic philosophy centers on these questions too. In “Verbal Disputes” I have argued for a deflationary view of the interest of these questions. Some component of these questions is inevitably verbal, and the non-verbal residue can be found without using ‘*X*’.

On the picture I favor, instead of asking “What is *X*”, one should focus on the roles one wants *X* to play, and see what can play that role. The roles in question here may in principle be properties of all sorts: so one focuses on the properties one wants *X* to have, and figures out what has those properties. But very frequently, they will be causal roles, normative roles, and especially explanatory roles.

For example, instead of asking “What is semantic content?” and expecting a determinate answer, one can instead focus on various explanatory roles one wants semantic content to play. One can then say, here are some interesting properties (of sentences or utterances): *S*₁ can play this role, *S*₂ can play these roles, *S*₃ can play these roles. Not much hangs on the residual verbal question of which is really semantic content.

Likewise, instead of asking “What is a belief? What is it to believe?” and expecting a determinate answer, one can instead focus on the various roles one wants belief to play. One can then say, here are some interesting states: *B*₁ can play these roles, *B*₂ can play these roles, *B*₃ can play these roles. Not much hangs on the residual verbal question of which is really belief.

The picture is also deflationary about the claims of some opponents of conceptual analysis. Such opponents often say: “I don’t care about the concept of *X*. I care about what *X* really is. Even though *X*₁ doesn’t mirror our concept of *X*, *X* really is *X*₁.” For example, when Ruth Millikan (1983, 73) gives her teleological theory of meaning, she says that she is not analyzing the concept of meaning, and that she is instead giving a theory about the natural nature of meaning. Likewise, when Hilary Kornblith (2002, 1-2) gives his naturalistic account of knowledge, he says that he is not analyzing the concept of knowledge, and is instead giving an account of knowledge itself.

I think these proposals about what *X* “really is” are often implausible, as the concept of *X* places constraints on what it picks out. But in any case, I think that these theorists’ point could be put much more plausibly and powerfully by saying: “I don’t care about what *X* is. I just care

about the associated explanatory role. And X_1 can play such-and-such crucial parts of that role.” On this way of putting things, it does not matter whether meaning really is teleological content (so the intuition that a swamp creature could speak meaningfully is irrelevant even if correct). What matters is whether teleological content can explain various phenomena that meaning was supposed to explain.

Of course roles themselves have to be expressed in language, and the danger of verbal disputes also arises in claims about roles. As always, one should always be prepared to give up any particular expression in stating these roles. Still, in practice, the expressions used in stating roles will often be relatively uncontested, where the expressions used in the original question were contested. So in practice, the move to roles can clarify matters greatly. In cases where potential verbal disputes also rise in stating the roles, one can repeat the process, in the hope that one will eventually find common terrain.

The model is not completely deflationary about conceptual analysis. On this model, the analysis of specific linguistic expressions and the associated concepts is relatively unimportant in understanding a first-order domain. But it is still interesting and important to analyze *conceptual spaces*. For example, it is especially interesting to determine which concepts are primitive, and to see how other concepts might be grounded in them. It is also interesting to investigate the space of concepts (and of the entities they pick out) that are relevant to a domain, determining which concepts can play which roles, what the relevant dimensions of variation are, and so on.¹

For example, the conceptual metaphysics that I discuss in the sixteenth excursus turns on the investigation of conceptual spaces rather than of specific concepts. Here the aim is to find out which concepts are fundamental—that is, serve as the primitives that ground the entire space of concepts—and thereby to investigate what is metaphysically fundamental. But nothing here turns on claims about which concepts are associated with which expressions. Likewise, the sort of conceivability arguments that I used in *The Conscious Mind* are sometimes described as conceptual analysis, but it is made clear that claims about the meaning of the word ‘consciousness’ are inessential to the main arguments there: what matters is that there is some property we have that

¹What about progress in conceptual analysis? One might take the discussion of ‘lie’ in “Verbal Disputes” to suggest that this progress is always broadly verbal, and therefore of interest largely if we are interested in metalinguistic and metaconceptual matters. I do not think this is quite right: progress in analyzing the concept of a limit was clearly substantive, for example. But I think this progress is substantive rather than verbal precisely because of the many roles that the concept of a limit plays for us, for example in characterizing continuous functions, derivatives, and so on. The nonmetalinguistic progress consists in determining just what sort of mathematical property can play these roles. Thanks to Kit Fine for discussion here.

is conceivably absent in a physical duplicate, not that that property is consciousness. Importantly, these claims in conceptual metaphysics are robust even when subject to the method of elimination, whereas many claims deriving from traditional conceptual analysis are not.

This view leads naturally to a sort of conceptual pluralism: there are multiple interesting concepts (corresponding to multiple interesting roles) in the vicinity of philosophical terms such as ‘semantic’, ‘justified’, ‘free’, and not much of substance depends on which concept goes with the term. The model also leads to a sort of pluralism about the properties that these concepts pick out. For example, it naturally leads to semantic pluralism: there are many interesting sorts of quasi-semantic properties of expressions, playing different roles. It leads to epistemic pluralism: there are many different epistemic relations, playing different roles. It leads to gene pluralism: there are many different things that deserve to be called “genes”, playing different roles. The same goes for confirmation pluralism, color pluralism, and so on.

I am inclined to think that pluralism should be the default view for most philosophical expressions, at least once we set aside expressions for bedrock concepts. Typically there will be no single privileged role associated with such an expression, and different roles will be played best by different properties. The same holds if we move to sets of roles: different speakers will associate the term with different sets of roles, and different sets of roles will be played best by different properties. It may be that in some cases, a term is uniformly associated with a single role (or set of roles) in our community, or that one role is much more important than all the others, or that one property plays every one of the relevant roles better than any other property. But I suspect that such cases are relatively rare. Setting such cases aside, we should expect there to be many interesting concepts and properties in the vicinity of a given expression.

The picture that emerges is somewhat deflationary about the role of ordinary language philosophy in its many manifestations: 1950s Oxford philosophy, Canberra-plan analyses of folk concepts, contextualism and related theses in epistemology, contemporary linguistics-based philosophy, and some parts of experimental philosophy. It is not entirely deflationary about these endeavors: one can certainly move from linguistic premises to substantive conclusions, if one is careful enough about bridging premises and arguments. Still, if one is not careful, one will end up making points that reflect the vicissitudes of one’s language rather than deeper philosophical truths.

For a more fine-grained analysis of the role of language, it is useful to focus on a specific sort of argument from ordinary language in philosophy. Perhaps the paradigmatic sort of argument here is a *disquotational argument*. Such an argument takes the following form:

Semantic premise: ‘X’ means Y.

Disquotational premise: ‘X’ means X

Conclusion: X is Y.

Here, the semantic premise is established through analysis of language and its usage, while the disquotational premise is a trivial premise about meaning, reference, or truth-conditions. The conclusion, however, appears to be a substantive first-order conclusion.²

For example, something like the following argument is given by Stanley and Williamson (19xx) and Stanley (2011):

(1) ‘S knows how to ϕ ’ is true iff there is a way w such that S knows that w is a way of ϕ ing. (Semantic premise)

(2) ‘S knows how to ϕ ’ is true iff S knows how to ϕ . (Disquotational premise)

——

(3) S knows how to ϕ iff there is a way w such that S knows that w is a way of ϕ ing.

Here, the first premise is offered as following from a standard linguistic analysis of “knowing-wh” constructions along the standard claim that “knowing how” is an instance of such a construction. The second premise is just an instance of Tarski’s T-schema. The conclusion, however, appears to be a substantive philosophical claim. It entails the “intellectualist” position that knowing how is a species of knowing that, as opposed to the “anti-intellectualist” position that knowing how differs from knowing that.

There is much to say about arguments of this sort. In effect, these arguments enable us to move from a premise about language to a conclusion about the world, providing a powerful response to the common charge that philosophy is about the world, and not just about language. Still, the arguments can be resisted in various ways. I will briefly discuss two ways of resisting that are

²A closely related sort of argument adds a third *truth premise* as follows: (1) ‘X’ means Y; (2) ‘X’ means X; (3) X; therefore (4) Y. For example: (1) ‘John believes something’ means ‘Exists p: John believes p’ (2) ‘John believes something’ is true iff John believes something; (3) John believes something; therefore (4) Exists p: John believes p. In this way one might use a linguistic argument to establish the existence of propositions. Related arguments with a truth premise have been used to argue for the existence of events (Davidson), possible worlds (Lewis), and the truth of the A-theory of time (Ludlow). Everything I say in what follows about disquotational arguments applies with appropriate changes to these arguments.

closely tied to semantics, before discussing a deeper worry about the upshot of these arguments tied to issues in “Verbal Disputes”.³

Most obviously, the semantic premise can be resisted by offering an alternative semantic analysis. Where semantic premises deriving from linguistics are concerned, it is especially relevant to observe that linguists’ semantic claims are often subject to fewer constraints than philosophers’ semantic claims. For example, linguists’ analyses often aim to capture inferences that are held to be valid by most or all competent users, whereas a philosopher might instead aim to capture inferences that are in fact valid. Likewise, a linguist might freely invoke abstract objects in their analyses, while some ontologically scrupulous philosophers might not. A philosopher might also give a role to empirical facts about science and naturalness in giving their analyses, while a linguist might not.

The linguists’ project (which we might call *semantics narrowly construed*) has the advantage that it allows the data of semantics to be relatively theory-neutral. But the cost is that it is far from obvious that this project yields genuine worldly truth-conditions. This point is recognized by many linguists. It has often been noted that on a natural model-theoretic account of the semantics of mass terms, derived from respecting speakers’ inferences, ‘water’ has distributive reference, which has the consequence that in all models, every part of water is water. Nevertheless, science suggests that some parts of water (hydrogen atoms, for example) are not water. Some conclude at this point that these model-theoretic analyses reveal more about speakers’ conceptions than about worldly truth-conditions.⁴

By contrast, the philosophers’ project (*semantics broadly construed*) might yield worldly truth-conditions. But here the data will not be theory-neutral: it may be that for a broad semantic analysis of ‘X’, one first needs a thorough philosophical or scientific analysis of X. To know what ‘water’ refers to, one has to engage in the science of water; to know what ‘knowledge’ refers to, one has to engage in the philosophy of knowledge. But to the extent that a semantic analysis (of ‘knowledge’, say) depends on a philosophical analysis, we will not be able to use it to support that philosophical analysis. Although a disquotational argument may be sound in such a case, one’s warrant for the semantic premise will rest on one’s warrant for the conclusion, so that one

³The next four paragraphs, on the semantic premise and the disquotational premise, are not directly tied to the issues about verbal disputes here and can be skipped, but I think the issues here are interesting in their own right.

⁴See for example Bunt 1979 and Lonning 1987, and Koslicki 1999 for criticism. Bach (1973) makes the case for “natural language metaphysics” as an area distinct from worldly metaphysics. Thanks to Zoltan Szabo and Brendan Balcerak Jackson for discussion of these matters.

cannot use the former warrant to generate warrant for the latter. More generally, an opponent of the first-order conclusion above, who holds that there are good reasons to hold that X is not Y , may hold that these are likewise good reasons to hold that ' X ' does not mean Y , at least when meaning invokes worldly truth-conditions. This is not to deny that claims deriving from semantic analysis may have some independent support and therefore some force, but it suggests that establishing their force is less straightforward than it might seem.⁵

One can also deny the disquotational premise by finding context-dependence in a key term, or by finding ambiguity or idiolectal variation. If different people use a key term in different ways, as arguably often happens in philosophy, a semantic analysis of one person's use may not apply to another's. Or following the worry above, one might hold that a semantic analysis yields something more like internal truth-conditions or model-theoretic truth-conditions than worldly truth-conditions, so that the disquotational premise does not apply.

The discussion in "Verbal Disputes" suggests another reaction to these arguments, though. An opponent might simply accept the conclusion, while denying that it affects their substantive position. For example, suppose a compatibilist makes a successful argument that 'free will' refers to the ability to do what one wants, and infers that free will is the ability to do what one wants. Rather than rejecting the conclusion, an incompatibilist might respond by accepting it, and simply denying that free will (so understood) is tied to moral responsibility. If so, their substantive views about the connection between determinism and moral responsibility will be unchanged, and their position will differ only verbally from their previous one.

Likewise, suppose an intellectualist makes a successful argument that 'S knows how to ϕ ' invokes a sort of propositional knowledge and infers that knowing how is a sort of propositional knowledge. Rather than rejecting the conclusion, an anti-intellectualist might react by accepting it,

⁵On the two-dimensional picture deriving from chapter 5 and the eleventh excursus, one may at least be able to establish claims about primary intensions (of 'water', say) independent of empirical scientific analysis (of water, say). But the idealization of primary intensions suggests that a priori philosophical analysis will still play a crucial role here. For example, the same sort of a priori philosophical analysis that goes into answering the question 'What is knowledge?' will also be relevant to the question 'What is the primary intension of "knowledge"'. Furthermore, while primary intensions are a sort of worldly truth-condition, for many or most expressions (in particular, for epistemically nonrigid expressions), they must be combined with empirical truths about the world to yield truths about reference. If one appeals to fine-grained "analytic" meanings to avoid idealization, the latter problem becomes worse: fine-grained meanings may need to be combined with substantive a priori truths about the world to yield substantive conclusions about reference. Perhaps for certain strong sorts of acquaintance concepts, reference will be determined in an especially direct way that does not require a priori reflection; but for these concepts, there may not be much role for semantic analysis to play in revealing reference.

and then saying that knowing how (so understood) is a relatively marginal phenomenon that does not play a central role in intelligent action, and that is grounded in abilities, which play the central role. If so, their substantive views about the character of intelligent action will be unchanged, and their position will differ only verbally from their previous one.

Generalizing the point: the mere fact that existing words like ‘know’ or ‘intentional’ or ‘see’ behave in a certain way does not suffice to settle substantive disputes about epistemology, action, or perception. After all, views based on these data may differ only verbally from views on which ‘know’ or ‘intentional’ or ‘see’ pick out something else, and which endorse apparently different first-order claims that use these words. For example, consider a contextualist view that says ‘know’ has two distinct referents in different contexts: $\text{knowledge}_{\text{low}}$, which we often have, and $\text{knowledge}_{\text{high}}$, which we never have. A proponent of this view may differ only verbally from a proponent of an apparently skeptical non-contextualist view on which ‘know’ always means $\text{know}_{\text{high}}$, or an apparently nonskeptical view on which ‘know’ always means know_{low} . Any substantive differences between the proponents will turn on what they say about further properties of the referents: for example, about the epistemic value and normative role of $\text{knowledge}_{\text{low}}$ and $\text{knowledge}_{\text{high}}$. Assessing these further properties requires going beyond the linguistic data.⁶

This brings out a crucial point about disquotational arguments. For their conclusions to be dialectically effective, there must be a tacit *role premise* about the role of X that is accepted by both parties to a dispute.⁷ In the case of free will, the role premise may be the claim that free will has a certain tie to moral responsibility. In the case of knowing how, the role premise may be the claim that knowing how plays a certain ubiquitous role as the grounds of intelligent action. In the case of knowledge, the role premise may be the claim that knowledge (or the variety of knowledge ascribed in ordinary contexts) has high epistemic value. If that role premise is fixed in place, then agreement on the conclusion of the disquotational argument (that X is Y) will yield agreement that Y (knowing propositions, be able to do what one wants, $\text{knowledge}_{\text{low}}$ plays the relevant role. But

⁶For more on applying the current methodology to contextualism, see the discussion linked at <http://consc.net/contextualism.html>. See also Sosa (2004), who uses considerations about verbal disputes to draw deflationary conclusions regarding certain aspects of contextualism. For related deflationary approaches to experimental philosophy, see Sosa (2007) and Scanlon (2010). In their “experimental philosophy manifesto”, Knobe and Nichols (2009) suggest that experimental philosophy aims not so much to discover the truth about intentional action, causation, and so on, but rather to discover how people think about these topics. I think that it is not out of the question for experimental philosophy to aim higher: perhaps it might help to establish truths about (what we call) intentional action, causation, and so on. But then the philosophical importance of the truths so established will be subject to the sort of critique given here.

if the role premise is not fixed in place, an opponent can hold on to their previous position simply by denying it. Knowing how or free will or knowledge (ascribed in ordinary contexts) turn out not to play the role in question, but everything else is unchanged.

There may be some cost to denying the role premise: after all, the claim that free will is a condition on moral responsibility or that knowing how plays the relevant role may seem like common sense, and denying it may involve ascribing error to many ordinary speakers. Still, commonsense is not sacred in these domains. And in any case, some element of common sense may be preserved by noting that while these claims are not literally true, there are nearby claims that are true: the incompatibilist may hold that a certain incompatibilist property (which we commonly take to be free will) is a condition on responsibility, and the anti-intellectualist may hold that certain abilities (which we commonly take to be knowing how) play the relevant role in action.

In such a case, the key substantive issue is whether Y (knowing propositions, being able to do what one wants) plays role R (a role in action or in moral responsibility), or whether some other property Y' (abilities, an incompatibilist property) plays that role. In there are strong first-order reasons to hold that Y' plays role R, it is implausible that the disquotational argument above can defeat these reasons: instead, if one accepts the conclusion that X is Y, the first-order reasons will then give one reason to deny that X plays R. For example, if semantics suggests that commonsense temporal claims (A happened before B) are committed to simultaneity but physics tells us that nothing plays the role of simultaneity, it is clear that the evidence from physics wins out: one will have to either deny the semantic analysis or deny the commonsense temporal claims. The same goes where strong philosophical arguments play the role of physics here. At best, the linguistic evidence deriving from the disquotational argument plus the role premise gives one some weak evidence that X plays R, evidence that may play a role when when other first-order evidence is weak.

This leaves some roles for ordinary language philosophy. First, following on the analysis above linguistic claims may play a weak evidential role with respect to nonlinguistic claims. The conclusion of a disquotational argument (say, that knowledge is knowledge₁) may not yet substantively differ from that of someone else who says 'Knowledge is knowledge₂'. But if we combine it

⁷The discussion in "Verbal Disputes" suggests that in the case of a bedrock expression (such as 'conscious', or 'right', perhaps) there need be no such role premise. But in these cases, it is not clear that semantic analyses do much to help give us a substantive grasp on reference. Furthermore, faced with a putative semantic analysis of such an expression, it will be open to an opponent to deny that the analysis is capturing the use of the expression to express the bedrock concept.

with a role premise saying that knowledge is valuable, it follows that knowledge₁ is valuable. Still, philosophical counterarguments, such as arguments that knowledge₁ is not valuable, may give one reason to question either the semantic premise or the role premise, as we have seen. Linguistic claims are just one sort of evidence among many here, and one always needs to be explicit about just how the evidence relations work.

Second, there is considerable intrinsic interest in understanding the meanings of the words we actually use in philosophy and in understanding the concepts and the patterns of thought that we actually deploy. The language and psychology of philosophy are important topics in their own right. But there is much more to philosophy than the language and psychology of philosophy, and great care is required in moving from the latter to the former. For example, if it turns out that freedom₁ and not freedom₂ is picked out by our ordinary concept of freedom, so that compatibilism is true of the ordinary concept, then that gives some extra interest to freedom₁ and perhaps is evidence that this is the sort of freedom that we really value. But to answer the hardest questions about freedom, one still needs to determine the nonlinguistic properties of freedom₁: is it truly valuable, what is its normative role, and so on. The fact that freedom₁ is picked out by our ordinary word or concept does not settle these questions.

Third, the practice of ordinary language philosophy can lead us to interesting concepts and important distinctions. This is especially likely if the doctrine of “the genius of our tongue” is correct: our language is a tool sharpened by a history of use, so that one can expect that if a concept or a distinction is important, our language may well have captured it already. Still, this point gives ordinary language more of a role in the context of discovery than the context of justification. After using ordinary language to find concepts and distinctions, one still needs to justify substantive claims involving those concepts and distinctions.

Fourth, there are important normative questions about what expressions *ought* to mean. These questions comprise what Peirce called “the ethics of terminology”. Ideal agents might be unaffected by which terms are used for which concepts, but for nonideal agents such as ourselves, the accepted meaning for a key term will make a difference to which concepts are highlighted, which questions can easily be raised, and which associations and inferences are naturally made. Following the “ameliorative” project of Haslanger (2005), one might argue that expressions such as ‘gender’ and ‘race’ play a certain practical role for us, and that role is played better by some conceptions than others, so ‘race’ and ‘gender’ ought to have certain meanings. The manifestly verbal dispute among astronomers about whether Pluto is a planet is best understood as a debate in the ethics of terminology: given the scientific and cultural roles that ‘planet’ plays, should ‘planet’

be used to include Pluto or exclude it? In philosophy, ‘meaning’ functions as something of an honorific (it attracts people to its study), so if one thinks that meaning₁ is more important than meaning₂, one might hold that ‘meaning’ ought to be used for meaning₁. Likewise, one might argue that ‘free will’ plays a useful practical role in judging the moral status of actions, or a pernicious role in supporting retributive punishment, so that we should use it for freedom₁ (so that we can attribute free will) or freedom₂ (so that we can deny it) respectively. The answers to these normative questions may depend on our purposes and values, but the questions themselves are certainly substantive.

So ordinary language philosophy is not unimportant. But one needs to be very clear about its use. If ordinary language philosophy is practiced in the material mode, as it often is, it is easy to move too quickly from linguistic data to substantive philosophical claims in a way that masks potential verbal disputes. Things work best when ordinary language philosophy is practiced in the formal mode, making claims about ‘freedom’ and not about freedom, for example, and being clear when these claims are descriptive or normative claims about actual usage. If one wants to draw non-linguistic conclusions while avoiding verbal disputes, one needs to be explicit about the bridge (a disquotational premise and a role premise, for example), and ideally one should be prepared to cast the conclusion without using the key expression. Doing things this way minimizes potential verbal disputes and maximizes clarity.