

4/30/2006

*Between Saying and Doing:
Towards an Analytic Pragmatism*

Lecture One (May 3, 2006)

Extending the Project of Analysis¹

My principal aim in these lectures is to present a new way of thinking about language, and about the relations between meaning and use, or between what is said and the activity of saying it. To that end, I will introduce a new metatheoretic conceptual apparatus, and develop it through applications to a number of sorts of locution that have properly been the focus of intense philosophical interest: logical and semantic vocabulary, indexical vocabulary, modal, normative, and intentional vocabularies. The concerns that animate this enterprise arise from a way of thinking about the nature of the general project pursued by analytic philosophy over the past century or so, and about its confrontation with Wittgensteinian pragmatism. Justifying that rendering of the tradition would take me far afield, but it will be well to begin with at least a sketch of that motivating picture.

Section 1: The Classical Project of Analysis

I think of analytic philosophy as having at its center a concern with semantic relations between what I will call ‘vocabularies’. (I use the term ‘*vocabulary*’ here in a very general and relaxed

¹ Throughout its gestation, this project has been generously supported by the Andrew W. Mellon Foundation, first through their funding of a research year at the Center for Advanced Study in the Behavioral Sciences, at Stanford University, and more recently, and even more lavishly, through their Distinguished Achievement in the Humanities Award. Whatever I have been able to accomplish here would not have been possible without that support, and the extended time for concentrated attention that it has afforded me.

fashion, to pick out, by what expresses it, what is expressed by any sort of language fragment or meaningful expression-type. In this usage, logically atomic sentences, semantic discourse, indexical and observational tokenings all count as vocabularies.²) Its characteristic form of question is whether and in what way one can make sense of the meanings expressed by one kind of locution in terms of the meanings expressed by another kind of locution. So, for instance, two early paradigmatic projects were to show that everything expressible in the vocabulary of number-theory, and again, everything expressible using definite descriptions, is expressible already in the vocabulary of first-order quantificational logic with identity.

The nature of the key kind of semantic relation between vocabularies has been variously characterized during the history of analytic philosophy: as analysis, definition, paraphrase, translation, reduction of different sorts, truth-making, and various kinds of supervenience—to name just a few contenders. In each case, however, it is characteristic of classical analytic philosophy that *logical* vocabulary is accorded a privileged role in specifying these semantic relations. It has always been taken at least to be *licit* to appeal to logical vocabulary in elaborating the relation between *analysandum* and *analysans*—target vocabulary and base vocabulary—and, according to stronger versions of this thesis, that may be the *only* vocabulary it is licit to employ in that capacity. I will refer to this aspect of the analytic project as its commitment to ‘*semantic logicism*’.³

² The intent of this capacious usage is in part to leave open what have sometimes been quite contentious issues about the proper specification of the objects of analysis: should what is to be analyzed be understood as expressions, meanings, concepts, or properties? I talk about vocabulary ‘*kinds*’ to indicate a certain *generality* of concern. “Bachelors are unmarried males,” may be offered as an *example* of an analysis (in these days of civil unions, now an obtrusively historically dated one), in order to highlight a particular kind of semantic relation, but the philosophical *interest* of that relation was always taken to lie in the possibility that it might hold not just for some particular locutions, but between what is expressed by whole stretches of discourse: between moral and naturalistic vocabularies, for instance.

³ In this usage, the logicism about mathematics characteristic of Frege’s *Grundgesetze* and Russell and Whitehead’s *Principia* is semantic logicism about the relations between mathematical and logical vocabularies.

All this is of course very rough. It is an attempt to produce a general formula, by varying the parameters of which one can encompass in one more-or-less smooth curve such otherwise disparate cardinal data-points of the classical project of analysis in the first half of the twentieth century as Russell and Moore, the Cambridge analysts of the twenties, Carnap and the Vienna Circle, Ayer, Ryle, C.I. Lewis, and Quine.

If we ask *which* were the vocabulary-kinds whose semantic relations it was during this period thought to be important to investigate, at least two *core programs* of classical analytic philosophy show up: *empiricism* and *naturalism*. These venerable modern philosophical traditions in epistemology and ontology respectively were transformed in the twentieth century first by being transposed into a *semantic* key, and second by the application of the newly available *logical* vocabulary to the self-consciously semantic programs they then became. By calling these ‘core programs’ of classical analytic philosophy I do not mean to be claiming that the index figures defining that tradition in every case subscribed to empiricist and naturalist conclusions. Moore, for instance, was anti-naturalist about moral normativity, and Neurath rejected empiricism in areas where it seemed to conflict with naturalism. Rather, these are core programs in the sense that philosophers in the tradition I am talking about typically felt obliged to *address* them, whether or not in the end they endorsed those programs.

As base vocabularies, different species of *empiricism* appealed to *phenomenal* vocabulary, expressing how things appear, or to *secondary-quality* vocabulary, or, less demandingly, to *observational* vocabulary. Typical target vocabularies include objective vocabulary formulating claims about how things actually *are* (as opposed to how they merely appear), primary-quality vocabulary, theoretical vocabulary, and modal, normative, and semantic vocabularies. The generic challenge is to show how what is expressed by the use of such target vocabularies can be reconstructed from what is expressed by the base vocabulary, when it is elaborated by the use of logical vocabulary. Carnap’s *Aufbau* and Ayer’s *Language, Truth, and Logic* may be thought of as paradigmatic attempts to carry out such logical empiricist programs.

As base vocabularies, different species of *naturalism* appealed to the vocabulary of *fundamental physics*, or to the vocabulary of the natural sciences (including the special sciences) more generally (loosening the reading of the unity-of-science commitment characteristic of the physicalist version of naturalism), or just to objective descriptive vocabulary, even when not regimented by incorporation into explicit scientific theories. Typical targets include normative, semantic, and intentional vocabularies. Once again, the generic challenge is to show how what is expressed by the use of those target vocabularies can be reconstructed from what is expressed by the base vocabulary, when it elaborated by the use of logical vocabulary. Here the ontological impetus for the program licenses less demanding semantic criteria of adequacy: mere co-reference, specification in naturalistic vocabulary of truth-makers for claims couched in the target vocabularies, or even just supervenience of the target vocabulary on the favored sort of naturalistic vocabulary have been thought to suffice for the success of a naturalistic analysis. Emotivism and its descendents about ethics, behaviorism and perhaps functionalism about intentionality, and causal-counterfactual, informational, and teleosemantic theories of semantics may be thought of as paradigmatic attempts to carry out such naturalist semantic programs.

Section 2: The Pragmatist Challenge

What I want to call the “classical project of analysis”, then, aims to exhibit the meanings expressed by various target vocabularies as intelligible by means of the logical elaboration of the meanings expressed by base vocabularies thought to be privileged in some important respects—epistemological, ontological, or semantic—relative to those others. This enterprise is visible in its purest form in what I have called the “core programs” of empiricism and naturalism, in their various forms. In my view the most significant conceptual development in this tradition—the

biggest thing that ever happened to it⁴—is the *pragmatist challenge* to it that was mounted during the middle years of the twentieth century. Generically, this movement of thought amounts to a displacement from the center of philosophical attention of the notion of meaning in favor of that of use: in suitably broad senses of those terms, replacing concern with *semantics* by concern with *pragmatics*. The towering figure behind this conceptual sea-change is, of course, Wittgenstein. In characterizing it, however, it will be useful to approach his radical and comprehensive critique by means of some more local semantically corrosive argumentative appeals to the practices of deploying various vocabularies rather than the meanings they express.

Wilfrid Sellars (one of my particular heroes) criticizes the empiricist core program of the classical project of analysis on the basis of what one must *do* in order to *use* various vocabularies. He argues that *none* of the various candidates for empiricist base vocabularies are practically autonomous, that is, could be deployed in a language-game one played though one played no other. For instance, no discursive practice can consist entirely of making non-inferential observation reports, for such reliably differentially elicited responses qualify as *conceptually* contentful or *cognitively* significant only insofar as they can serve as *premises* from which it is appropriate to draw *conclusions*, that is, as *reasons* for other judgments. Drawing such conclusions is applying concepts *inferentially*—that is, precisely *not* making non-inferential *observational* use of them.⁵

⁴ The subsequent *modal revolution* might be the next most important event in the development of the analytic tradition. It will be discussed beginning in Lecture Four.

⁵ This argument occupies roughly the first half of his classic “Empiricism and the Philosophy of Mind” [reprinted by Harvard University Press, 1997]. His critique of the phenomenalist version of empiricism can be found in “Phenomenalism”, in his collection *Science, Perception, and Reality* [Routledge Kegan Paul 1963].

Further, he argues, phenomenalist vocabulary, which says how things merely *look* or *appear* to be is likewise not semantically autonomous, because not pragmatically autonomous. On his account, what one is *doing* in saying that something *looks* red essentially includes evincing one's other-things-being-equal disposition to respond to it by claiming that it *is* red, while simultaneously explicitly *withholding* that endorsement (perhaps on the basis of a suspicion that other things are in this case *not* equal). If that is right about the practice of deploying phenomenal vocabulary, it follows that the phenomenal vocabulary used to say how things appear to one is also not pragmatically autonomous. One can use 'looks red'-talk only as part of a larger game that includes also using 'is red'-talk. As for secondary-quality vocabulary, at least a necessary condition of using ϕ in that way is that one cannot count as having mastered the concept ϕ unless one also can use 'looks- ϕ '. But if that is so, then the combination of the semantic (because pragmatic) dependence of 'looks'-talk on 'is'-talk and of non-inferential on inferential uses of expressions will, under plausible further assumptions, have the consequence that the ability to use secondary-quality vocabulary requires also the ability to use primary-quality vocabulary. Sellars takes it that every plausible philosophical motivation for *caring* about the possibility of an empiricist semantics for some other region of discourse depends on the assumption that the empiricist base vocabulary is semantically autonomous. In any case, the interest of a phenomenalist account of objective talk, of a secondary-quality account of primary-quality talk, and of an account in observational terms of non-observational applications of concepts, lapses if it turns out the capacity to use those target vocabularies is already presupposed by the capacity to use the empiricist base vocabularies.

My point here is not to claim that these arguments are devastating to the empiricist program in both its traditional and its twentieth-century logical forms—though I certainly think there is something to these arguments. It is rather to point out how Sellars appeals to considerations concerning the *practices of using* expressions, the *abilities* required to count as deploying vocabulary with the sorts of meaning in question, the way in which the practices necessary to use one sort of locution can depend practically on practices sufficient to use another, and the notion of an autonomous discursive practice, in criticizing a core *semantic* program of analysis. In a similar pragmatist vein, Sellars argues that fundamental criteria of adequacy for semantic accounts of talk of how things objectively are in terms of how they subjectively seem, of primary-quality talk in terms of secondary-quality talk, or theoretical talk in terms of observational talk, must make essential use of *modal* vocabulary specifying how things

would look if..., what secondary qualities *would* present themselves under various circumstances, and what *would* be observed if an electron-flow were present. If only the practical abilities made explicit by modal vocabulary make possible the semantic elaboration of empiricist base vocabularies into their target vocabularies, then the empiricist program is globally and in principle hostage to solving Hume's problem about the empiricist semantics of modal vocabulary.⁶ Again, Sellars argues that the *antecedents* of those counterfactually robust conditionals will not in general be specifiable in the empiricist base vocabularies, without using some of the target vocabularies. The phenomenalist must appeal to what would happen if I *actually* turn my head, or go in the next room, the observationalist to what would be observed if the apparatus is actually working properly—something specifiable only in theoretical terms—and so on.

Quine offers an even broader pragmatist objection, not only to the empiricist program, but to essential aspects of the whole analytic semantic project. For he attacks the very notion of meaning it presupposes. Quine is what I have elsewhere called a “methodological” pragmatist.⁷ That is, he takes it that the whole point of a theory of meaning is to explain, codify, or illuminate features of the *use* of linguistic expressions. He, like Dummett, endorses the analogy: *meaning is to use as theory is to observation*. And he argues that postulating meanings associated with bits of vocabulary yields a *bad* theory of discursive practice. He takes the notion of meaning to presuppose an essentialist view of language-use. “Meaning is what essence becomes, when it is detached from the thing and attached to the word.”⁸ The idea of meanings as entities—of whatever sort—somehow associated with expressions and determining how it is correct to use them in a way that is not in turn reciprocally sensitive to how we *actually* use them is a philosopher's fiction, a metaphysical reification, the “myth of the museum” in which those meanings are pinned, labeled, and laid out for our inspection.

⁶ This issue is addressed in my fourth lecture.

⁷ “Pragmatics and Pragmatism” in James Conant and Urszula M. Zeglen (eds.) *Hilary Putnam: Pragmatism and Realism* [Routledge 2002], translated as “Pragmatik und Pragmatismus”, pp. 29-58 in Sandbothe, M (ed.) *Die Renaissance des Pragmatismus* [Velbrück Wissenschaft, 2000].

⁸ “Two Dogmas of Empiricism” *The Philosophical Review* **60** (1) January 1951 pp. 20-43. Reprinted in revised form in Quine's *From a Logical Point of View* [Cambridge, MA, Harvard University Press, 1953].

If there were such things as meanings that determine how it would be correct to use our expressions, then those meanings would at least have to determine the inferential roles of those expressions: what follows from applying them, what applying them rules out, what is good evidence for or against doing so. But what follows from what depends on what else is true—on laws of nature and obscure contingent facts—that is, on what claims can serve as auxiliary hypotheses or collateral premises in those inferences. If we look at what practical abilities are required to deploy various bits of vocabulary—at what one has to be able to *do* in order to count as *saying* something with them—we do not find any special set of these whose practical significance can be understood as *pragmatically* distinctive of *semantically* necessary or sufficient conditions.⁹

Quine thought one could save at least the naturalist program by retreating semantically to the level of reference and truth-conditions. James and Dewey appeal to the same sort of methodological pragmatism in support of more sweeping sorts of semantic revisionism—pursuing programs that Rorty, for instance, argues should be understood as more rejectionist than properly revisionist. And under the banner “Don’t look to the meaning, look to the use,” Wittgenstein further radicalizes the pragmatist critique of semantics. Pointing out to begin with that one cannot assume that uses of singular terms have the job of picking out objects, nor that declarative sentences purport to state facts, he goes on to deny, in effect, that such uses even form a privileged center, on the basis of which one can understand more peripheral ones. (“Language has no downtown.”) Of course there is a *lot* going on in the later Wittgenstein, but I am going to engage in the doubly risky enterprise of both focusing on a few lessons that are of particular importance to my story, and trying to formulate them as explicit principles—not a kind of undertaking he either indulges in himself or encourages in others.

⁹ Ibid.

I take it that Wittgenstein, like the other methodological pragmatists, understands the home language-game of the concept of meaning to be explanation of how expressions are correctly *used*. And he is profoundly skeptical about the utility or applicability of the model of postulation, explanation, and systematization in the case of discursive practices—about the possibility of systematically *deriving* aspects of use from assigned meanings. Seen from this perspective, the idea of the classical project of analysis is to codify, using logical vocabulary, the meanings expressed by one vocabulary—from which we are to derive proprieties of its use—in terms of the meanings expressed by some *other* vocabulary—from which we can derive proprieties of *its* use. One idea, I think, is that this makes sense as an enterprise only if we think of the uses as species of a genus—of them all being the same general *kind* of use, say stating facts, or representing states of affairs. This may seem plausible if we focus on a very restricted set of uses—just as, in the case of tools, we might be impressed to notice that hammer and nails, screwdriver and screw, glue, gluepot and brush all have the function of attaching more-or-less flat things to one another. So we can think of declarative sentences as stating empirical, physical, normative, modal, and intentional facts, making *claims* about such states of affairs. But if we think of the uses as *very* different—if we think also about the carpenter’s pencil, ruler, level, and tool-belt—if we think of linguistic practice as a *motley*, of uses as not coming in a simple or systematic or even determinate variety, then the very idea that there is such a thing as *meanings* that permit the codification of proprieties of quite disparate kinds of use—even with liberal use of logical *elaboration* of the meanings—becomes contentious and in need of justification both in general and in each particular case.

More specifically, Wittgenstein uses the image of “family resemblances” to urge that the *kinds* into which linguistic practices and the vocabularies caught up in them are functionally

sorted—what belong together in boxes labeled ‘game’, ‘name’, ‘assertion’, ‘observation’ and so on—do not typically admit of specification in terms of underlying principles specifiable in other vocabularies, whether by genus and differentiae or any other kind of explicit rule or definition. It is easy to understand this line of thought as entailing a straightforward denial of the possibility of semantic analysis in the classical sense.

I think that the thought underlying these observations about the unsystematic, unsurveyable variety of kinds of uses of expressions and about the uncodifiable character of those kinds concerns the essentially *dynamic* character of linguistic practice. I think Wittgenstein thinks that an absolutely fundamental discursive phenomenon is the way in which the abilities required to deploy one vocabulary can be practically *extended*, elaborated, or developed so as to constitute the ability to deploy some further vocabulary, or to deploy the old vocabulary in quite different ways. Many of his thought-experiments concern this sort of process of *pragmatic projection* of one practice into another. We are asked to imagine a community that uses proper names only for people, but then extends the practice to include rivers. There is no guarantee that interlocutors can master the extended practice, building on what they can already do. But if they can, then they will have changed the only ‘essence’ proper-name usage could be taken to have had. In the old practice it always made sense to ask for the identity of the *mother and father* of the named item; in the new practice, that question is often senseless. Again, we are asked to imagine a community that talked about having gold or silver in one’s tooth, and extends that practice to talk about having pain in one’s tooth. If as a matter of contingent fact the practitioners can learn to use the expression ‘in’ in the new way, building on but adapting the old, they will have fundamentally changed the ‘meaning’ of ‘in’. In the old practice it made sense to ask where the gold was *before* it was in one’s tooth; in the new practice

asking where the pain was before it was in the tooth can lead only to a distinctively philosophical kind of puzzlement (a permanent, unavoidable phenomenon because the pragmatic projection and extension of one set of discursive practices into another that generates it is ubiquitous, and its possibility is fundamental to the discursive—and possibly to the practical—as such¹⁰).

At every stage, what practical extensions of a given practice are possible for the practitioners can turn on features of their embodiment, lives, environment, and history that are contingent and wholly particular to them. And which of those developments actually took place, and in what order, can turn on any obscure fact. The reason vocabulary-kinds resist specification by rules, principles, definitions, or meanings expressed in other vocabularies is that they are the current time-slices of processes of development of practices that have this character—and that is why the collection of uses that is the cumulative and collective result of such developments-by-practical-projection is a motley.¹¹ If that is right, then any codification or theoretical systematization of the uses of those vocabulary-kinds by associating with them meanings, specifiable in other vocabularies, which determine which uses are correct will, if at all successful, be successful only contingently, locally, and temporarily. Semantics on this view is an inherently Procrustean enterprise, which can proceed only by theoretically privileging some aspects of the use of a vocabulary that are not at all practically privileged, and spawning philosophical puzzlement about the intelligibility of the rest.¹² On this conception, the classical project of analysis is a disease that rests on a fundamental and perennial

¹⁰ I learned to think about Wittgenstein this way from Hans-Julius Schneider's chapter on Wittgenstein in *Phantasie und Kalkul: Über die Polarität von Handlung und Struktur in der Sprache* [Frankfurt am Main, Suhrkamp Verlag, 1992].

¹¹ A patient and detailed investigation of the mechanisms of this phenomenon in basic descriptive and scientific concepts, and an extended argument for its ubiquity can be found in Mark Wilson's exciting and original *Wandering Significance* [Oxford University Press, 2006]. I return to a consideration of the practical, non-rule-governed extension of usage in my third lecture.

¹² I would be happy if those who dance with his texts find affinities here with Hegel's insistence that the metaconceptual categories of *Verstand* must be replaced by those of *Vernunft*. It is characteristic of his philosophical ambition that draws the opposite of Wittgenstein's conclusions from an appreciation of the dynamics of conceptual development and its sensitivity to arbitrary contingent features of the practitioners, devoting himself to elaborating what he insists is the *logic* of such processes and the conceptual contents they shape.

kind of misunderstanding—one that can be removed or ameliorated only by heeding the advice to replace concern with *meaning* by concern with *use*. The recommended philosophical attitude to discursive practice is accordingly *descriptive particularism*, *theoretical quietism*, and *semantic pessimism*.

Section 3: Extending the Project of Analysis: Pragmatically Mediated Semantic Relations

On this account Wittgenstein is putting in place a picture of discursive meaningfulness or significance that is very different from that on which the classical project of analysis is predicated. In place of *semantics*, we are encouraged to do *pragmatics*—not in the sense of Kaplan and Stalnaker, which is really the semantics of token-reflexive expressions, nor again in the sense of Grice, which addresses conversational heuristics in terms that presuppose a prior, independent, classical semantics—but ‘pragmatics’ in the sense of the study of the *use* of expressions in virtue of which they are meaningful at all. To the formal, mathematically inspired tradition of Frege, Russell, Carnap, and Tarski, culminating in model-theoretic semantics, is opposed an anthropological, natural-historical, social-practical inquiry aimed both at demystifying our discursive doings, and at deflating philosophers’ systematic and theoretical ambitions regarding them. I think that philosophers of language have tended to draw this opposition in the starkest possible terms, treating these approaches as mutually exclusive, hence as requiring that a choice be made between them. Those who are moved by the pragmatist picture generally accept the particularist, quietist, anti-theoretical conclusions Wittgenstein seems to have drawn from it. And those committed to some version of the project of semantic analysis have felt obliged to deny the significance of pragmatics in this sense, or at the least to dismiss it as irrelevant to properly semantic concerns. In the most extreme cases, the attitudes of anti-pragmatist philosophers of

language to Wittgenstein's picture verges on that of the Victorian lady to Darwin's theory of evolution: One hopes that it is not true, and that if should turn out to be true, at least that it not become generally known.

But I do not think we are obliged to choose between these approaches. They should be seen as complementing rather than competing with one another. Semantics and pragmatics, concern with meaning and concern with use, ought to be understood as aspects of one more comprehensive picture of the discursive. Pragmatist considerations do not oblige us to focus on pragmatics to the exclusion of semantics; we can deepen our semantics by the addition of pragmatics. If we extract consequences from the pragmatists' observations somewhat more modestly and construe the analytic project somewhat more broadly, the two will be seen not only as compatible, but as mutually illuminating. If we approach the pragmatists' observations in an analytic spirit, we can understand pragmatics as providing special resources for extending and expanding the analytic project, from exclusive concern with relations among meanings to encompass also relations between meaning and use. In its most ambitious form, such an enterprise would aspire to articulate something like a *logic* of the relations between meaning and use.

If we leave open the possibility that the use of some vocabulary may be illuminated by taking it to express some sort of meaning or content—that is, if we do not from the beginning embrace semantic nihilism—then the most important positive pragmatist insight will be one complementary to the methodological pragmatism I have already identified. The thought underlying the pragmatist line of thought is that what makes some bit of vocabulary mean what it does is how it is used. What we could call *semantic* pragmatism is the view that the only

explanation there could be for how a given *meaning* gets associated with a vocabulary is to be found in the *use* of that vocabulary: the practices by which that meaning is conferred or the abilities whose exercise constitutes deploying a vocabulary with that meaning. To broaden the classical project of analysis in the light of the pragmatists' insistence on the centrality of pragmatics we can focus on this fundamental relation between use and meaning, between practices or practical abilities and vocabularies. We must look at what it is to use locutions *as* expressing meanings—that is, at what one must *do* in order to count as *saying* what the vocabulary lets practitioners express. I am going to call this kind of relation “practice-vocabulary sufficiency”—or usually, “PV-sufficiency,” for short. It obtains when engaging in a specified set of practices or exercising a specified set of abilities is sufficient for someone to count as *deploying* a specified vocabulary.

The labels I want to use for the two terms of this basic relation between meaning and use are ‘vocabulary’ and ‘practice or ability’. I said something early on about the relaxed and elastic way in which I want to use the word ‘vocabulary’. Aiming at a potentially fruitful generality, I will adopt a similar policy governing the phrase ‘practice or ability’. One rationale for this phrase is that I want to be noncommittal concerning the important strategic issue of whether to focus on *social* practices, as communally *conferring* the meanings or *instituting* the cognitive significances in virtue of which some noises, marks, or performances count as uses of a *vocabulary*, or on *individually* manifestable abilities that count as deploying those noises, marks, or performances *as* a vocabulary. Talking about use—the subject of pragmatics in the sense in which each of methodological and semantic pragmatism in its own way asserts the priority of pragmatics over semantics—in terms of “practices or abilities” is intended to make what I go on to say here apply whether one takes a social or an individualistic approach to the activities in virtue of which otherwise semantically inert sign-designs (for instance, the sign-post thought of just as a piece of wood) qualify as constituting a vocabulary.


Of course it matters a lot how we think about these content-conferring, vocabulary-deploying practices or abilities. The semantic pragmatist's claim that use confers meaning (so talk of practices or the exercise of abilities as deploying vocabularies) reverts to triviality if we are allowed to talk about "using the tilde to express negation," "the ability to mean red by the word 'red'," or "the capacity to refer to electrons by the word 'electron'." And that is to say that the interest of the PV-sufficiency of some set of practices or abilities for the deploying of a vocabulary is quite sensitive to the *vocabulary* in which we *specify* those practices-or-abilities. Semantic pragmatism would have the greatest explanatory power if it were possible to specify practices or abilities PV-sufficient to deploy some substantial vocabulary in non-semantic, non-intentional terms—that is, in a vocabulary that does *not* include such locutions as 'express', 'mean', and 'refer', and which did not itself appeal to the concepts mastery of which is being characterized. That may or may not be possible.¹³ What is clear is that talk of practices-or-abilities has a definite sense only insofar as it is relativized to the vocabulary in which those practices-or-abilities are specified. And that means that besides PV-sufficiency, we should admit a second basic meaning-use relation: "vocabulary-practice sufficiency," or just "*VP-sufficiency*," is the relation that holds between a vocabulary and a set of practices-or-abilities when that vocabulary is sufficient to *specify* those practices-or-abilities. Specifying PV-sufficient practices is *saying* what one must *do* in order to count as *saying* something, deploying a vocabulary. VP-sufficient vocabularies let one *say* what it is one must *do* to be engaging in those practices or exercising those abilities.

PV-sufficiency and VP-sufficiency are two basic *meaning-use* relations (MURs). In terms of those basic relations, we can define a more complex relation: the relation that holds between vocabulary V' and vocabulary V when V' is VP-sufficient to specify practices-or-

¹³ See John McDowell's "In Defense of Modesty" and "Another Plea for Modesty", in his *Meaning, Knowledge, and Reality* [Harvard University Press, 2001].

abilities P that are PV-sufficient to deploy vocabulary V. This VV-relation is the *composition* of the two basic MURs. When it obtains I will say that V' is a *pragmatic metavocabulary* for V. It allows one to *say* what one must *do* in order to count as *saying* the things expressed by vocabulary V. We can present this relation graphically in a *meaning-use diagram* (MUD):

**Meaning-Use Diagram #1:
Pragmatic
Metavocabulary**



The conventions of this diagram are:

- Vocabularies are shown as ovals, practices-or-abilities as (rounded) rectangles.
- Basic meaning-use relations are indicated by solid arrows, numbered, and labeled as to kind of relation.
- Resultant meaning-use relations are indicated by dotted arrows, numbered, and labeled as to kind and the basic MURs from which they result.

The idea is that a resultant MUR is the relation that obtains when all of the basic MURs listed on its label obtain.

Being a pragmatic metavocabulary is the simplest species of the genus I want to introduce here. It is a *pragmatically mediated semantic relation* between vocabularies. It is pragmatically mediated by the practices-or-abilities that are *specified* by one of the vocabularies

(which say what counts as doing that) and that *deploy* or are the *use* of the other vocabulary (what one says by doing that). The relation that is established thereby between the two vocabularies is of a distinctive sort, quite different, for instance, from the sort of *semantic* equivalence aimed at by the relations of definability, translatability, reducibility, and supervenience that have occupied the analytic tradition. My basic suggestion for extending the classical project of analysis so as to incorporate as essential positive elements the insights that animate the pragmatist critique of that project is that alongside these classical semantic relations between vocabularies we consider pragmatically mediated ones, of which the relation of being a pragmatic metavocabulary can serve as a paradigm. I will introduce an apparatus that recursively generates an infinite set of such pragmatically mediated semantic relations. In fact I will eventually argue that unless we take steps along these lines, we cannot understand the expressive roles played by some of the kinds of vocabulary with which the analytic tradition has been most centrally concerned: *logical*, *modal*, *normative*, and *intentional* vocabularies.

Under what circumstances would this simplest pragmatically mediated semantic relation—being a pragmatic metavocabulary—be philosophically interesting, when considered in connection with the sorts of vocabularies that have been of most interest to classical analysis? At least one sort of result that could be of considerable potential significance is if it turned out that in some cases pragmatic metavocabularies exist that differ significantly in their expressive power from the vocabularies for the deployment of which they specify sufficient practices-or-abilities. I will call that phenomenon “*pragmatic expressive bootstrapping*.” If one vocabulary is strictly weaker in expressive power than the other, I will call that *strict* expressive bootstrapping. We are familiar with this sort of phenomenon in ordinary semantics, where sometimes a semantic metalanguage differs substantially in expressive power from its object language—for instance, where we can produce an extensional metalanguage for intensional languages, as in the case of possible worlds semantics for modality. But in the case of semantic

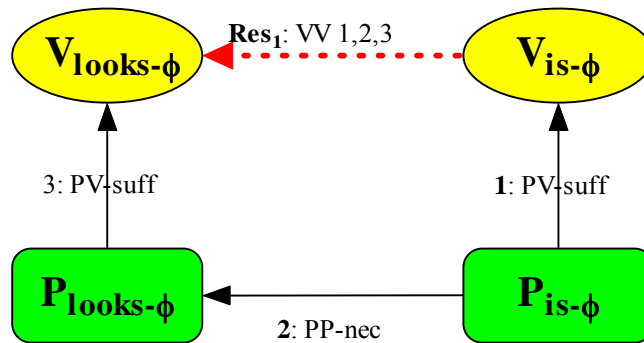
metalanguages, as Tarski forcibly reminds us, we typically need a metalanguage that is *more* expressively powerful than the object language to which it is addressed. One example of a claim of this shape in the case of pragmatically mediated semantic relations—though of course it is not expressed in terms of the machinery I have been introducing—is Huw Price’s pragmatic normative naturalism.¹⁴ Price argues, in effect, that although normative vocabulary is not *reducible to* naturalistic vocabulary, it is possible to *say* in wholly naturalistic vocabulary what one must *do* in order thereby to be *using* normative vocabulary. If such a claim about the existence of an expressively bootstrapping naturalistic pragmatic metavocabulary for normative vocabulary could be made out, it would evidently be an important chapter in the development of the naturalist core program of the classical project of philosophical analysis. It would be a paradigm of the sort of payoff we could expect from extending that analytic project by including pragmatically mediated semantic relations.

The meaning-use diagram of the pragmatically mediated semantic relation of being a pragmatic metavocabulary illustrates a distinctive kind of *analysis* of that relation. It exhibits that relation as the resultant, by composition, of the two basic meaning-use relations of PV-sufficiency and VP-sufficiency. A complex MUR is analyzed as the product of operations applied to basic MURs. This is *meaning-use analysis*. The same analytic apparatus applies also to more complex pragmatically mediated semantic relations. Consider one of the pragmatist criticisms that Sellars addresses to the empiricist core program of the classical analytic project. It turns on the assertion of the *pragmatic dependence* of one set of practices-or-abilities on another. Because one cannot *withhold* a commitment that one cannot *undertake*, given that he thinks part of what one is *doing* in saying how things merely appear is withholding a

¹⁴ “Naturalism without Representationalism” in Mario de Caro and David Macarthur (eds.) *Naturalism in Question* [Harvard University Press, 2004], pp. 71-90.

commitment to their actually being that way, Sellars concludes that one cannot have the ability to do that unless one also has the ability to say how things actually are. In effect, this Sellarsian pragmatist critique of the phenomenalist form of empiricism consists in the claim that the practices that are PV-sufficient for ‘is’- ϕ talk are PP-necessary for the practices that are PV-sufficient for ‘looks’- ϕ talk. That pragmatic dependence of practices-or-abilities then induces a resultant pragmatically mediated semantic relation between the vocabularies. The meaning-use diagram for this claim is:

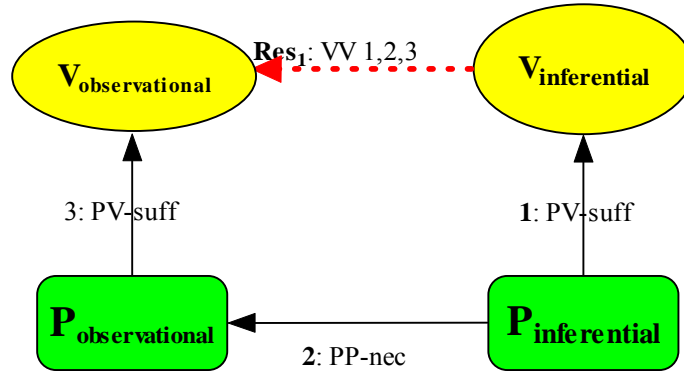
**Meaning-Use Diagram #2:
Pragmatically Mediated
Semantic Presupposition**



The resultant MUR here is a kind of complex, pragmatically mediated VV-necessity, or presupposition.


In fact, although Sellars’s argument for the crucial PP-necessity relation of pragmatic dependence of one set of practices-or-abilities on another is different, his argument against the observational version of empiricism—the claim that purely non-inferential, observational uses do not form an autonomous discursive practice, but presuppose inferential uses—has exactly the same form:

**Meaning-Use Diagram #3:
Pragmatically Mediated
Semantic Presupposition**




For these cases, we can say something further about the nature of the pragmatically mediated semantic relation that is analyzed as the resultant MUR in these diagrams. For instead of jumping directly to this VV resultant MUR, we could have put in the composition of the PP-necessity and second PV-sufficiency relation, yielding a kind of complex pragmatic presupposition:

**Meaning-Use Diagram #4:
Composition**



If this diagram were completed by an arrow from $V_{is-\phi}$ to $V_{looks-\phi}$ such that the *same* diagonal resultant arrow could represent *both* the composition of relations 2 and 3 *and* the composition of relation 1 and the newly supplied one, then category theorists would say that the diagram *commutes*. And the arrow that needs to be supplied to make the diagram commute they call the *retraction* of relation 1 through the composition Res_2 :

**Meaning-Use Diagram #5:
Composition and
Retraction**



After composition, then, the next most complex form of resultant MUR is retraction. Analyzing the structure of Sellars’s pragmatist arguments against empiricism requires recognizing the pragmatically mediated semantic relation that he claims holds between phenomenal and objective vocabulary as the retraction of a constellation of more basic meaning-use relations.

Section 4: Automata: Syntactic PV-sufficiency and VP-sufficiency

Now this is all extremely abstract. To make it more definite, we need to fill in (at least) the notions of vocabulary, practice-or-ability, PV-sufficiency, and VP-sufficiency¹⁵, which are the fundamental elements that articulate what I am calling the “meaning-use analysis” of resultant meaning-use relations—in particular, the pragmatically mediated semantic relations between vocabularies that I am claiming we must acknowledge in order to pursue the classical project of philosophical analysis. We can begin to do that by looking at a special case in which it is possible to be unusually clear and precise about the things and relations that play these metatheoretic roles. This is the case where ‘vocabulary’ takes a purely *syntactic* sense. Of

¹⁵ PP-sufficiency and PP-necessity will be discussed further in the next lecture.

course, the cases we eventually care about—and will be discussing in the remaining lectures—involve vocabularies understood in a sense that includes their *semantic* significance. But besides the advantages of clarity and simplicity, we will find that some important lessons carry over from the syntactic to the semantic case.

The restriction to vocabularies understood in a restricted syntactic sense leads to correspondingly restricted notions of what it is to *deploy* such a vocabulary, and what it is to *specify* practices-or-abilities sufficient to deploy one. Suppose we are given an *alphabet*, which is a finite set of primitive sign types—for instance, the letters of the English alphabet. The *universe* generated by that alphabet then consists of all the finite strings that can be formed by concatenating elements drawn from the alphabet. So long as no upper bound is put on the length of the finite strings, the universe generated by any finite alphabet, even one consisting of a single character, contains an infinite number of strings: a, aa, aaa, and so on. A *vocabulary* over such an alphabet—in the syntactic sense I am now after—is then any subset of the universe of strings that alphabet generates. If the generating alphabet is the English alphabet with punctuation including spaces, then the vocabulary might consist of all English sentences, or all possible English texts.¹⁶


What can we say about the abilities that count as *deploying* a vocabulary in this spare syntactic sense?¹⁷ The abilities in question are the capacity to *read* and *write* the vocabulary. In this purely syntactic sense, ‘reading’ it means being able practically to *distinguish*, within the universe generated by the vocabulary, strings that do, from those that do not, belong to the

¹⁶ Computational linguists, who worry about vocabularies in this sense, have developed metalanguages for specifying important classes of such vocabularies: the syntactic analogues of semantic metalanguages in the cases we will eventually address. So, for instance, for the alphabet {a,b}, ‘aⁿbⁿ’ characterizes the vocabulary that comprises all strings of some finite number of ‘a’s followed by the *same* number of ‘b’s. ‘a(ba)^{*}b’ characterizes the vocabulary that comprises all strings beginning with an ‘a’, ending with a ‘b’, and having *any* number of repetitions of the sub-string ‘ba’ in between.

¹⁷ Here we can safely just talk about *abilities*, without danger of restricting the generality of the analysis.

specified vocabulary. And ‘writing’ it means practically being able to *produce* all and only the strings in the alphabetic universe that do belong to the vocabulary.

We assume as *primitive* abilities the capacities to read and write, in this sense, the *alphabet* from whose universe the vocabulary is drawn—that is, the capacity to respond differentially to alphabetic tokens according to their type, and to produce tokens of specified alphabetic types. Then the abilities that are PV-sufficient to deploy some vocabularies can be specified in a particularly simple form. They are *finite-state automata* (FSAs). As an example, suppose we begin with the alphabet {a, h, o, !}. Then we can consider the *laughing Santa vocabulary*, which consists of strings such as ‘hahaha!’, ‘hohoho!’, ‘hahahoho!’ ‘hohoha!’, and so on.¹⁸ Here is a graphical representation of a *laughing Santa finite-state automaton*, which can read and write the laughing Santa vocabulary:



The numbered nodes represent the *states* of the automaton, and the alphabetically labeled arcs represent *state-transitions*. By convention, the starting state is represented by a square (State 1), and the final state by a circle with a thick border (State 4).

As a *reader* of the laughing Santa vocabulary, the task of this automaton is to process a finite string, and determine whether or not it is a licit string of the vocabulary. It processes the string one alphabetic character at a time, beginning in State 1. It recognizes the string if and only

¹⁸ In the syntactic metalanguage for specifying vocabularies that I mentioned in the note above, this is the vocabulary (ha/ho)*!

if (when and only when) it arrives at its final state, State 4. If the first character of the string is not an ‘h’, it remains stuck in State 1, and rejects the string. If the first character is an ‘h’, it moves to State 2, and processes the next character. If that character is not an ‘a’ or an ‘o’, it remains stuck in State 2, and rejects the string. If the character is an ‘a’ or an ‘o’, it moves to State 3. If the next character is an exclamation point, it moves to State 4, and recognizes the string ‘ha!’ or ‘ho!’—the shortest ones in the laughing Santa vocabulary. If instead the next character is an ‘h’, it goes back to State 2, and repeats itself in loops of ‘ha’s and ‘ho’s any number of times until an exclamation point is finally reached, or it is fed a discordant character.

As a *writer* of the laughing Santa vocabulary, the task of the automaton is to *produce* only licit strings of that vocabulary, by a process that can produce any and all such strings. It begins in its initial state, State 1, and emits an ‘h’ (its only available move), changing to State 2. In this state, it can produce either an ‘a’ or an ‘o’—it selects one at random¹⁹—and goes into State 3. In this state, it can either tack on an exclamation point, and move into its final state, State 4, finishing the process, or emit another ‘h’ and return to State 2 to repeat the process. In any case, whenever it reaches State 4 and halts, the string it has constructed will be a member of the laughing Santa vocabulary.

I hope this brief rehearsal makes it clear how the constellation of nodes and arrows that makes up this directed graph represents the abilities to read and write (recognize and produce arbitrary strings of) the laughing Santa vocabulary.²⁰ What it represents is abilities that are *PV*-

¹⁹ As a matter of fact, it can be shown that every vocabulary readable/writeable by a non-deterministic finite-state automaton—such as the laughing Santa automaton—is also readable/writeable by a deterministic one. [ref.]


²⁰ For practice, or to test one’s grip on the digraph specification of FSAs, consider what vocabulary over the same alphabet that produces the laughing Santa is recognized/produced by this automaton:

sufficient to deploy that vocabulary—in the attenuated sense appropriate to this purely syntactic case. And the digraph representation is itself a *vocabulary* that is *VP-sufficient to specify* those vocabulary-deploying abilities. That is, the digraph representation of this finite-state automaton is a *pragmatic metavocabulary* for the laughing Santa vocabulary. The relation between the digraph vocabulary and the laughing Santa vocabulary is, then, a *pragmatically mediated*—not now *semantic*, but *syntactic*—relation between vocabularies.

It may seem that I am stretching things by calling the digraph form of representation a ‘vocabulary’. It will be useful, as a way of introducing my final point, to consider a different form of pragmatic metavocabulary for the laughing Santa vocabulary. Besides the digraph representation of a finite-state automaton, we can also use a *state-table* representation. For the laughing Santa automaton this is:

	State 1	State 2	State 3
--	---------	---------	---------

The "I'll Have What She's Having" Automaton



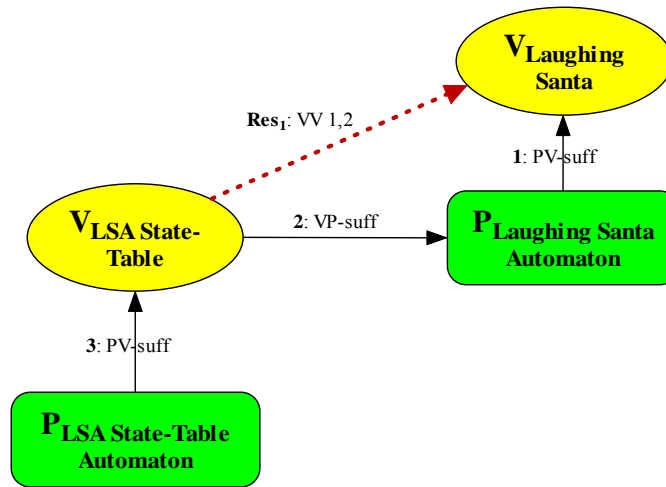
a	Halt	3	Halt
h	2	Halt	2
o	Halt	3	Halt
!	Halt	Halt	4

In read mode, the automaton starts in State 1. To see what it will do if fed a particular character, we look at the row labeled with that character. The LSA will Halt if the input string starts with anything other than an ‘h’, in which case it will change to State 2. In that state, the automaton specified by the table will halt unless the next character is an ‘a’ or an ‘o’, in which case it changes to State 3, and so on. (There is no column for State 4, since it is the final state, and accepts/produces no further characters.) Clearly there is a tabular representation corresponding to any digraph representation of a finite state automaton, and vice versa. Notice further that we need not use a two-dimensional table to convey this information. We could put the rows one after another, in the form:

aHalt3Halth2Halt2oHalt3Halt!HaltHalt4.

This is just a string, drawn from a universe generated by the alphabet of the LSA, together with ‘Halt’ and the designations of the states of that automaton. The strings that specify finite-state automata that deploy vocabularies defined over the same basic alphabet as the LSA then form a vocabulary in the technical syntactic sense we have been considering. And that means we can ask about the automata that can read and write *those* state-table encoding vocabularies. The meaning-use diagram for this situation is then:

Meaning-Use Diagram #6:
Specifying the Automaton
that Deploys the Laughing
Santa Vocabulary



Section 5: The Chomsky Hierarchy and a Syntactic Example of Pragmatic Expressive Bootstrapping

Restricting ourselves to a purely syntactic notion of a vocabulary yields a clear sense of ‘pragmatic metavocabulary’: both the digraph and the state-table vocabularies are VP-sufficient to specify practical abilities articulated as a finite-state automaton that is PV-sufficient to deploy—in the sense of recognizing and producing—the laughing Santa vocabulary, as well as many others. Perhaps surprisingly, it also offers a prime example of *strict pragmatic expressive bootstrapping*. For in this setting we can *prove* that one vocabulary that is expressively weaker than another can nonetheless serve as an adequate *pragmatic* metavocabulary for that stronger vocabulary. That is, even though one cannot *say* in the weaker vocabulary everything that can be *said* in the stronger one, one can still *say* in the weaker one everything that one needs to be able to *do* in order to deploy the stronger one.

Here the relevant notion of the relative expressive power of a vocabulary is also a purely syntactic one. Already in the 1950's, Chomsky offered mathematical characterizations of the different sets of strings of characters that could be generated by different classes of grammars (that is, in my terms, characterized by different kinds of syntactic metavocabularies) and computed by different kinds of automata. The kinds of vocabulary, grammar, and automata line up with one another, and can be arranged in a list of strictly increasing expressive power: what is now called the “Chomsky hierarchy.” It is summarized in the following table:

<u>Vocabulary</u>	<u>Grammar</u>	<u>Automaton</u>
Regular	$A \rightarrow aB$ $A \rightarrow a$	Finite State Automaton
Context-Free	$A \rightarrow \langle \text{anything} \rangle$	Push-Down Automaton
Context-Sensitive	$c_1 A c_2 \rightarrow c_1 \langle \text{anything} \rangle c_2$	Linear Bounded Automaton
Recursively Enumerable	No Restrictions on Rules	Turing Machine (= 2 Stack PDA)

The point I want to make fortunately does not require us to delve very deeply into the information summarized in this table. A few basic points will suffice. The first thing to realize is that not all vocabularies in the syntactic sense we have been pursuing can be read and written by finite-state automata. For instance, it can be shown that no finite-state automaton is PV-sufficient to deploy the vocabulary $a^n b^n$, defined over the alphabet $\{a,b\}$, which consists of all strings of any arbitrary number of ‘a’s followed by the same number of ‘b’s. The idea behind the proof is that in order to tell whether the right number of ‘b’s follow the ‘a’s (when reading) or to produce the right number of ‘b’s (when writing), the automaton must somehow keep track of how many ‘a’s have been processed (read or written). The only way an FSA can store information is by being in one state rather than another. So, it could be in one state—or in one of

a class of states—if one ‘a’ has been processed, another if two have, and so on. But by definition, a finite-state automaton only has a finite number of states, *and that number is fixed in advance* of receiving its input or producing its output. Whatever that number of states is, and whatever system it uses to code numbers into states (it need not be one-to-one—it could use a decimal coding, for instance), there will be some number of ‘a’s that is so large that the automaton runs out of states before it finishes counting. But the vocabulary in question consists of arbitrarily long strings of ‘a’s and ‘b’s. In fact, it is possible to say exactly which vocabularies finite-state automata (specifiable by digraphs and state-tables of the sort illustrated above) *are* capable of deploying. These are called the ‘regular’ vocabularies (or languages).

The next point is that slightly more complex automata *are* capable of deploying vocabularies, such as $a^n b^n$, that are not regular, and hence cannot be read or written by finite-state automata. As my brief discussion indicated, intuitively the problem FSAs have with languages like $a^n b^n$ is that they lack *memory*. If we give them a memory, we get a new class of machines: (non-deterministic²¹) *push-down automata* (PDAs). In addition to being able to respond differentially to and produce tokenings of the alphabetic types, and being able to change state, PDAs can *push* alphabetic values to the top of a *memory-stack*, and *pull* such values from the top of that stack. PDAs can do everything that finite-state automata can do, but they can also read and write many vocabularies, such as $a^n b^n$, that are *not* regular, and so cannot be read and written by FSAs. The vocabularies they can deploy are called “context-free”. All regular vocabularies are context-free, but not *vice versa*. This proper containment of classes of vocabularies provides a clear sense, suitable to this purely syntactic setting, in which one vocabulary can be thought of as “expressively more powerful” than another: the different kinds of grammar can specify, and the

²¹ By contrast to FSA’s, there need not in general be for every vocabulary computable by a non-deterministic PDA, some deterministic PDA that reads and writes the same vocabulary.

different kinds of automata can compute, ever larger classes of vocabularies. Context-free vocabularies that are not regular require more powerful grammars to specify them, as well as more powerful automata to deploy them.²² FSAs are special kinds of PDAs, and all the automata are special kinds of Turing Machines. Recursively enumerable vocabularies are not in general syntactically reducible to context-sensitive, context-free, or regular ones. And the less capable automata cannot read and write the all the vocabularies that can be read and written by Turing Machines.

Nonetheless, if we look at *pragmatically mediated* relations between these syntactically characterized vocabularies, we find that they make possible a kind of *strict expressive bootstrapping* that permits us in a certain sense to evade the strict restrictions on expressive power enforced for purely syntactic relations between vocabularies. The hierarchy dictates that only the abilities codified in Turing Machines—two-stack push-down automata—are *PV-sufficient* to *deploy* recursively enumerable vocabularies in general. But now we can ask: what class of languages is *VP-sufficient* to *specify* Turing Machines, and hence to serve as sufficient *pragmatic* metavocabularies for recursively enumerable vocabularies in general? The surprising fact is that **the abilities codified in Turing Machines—the abilities to recognize and produce recursively enumerable vocabularies—can quite generally be specified in context-free vocabularies**. It is demonstrable that context-free languages are strictly weaker in syntactic expressive resources than recursively enumerable languages. The push-down automata that can read and write only context-free languages cannot read and write recursively enumerable

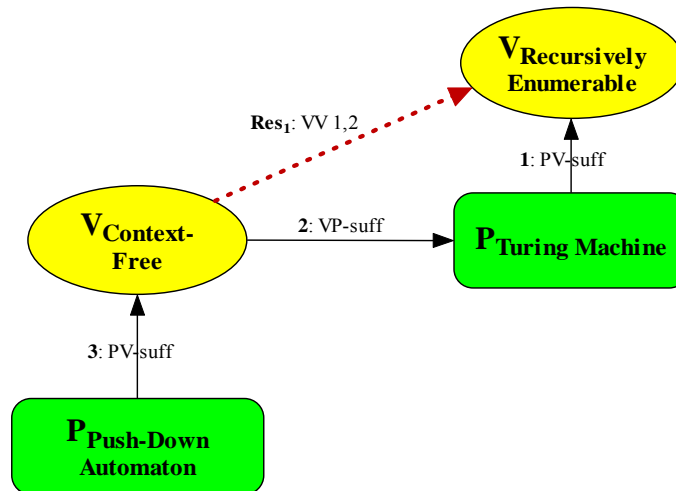
²² Regular languages are those that can be specified entirely by production rules of the form: $A \rightarrow aB$ and $A \rightarrow a$. Context-free languages relax these restrictions, allowing production rules of the form: $A \rightarrow \langle \text{anything} \rangle$.

languages in general. But it is possible to *say* in a context-free language what one needs to *do* in order to deploy recursively enumerable languages in general.

The proof of this claim is tedious, but not difficult, and the claim itself is not at all controversial (my introductory textbook leaves the proof as an exercise to the reader²³)—though computational linguists make nothing of it, having theoretical concerns very different from those that lead me to underline this fact. General-purpose computer languages such as Pascal and C++ can specify the algorithms a Turing Machine, or any other universal computer, uses to compute any recursively enumerable function, hence to recognize or produce any recursively enumerable vocabulary. And they are invariably context-free languages—in no small part just because the simplicity of this type of grammar makes it easy to write parsers for them. Yet they suffice to specify the state-table, contents of the tape (or of the dual stacks), and primitive operations of any and every Turing Machine. Here is the MUD characterizing this pragmatically mediated relation between syntactically characterized vocabularies:

Meaning-Use Diagram #7:

**Syntactic Pragmatic
Expressive Bootstrapping**



I called the fact that context-free vocabularies can be adequate pragmatic metavocabularies for recursively enumerable vocabularies in general ‘surprising’, because of the provable syntactic irreducibility of the one class of

²³ Thomas A. Sudkamp *Languages and Machines* 2nd edition [Addison Wesley Longman 1998], Chapter 10.

vocabularies to the other. But if we step back from the context provided by the Chomsky hierarchy, we can see why the possibility of such pragmatic expressive bootstrapping should not, in the end, be surprising. For all the result really means is that context-free vocabularies let one *say* what it is one must *do* in order to say things they cannot themselves say, because the ability to deploy those context-free vocabularies does not include the abilities those vocabularies let one specify. Thus, for instance, there is no reason that an FSA could not read and write a vocabulary that included commands such as “Push an ‘a’ onto the stack,”—and thus specify the program of a PDA—even though it itself has no stack, and could not *do* what the vocabulary it is deploying specifies. More generally, a coach might be able to tell an athlete exactly what to do, and even how to do it, even though the coach cannot himself do what he is telling the athlete to do, does not have the abilities he is specifying. We ought not to boggle at the possibility of an expressively weaker pragmatic metavocabulary having the capacity to say what one must do in order to deploy an expressively stronger one. We should just look to see where this seems in fact to be possible for vocabularies we care about, and what we can learn from such relations when they do obtain.

Section 6: Looking Ahead

Let us recall what motivated this rehearsal of some elements of automaton theory and introductory computational linguistics. I suggested that a way to extend the classical project of semantic analysis so as to take account of the insights of its pragmatist critics is to look analytically at relations between meaning and use. More specifically, I suggested focusing to begin with on two in some sense complementary relations: the one that holds when some set of practices-or-abilities is PV-sufficient to *deploy* a given vocabulary, and the one that holds when some vocabulary is VP-sufficient to *specify* a given set of practices-or-abilities. The composition of these is the simplest *pragmatically mediated semantic relation* between vocabularies: the relation that holds when one vocabulary is a sufficient pragmatic metavocabulary for another. It is a paradigm of the infinite, recursively generable class of complex, pragmatically mediated

semantic relations that I propose to lay alongside the other semantic relations between vocabularies that have been investigated by analytic philosophers (for instance those who address the core programs of empiricism and naturalism): relations such as analyzability, definition, translation, reduction, truth-making, and supervenience. I suggested further that pragmatic metavocabularies might be of particular interest in case they exhibited what I called “expressive bootstrapping”—cases, that is, in which the expressive power of the pragmatic metavocabulary differs markedly from that of the target vocabulary, most strikingly, when the metavocabulary is substantially expressively weaker—a phenomenon Tarski has led us *not* to expect for semantic metavocabularies.

We have now seen that all of these notions can be illustrated with particular clarity for the special case of purely syntactically characterized vocabularies. The abilities that are PV-sufficient to deploy those vocabularies, in the sense of the capacity to recognize and produce them, can be thought of as various sorts of automata. There are several well-established, different-but-equivalent vocabularies that are known to be VP-sufficient to specify those automata. In this special syntactic case we can accordingly investigate the properties of pragmatic metavocabularies, and when we do, we find a striking instance of *strict expressive bootstrapping* in a pragmatically mediated syntactic relation between vocabularies.

Of course, the cases we really care about involve *semantically* significant vocabularies. Are there any interesting instances of these phenomena in such cases? I have indicated briefly how some of Sellars’s pragmatist criticisms of various ways of pursuing the empiricist program can be understood to turn on pragmatically mediated semantic relations. And I mentioned Huw Price’s idea that although normative vocabulary is not semantically reducible to naturalistic

vocabulary, naturalistic vocabulary may suffice to specify what one must *do*—the practices-or-abilities one must engage in or exercise—in order to deploy normative vocabulary. Here is another example that I want to point to, though I cannot develop the claim here. For roughly the first three-quarters of the twentieth century, philosophers who thought about indexical vocabulary took for granted some version of the doctrine that a tokening *n* of an expression of the type ‘now’ was *synonymous* with, definable or semantically analyzable as, ‘the time of utterance of *n*,’ and similarly for ‘here’ and ‘the place of utterance of *h*,’ and so on. During the 1970’s philosophers such as John Perry, David Lewis, and G. E. M. Anscombe, by focusing on the use of indexicals in modal and epistemic contexts, showed decisively that this cannot be right: what is expressed by indexical vocabulary cannot be expressed equivalently by non-indexical vocabulary. This fact seems so obvious to us now that we may wonder what philosophers such as Russell, Carnap, and Reichenbach could have been thinking. I want to suggest that the genuine phenomenon in the vicinity is a *pragmatically mediated* semantic relation between these vocabularies. Specifically, in spite of the semantic irreducibility of indexical to nonindexical vocabulary, it is possible to *say*, entirely in non-indexical terms, what one must *do* in order to be deploying indexical vocabulary correctly: to be saying essentially indexical things. For we can formulate practical rules such as:

- If, at time *t* and place $\langle x,y,z \rangle$, speaker *s* wants to assert that some property *P* holds of $\langle x,y,z,t,s \rangle$, it is correct to say “*P* holds of *me here and now*.”

In fact, though this is much less obvious, one can also say in wholly non-indexical terms what one must do in order to recruit other vocabulary to play the distinctive role in the aetiology of behavior that is characteristic of “essential” indexicals, in virtue of which they are not semantically equivalent to any non-indexical expressions. That is to say that non-indexical vocabulary can serve as an adequate *pragmatic* metavocabulary for indexical vocabulary. The fact that one nonetheless cannot *say* in non-indexical terms everything that one

can *say* with indexical vocabulary just shows that these vocabularies have different expressive powers, so that the pragmatically mediated semantic relation between them is a case of pragmatic expressive bootstrapping.

In the lectures to come, I will be doing three things:

- further developing the conceptual apparatus of meaning-use analysis, by introducing both new basic meaning-use relations and new combinations of them;
- applying that apparatus to vocabularies of ongoing philosophical interest (logical, modal, normative, intentional); and
- seeing what new pragmatically mediated semantic relations become visible in that way.

Each subsequent lecture will report some unexpected, suggestive results, which fit together cumulatively to constitute a distinctive, novel picture of what we would previously have thought was familiar terrain.

Besides pragmatically mediated semantic relations between vocabularies, there is another sort of pragmatic analysis, which relates one constellation of practices-or-abilities to another. It corresponds to another basic meaning-use relation: the kind of PP-sufficiency that holds when having acquired one set of abilities means one can already do everything one needs to do, *in principle*, to be able to do something else. One concrete way of filling in a definite sense of “in principle” is by *algorithmic elaboration*, where exercising the target ability just is exercising the right basic abilities in the right order and under the right circumstances. As an example, the ability to do long division just consists in exercising the abilities to do multiplication and subtraction and write down the results of those calculations, according to a particular conditional

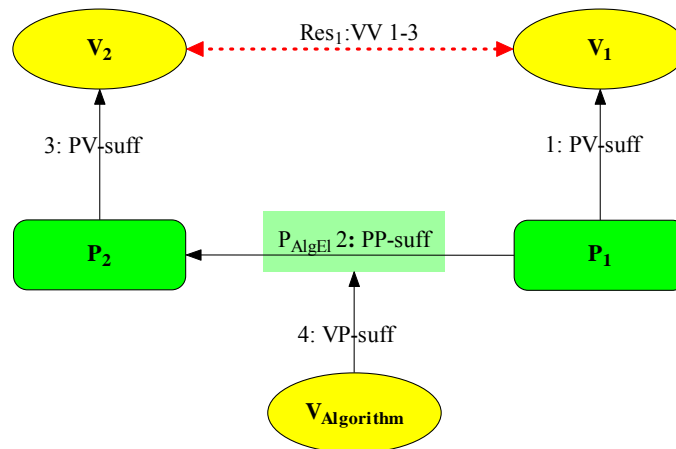
branched-schedule algorithm. The practical abilities that *implement* such an algorithmic PP-sufficiency relation are just those exercised by a finite-state *automaton*. Indeed, automata are defined by a definite set of *meta-abilities*: abilities to *elaborate* a set of primitive abilities into a set of more complex ones, which can accordingly be pragmatically *analyzed* in terms of or *decomposed* into the other.

To get a usefully general concept of the PP-sufficiency of a set of basic abilities for a set of more complex ones, we need to move beyond the purely *syntactic* automata I have described so far. One way to do that is to replace their specialized capacities to read and write symbols—in the minimal sense of classifying tokens as to types and producing tokens of specified types—by more general recognitional and productive capacities. These are abilities to respond differentially to various non-symbolic stimuli (for instance, the visible presence of red things), corresponding to reading, and to respond by producing performances of various non-symbolic kinds (for instance, walking north for a mile), corresponding to writing. What practically implements the algorithmic elaboration of such a set of basic differential responsive abilities is a finite state *transducing* automaton (and its more sophisticated push-down brethren).

In my third lecture, I will argue that the notion of the algorithmic decomposability of some practices-or-abilities into others that results suggests in turn a pragmatic generalization of the classical program of artificial intelligence functionalism—which, though a latecomer in the twentieth century, deserves, I think, to count as a *third* core program of classical semantic analysis. AI functionalism traditionally held itself hostage to a commitment to the purely *symbolic* character of intelligence in the sense of sapience. But broadening our concern from

automata as purely syntactic engines to the realm of *transducing* automata puts us in a position to see AI functionalism as properly concerned with the algorithmic decomposability of discursive (that is, vocabulary-deploying) practices-and-abilities. What I will call the ‘pragmatic’ thesis of artificial intelligence is the claim that the ability to engage in some autonomous discursive practice—a language game one could play though one played no other—can be *algorithmically decomposed* into non-discursive abilities. (By non-discursive abilities, I mean abilities each of which can in principle be exhibited by something that does *not* engage in any autonomous discursive practice.)²⁴ The arguments for and against this pragmatic version of AI-functionalism look quite different from those arrayed on the opposing sides of the debate about the prospects of symbolic AI.

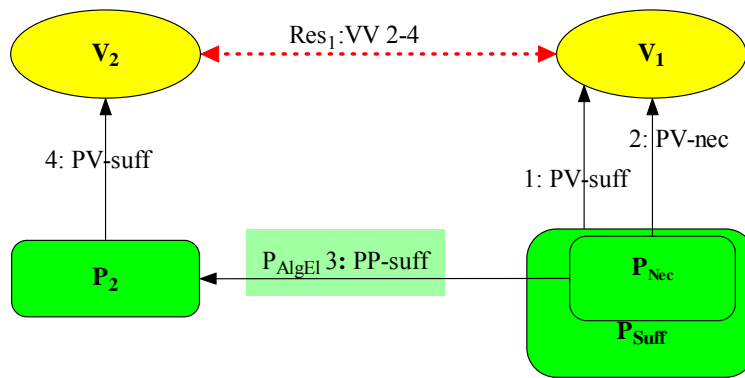
The notion of PP-sufficiency brings into view a slightly more complicated pragmatically mediated semantic relation between vocabularies: that which obtains when practices PV-sufficient for V_1 are PP-sufficient (in the sense that they can be algorithmically elaborated into) practices PV-sufficient for V_2 . A meaning-use diagram for this is [NOT IN HANDOUT]:



Another basic meaning-use relation of the kind we have been considering is PV-*necessity*, the converse of PV-sufficiency. It obtains when one cannot deploy a certain vocabulary without


²⁴ Without that restriction on the primitive abilities out of which discursive ones are to be algorithmically elaborated, the claim would be trivial, since the null algorithmic decomposition is also a decomposition.

engaging in the specified practice, or exercising the specified ability. For example, I have argued elsewhere that nothing could count as engaging in an autonomous discursive practice (hence using a vocabulary one could use though one used no other) that did not include *asserting* and *inferring*. Considering that basic MUR permits the formulation of a complex resultant MUR that is a variant on the prior one: a relation that obtains where practices PV-necessary for V₁ are PP-sufficient for practices-or-abilities PV-sufficient for V₂. The corresponding MUD is [NOT IN HANDOUT]:



It can happen, I will argue, that such a V₂ is also VP-sufficient to *specify* the practices-or-abilities that are PV-sufficient to deploy V₁. A MUD for this is:

LX: V₂ is Elaborated from and Expliative of Practices PV-Necessary for V₁



In my next lecture, I will introduce a version of this complex resultant pragmatically mediated semantic relation (what I call for short being “universally LX”), and argue that it constitutes the genus of which *logical* vocabulary is a species. More specifically, I will argue that logical vocabulary both can be algorithmically elaborated from and is explicative of practices that are PV-necessary for the autonomous deployment of any vocabulary at all. And I will argue that the most illuminating way to explain and justify the distinctive privileged role accorded to logical vocabulary by the classical project of philosophical analysis—what I have here called “semantic logicism”—is by appeal to this whole constellation of basic meaning-use relations, and the complex pragmatically mediated semantic relation that results from it.

My last three lectures will address modal vocabulary, normative vocabulary, and the pragmatically mediated semantic relations they stand in to ordinary objective, empirical, and naturalistic vocabularies, and to each other. The modal revolution in the last third of the 20th century breathed new life into semantic logicism, providing powerful new expressive tools, which have been of great use to those pursuing naturalistic programs, for instance. But this successor version raises the same question of vindication that I consider for semantic logicism in my second lecture: what justifies according modal concepts this special, privileged role in our semantic analytic enterprise? This question is particularly urgent since the empiricist program had always been—traditionally with Hume, and in the 20th century logical form, with Quine, particularly and specifically hostile to and critical of this vocabulary.

I will begin my consideration of modality, in my fourth lecture, with a consideration of this question, and with a vindication of the role of modal vocabulary that parallels the one I will already have offered for ordinary logical vocabulary: modal vocabulary, too, can be elaborated

from and is explicative of, features integral to every autonomous discursive practice—features intimately related to, but distinct from, those made explicit by ordinary logical vocabulary. I will then enter into an extended treatment of the relation between *alethic* and *deontic* (modal and normative) vocabularies (where the norms in question are those that govern the use of vocabulary, that is, *conceptual* norms rather than *moral* ones). When we look at those vocabularies through the lens of meaning-use analysis, a sequence of startling relations between them emerges.

For a start, I argue that deontic normative vocabulary is *also* universally LX (that it is VP-sufficient to specify practices-or-abilities that are both PV-necessary for deploying any autonomous vocabulary, and PP-sufficient for practices-or-abilities PV-sufficient for deploying the deontic normative vocabulary that explicates them). Although in this regard it belongs in a box with alethic modal vocabulary, the features of autonomous vocabulary use that it explicates are quite different from those explicated by modal vocabulary. I then argue that what lies behind Sellars's dark and pregnant claim that "The language of modality is a transposed language of norms" is the fact that deontic normative vocabulary can serve as a *pragmatic metavocabulary* for alethic modal vocabulary. In my fifth lecture, I will show how exploiting that relation makes possible a new kind of directly modal formal semantics: incompatibility semantics. It in turn gives us a new semantic perspective both on traditional logical vocabulary, and on modal vocabulary. The final lecture will then weave all these strands into a meaning-use analysis of intentionality itself (what is expressed by intentional vocabulary) as a pragmatically mediated semantic relation essentially involving both what is expressed by modal and what is expressed by normative vocabulary.

The substantive cumulative result of this sequence of revelations about modal and normative vocabulary is to put new flesh on the bones of ideas that originate with Kant, and are developed by his tradition up through the traditional American pragmatists, and are reinterpreted by Sellars in the middle years of the 20th century. And the methodological result of this development and application of meaning-use analysis is a new synthesis of pragmatism and analytic philosophy—one that shows how concerns and considerations at the heart of the pragmatist critique of semantic analysis can be seen to have been implicitly at work within the analytic tradition all along.

The title of this lecture series, “Between Saying and Doing,” evidently refers to my aspiration to present a new way of thinking about the relations between meaning and use that arises when we think systematically about *saying* what we are *doing* when we are *saying* something. But the phrase itself is taken from an Italian proverb: “Between saying and doing, many a pair of shoes is worn out.” Following the argumentative and constructive path I am proposing for exploring the intricate and revealing ways in which semantics and pragmatics interdigitate will require wearing out a few.

END

[4-9 a: 8410 words (= 28.0 pages) in large print]