Towards an Analytic Pragmatism:

Meaning-Use Analysis

Within the Anglophone tradition, pragmatism has often appeared as a current of thought that stands apart from, and indeed runs in opposition to, the mainstream of analytic philosophy. This is true whether one uses 'pragmatist' in a narrow sense tailored to the triumvirate of Peirce, James, and Dewey (here one might think of Russell's dismissive responses to the latter two), or in a more capacious sense that includes the early Heidegger, the later Wittgenstein, and, more recently, neo-pragmatists such as Rorty and Putnam. There are good reasons on both sides for adopting somewhat adversarial stances, but I think that when we examine them more closely it becomes possible to see the outlines of a common project, in the service of which the two camps might find themselves joining forces. In my 2006 John Locke lectures, entitled *Between Saying and Doing: Towards an Analytic Pragmatism[1],* I explore in more detail one way of pursuing such a project. In this essay I want to offer a sketch of the basic understanding of the principal aims of the two movements that motivates that more extended discussion, and to indicate in general terms the sort of *pragmatic* semantic analysis (not, I will be insisting, an oxymoron) that might emerge from unifying their only apparently disparate concerns. The intended spirit is irenic, synthetic, and constructive.

I. 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'. 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. I will refer to this aspect of the analytic project as its commitment to '*semantic logicism*'.[2]

If we ask *which* were the vocabulary-kinds whose semantic relations it was thought to be important to investigate during this period, 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.

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.

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, or just to objective descriptive vocabulary, even when not regimented by incorporation into explicit scientific theories. Typical targets include normative, semantic, and intentional vocabularies.

II. 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 it—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 <u>meaning</u> in favor of that of <u>use</u>: 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, and so to count as *saying* or *thinking* various things. 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.[3]

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 in Chapter Two I called a "methodological" pragmatist. 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.

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.

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 are in the business of stating 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," he says, "has no downtown.")

I take it that Wittgenstein also takes the home language game of the concept of <u>meaning</u> 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 theoretical systematization in the case of discursive practices—about the possibility of systematically *deriving* aspects of correct 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—from the meanings expressed by some *other* vocabulary—from which we can derive proprieties of *its* use. One idea, I think, is that this enterprise makes sense only if we think of the uses as species of a genus—of them all being the same general *kind* of use, say describing, 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 nails and hammer, screws and screwdriver, glue 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 (even if we then find ourselves metaphysically puzzled about the nature of the fact-kinds to which we have thereby committed ourselves). But if we think of the uses as *very* different, if we think also about the carpenter's level, pencil, 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', 'description', 'assertion', 'observation' and so on—do not typically admit of specification in terms of underlying principles specifiable in other vocabularies, whether by genus and differentia(e) 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 one 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 ˢessenceˢ proper-name usage could be taken to have had.[4] 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 teeth, and extends that practice to talk about having pain in one's teeth. 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 ˢmeaningˢ 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.[5]

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 dynamic character—and that is why the collection of uses that is the current cumulative and collective result of such developments-by-practical-projection is a

motley.[6] If that is right, then any codification or theoretical systematization of the uses of those vocabulary-kinds by associating with them meanings that 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.[7] On this conception, the classical project of analysis is disease that rests on a fundamental, if perennial, 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*.

III. 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 and possible worlds 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 contemporary 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, thereby marking out a substantial sociological faultline in the discipline. Those who are moved by the pragmatist picture generally accept the particularist, quietist 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: One hopes that it is not true, and that if it is 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 surely 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 semantic project, from exclusive concern with relations among meanings to encompass also relations between meaning and use. In its most ambitious form, as in the present project, 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 theoretical 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[8] is sufficient for someone to count as *deploying* a specified 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 <u>negation</u>," "the ability to mean <u>red</u> by the word 'red'," or "the capacity to refer to electrons by the word 'electron'," (or, I think, even *intentions* so to refer). 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. 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 consider 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.[9] VP-sufficient vocabularies that *specify* PV-sufficient practices let one *say*

what it is one must *do* to count as engaging in those practices or exercising those abilities, and so to deploy a vocabulary to *say* something.

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-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
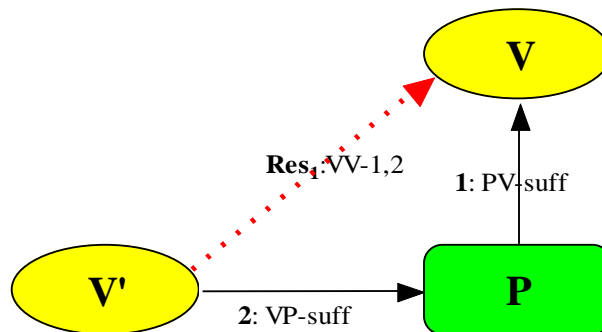
The conventions of this diagram are:

- 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 semantic relation that is established thereby between the two vocabularies is of a distinctive sort, quite different from, for instance, definability, translatability, reducibility, and supervenience. **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 the classical semantic relations between vocabularies that project has traditionally appealed to, we consider also *pragmatically mediated* ones—of which the relation of being a pragmatic metavocabulary is a paradigm.**

Under what circumstances would this simplest pragmatically mediated semantic relation be philosophically interesting, when considered in connection with the sorts of vocabularies that have traditionally been of most interest to classical analysis? At least one sort of result that could be of considerable potential significance, I think, 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. 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.[10] He argues, in effect, that although normative vocabulary is not *reducible to* naturalistic vocabulary, it might still be possible to *say* in wholly naturalistic vocabulary what one must *do* in order 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. (Later on I'll discuss briefly a claim of this shape concerning *indexical* vocabulary.)

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 vocabulary-deploying

practices-or-abilities on another.

Because he thinks part of what one is *doing* in saying how things merely appear is withholding a commitment to their actually being that way, and because one cannot be understood as *withholding* a commitment that one cannot *undertake*, Sellars concludes that one cannot have the ability to say or think how things *seem* or *appear* unless one also has the ability to make claims about 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.[11] 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**
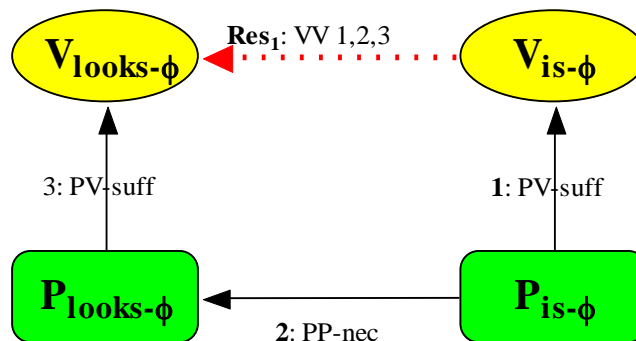
The resultant MUR here is a kind of complex, pragmatically mediated, VV-necessity, or semantic presupposition.

In fact, although Sellars's argument for the crucial PP-necessity relation of pragmatic dependence of one set of vocabulary-deploying 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:
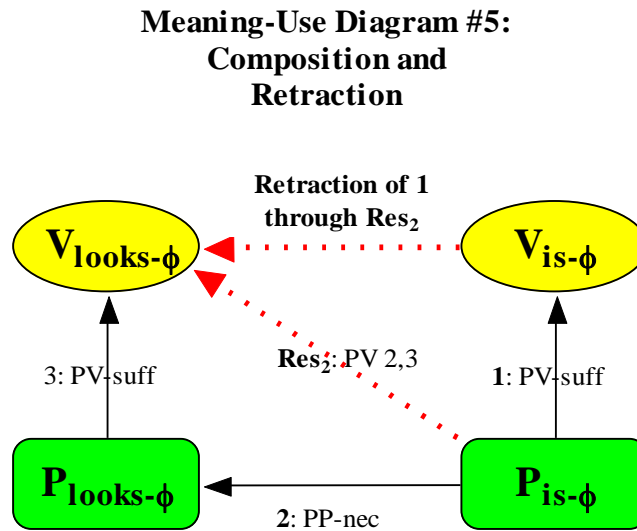**Pragmatically Mediated**
Semantic Presupposition
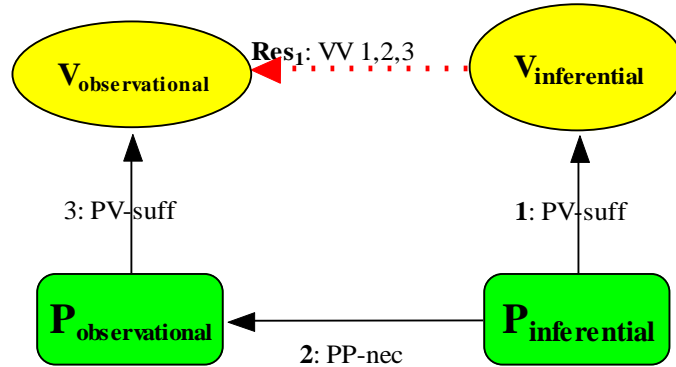
For these cases, we can say something further about the nature of the pragmatically mediated semantic relation that is analyzed as the resultant MUR in these diagrams. For instead of jumping directly to this VV resultant MUR, we could have put in the composition of the PP-necessity and second PV-sufficiency relation, yielding a kind of complex pragmatic presupposition:

Meaning-Use Diagram #4:
Composition
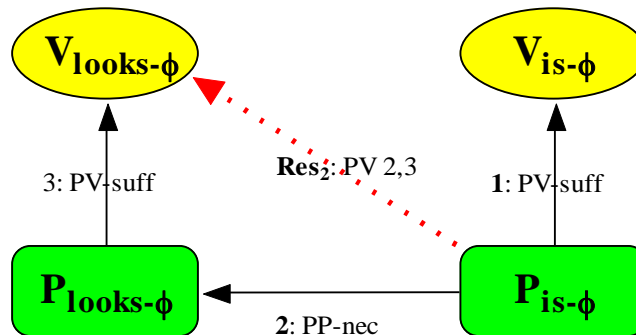
If this diagram were completed by an arrow from $V_{is-\varphi}$ to $V_{looks-\varphi}$ such that the *same* diagonal resultant arrow could represent *both* the composition of relations 2 and 3 *and* the composition of relation 1 and the newly supplied one, then category theorists would say that the 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 he claims holds between phenomenal and objective vocabulary as the retraction of a constellation of more basic meaning-use relations.

IV. 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-sufficiency, 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 in the light of what is right about the pragmatist critique of it. 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 course, the cases we eventually care about 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 spare 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. 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, then the vocabulary might consist of all English sentences, all possible English texts, or all and only the sentences of *Making It Explicit*.[12]

What can we say about the abilities that count as *deploying* a vocabulary in this spare syntactic sense?[13] 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 alphabet, strings that do, from those that do not, belong to the 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 antecedently 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.[14]   Here is a graphical representation of a *laughing Santa finite-state automaton*, which can read and write the laughing Santa vocabulary:

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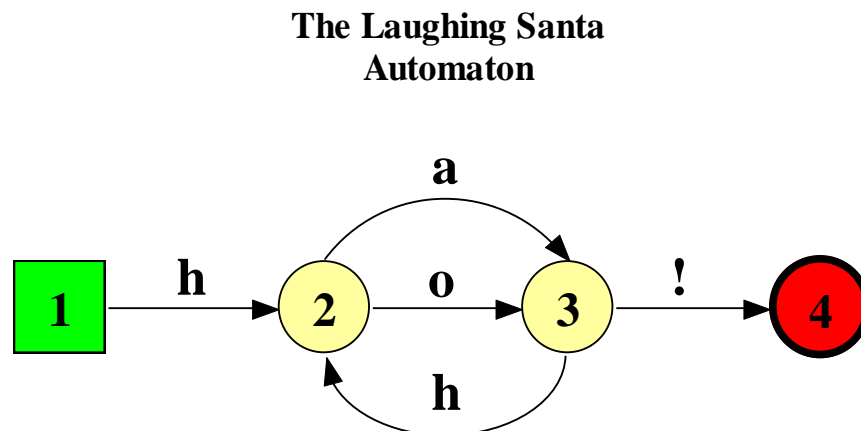
The numbered nodes represent the *states* of the automaton, and the alphabetically labeled arcs represent *state-transitions*. By convention, the starting state is represented by a square (State 1), and the final state by a circle with a thick border (State 4).

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 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 random[15]—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.[16] What it represents is abilities that are *PV-sufficient* to *deploy* that vocabulary—that is, read and write it, 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 in the vicinity, 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**
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 an FSA, 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**
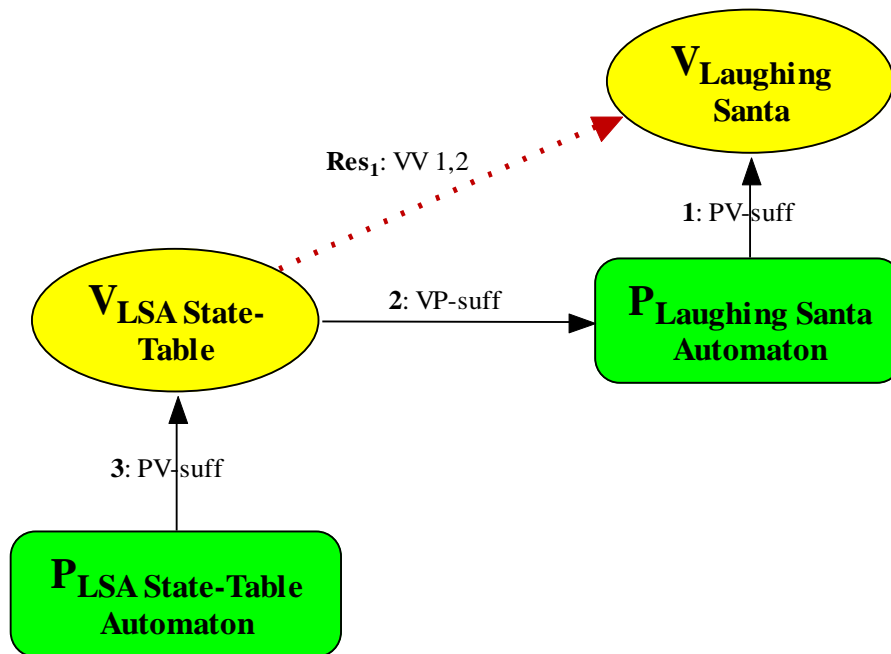
I. The Chomsky Hierarchy:
A Syntactic Example of Pragmatic Expressive Bootstrapping

Restricting ourselves to a purely syntactic notion of a <u>vocabulary</u> 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. (Of course, it does that only against the background of a set of abilities PV-sufficient to deploy *those* vocabularies.) 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 can*not 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 <u>relative expressive power</u> of vocabularies 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 lined up with one another, and could be arranged in a strict expressive hierarchy: the Chomsky hierarchy. It is summarized in the following table:

Vocabulary	**Grammar**	**Automaton**
Regular	A→aB	Finite State

	A→a	Automaton
Context-Free	A→<anything>	Push-Down Automaton
Context-Sensitive	c_1Ac_2→c1<anything>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^nb^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 our 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-deterministic[17]) *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 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. 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 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 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 arbitrary recursively enumerable vocabularies—can quite generally be specified in *context-free* vocabularies**. It is demonstrable that context-free vocabularies are strictly weaker in syntactic expressive resources than recursively enumerable vocabularies. The push-down automata that can read and write only context-free vocabularies cannot read and write recursively enumerable vocabularies in general. But it is possible to *say* in a context-free vocabulary what one needs to be able to *do* in order to deploy recursively enumerable vocabularies in general.

The proof of this claim is tedious, but not difficult, and the claim itself is not at all controversial—though computational linguists make nothing of it, having theoretical concerns very different from those that lead me to underline this fact. (My introductory textbook leaves the proof as an exercise to the reader.[18]) 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[19]—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**
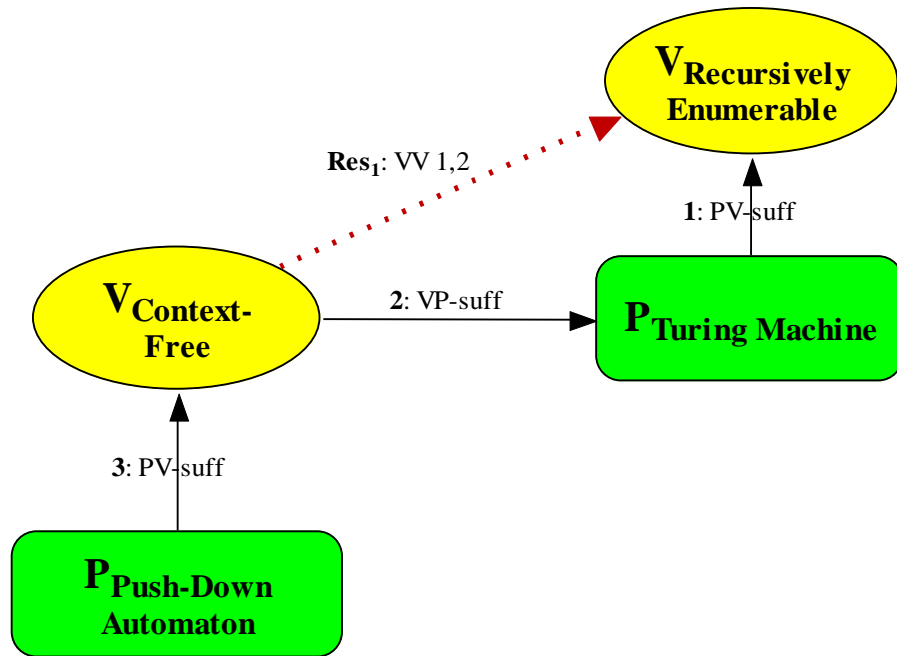
I called the fact that context-free vocabularies can be adequate pragmatic metavocabularies for

recursively enumerable vocabularies in general 'surprising', because of the provable syntactic

irreducibility of the one class of 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. 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.

II.   Semantic Examples of Pragmatic Expressive Bootstrapping and Further Basic and Resultant Meaning-Use Relations

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, which in general must be expressively *stronger* than the vocabularies they address.

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 might 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 might be led to wonder what philosophers such as Russell, Carnap, and Reichenbach could have been thinking for all those years. 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 and irreducibly indexical things. For we can formulate practical rules such as:

1. If, at time t and place $<x,y,z>$, speaker s wants to assert that some property P holds of $<x,y,z,t,s>$, it is correct to say "P holds of *me*, *here* and *now*."

And

2. If a speaker s at time t and place $<x,y,z>$ asserts "P holds of *me*, *here* and *now*," the speaker is committed to the property P holding of $<x,y,z,t,s>$.

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 strict pragmatic expressive bootstrapping.

Here is another example. 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

262

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. (Of course, this is just *one* species of the genus of practical projection that Wittgenstein brings to our attention.) As an example, the ability to do long division just consists in exercising the abilities to do multiplication and subtraction according to a particular conditional branched-schedule algorithm. The practical abilities that *implement* such an algorithmic PP-sufficiency relation are just those exercised by finite-state *automata* (in general, Turing Machines). Indeed, automata should be thought of as consisting in 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.[20]

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 in general non-symbolic stimuli (for instance, the visible presence of red things), corresponding to reading, and to respond by producing performances of various in general 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).

A Finite-State Transducing Automaton
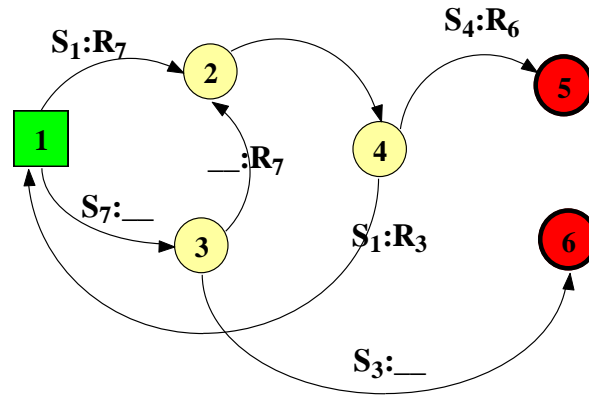
This is a diagram of an FSTA that has an initial set of stimuli to which it can respond differentially, and an initial set of responses it can differentially produce. And the diagram indicates that in its initial state, if presented with a stimulus of kind 1, it will produce a response of kind 7 and shift to state 2, and if presented instead with a stimulus of kind 7 it will produce no response, but will shift to state 3. It is important to note that although the recognitive and performative abilities that such an automaton algorithmically elaborates are to be considered as 'primitive' or 'basic' with respect to such elaboration, this does not mean that they are so in any absolute sense. The stimulus-response formulation by itself does not keep us from considering as 'primitive' capacities the abilities to keep ourselves at a suitable distance from a conversational partner, distinguish cubist paintings done by Braque from those done by Picasso, drive from New York to San Francisco, or build a house.

The notion of the algorithmic decomposability of some practices-or-abilities into others 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, alongside empiricism and naturalism. AI

264

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, we are now in a position to see automaton 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—where 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. (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.) The capacity to talk-and-think as I am addressing it is the capacity to deploy an autonomous vocabulary. But unlike classical symbolic AI, the pragmatic thesis of artificial intelligence does not presume that the practical capacities from which some transducing automaton can algorithmically elaborate the ability to engage in an autonomous discursive practice must themselves consist exclusively of symbol-manipulating abilities, never mind ultimately syntactic ones.[21]

The algorithmic practical elaboration model of AI gives a relatively precise shape to the pragmatist program of explaining knowing-*that* in terms of knowing-*how*: specifying in a non-intentional, non-semantic vocabulary what it is one must *do* in order to count as deploying some vocabulary to *say* something, hence as making intentional and semantic vocabulary applicable to the performances one produces. In particular, it offers a construal of the basic claim of AI-

functionalism as a pragmatic *expressive bootstrapping* claim about computer languages as pragmatic metavocabularies for much more expressively powerful autonomous vocabularies, namely natural languages. The arguments for and against this pragmatic version of AI-functionalism accordingly look quite different from those arrayed on the opposing sides of the debate about the prospects of symbolic AI.

Combining the notion of PP-sufficiency that holds between two constellations of practices-or-abilities when one can be algorithmically elaborated from the other with the two sorts of basic meaning-use relations out of which I previously constructed the notion of expressively bootstrapping pragmatic metavocabularies—namely, a set of practices-or-abilities being PV-sufficient to *deploy* a vocabulary and a vocabulary being VP-sufficient to *specify* a set of practices-or-abilities—makes it possible to define further kinds of pragmatically mediated semantic relations. As my final example, consider the relation between *logical* vocabulary—paradigmatically, conditionals—and ordinary, non-logical, empirical descriptive vocabulary. I take it that every autonomous discursive practice must include performances that have the pragmatic significance of *assertions* and *inferences* (which I would argue come as an indissoluble package). I actually think this PP-necessary condition on any practices PV-sufficient for autonomously deploying a vocabulary can usefully be treated as sufficient as well—that is, as what distinguishes *discursive* practices as such. But nothing in what follows turns on that further commitment. To count as engaging in such practices, practitioners must exercise an ability, however fallible, to assess the goodness of material inferences: to sort them into those they accept and those they reject. This is part of what one must *do* in order to *say* anything. But it is easy to say how those recognitional and performative abilities, for these

purposes counted as primitive, can be algorithmically elaborated into the capacity to use conditionals. An algorithm VP-sufficient to specify an automaton that practically implements such a pragmatic elaboration or PP-sufficiency relation is the following:

3. Assert the conditional '*if p then q*' if one endorses the inference from *p* to *q*;

4. Endorse the inference from *p* to *q* if one asserts the conditional '*if p then q*'.

These rules of usage codify introduction and elimination rules for the conditional. So the capacity to use conditionals can be algorithmically elaborated from the capacities to make assertions and assess inferences. This is the composition of a PP-sufficiency relation with a PV-sufficiency relation, and is expressed in the following meaning-use diagram[22]:

**Elaborating
Conditionals**
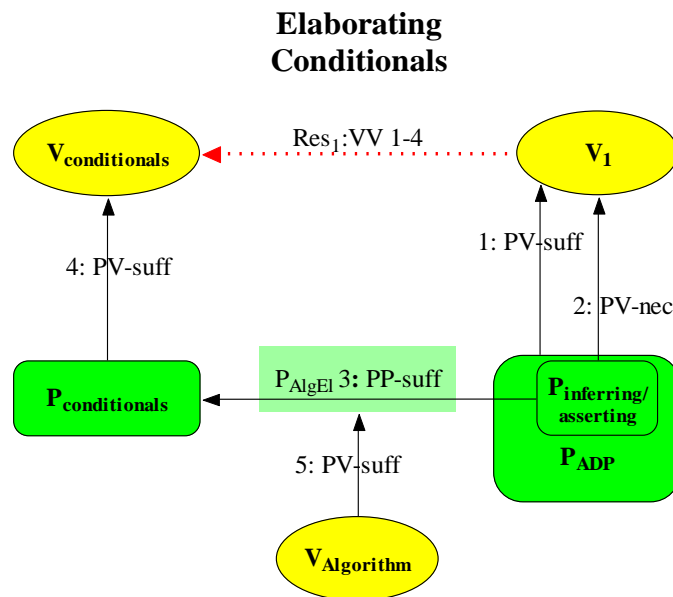
The complex resultant meaning-use relation indicated by the dotted arrow at the top of the diagram is a further pragmatically mediated semantic relation. The diagram indicates *exactly* what constellation of sub-claims about *basic* meaning-use relations must be justified in order to justify the claim that this relation obtains between two vocabularies, and hence the diagram graphically presents a distinctive kind of meaning-use *analysis* of that semantic relation.

In fact, if we think further about this example, by filling in another basic meaning-use relation that obtains in this case, we can define an even more articulated pragmatically mediated semantic relation between vocabularies. For when conditionals are deployed with the practical circumstances and consequences of application specified in the algorithm stated above, they let practitioners *say* what otherwise they could only *do*; that is, they express *explicitly*, in the form of a claimable, hence propositional, content, what practitioners are *implicitly* doing in endorsing some material inferences and rejecting others. This is a VP-sufficiency relation: conditionals let one *specify* the practices of taking-or-treating inferences as materially good or bad. Adding in this *explicating* relation between conditionals and the practices-or-abilities they make explicit yields a new pragmatically mediated semantic relation that conditionals stand in to every autonomously deployable vocabulary. Its meaning-use diagram is this:

Elaborated-Explicating (LX) Conditionals
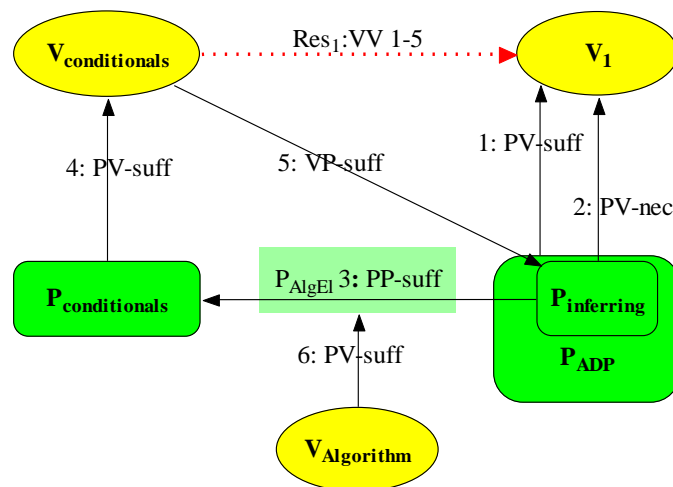
The practical capacity to deploy conditionals (that is, something PV-sufficient for their use) both can be *elaborated from* practices PP-*necessary* for every ADP, and *explicates* those practices (in the sense of being VP-sufficient for them). It is *elaborated-explicative* relative to every autonomous vocabulary. We say, it is LX for every AV, hence for every vocabulary (since the use of

any vocabulary presupposes, and in that sense is parasitic on, the capacity to use some autonomous vocabulary).

I believe that this complex resultant pragmatically mediated semantic relation is important for understanding the distinctive semantic role played by *logical* vocabulary generally: not just conditionals, but also negation (which makes explicit a central feature of our practice of treating claims as materially *incompatible*), and even modal vocabulary (which makes explicit a central feature of our practice of associating *ranges of counterfactual robustness* with material inferences). In my initial characterization of the classical semantic project of philosophical analysis, I pointed to the special status that is accorded to logical vocabulary in that project. What I called "semantic logicism" is its commitment to the legitimacy of the strategy of using logical vocabulary to articulate the semantic relations between vocabularies that is its goal— paradigmatically in connection with the core projects of empiricism, naturalism, and functionalism. One interesting way to vindicate that commitment (that is, at once to explain and to justify it) would be to appeal to the fact that logical vocabulary is elaborated from and explicating of every autonomously deployable vocabulary whatsoever. For that means that the capacity to use logical vocabulary is both in this very clear and specific sense *implicit* in the capacity to use any vocabulary, and has the expressive function of making *explicit* something already present in the use of any vocabulary.

I won't say anything more here about how such a vindication might proceed, contenting myself with the observation that insofar as there is anything to an account along these lines, supplementing the traditional philosophical analytical concern with semantic relations between the *meanings* expressed by different kinds of vocabulary by worrying also about the pragmatic

relations between those meanings and the *use* of those vocabularies in virtue of which they express those meanings is not so much *extending* the classical project of analysis as *unpacking* it, to reveal *explicitly* a pragmatic structure that turns out to have been *implicit* in the analytic semantic project all along. For the conclusion will be that it is *because* some vocabularies are universal pragmatically elaborated and explicitating vocabularies that semantic analysis of the logicist sort is both possible and legitimate at all. I don't claim to have entitled myself to that conclusion here, only to have introduced some conceptual machinery that might make it possible to do so—and so at least to have sketched a way in which the insights of the *pragmatist* tradition can be assembled and developed so as to be constructively helpful to, rather than destructively critical of, the classical project of philosophical semantic *analysis*, and so to open the way to extending that project in promising new directions.

Notes to Chapter Six:

[1] Oxford: Oxford University Press, 2008.

[2] 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.

[3] This argument occupies roughly the first half of his classic *Empiricism and the Philosophy of Mind,* reprinted with an introduction by Richard Rorty and a study guide by Robert Brandom (Cambridge, MA: Harvard University Press, 1997). His critique of the phenomenalist version of empiricism can be found in "Phenomenalism," in *In the Space of Reasons,* ed. Kevin Scharp and Robert Brandom (Cambridge, MA: Harvard University Press, 2007), p. 303-49.

[4] Cf. Quine's remark (in "Two Dogmas of Empiricism"): "Meaning is what essence becomes when it is detached from the thing and attached to the word." I distinguish ˢscare quotesˢ from regular quotes by using superscripted 's's. My official theory of such quotation, as the conceptual converse of *de re* ascriptions of propositional attitude, can be found in *Making It Explicit,* pp. 545-7 and 588-90.

[5] I am indebted for this way of thinking of Wittgenstein's point to Hans Julius Schneider's penetrating discussion in his *Phantasie und Kalkul* (Frankfurt: Suhrkamp, 1992].

[6] 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: Oxford University Press, 2006).

[7] 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 Hegel 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.

[8] For the purposes of the present project, I will maintain a studied neutrality between these options. The apparatus I am introducing can be noncommittal as to whether we understand content-conferring *uses* of expressions in terms of social practices or individual abilities.

[9] Somewhat more precisely: some theory (a set of sentences), formulable in the vocabulary in question, is such that if all those sentences are true of some interlocutor, then it thereby counts as exercising the relevant ability, or engaging in the relevant practices.

[10] See his "Naturalism without Representationalism," in *Naturalism in Question,* ed. Mario de Caro and David Macarthur, (Cambridge, MA: Harvard University Press, 2004), pp. 71-90. Price calls his view "subject naturalism," as opposed to the more traditional (and more metaphysical) "object naturalism." I discuss it further in the next chapter.

[11] I discuss this argument in greater detail in the final chapter of *Tales of the Mighty Dead* (Cambridge, MA: Harvard University Press 2004).

[12]   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}, '$a^n b^n$' 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.

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

[14]   In the syntactic metalanguage for specifying vocabularies that I mentioned in the note above, this is the vocabulary (ha|ho)*!

[15]   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. M. O. Rabin and D. Scott, "Finite Automata and their Decision Problems," *IBM Journal of Research and Development* 3, no. 2 (1959), pp. 115-25.

[16]   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:

The "I'll Have What She's Having" Automaton
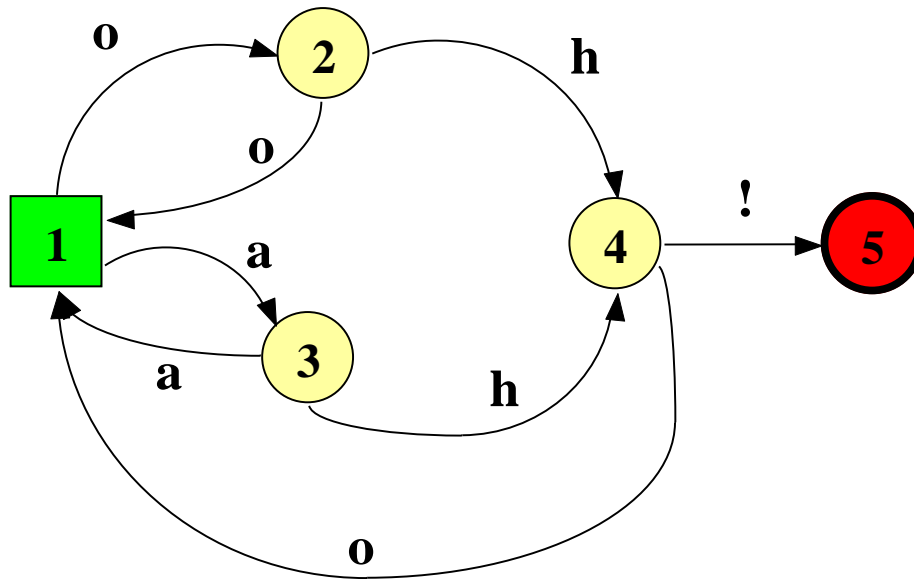
[17] By contrast to FSA's, there need not in general be, for every vocabulary computable by a non-deterministic PDA, some deterministic PDA that reads and writes the same vocabulary.

[18] Thomas Sudkamp, *Languages and Machines,* 2nd ed. (Reading, MA: Addison-Wesley, 1997), Chapter 10.

[19] In principle. There are subtleties that arise when we look at the details of actual implementations of particular computer languages, which can remove them from qualifying as strictly context-free.

[20] There are various vocabularies that are VP-sufficient for specifying those meta-abilities. Specifying them in terms of the differentially elicitable capacities to change state and to store and retrieve symbols is just one of them.

21   For this reason, the frame problem, as it is often formulated, does not *immediately* arise for the pragmatic version of AI-functionalism. But as is explored in the third chapter of *Between Saying and Doing,* it does get a grip, at a different point.

22   I indicate PP-necessity relations by including the rounded rectangle for one set of practices-or-abilities in another.