Metaphilosophical Reflections on the Idea of Metaphysics[1]

Robert Brandom

1

One constant in Wittgenstein's thought, early and late, is his denial not only of metaphysical naturalism, but of methodologically monistic scientism, a broadly epistemological view that is often taken to be a consequence of such a metaphysical view. This is the claim, roughly, that scientific knowledge is *the* form of knowledge, and scientific understanding is the only kind of understanding that deserves the name.[2] "Philosophy is not one of the natural sciences," he says in the *Tractatus*[3], and this view seems to be part of what lies behind the theoretical quietism of the later work. In fact, I think Wittgenstein thinks that if systematic philosophical theorizing *were* possible, it would mean that philosophy *is* an empirical science. Since it is not, philosophers must eschew theorizing, restricting themselves instead to light, local descriptions of discursive practices, where such descriptions might provide helpful reminders in freeing ourselves from the sort of misunderstandings and puzzlements that arise precisely from the theories implicit in inherited pictures of what is going on when we think and talk. Whether or not Wittgenstein himself reasoned in this way, I take it that it is common for his admirers to see him as presenting us with a forced choice: either embrace scientism about philosophy of the methodologically monistic sort—that is, take philosophy to be an empirical, scientific discipline—or give up the idea of systematic philosophical theorizing once and for all.

I think this is a false choice. Rejecting scientism of the methodological monist sort does not entail giving up the possibility of systematic philosophical theorizing about discursive practice. One of the most powerful methodological features of the natural sciences is the postulation of unobservable theoretical entities, and their deployment in constructions aimed at explaining what is observable. Theoretical entities are those about which we can make only theoretical, and not observational claims. Theoretical claims are ones that we can only become entitled to as the conclusions of *inferences* from other claims, not non-inferentially, as the results of exercising reliable dispositions to respond differentially to environing states of affairs by making observation reports of them. A generalization of this method would have the role played by observational vocabulary played by any antecedently available vocabulary, whether observational or not. So for instance one might postulate *meanings* to explain proprieties of *use*, where the latter are expressed in a non-semantic vocabulary, whether or not our access to claims about correct usage are made observationally or themselves inferentially.[4] The claim that theorizing of this sort could be legitimate in philosophy does *not* commit one to the claim that this method is the *only* legitimate method of acquiring philosophical understanding—which is what methodologically monistic scientism claims. The generalized method of postulation and construction might be *one* form of philosophical understanding among others. I want to claim that what is objectionable about the methodologically monistic form of scientism is its exclusivity. Rejecting that at least leaves open the question of whether, and which, features of natural scientific investigation, explanation, knowledge, and understanding ought also to be counted among those useful and appropriate in philosophy. After all, *description* is also a central and essential

[2] It is important to distinguish this very strong claim about understanding from the weaker (though already substantially committive) *scientia mensura* of Sellars: "In the dimension of describing and explaining the world, science is the measure of all things, of what is, that it is, and of what is not, that it is not." §42 of *Empiricism and the Philosophy of Mind* edited, with a Study Guide by Robert Brandom [Harvard University Press, 1997]. The latter claim is compatible with claiming also, as Sellars does, that distinctive forms of understanding are involved in the use of vocabulary that is *not* principally in the business of describing and explaining: for instance, normative vocabulary (and therefore, according to Sellars, also semantic and intentional vocabularies).
[3] Proposition 4.111.
[4] The analogous postulation of intentional states to explain behavior Sellars calls "philosophical behaviorism," by contrast to the "logical behaviorism" that is committed to *defining* the states in terms of behavior. In the case of meaning and use, the corresponding non-theoretical move is a semantic instrumentalism that insists, as Dummett used to do, that every aspect of meaning be *manifestable* in use.

element of scientific methodology, and even the most rigorous versions of Wittgensteinean quietism allow philosophers to describe features of our linguistic practice.

2

One reason to answer that question in the negative is provided by a positive view that is something like a converse of methodologically monistic scientism, and which is also often associated with the later Wittgenstein (though by no means exclusively with him). That is that the subject-matter itself settles that discursive practices—the capacity to use or deploy vocabularies—requires a distinctive kind of understanding. The reason we should not be methodological monists is that understanding talking and thinking, concept use, vocabularies, natural language utterances and texts, is a distinctive sort of achievement. This kind of understanding, what we might call "*hermeneutic* understanding," is not expressible in explicit rules, formalizable in regimented technical or artificial languages. The mathematized mature natural sciences have had great success in achieving what we might call "*algebraic* understanding" of great swathes of the inanimate natural world. (Whether the animate biological world, including sentient-but-not-sapient creatures and their activities, itself already calls for further special sorts of understanding remains a lively and controverted question.[5]) But when the topic is *culture* rather than *nature*, another sort of approach is called for. Here the paradigm of understanding is that exhibited by competent native speakers of natural languages when confronted by everyday utterances expressed in familiar vocabulary. This sort of practical grasp of meanings (the medium of the cultural) is not in the most fundamental cases a matter of explicit theorizing at all. And it is not a matter of mapping or translating the utterance into some other vocabulary (perhaps with the use of auxiliary logical vocabulary) either. (In the sense that matters for this point, the language of my thought is just my language: a language I speak.) More sophisticated forms of hermeneutic understanding, of the sort exercised by the literary critic, jurisprudential interpreters, and reader of philosophical texts, are possible, but they are both rooted in the basic one and do not come closer to having the structure of algebraic understanding.

A pragmatist line of thought common common to the Dewey of *Experience and Nature* and *Art and Experience*, the Heidegger of *Being and Time*, and the Wittgenstein of the *Philosophical Investigations* is that there is such a thing as hermeneutic understanding in this sense, it is a genuine and distinctive kind of understanding, and it is the most basic kind of understanding, in the sense that all other sorts of understanding are parasitic on it and develop out of it. It is the primordial sort of practical discursive know-how: the capacity to engage in an autonomous discursive practice. In particular, they are concerned to insist that the sort of algebraic understanding characteristic of mature mathematized sciences—the sort for which analytic philosophers long—is pragmatically dependent on everyday hermeneutic understanding, which accordingly cannot be replaced by or reduced to the more technical kind.

I accept all of these pragmatist claims about the distinctiveness and basicness of ordinary hermeneutic understanding of discursive performances and their products. But this pragmatist line of thought does not entail that many aspects of discursive practice might not *also* be susceptible to understanding of the sort I have called "algebraic." And where it is possible, broadly algebraic understanding has distinctive virtues, which adherents of the project of analysis are right to esteem and treasure. The slogan of the analytic project is "Faith, hope, and clarity—and the greatest of these is clarity." The clarity in question is specifically *conceptual* clarity. It would seem to have two dimensions: definiteness and perspicuity. From a pragmatic point of view, the significance of a speech act is definite insofar as its normative significance is settled. From the point of view of semantic inferentialism, this means that concepts are definite insofar as their circumstances and consequences of application are settled: when one is committed and entitled to apply them, and what such application commits and entitles one to. Perspicuity is then epistemic or psychological ease of access to those circumstances and consequences of application. On this line, thinking clearly is both formulating one's claims (claimables) so as to fix what one would be committing oneself to by endorsing them and what would entitle one to do so, and being aware of those definite consequences and circumstances of application of the concepts that articulate the contents of the concepts one is applying. Writing clearly is choosing one's words so as both to determine the inferential boundaries (or, one could equally well say, truth conditions) of one's claims and to convey them to the reader.

---

[5]   See the discussion in Michael Thompson's astonishing, original, pathbreaking book *Life and Action* [Harvard University Press, 2008].

What I'm calling the "algebraic" form of understanding achieves clarity along both the dimensions of definiteness and perspicuity by *constructing* the conceptual contents expressed by a target vocabulary It does that by exhibiting them as complexes formed as the products of applying explicit algorithms to the conceptual contents expressed by a base vocabulary (treated for this purpose and relative to this construction, as simple). What corresponds in this semantic-analytic project to the postulation of unobservables in empirical scientific theorizing is the employment in the algorithmic construction also of some further auxiliary vocabulary, whose use is not governed by antecedent norms but is determined instead by stipulated inferential connections to both base and target vocabulary.[6] This algorithmic-constructional method (building complex things by applying well-defined operations to simpler things) is a very good, perhaps superlative, way of securing clarity of understanding. I have elsewhere called it for this reason the "gold standard" of understanding generally—by which I mean that *when and to the extent it is available*, it is the very best sort of understanding to have. For it takes the issue of what one means (what one is committed to by a claim, what is incompatible with it, what would count as evidence for or against it, and so on) out of the hands, out from under the authority, of the one making the claims. It establishes a fact of the matter about the inferential relations that articulate the contents of the concepts expressed by the target vocabulary that swings free of the beliefs and preferences of the concept-user: what she would *like* to be committed to or treat as evidence for those claims. If a dispute arises, those who are assessing the claim in question can say, with Leibniz, "Let us calculate." This aspiration to develop "…a general method in which all truths of reason would be reduced to a kind of calculation…and errors—except those of fact—would be mere mistakes in calculation…"[7] (I take it that "truths of reason" here stands in for inferential relations that articulate the contents of the concepts involved) is one of the reasons Leibniz was a hero for Russell in the latter's attempt to develop a notion of philosophical analysis. This sort of clarity of understanding is a pearl without price—all the more to be prized where the target vocabulary it concerns is weightier and more difficult, as is the case with many of those either used or addressed by philosophers. This sort of clarity facilitates communication—scientism in the sociological sense.

Appreciating this cardinal virtue of the algebraic form of understanding does not require taking issue with the pragmatist point that it is in principle parasitic on and intelligible in principle only against the background of a more basic sort of practical discursive understanding that does not at all have this explicit theoretical form. It is useless—for instance, in settling disputes about what someone is committed to by a claim couched in the target vocabulary being (re)constructed—unless there is a shared base vocabulary about whose proper use all parties can agree in their practice. We are not in a position to calculate unless we can all practically go on in the same way in counting and adding—as Wittgenstein is at pains to remind us in many different ways and many different contexts. And the same is true of algebraically computing the inferential roles or truth conditions of complex expressions from those of simpler ones. Algorithmic elaboration is a way of *leveraging* practical agreement in the use of one vocabulary into practical agreement in the use of another. It is true that what plays the role of a base vocabulary for one such constructive enterprise may be the target vocabulary whose proper use is algorithmically reconstructed by another. But the point Wittgenstein was after here is that it cannot be algorithmic elaboration all the way down. At some point each such chain must be anchored in practical agreement about what it is and is not correct to do with a vocabulary that is *not* settled by being algorithmically handed off to some prior one. And that is to say that we should *not* make the jump from the legitimate *local* aspiration to be able to settle some semantic-inferential disputes in the "Calculemus" way to Leibniz's dream of a *global lingua characteristic*, *all* of whose concepts are governed by a *calculus ratiocinator* that is in this sense *universal*.

Acknowledging the value of the unique clarity afforded by algebraic understanding accordingly does not entail commitment to this sort of understanding being available in every case, even in principle. It does not oblige one to embrace the shaky method of the drunk who looks for his keys under the streetlamp, not

---

[6]   Thought of in this framework, in the case of empirical scientific theorizing, the base claims and the target claims are formulated in the same antecedent vocabulary—which may be, and must include, observational claims in the strict sense of those elicited by the exercise of reliable noninferential differential responsive dispositions, but which also include statements couched in the vocabulary (including theoretical, that is, only inferentially applicable vocabulary) of other scientific disciplines, for instance, those that address the workings of measuring instruments and the ranges of counterfactual robustness of various collateral premises.
[7]   In his 1685 ms. "Machina arithmetica in qua non additio tantum et subtractio sed et multiplicatio nullo, divisio vero paene nullo animi labore peragantur."

because they are likely to be there, but just because the light is better there. We should admit that sometimes algebraic understanding is not available—indeed, that *every* context in which it *is* available contains an appeal to a base vocabulary whose use is *not* held in place algebraically, but depends on *another* sort of practical mastery and understanding. Algebraic understanding can be no more legitimate than the hermeneutic understanding on which it depends and which it leverages, amplifies, and concentrates. It follows that philosophy cannot be *identified with* analysis, thought of as comprising the tasks of understanding algebraic understanding and applying it in semantics. Even under the broad heading of trying to understand discursive practice, there is a more basic sort of hermeneutic understanding, both whose implicit, practical, everyday species and whose explicit, theoretical, sophisticated species must both be studied and exercised by philosophers. Thinking through the presuppositions of its project shows that analytic philosophy can aspire at most to being one species of the genus.[8] (In the third, methodological, chapter of *Tales of the Mighty Dead*, and again in "Hermeneutic Practice and Theories of Meaning"[i] I try to say something specific and systematic about how the different aspects of discourse addressed by these two sorts of understanding and their associated disciplines complement one another.)

3

There is, then, a lot more to be understood about discursiveness than can be understood algebraically. This is obviously true *de facto*, and I have just rehearsed an argument that it is true also *de jure*. But can we know *in advance* that the algebraic sort of understanding is not available at all for some subject matters? Might it not be the case that the very nature of discursive practice makes it unsuitable for this sort of account? Perhaps algebraic understanding must inevitably "murder to dissect," the very method it employs making it impossible for it ever to grasp the essence of the phenomenon it addresses.

It is at any rate important to keep in mind that the claim that there are some vocabularies, some discursive practices-or-abilities, that are by their very nature not amenable to analytic algebraic reconstruction does not *follow* just from the observation made above (in denying methodologically monistic scientism) that every analysis or algebraic reconstruction of a target vocabulary must make use of and so depend on the prior semantic determinateness and understanding of what is expressed by some base vocabulary. That is, it does not follow that there is some order of, as it were *natural* basicness among vocabularies, which must have unexplained unexplainers (base vocabularies that do not admit of analytic algebraic reconstruction in terms of others) as its most basic elements. It might well be that although each analytic-algebraic account of the use of any vocabulary must appeal to some base vocabulary whose use is not explicated in that account, every vocabulary that plays that role of base vocabulary in some analyses plays the role of target vocabulary in some other successful analysis. A claim of the form $\forall x \exists y [Rxy]$ does not entail one of the form $\exists y \forall x [Rxy]$. (It is true that the world has a population problem because during every minute there is a woman somewhere in the world having a baby. But it is not a productive way to address the problem to look for the woman who is having all those babies and make her stop doing what she is doing.) The sense in which algebraic understanding rests *de jure* on hermeneutic understanding may be merely of the local, $\forall x \exists y$ sort, not the global $\exists y \forall x$.[ii]

The algebraic form of understanding requires distinguishing between *base* vocabularies and *target* vocabularies: between those employed in an account and those whose use is accounted for. But we can understand that distinction in two different ways. We can take it to be local and relative to particular expressive-explanatory undertakings, on the one hand, or we can take it to be global and absolute, on the other—take it to be a matter of cognitive convenience or taste, or take it to be something we could get substantively wrong because of how things anyway are. That thought brings into view the notion of *universal* base vocabularies, and that, in turn will bring us to the idea of metaphysics.

What is distinctive of empiricism and naturalism, considered abstractly, is that they each see some one vocabulary (or vocabulary-kind) as uniquely privileged with respect to all other vocabularies. Empiricism takes its favored vocabulary (whether it be phenomenal, secondary-quality, or observational) to be *epistemologically* privileged relative to all the rest. In what I think of as its most sophisticated forms, the

---

[8] In the third, methodological, chapter of *Tales of the Mighty Dead*, and again in "Hermeneutic Practice and Theories of Meaning" [*SATS* Nordic Journal of Philosophy Vol 5 No. 1 2004] I offer some more specific and systematic ideas about how the different aspects of discourse addressed by these two sorts of understanding and their associated disciplines complement one another.

privilege is understood more fundamentally to be *semantic*, and only derivatively and consequentially epistemological. Naturalism takes its favored vocabulary (whether it be that of fundamental physics, the special sciences, or just descriptive) to be *ontologically* privileged relative to all the rest. In both cases, what motivates and gives weight and significance to the question of whether, to what extent, and how a given target vocabulary can be logically or algorithmically elaborated from the favored base vocabulary is the philosophical argument for epistemologically, semantically, or ontologically privileging that base vocabulary. These are arguments to the effect that everything that can be known, said or thought, every fact must in principle be expressible in the base vocabulary in question. It is in this sense (epistemological, semantic, or ontological) a *universal* vocabulary. What it can*not* express is fatally defective: unknowable, unintelligible, or unreal. One clear thing to mean by 'metaphysics' is the making of claims of this sort about the universal expressive power of some vocabulary.

4

One might think that metaphysics in this sense is in principle a defective enterprise. I am not really sure how arguments for such a claim at this level of generality go. I have heard four sorts, and am somewhat moved by a fifth. Empiricists rejected metaphysical claims because they want to make claims there can be empirical evidence for, because they take that to be a necessary condition of those claims being candidates for expressing *knowledge*, or, indeed, *meaning* anything. Of course, that this methodological principle contradicted their own empiricist metaphysical principles, including this one, was forcefully pointed out by Hempel in his masterful "The Empiricist Criteria of Cognitive Significance."[9] Corresponding considerations, of course, show that one who rejects empiricism on the grounds that it is unacceptably metaphysical cannot offer these empiricist reasons for rejecting metaphysics. Naturalists often reject metaphysics on the grounds that it is not a product of natural science. When we have real physics, why should we want, and how would we be justified in adding, *meta*physics? From this point of view it seems like the attempt to add an otiose layer of *hyper*-physics. Or, it is an attitude toward the results of science that is itself not science, but a kind of scientism that itself has no scientific credentials. Arthur Fine's rejection of scientific realism in favor of the "natural ontological attitude" is a sophisticated version of this thought. Again, though, someone who rejects naturalism as objectionable metaphysics cannot do so for these naturalistic reasons.

Pragmatists such as Rorty object to the privileging of some vocabularies as universal base vocabularies characteristic of metaphysics on the grounds that it depends on a false, indeed, ultimately magical, understanding of the nature of the sorts of privilege or authority involved. All normative statuses, including these, are instituted by social practice. There are no such normative statuses apart from our practical normative attitudes of taking or treating some things as privileged or authoritative. And for instrumentalist pragmatists like Rorty's hero Dewey (whom he *sometimes* follows in this respect), the ultimate source of those attitudes is our own needs, wants, and convenience. What is wrong with the metaphysical sort of privileging of vocabularies is that it requires the idea of some vocabulary being necessarily privileged by *how things are*—God's vocabulary, or Nature's, or even Mind's, or Meaning's vocabulary—quite apart from our contingent projects and attitudes. Once again there would seem to be an issue about the self-referential stability of this view: is the social nature of normativity, and its normative capacity to trump metaphysical projects, just a feature of how things anyway are? Rorty's response is that this commitment, too, is indeed to be assessed relative to our needs, welfare, and projects. We learned (well, we were supposed to have all learned, though current events make it dubious that we did) from the Enlightenment that it was bad for us in our development as mature humans in charge of our lives and institutions, to understand *moral* normativity as simply reflecting how things were with a non-human (albeit divine) reality. And Rorty's practical proposal for a second Enlightment, completing the work of the first, is to extend that lesson of self-reliance from the practical to the theoretical sphere, for reasons analogous to those that warrant the first move.[10] This is radical and controversial.

---

[9] [ref.]

[10] See for instance "Pragmatism as Anti-authoritarianism." *Revue Internationale de Philosophie* 53:1 (207) (1999): 7–20, and "Universality and Truth" Chapter One of Robert Brandom (ed.) *Rorty and His Critics* [Blackwell's Publishers, 2000]. I offer an assessment of this argument in Section III of "When Philosophy Paints its Blue on Gray: Irony and the Pragmatist Enlightenment" *boundary2* Vol 29 No 2, Summer 2002, pp. 1-28.

A more Wittgensteinean pragmatism addresses metaphysical programs more in a retail than a wholesale spirit. It addresses empiricism and naturalism, and any successor projects one by one, seeking to undermine the specific claims of privilege they put forward. (Here the various criticisms Sellars addresses to empiricism, adverted to in my opening lecture, as well as those of Quine and Austin can serve as paradigms.) But it also expresses a more general suspicion that any such program will turn out, upon examination, to have been motivated by a philosophical anxiety that can be traced to some relatively specific misleading philosophical picture of what knowledge, mind, meaning, or reality *must* be like—on pain of some Bad Consequence. The best anti-metaphysical strategy is then to diagnose and dissolve that underlying misconception, thereby relieving the felt pressure that had made a metaphysical response seem possible because necessary. McDowell reads Kant and Hegel as already engaged in enterprises with this diagnostic-therapeutical, antimetaphysical shape. As far as the general issue is concerned, I think this is an anti-metaphysical attitude, and a template for arguing against metaphysical programs, rather than an argument as such.

5

What most gives me pause about the commitments underlying programs of the sort I am calling 'metaphysical' is that they essentially require us to quantify over all possible vocabularies. Universal base languages are base languages from which *every* vocabulary that is legitimate in some sense (specific to the metaphysical program) can be elaborated as a target vocabulary. I have my doubts about that notion. It is not that I am confident that no sense can be made of the notion of all possible vocabularies. It is rather that I do not think it comes with a clear sense. If it is to make sense, we must give it a sense. And I don't know how to do that. As a graduate student, I attended a seminar offered by the great anthropologist, Clifford Geertz. In the first session, in the course of defining 'culture,' the topic of anthropology, he defined a *language* as "a symbol system in which everything can be said." In a question I expressed my natural curiosity as to what he meant by that "everything." He declined further comment—wisely, perhaps. I'd still like to know what one might mean by it. The worry is that no definite or determinate totality is being delineated. Maybe new vocabularies become possible all the time. This issue arises equally, and for the same reason for talk of all possible *facts*. Following Frege, I understand facts as true thoughts—in the sense of true thinkables, rather than true thinkings, of course. And I understand what is thinkable to be what is claimable, what is expressible in some vocabulary. So there are as many facts as there are true statements in any vocabulary. (I don't object to people using 'fact' so as to allow for the possibility of facts not expressible in any vocabulary. But they must undertake the labor of making sense of that more capacious notion, and showing that there is real work for it to do.) So talking about "all the facts" and talking about "all possible vocabularies" involve a common set of commitments, ones about which I am uneasy. (I am equally uncomfortable with quantification over all "objects." objects : sortals :: facts : claims, all depending on what possible vocabularies there are.)

One might think in this connection about the third proposition of the *Tractatus*: "The world is determined by the facts, and by their being all the facts." (1.11) Wittgenstein rejects the idea of facts that cannot be expressed: for any fact, there is some vocabulary in which it can be stated (clearly). To talk about "all the facts" is then to talk about everything expressible in *any* vocabulary. The *Tractatus* is predicated on there being some *one* vocabulary in which everything sayable at all can be said. It is in this sense a work of metaphysics. And the claim I've cited is a metaphysical claim. Nothing but embarrassment results if we ask about the official status of the second conjunct of this proposition in the botanization of the *Tractatus*. That a specified collection of facts is *all* the facts is not itself one of those facts, nor a fact of that kind. But facts are what can be said. Wittgenstein officially denies that we can even say that there *are* facts, or how many there are (4.1272). Is it, then, something that is *shown*? But how? The awkwardness here shows how the issue of the intelligibility of quantifying over all vocabularies arises in this particular setting. But cognate difficulties attend claims such as that "All the facts are physical facts." Putting aside worries about the boundaries of "physical facts" (it is a very implausible claim if one means *contemporary* physics, and it is hard to say what one means by something like "eventual" or "ideal" physics in a way that remains plausible without circularity) and the sense of "are," the question of how to give a definite sense to "all the facts" in such a way as to make the claim plausible without being question-begging remains. I just don't know what we are saying when we talk this way.

The main point I want to make is that one need not endorse any commitment to *globally* privileging some vocabularies as base vocabularies with respect to *all* vocabularies—which are then taken to be legitimate in one sense or another only if and insofar as they can be reconstructed as target vocabularies elaborated from those base vocabularies—in order to vindicate the project of *locally* identifying *particular* cases where the base-vocabulary/target-vocabulary relation can be shown to obtain. For there is a distinctive kind of[11] philosophical understanding that consists in practical mastery of the relations among vocabularies that become visible when we trace various possible base→target relations through the (admittedly in general ill-defined) field of vocabularies. As an example, functionalism, and its specifically computational AI species, is an attempted *local* privileging of one vocabulary over another: functional vocabulary over intentional vocabulary. It has no pretensions at all to *global* privileging of functional vocabulary, to its being a *universal* base vocabulary.

Though elsewhere I have offered some reasons for skepticism about even the most plausible, pragmatist, version of AI-functionalism about intentionality, I think that a suitably broadly construed functionalism is something like the only straw floating as a prospect for an account, in other terms, of intentionality.[12] Of course, it may well be, as some Wittgensteinians, for instance, think, that intentional discourse can only be understood in its own terms, and not in terms of some other vocabulary. For reasons I discuss above, that does not mean that understanding would totter. Analytic (algorithmic, algebraic) understanding is not the only kind of genuine philosophical understanding, and it is not always available.

David Lewis propounded a view of philosophy that was inspiring to me when I was his student, and inspires me still. He thought what philosophers should do is lay down a set of premises concerning some topic of interest as clearly as possible, and extract consequences from them as rigorously as possible. Having done that, one should lay down another, perhaps quite different set of premises, and extract consequences from them as rigorously as possible. The point was not in the first instance to endorse the conclusions of any of these chains of reasoning, but to learn our way about in the inferential field they all defined, by tracing many overlapping, intersecting, and diverging paths through the terrain. That is how we would learn what difference it would make, in various contexts, if we were to endorse some claim that figures as a premise in many of the inferences, and what might entitle us to a claim that shows up as a consequence in many of the inferences. Actually plumping for and defending any of these theses is then a subsequent, parasitic, and substantially less important stage of the process. The principal aim is not belief, but understanding.

One thing that was liberating and exhilarating about this metaphilosophical attitude is that Lewis accordingly didn't care much what reasons one had for starting with one set of premises rather than another. He was entirely open to and indeed eager to turn his awe-inspiring intellect to following out the consequences of even the wackiest of claims. Pavel Tichy visited while I was still in graduate school at Princeton, and he was then gripped by just such an obsession. He had somehow gotten the idea that there was some finite number *n* such that it was a necessary truth that there were exactly *n* things in each possible world, and that there were no further constraints on trans-world identification of objects or counterparts. I don't know how he thought he could count objects (Lewis later had what is at least a responsive answer to this question from a physicalist point of view). And I can't imagine what would make one think his axiom true. But Lewis was perfectly willing to figure out the sort of modal logic and metaphysics one would get on that assumption. You never know where you might learn something.

I think this is the spirit in which we should think about semantic relations between different vocabularies. It is worth seeing how and to what extent different target vocabularies can be elaborated from various base vocabularies—including, and perhaps especially with pragmatic detours through the specifications of practices-or-abilities necessary or sufficient to deploy those vocabularies, as I recommend in *Between Saying and Doing*—because that is a way of coming to practically to know our way around those vocabularies, our discursive practices, and the subject-matters they make it possible for us to talk and think about. Exploring all the available paths between landmarks is a way, perhaps the only way, of learning to find our way around these woods, acquiring a practical conceptual mastery of the many aspects of discursive practice, and their relations to one another. Let me repeat that I am not claiming this is all there is to philosophy, or that this is the only way to do philosophy. But this sort of semantic analysis yields one valuable kind of understanding that is potentially of great philosophical value.

---

[11] In Chapter Three of *Between Saying and Doing*.
[12] I say something about this in the first chapter of *Tales of the Mighty Dead*.

Notice that on this view, one's exploration of semantic relations (including pragmatically mediated ones) among vocabularies of antecedent philosophical interest, need not be motivated by some global, monolithic program, such as empiricism or naturalism.

The distinctive kind of semantic understanding I am suggesting is the genus of which metaphysical understanding is a species is not hostage to such programs. That understanding can be well served by accumulating particular, local connections that support no antecedent global program and perhaps could be predicted by none. Nor must the search for such semantic relations among vocabularies and the discursive practices-or-abilities they specify or that deploy them be motivated by some deep-seated philosophical anxiety or puzzlement, the proper deflating diagnosis of which then exhibits or renders the task of exploring those relations otiose. Simple curiosity, the desire to deepen our understanding can suffice, for this sort of philosophical theorizing as for the empirical scientific variety. Indeed, as Kuhn has taught us that in the latter case, it really doesn't matter *why* the scientists do what they do, since the institution can ensure that so long as they act professionally, the result will be to extend our knowledge and deepen our understanding, so we might strive to make it be in philosophy—a light and harmless sort of motivational scientism.

6

With this thought on the table, I want to return briefly to the issue of metaphysics. Taking our cue from Geertz, we might think of metaphysics as the enterprise of crafting a vocabulary in which everything can be said. Now he might be right that, in the sense I am after, natural languages are autonomous vocabularies in which everything can be said. But "craft" is doing some work in this definition. The metaphysician aims to *construct* a *technical*, *artificial* vocabulary with that same expressive power. Why? The greater control that regimentation gives vocabularies whose basic semantics is stipulated—in some other vocabulary, perhaps a natural language (no escaping the need for hermeneutic understanding)—and the rest of whose semantics is computed algorithmically. For we have a distinctively clear sort of understanding of whatever other vocabularies can be elaborated as target vocabularies from a base vocabulary constructed so as to exhibit this structure. We can get around my earlier worries about the concept of <u>all possible vocabularies</u> by understanding the "everything" regulatively, rather than constitutively. That is, for every vocabulary anyone comes up with, the metaphysician is committed to the favored base vocabulary being adequate, when suitably elaborated, to express what it expresses. We start by trying to codify the vocabularies we have, but acknowledge the commitment to address any more that may come along.

I think metaphysics in this sense is a perfectly reasonable undertaking, and that we potentially have a lot to learn from pursuing it. It is, perhaps, somewhat Quixotic—but that is a practical, not a theoretical drawback. If we are to reap the rewards in understanding that engaging in this kind of metaphysical enterprise promises, however, I think it is crucial that it be pursued in the open-minded, pluralistic spirit of Lewis, and not in a more small-minded and exclusionary one. The distinction arises when the metaphysician fails to reconstruct in the favored terms *all* the antecedent uses of *all* the vocabularies it aspires to codify. I take it that such partial failures are inevitable. The task is just too hard, both for practical reasons and for principled ones. It is probably too much to ask even that for every target vocabulary one find some other regimented base vocabulary from which it can be elaborated. Finding some one regimented base vocabulary in terms of which every such target vocabulary can be reconstructed is far more difficult and unlikely. That is why it is to some extent a Quixotic quest. But for the same sort of reasons that led to Popper's methodological recommendation to endorse the strongest, most easily falsifiable theory not already falsified by the evidence, it can make sense to pursue the quantificationally more difficult goal. For to the limited extent that one does succeed, one finds out more both about the metaphysical base vocabulary, and about the target vocabularies to which it turns out either to be expressively adequate, or not to be expressively adequate. Traditional metaphysics treats the distinction as invidious. It denigrates and dismisses what resists formulation in its favored terms as ontologically second class: as unreal, as mere appearance. So for Leibniz relations and evil are unreal, relegated to the phenomenal realm of appearance. Later metaphysicians found themselves similarly rejecting as unreal such phenomena as time. (One of my favorite mind-benders is the attempt of some British idealists to dismiss finitude as an illusion, the effect produced by the infinite Absolute—what there really is—on poor incapable…merely finite minds!) In a more contemporary semantic key, the term of disapprobation may be 'unintelligible' rather than the ontological 'unreal.' That, at any rate, is the way the logical empiricists talk, and admirers of theirs such as

Quine continued the practice (modal logic merely "engenders an illusion of understanding" of modal vocabulary, and intentional vocabulary is merely apparently coherent). (Nineteenth century Idealism, which more or less equates the unreal with the unintelligible, may be thought of in this respect as a transitional phase.)

But such a mean-spirited, suspicious, begrudging, exclusionary attitude is not the only one possible. One might instead take the great positive payoff of a particular metaphysical effort to consist in no small part in the particular line it draws between what target vocabularies (and which bits of each) can and which cannot be captured with the expressive resources of the base vocabulary whose metaphysical credentials are being explored. Each regimented base vocabulary, we might suppose, will determine a different boundary between the (relatively) expressible and inexpressible. Suppose we found out (it wouldn't surprise me) that there is simply no way to say in the language of fundamental physics, no matter how it is elaborated with the resources of the most advanced logic, what Beckett said when he said "I can't go on. I'll go on," or what Nixon said when he said "I am not a crook," or what the Buddha said when he told the hot-dog man "Make me one with everything," or… . We don't need to say that they didn't say anything, or to pretend that we can't understand what they did say, we don't need to deny that there is such a thing as going on, being a crook, and so on, to learn something about saying things from the relation between their vocabularies and that of fundamental physics that shows up in this expressive mismatch. Indeed, I think we learn more if we do not go on to adopt the wholly optional dismissive attitude. Our slogan should be "Metaphysical discrimination without denigration." And just as Lewis thought it essential that we draw consequences rigorously from *many* sets of premises, so as to learn our way around by taking many different paths through the terrain, so the virtues of the metaphysical enterprise will manifest themselves most fully if we try out *many* different possible metaphysical base vocabularies. (A side benefit of adopting this plan is that we then need not resign ourselves to living out our lives oppressed by the steady drip, drip, drip… of naturalistic semantics and ontology.)

My characterization of metaphysics transposes what is normally thought of in *ontological* terms into a *semantic* key. These versions can be thought of as related to one another on the model of Carnap's *material* and *formal* modes. I've described the metaphysical project in metalinguistic terms. It might seem that the translation back and forth between these two ways of talking is so straightforward that it is perverse to flout ordinary philosophical usage by insisting on the metalinguistic version here. But in this case there is a significant asymmetry between them. Indeed, I think the asymmetry here reveals something important about the Carnapian dyad that we might otherwise not have seen: a new justification, from the side of pragmatism, for the characteristically analytic preference for the formal mode.

For thinking of the metaphysical enterprise in semantic terms, as seeking to establish distinctive sorts of relations among vocabularies, opens up the possibility of considering in this case, too, *pragmatically mediated* semantic relations between vocabularies, in addition to the traditional kind. In particular, we can lay alongside the aspiration to find a vocabulary in which everything can be said, the aspiration to find one in which one can say everything one must be able to *do* in order to say anything, that is to *use* any vocabulary whatsoever. This is just the idea of a <u>universal *pragmatic* metavocabulary</u>.

The sort of illumination one would get from succeeding at the task of constructing a regimented *de facto* universal pragmatic metavocabulary is not exactly the same as what one would get from succeeding at the task of constructing a regimented vocabulary whose expressive power encompassed that of all the vocabularies we could come up with to try it out on. But they would evidently be complementary forms of understanding: one telling us what we can *say*, and the other what we must *do* to say it. Further, any adequate pragmatic metavocabulary for a semantically adequate metaphysical vocabulary would *be* a universal pragmatic metavocabulary. I have already suggested, however, that the real payoff from the metaphysical enterprise should not be thought of as consequent upon the anticipation of complete success at producing a regimented semantically expressively universal vocabulary. In place of such a wholesale cognitive reward, we should think of the accumulation of retail rewards. Each only partially successful try at a universal metaphysical vocabulary draws a line between those antecedent vocabularies it can reconstruct, and those it cannot. And each such endeavor will draw a *different* line. The lesson I drew from the young David Lewis's methodologically principled polymorphous theoretical promiscuity is that a valuable kind of understanding consists in the sort of knowing our way about secured only by multiplying the criss-crossing of concrete ways of drawing the boundary between the expressible and the inexpressible, not globally and absolutely, but locally and relative to specific base vocabularies. The same will hold true of attempts to construct regimented universal pragmatic metavocabularies: their value lies in the details of their only partial successes, in where, specifically, they fail, and in how the line between partial success and

partial failure varies as we try out quite different candidate base vocabularies. Here one thinks of the parable with which Hempel closes "Empiricist Criteria of Cognitive Significance," in which a dying father tells his sons a vast treasure is buried in their vineyard. Only many years later do they realize that their energetic but fruitless digging in search of the imaginary trove has led to their turning over the soil in just the way needed to ensure that their vines flourish. The romantic dream of total transformation by a single, magical find inspires the hard, unexciting daily work that gradually lays up the real treasure.

The parallel between the metaphysical goal of crafting a regimented universally expressive vocabulary and that of constructing a universal pragmatic metavocabulary, the genus of which these two tasks are species, is invisible if we think of metaphysics exclusively in ontological terms. The wider perspective is available only if we construe it semantically, in terms of relations between vocabularies. So viewed, they show up as complementary, corresponding to elements not only of the dimension defined by the semantic/pragmatic distinction, but also of the objective/subjective dimension: what is talked about and talking about it. They address objective-ontological and subjective-practical sides of the coin of discursiveness. Although both these sorts of totalizing project come into view from the vantage-point of meaning-use analysis, what I am doing in *Between Saying and Doing* is neither metaphysics nor an attempt to construct a universal pragmatic metavocabulary. Rather, I am looking for a different kind of metavocabulary. It is at a higher level, making it possible to express crucial structural features of the *relations* between the dimensions of discursiveness they take as their targets, relations between what is said and what is done. But out of that project grows both a new way of thinking about the project of metaphysics, and about the possibility of a new, pragmatic, kind of universal metavocabulary. Perhaps the availability of this new chart in the metaphilosophical atlas will yield other worthwhile dividends as well.