

May 2, 2007

*Between Saying and Doing:
Towards an Analytic Pragmatism*

Lecture 2:

Elaborating Abilities: The Expressive Role of Logic¹

Section 1: Automata as Algorithmically Elaborating Abilities

In the previous lecture I began to set out my enterprise against the background of an understanding of the classical project of philosophical analysis as having the task of exhibiting what is expressed semantically by one vocabulary (one sort of meaning) as the logical elaboration of what is expressed by another. From this point of view, empiricism and naturalism show up as generic core analytic programs, with their species distinguished in part by the vocabularies they treat as basic and by those they seek to elaborate on those bases, and in part by the sort of logically articulated elaboration they see as relating the basic and target vocabularies. Pragmatism can be thought of as challenging these analytic undertakings, by insisting that appeal to *meaning* is just *one*, optional, theoretically laden way of dealing with *use*, one that is based on a defective, scientific, uniformitarian theory of what is required to understand what one has to *do* in order to count as *saying* various kinds of thing—as deploying various vocabularies. But if we accept the methodological pragmatism that tells us that the point of *semantic* theory is to make sense of *pragmatics*—that *meanings* should be thought of as theoretical entities postulated to explain or at least to codify proprieties of *use*—and supplement it with the *semantic* pragmatism that tells us that only its use can explain the association of meaning with, or its

¹ The work reported here was conducted with the support of the A.W. Mellon Foundation, through their Distinguished Achievement in the Humanities Award, and the Center for Advanced Study in the Behavioral Sciences at Stanford University (where I was also supported by the Mellon Foundation).

expression by, a vocabulary, we see that we are not forced to choose between thinking in terms of the *meanings* expressed by vocabularies and thinking in terms of their *use*. I want to show how pragmatism can be turned from a pessimistic, even nihilistic, counsel of theoretical despair into a definite, substantive, progressive and promising program in the philosophy of language: indeed, how it can be understood simply as the latest phase of the analytic project.

A central task must be to get clearer about the notion of *use* that pragmatism seeks (irenicly) to lay alongside semantics, or with which it seeks (martially) to displace semantics. The starting-place I suggested last time is the way in which automaton theory lets us specify the abilities needed to *deploy* syntactically characterized vocabularies, in the sense of being able both to distinguish and to generate them. This idiom illustrated a number of basic meaning-use relations: relations between (as I suggested we label them) vocabularies and practices-or-abilities. My discussion of them began with vocabularies that are *syntactically* specified, rather than *semantically* specified—a restriction that pays off in the clarity and definiteness of the relations involved, but only at the high price of abstracting from the dimension of *semantic* expressiveness that makes us care philosophically about vocabularies in the first place. (After all, it is *meaning*-use relations we are ultimately after.) Nonetheless, we can see already at this point that the metavocabulary we use to characterize a vocabulary makes a significant difference to what sort of practices-or-abilities we can count as deploying it. We also saw that the practices-or-abilities to deploy those metavocabularies—the ones that are, in my terms, PV-sufficient for them—must be thought of in turn in terms of the vocabularies in which *they* are specified.

I promised that we could build on this initial, simplified syntactic model of basic and resultant meaning-use relations—especially PV-sufficiency, VP-sufficiency, the pragmatically mediated VV-relation that is their composition, and the sort of pragmatic expressive bootstrapping exhibited by such relations even in the syntactic case—to yield insights into corresponding relations between genuine meanings and uses for vocabularies characterized in full-blooded, *semantic* meta-vocabularies. So far I have presented automata as merely syntactic engines, which is to say as sets of primitive PV-sufficient abilities to deploy purely syntactically characterized vocabularies. But there is another way to think about the abilities that the automaton-theoretic VP-sufficient vocabulary specifies. Automata put together primitive abilities so that they add up to more complex ones. Automata are the practical embodiments of *algorithms*. And algorithms generally say how some set of primitive abilities can be so exercised as to constitute more complex abilities. For instance, an algorithm implemented by an automaton tells us how to put together the capacities to multiply and subtract so as to amount to the capacity to do long division.

Thought of this way, 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. In the metavocabulary for meaning-use relations I am introducing here, they *implement* PP-sufficiency relations: the kind of relation that obtains when the capacity to engage in one sort of practice or exercise one sort of ability is *in principle sufficient* for the capacity to engage in other practices, or exercise other abilities. Putting the point the other way around, an automaton-theoretic specification of how a set of primitive practices-or-abilities can be algorithmically elaborated into a more complex set of practices-or-abilities shows how the latter can be *pragmatically analyzed* into the former.

Semantic analysis in the most straightforward case—for instance Russell’s analysis of definite descriptions—shows how some *meanings* can be exhibited as complex, in the sense that they can be understood as compounded by definite means out of more basic meanings. The semantic logicism characteristic of the classical project of philosophical analysis in the twentieth century insists that it is *logical* vocabulary that articulates or makes explicit such semantic compounding. We can lay alongside this conception a notion of *pragmatic analysis*. This is the analysis of some *uses* as complex, in the sense that they can be understood as compounded by definite means out of simpler uses, that is, in the regimentation I am employing, practices-or-abilities.² My claim is that automata can be thought of as one way of implementing such pragmatic analyses. What thinking about automata in this broad sense will do is teach us that *algorithmic elaboration* of primitive abilities into complex ones plays the same role in pragmatic analysis that *logic* does in semantic analysis. Algorithmic elaboration is a logic of practical abilities.

My aim is to weave these two sorts of analysis together as essential elements of a more comprehensive kind of meaning-use analysis. In today’s lecture I will begin to explore what happens when we add the PP-sufficiency relation of algorithmic elaboration, implemented by automata, to the conceptual apparatus of the PV-sufficiency of practices-or-abilities to deploy vocabularies and the VP-sufficiency of vocabularies to specify practices-or-abilities. Doing that is the first step in the move to understanding the practices-or-abilities involved in deploying *semantically* characterized vocabularies.

² Of course, the issue of how some *doings* can be presented as products, compounded in antecedently specifiable ways, out of more primitive doings, is a *very* general one. But we have the advantage that the doings we care most about are sayings: practices-or-abilities that count as deploying vocabularies, as conferring or applying meanings: *discursive* practices-or-abilities. This gives our approach to the general question a more particular focus.

Section 2: Transducing Automata

Automata, in the general sense in which I want to think about them, are constellations of practices-or-abilities that algorithmically elaborate sets of primitive practices-or-abilities into more complex ones. *Transducing* automata³ generalize the primitive reading-and-writing abilities of finite state automata to include *discriminating stimuli* of any kind, on the input side, and *differentially responding* in any way, on the output side. That is, instead of an alphabet of character-types, tokenings of which can be indifferently read or written, these automata are defined over a pair $\langle S, R \rangle$ of *stimulus*-types and *response*-types. The stimulus-types are any circumstances to which they system is able to respond differentially. The response-types are any kinds of performance that can be differentially elicited from the system—that is, which it produces in some circumstances and not others. A *single-state transducing automaton* (SSTA) can be specified by a state-table that is just a set of pairs of stimulus-kinds and the kinds of responses they would elicit. It is just a set of reliable differential responsive dispositions.

The stimulus-response model might seem to impose a formal, narrowly behaviorist strait-jacket on what counts as a primitive ability. Behaviorism is lurking in the vicinity, but I think it is important to see that its proximity does *not* arise from this characterization of primitive abilities. For how restrictive the stimulus-response model of such abilities is depends on the VP-sufficient vocabulary that we are allowed to use in specifying the discriminable stimuli and differentially elicitable performances. If *no* restrictions are imposed, then candidate stimuli one might have the ability to discriminate can include such species as poetry that qualifies as lyrical,


³ In the case of finite-state machines, these are Mealy machines (or—only slightly different, but basically equivalent—Moore machines).

actions that are cruel, remarks that are witty or telling, historical events that illustrate the superiority of liberal political arrangements, and so on. And the responsively elicitable performances could include anything one had the ability to do: painting well-composed pictures, toeing the Party line, riding a bike, standing the right conversational distance from someone, and so on. So the practices-or-abilities counted as ‘primitive’ in the sense that they are regarded as inputs to the process of algorithmic elaboration need not be crude, easily acquired, or shareable with lab rodents.

What *is* crudely behavioristic about anything specifiable as a *single-state* transducing automaton—whether it became so specifiable by a process of learning or came that way innately—is the *inflexibility* of its behavior. Though it can, by definition, respond differently to different situations, it always responds the same way to the same kind of situation. Its behavior is governed by a set of reliable differential responsive dispositions. From an algorithmic point of view, the different stimulus-kinds can be thought of as *instructions* to produce performances of the kind associated with those stimulus kinds in the state-table that specifies the automaton. So if we are in a position to produce stimuli of desired kinds at will, we can program the SSTA to execute arbitrary *straight-schedule algorithms* over its specified behavioral repertoire. That is, we can instruct it to do any of the things it can do, to produce performances of any of the kinds in its responsive repertoire R, sequentially, in any order we like. It can follow a *list* of instructions. That is one kind of practical elaboration of basic abilities into more complex ones, but it is a poor one.

Much greater flexibility is exhibited by *finite-state transducing automata* (FSTAs). Besides responding differentially to stimuli by producing performances from its responsive repertoire, an FSTA can respond differentially by changing state.

A Finite-State Transducing Automaton




This means it can combine its primitive abilities according to *conditional branched-schedule algorithms*, which specify how it *alters* its differential responsive dispositions in response to the actual outcome of something it has done—for instance, its success or failure at achieving some recognizable result. These multi-state *functional* systems are accordingly much more capable than the corresponding single-state *behavioral* systems definable over the same set of primitive discriminable situation-kinds and elicitable performance-kinds. They can elaborate much more complex practices-or-abilities on the basis of the same primitive abilities. The advance from *behaviorism* to *functionalism* in the philosophy of mind corresponds to the move from a single-state to a multi-state model. I will have more to say about functionalism in this connection later on.


Transducing automata are more than merely syntactic elaborating engines because the stimuli they can respond to and the responses they can produce are not limited to symbol-types. Depending on the ‘alphabet’ of stimulus- and response-kinds they elaborate, however, they can also manipulate symbols. But they also allow us to think about symbols in a new way: still not representationally, nor yet semantically, but not just syntactically either. For we can think of *symbols* generically as anything that can both be read and written, that is, recognized and produced. In this broad sense, anything in the intersection $S \cap R$ of S and R can be used as a symbol: any stimulus-kind instances of which the system can produce as responses.

Algorithmic elaborating abilities—paradigmatically abilities to respond differentially to stimuli by changing state—implement PP-sufficiency relations between primitive abilities and more complex ones. They accordingly delineate a notion of one set of practices-or-abilities being *in principle sufficient* for another, of what it is for the elaborated practices-or-abilities to be *practically implicit in* the primitive ones from which they are algorithmically elaborated. They give a definite sense to saying that *if* a system has the abilities $A_1 \dots A_n$, it already knows how to do everything it needs *in principle* to know how to do in order to do A. There is a kind of idealization involved in any such claim. What is required to define the basic relation of PP-sufficiency is to be clear about the nature of that idealization. Not every system may in fact have the algorithmic elaborating abilities. Algorithmic PP-sufficiency is what holds in case *all* it needs to elaborate its basic abilities into the complex one in question is those algorithmic abilities.

In this sense, we can diagram the meaning-use relations for two of the examples already considered so as to show the automata that implement the algorithmic elaboration PP-sufficiency relations like this:



This is an automaton-implemented, algorithmically elaborated, pragmatically mediated *syntactic* relation between vocabularies.



The resultant meaning-use relation in this latter MUD is an automaton-implemented, algorithmically elaborated, pragmatically mediated *semantic* relation.

The automaton-theoretic metavocabulary for specifying abilities that implement PP-sufficiency relations typically involves further idealizations as well. These idealizations have not, so far as I can determine, been much discussed, but they are substantial and noteworthy in connection with the present enterprise. The first idealizing assumption is that any stimulus to which a system can respond differentially can be connected to any response it is capable of differentially producing. That is, there is no restriction on which elements of S and of R can be linked together as elements of the state table of an SSTA, FSTA, or PDTA. I will call this assumption “*response substitution*”: if a system is capable of producing some response differentially to Ss, and can respond differentially to some kind of stimulus by producing Rs, then it is in this sense in principle capable of responding differentially to Ss by producing Rs.

For many actual systems, this is not a realistic assumption. I can respond with delight to the baby's chortle, the kitten's antics, and an unexpected glimpse of a beautiful orchid. I can distinguish disgusting smells, tastes, and sights. Could I really be trained to respond with delight to those disgusting stimuli? I am not, of course, claiming that I could. I am delineating a sense of "in-principle capability" in which I *in principle* already have all the basic discriminative and performative abilities needed for such a responsive connection. Insofar as I cannot in fact responsively hook those basic abilities together in arbitrary combinations, those restrictions count—relative to the idealization—as *psychological* restrictions, in a broad sense. (In this sense, automata have no psychology, though physical implementations of them do.)

The second idealization is that the stimulus-response connections of which a system is capable—either ideally, according to response substitution, or really psychologically—can be arbitrarily combined into *states*. If I can respond to stimuli of kind S_i with responses of kind R_j and to stimuli of kind S_k with responses of kind R_l , and if S_i and S_k are compatible and so are R_j and R_l , then I can be in a state in which I am disposed to respond to S_i with R_j *and* to S_k with R_l . I will call this idealizing assumption "*arbitrary state formation*." Once again actual failures to form state-tables arbitrarily will be counted as due to *psychological* restrictions of the system, in the broad sense defined by contrast to this idealization.

The final idealization of the transducing-automaton model of PP-sufficiency is then "*arbitrary state permutation*," which is arbitrary formation of state *tables*. If a system is capable of entering into responsive $State_1$ and is capable of entering into responsive $State_2$, defined over

the stimulus-response-kind universe $\langle S, R \rangle$, then it is capable in principle (in this sense) of responding to stimuli of any kind $S_i \in S$ by changing from $State_1$ to $State_2$ and *vice versa*.

The algorithmic abilities to form states as constellations of differential responsive dispositions, and to change them in response to already-discriminable stimuli, together with these idealizations concerning those abilities give a well-defined sense to the concept of one set of practices-or-abilities being PP-sufficient for another, in the sense that the latter can be exhibited as the result of *algorithmically elaborating* the former—or equivalently, that the latter can be *algorithmically decomposed or analyzed* into the latter.

Section 3: Autonomous Discursive Practices, Asserting, and Inferring

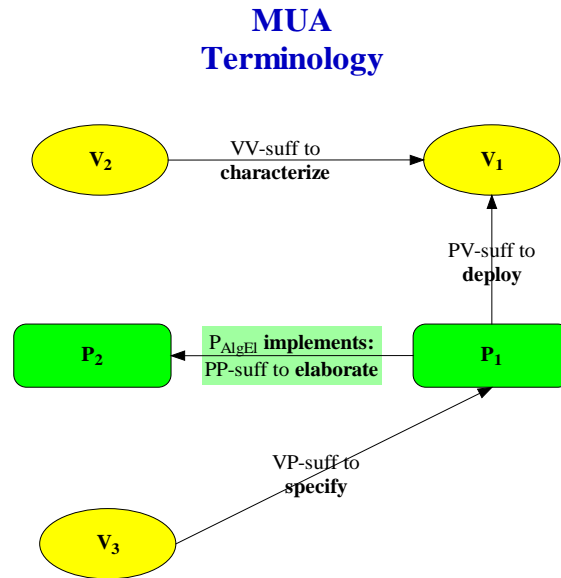
So far I have introduced three basic meaning-use relations, and used automaton-theoretic considerations to illustrate and clarify them. These basic MURs are:

- the PV-sufficiency of practices-or-abilities to *deploy* a vocabulary,
- the VP-sufficiency of a vocabulary to *specify* a set of practices-or-abilities, and
- the PP-sufficiency of a set of practices-or-abilities that can be *elaborated* into another, by a set of algorithmic abilities that *implement* that practical elaboration.

I have also appealed to a fourth:

- the VV-sufficiency of one vocabulary to *characterize* another (the relation of being a direct or immediate semantic or syntactic metavocabulary).

This terminology lines up with the conventions I have offered for meaning-use diagrams like this:



All of these are *sufficiency* relations. Now I want to discuss briefly some corresponding *necessity* relations. In the previous lecture I mentioned one sense of pragmatic presupposition: the PP-necessity relation that obtains when it is not possible to engage in or exercise one set of practices-or-abilities unless one also engages in or exercises another.

The sense of pragmatic presupposition I want to focus on here, however, is PV-necessity: the sense in which the capacity to *say* something of a certain kind, to deploy a particular vocabulary, can require being able to *do* something of a specifiable kind. For example, using *observational* vocabulary—using expressions in the observational way, making noninferential reports, for instance of the visible presence of red things—requires being able reliably differentially to respond to stimuli of the relevant kind, for instance, visible red things. This is a capacity those who *can* deploy such vocabularies might share with those, such as pigeons, which *cannot* do so.

John Dewey bequeathed to us a notion of *pragmatism* as opposed to *intellectualism* or *platonism*. These two broad currents of philosophical thought are identified with different orders of explanation: pragmatists appeal to knowing-how in order to explain knowing-that (or, more carefully, saying- or believing-that), and their intellectualist opponents (virtually the entire prior philosophical tradition, they thought) go the other way around, finding principles standing behind every propriety of practice, and rules grounding every practical ability. We are in a position to subdivide pragmatism in this traditional sense into two subsidiary claims: that for any vocabulary (any kind of saying-*that*) there are some practical abilities (some bits of know-how) that are PV-*necessary* to count as deploying it, and that there are some practices-or-abilities that are PV-*sufficient* to confer those contents or count as deploying that vocabulary. I have already called the latter ‘semantic’ pragmatism. I will call the former, the PV-necessity claim, ‘*fundamental*’ pragmatism. Together they articulate the sense in which pragmatists take practices-or-abilities to be privileged with respect to, to “stand behind”, the capacity to say, mean, or believe (hence to know) anything discursively.

Besides asking what practical abilities are necessary in order to deploy a *particular* vocabulary—such as observational, logical, indexical, or normative vocabulary—we can ask whether there is some kind of thing one must be able to do in order to deploy any vocabulary whatever, no matter what it is. Are there any practical abilities that are *universally* PV-necessary? This is a way of asking what sorts of *doings* deserve to count as *sayings*. In the very general way in which I have been using the term ‘vocabulary’, many of the vocabularies of most concern are language *fragments*: expressions whose use is not an *autonomous discursive practice* (ADP), in the sense of a language game one could play though one played no other, or a set of discursive abilities one could have though one had no other specifically discursive abilities. Singular terms, for instance, constitute a vocabulary in my sense. But no-one could count as using singular terms unless she could also use sentences containing them, and hence predicates

as well. Vocabularies deployed by autonomous discursive practices may be called ‘*autonomous*’ vocabularies, or just ‘languages’. (In this sense, there is no “language of physics”, though there is a *vocabulary* of physics.) Any vocabulary that is fragmentary (that is, not autonomous) pragmatically presupposes, in the PV-necessity sense, some set of autonomous discursive practices, which *are* PV-sufficient for deploying the autonomous vocabulary of which the vocabulary in question is a fragment. If that is right, then any practices PV-necessary for every *autonomous* vocabulary would be PV-necessary for every vocabulary whatsoever.

It certainly can coherently be denied that there is any core of practices-or-abilities common to all autonomous discursive practices. Wittgenstein seems to do so in thinking of such practices as language-*games* (Sprachspiele), given his insistence that the concept game does not have an essence or a definition, but is structured rather by family resemblances. That is why he does not see language as having a ‘downtown’, by more or less peripheral relation to which something else can count as linguistic. But I think there is a relatively bright line to be drawn in the vicinity, marking a good thing to mean by ‘linguistic’ or ‘discursive’. Specifically *linguistic* practices are those in which some *doings* have the practical significance of *sayings*. The core case of *saying* something is making a *claim*, *asserting* something. The practices I will call ‘linguistic’ or ‘discursive’ are those in which it is possible to make assertions or claims. Although, as Wittgenstein is concerned to point out, their occurrence can have other kinds of import, the home language game of what are for that very reason called ‘declarative’ sentences is their free-standing use in asserting.

By this assertional pragmatic criterion of demarcation of the discursive, many of Wittgenstein's *Sprachspiele* are not really *Sprachspiele*. They are *vocal* practices, but not *verbal* ones. For instance, the practices introduced in Section 2 of the *Philosophical Investigations* are specified as:

...a language consisting of the words "block", "pillar", "slab", "beam". A calls them out;
 B brings the stone which he has learnt to bring at such-and-such a call. Conceive this as a complete primitive language.

These 'calls'⁴ are properly so-called. They are *signals*, appropriately responded to, according to the practice, in one way rather than another. But they are not *orders*. For an order specifies *how* it is appropriately responded to by *saying* what it is one must *do* in order to comply. "Shut the door!" can be a saying of the imperative kind only as part of a larger practice in which "The door is shut," can be a saying of the declarative kind.

I am suggesting that we treat assertional practices-or-abilities as PV-necessary for any autonomous practice we count as 'discursive'—that we think of asserting as the minimal kind of *doing* that counts as a *saying*. But what is asserting?⁵ I think the beginning of wisdom in answering this question is to see that *asserting* and *inferring* are internally related practices, in the sense that each is PP-necessary for the other. Assertions are essentially, and not just accidentally, speech acts that can play the role both of premises and as conclusions of inferences. In my book *Making It Explicit*, I pursue the explanatory strategy of treating the connections in both directions also as PP-*sufficiency* relations. That is a bold and deservedly contentious theoretical approach. Here I am making only the much weaker and less controversial claim that

⁴ Wittgenstein's verb is 'ausrufen'.

⁵ [snarky footnote on Searle in *Speech Acts*: "an undertaking to the effect that p." But look up the exact quote and give the page reference.]

only doings that can serve as both kinds of termini for inferences should be thought of as assertions (when the point is put pragmatically), and so that only what can stand in inferential relations should be thought of as the contents expressed by declarative sentences (when the point is put semantically). According to this way of thinking, *inferential* practices are PP-necessary components of every autonomous discursive practice, hence PV-necessary for the deployment of every autonomous vocabulary, hence PV-necessary for the deployment of every vocabulary whatsoever. They are *universally* PV-necessary.

The combination of this assertional pragmatic criterion of demarcation of the discursive and the inferential necessary condition on assertion is a distinctive kind of pragmatic *rationalism* about the discursive. *Language games*, *Sprachspiele*, must, according to this line of thought, include practices of giving and asking for *reasons*—because *assertions*, the most basic kind of *sayings*, must be capable of both serving as and standing in need of reasons. Pragmatic rationalism is the view that language *does* have a ‘downtown’, and it comprises the practices of making claims and giving and asking for reasons for them.

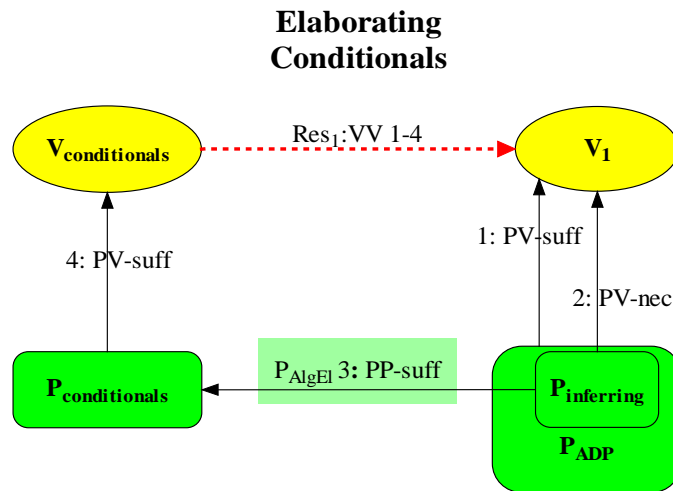
It is perhaps worth noticing in passing that this view permits a substantive response to Derrida’s charge that philosophers have self-servingly fetishized reason-giving, which he sees as simply *one* game one can play with language, deserving no privilege of any kind relative to the myriad of others. Rather than simply ignoring him, or demonizing him as a dangerous irrationalist just for raising the challenge, on the one hand, or acquiescing in the radical conclusion he draws from what he takes to be the unanswerability of his challenge to justify the privileging he calls ‘logocentrism’, on the other, the pragmatic rationalist offers a responsive

answer to that challenge: that our expressions play a suitable role in reasoning is an essential, necessary element of *our saying*, and *their meaning* anything at all. Apart from playing such a role in justification, inference, criticism, and argument, sentences and other locutions would not have the meanings appealed to and played with by all the other games we can play with language. We philosophers should be proud to acknowledge and affirm our logocentrism, but should also justify it by an account of the relations between meaning and use, conceptual content and discursive practice.

Section 4: Introducing Conditionals

If these PV-necessity claims are correct, then there are two abilities that must be had by any system that can deploy any vocabulary, as part of the autonomous discursive practice of which the use of that vocabulary is a part: the ability to respond differentially to some sentence-tokenings as expressing *claims* the system is disposed to *assert*, and the ability to respond differentially to moves relating one set of such sentence-tokenings to another as *inferences* the system is disposed to *endorse*. These may be treated as primitive abilities for the purpose of algorithmic elaboration of further abilities, for which they are in that sense PP-sufficient. If we do that, we see that the algorithmically primitive abilities to make assertions and to sort inferences into those that are and those that are not materially good ones—in the sense of making such a distinction in practice, regardless of whether or not one gets the distinction *right* by some further standard—suffices to elaborate practices that are PV-sufficient to deploy a further vocabulary, namely *conditional* locutions. That constellation of basic meaning-use relations then institutes a pragmatically mediated resultant semantic relation between the original vocabulary

and a more complex vocabulary that involves conditionals relating sentences deployed in the underlying ADP.



In this meaning-use diagram, the fact that the rectangle representing one set of practices-or-abilities is contained in another indicates that the smaller one is PP-necessary for the one it is contained in.

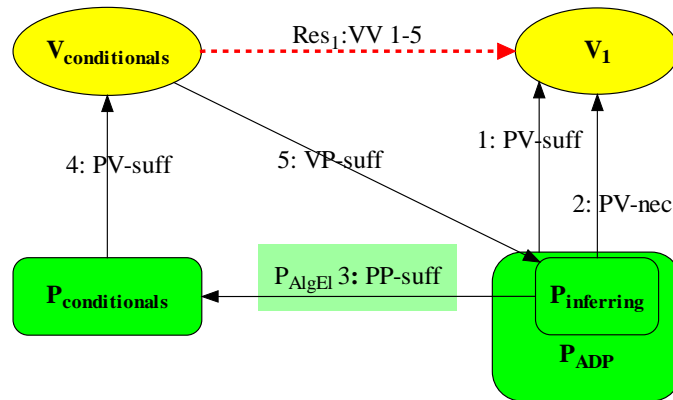
Here is a very rough sketch of how the algorithmic elaboration in question might work. By hypothesis, the system has the ability to respond differentially to the inference from p to q by accepting or rejecting it. It also must have the ability to produce tokenings of p and q in the form of assertings. We assume that since it can produce those assertings, we can teach it also to produce-assertively tokenings of the new form “if p then q .” What is required, then, is first that this new sort of response be hooked up responsively to the previously discriminable stimulus, so that it is asserted just in those cases where the inference from p to q would have been responded to as a good one. This is an exercise of the algorithmic elaborative ability I earlier called “response substitution”: responsively connecting a previously distinguishable stimulus-kind to an already elicitable performance-kind. This rule codifies the *circumstances* of appropriate

application of the newly-introduced conditionals relating sentences deployed by the ADP. For the *consequences* of application, we need another bit of response substitution. The system can already, by hypothesis, respond to some stimuli by treating an inference as good or bad. We must now hook up that response to a new stimulus-kind. The system must respond to its *assertion* of the *conditional* “if p then q ” by treating the *inference* from p to q as a good one—for instance, by being disposed to endorse q assertively if it is disposed to endorse p assertively. These new differential responsive abilities, achieved by reshuffling prior ones, then settle the state-table that specifies how the system is able to respond to different presented stimuli: non-logical sentences and inferences involving them, and now also conditional sentences and inferences involving them—paradigmatically, *modus ponens*. In a clear sense, then, the capacity to distinguish good from bad inferences involving *non-logical* sentences is (PP-)sufficient for the ability to deploy conditionals involving those sentences. That is what is represented by the MUD above.

But the pragmatically mediated semantic relation between the vocabulary of conditionals and any autonomously deployable vocabulary that obtains in virtue of the PV-necessity of material inferential practices-or-abilities for the deployment of any vocabulary includes another crucial element as well. For conditionals let one *say* something, where before one could only *do* something. Saying that if something is copper, then it conducts electricity is a new way of doing—by saying—what one was doing before by endorsing the material inference from “That is copper,” to “That conducts electricity.” Conditionals make *explicit* something that otherwise was *implicit* in the practical sorting of non-logical inferences into good and bad. Where before one could only in practice *take* or *treat* inferences *as* good or bad, after the algorithmic

introduction of conditionals one can endorse or reject the inference by explicitly *saying* something, by asserting or denying the corresponding conditionals. What the conditional says explicitly is what one endorsed implicitly by doing what one did. The expressive role distinctive of conditionals whose use is elaborated in the way I have just specified is to *codify* inferences, to *specify* inferential practices-or-abilities, to *explicate* them, in the sense of making explicit something that was implicit in them.⁶ That is to say that the MUD for the resultant MUR that is put in play by the introduction of conditionals is really this one:

**Elaborated-Explicating (LX)
Conditionals**



Conditionals are both *elaborated* from and *explicative* of inferential practices. As shorthand, I will call this important kind of pragmatically mediated semantic VV-relation an “LX” relation.

⁶ The effect of asserting the new sayable that is a conditional (the practices for deploying which can be elaborated from inferential practices) is not to *say that* an act of inferring is permissible. For that one needs *normative* vocabulary, and the concepts of *saying* and *inferring*—expressed by vocabulary one need *not* master in order to master the use of conditionals. Rather, conditionals assert explicitly that one thing that can be said *follows from* another thing that can be said, that the one is a *consequence* of the other. In Lecture Four (and further, early in Five) I discuss various sorts of semantic inferential relations among contents and their relation to pragmatic relations among deontic statuses, and say something about how to introduce the normative vocabulary that is VP-sufficient to specify this aspect of inferential practices. What I mean to be introducing here is the notion of a *genus* of relations of VP-sufficiency to express different aspects of practices PV-sufficient to deploy vocabularies. The relations between the aspects of practice made explicit by logical and modal vocabulary (those pertaining to the *content* expressed), on the one hand, and the different but complementary aspects made explicit by normative vocabulary (those pertaining to the *act* of expressing) on the other hand, are the topic of Lecture Six. Detailed consideration of the sense of “making explicit what is implicit” that is invoked here by the VP-sufficiency relation that constitutes the second half of the pragmatically mediated semantic relation of one’s vocabulary being LX for another must accordingly await clarification until then.

It obtains when some practices-or-abilities PV-necessary for the deployment of one vocabulary can be algorithmically *elaborated* (the ‘L’ part) into a set that is PV-sufficient to deploy a vocabulary that is VP-sufficient *explicitly* to specify or codify the original set of practices (the ‘X’ part). Just in virtue of being able to assert and to sort inferences into materially good and materially bad ones (once again, whether or not correctly), one already knows how to do everything necessary in principle to introduce vocabulary that will let one *say* something, the saying of which *is* taking some inferences to be good and other ones not good.

Conditionals are a paradigm of *logical* vocabulary. For instance, the conditional is the very first bit of logical vocabulary Frege introduces in his *Begriffsschrift*, the founding document of modern logic. I want to suggest that the meaning-use analysis just offered of conditionals specifies the genus of which logical vocabulary is a species. That genus is distinguished by three characteristics:

- i) being deployed by a set of practices-or-abilities that are algorithmically elaborated from
- ii) practices-or-abilities that are PV-necessary for *every* autonomous vocabulary (and hence every vocabulary whatsoever) and that
- iii) suffice to specify explicitly those PV-necessary practices-or-abilities.

Any vocabulary meeting these conditions I will call a “universal LX-vocabulary.” In a later lecture I will offer more details about how one can think of other logical connectives as exhibiting this pattern of pragmatically mediated semantic relation to non-logical vocabularies generally.

By way of anticipatory illustration, I can say that I take it that just as every autonomous discursive practice must involve distinguishing some inferences as materially good, so it must involve distinguishing some claims as materially *incompatible* with others. That a monochromatic patch is red rules out its being blue. Only algorithmic elaboration is required to turn the ability to distinguish material incompatibility of claims into the ability to deploy logical *negation*. And once that bit of logical vocabulary is deployed, it (together with the conditional) lets one *say that* two claimables are materially incompatible: “If a monochromatic patch is red, then it is *not* blue.” That is, negation lets one make explicit, in the form of claims—something that can be said and (so) thought—a relation that otherwise remained implicit in what one practically did, namely treat two claims as materially incompatible. So negation also qualifies as being LX (elaborated-explicitating) for every autonomous vocabulary, and hence as belonging to the pragmatically specified semantic genus within which I am suggesting we locate *logical* vocabulary.

Section 5: Characterizing Logic: the Logician’s Dilemma

I am suggesting that standing in this complex, resultant meaning-use relation to every autonomously deployable vocabulary can serve as a partial answer to a central question in the philosophy of logic: the demarcation question. That question is, roughly, “What is logic?”—or, somewhat more carefully, “What is logical vocabulary?”, that is, “What features should be taken as distinguishing some bit of vocabulary as distinctively *logical* vocabulary?”

Even when we envisage a complete answer as taking the form of necessary and sufficient conditions, the general demarcational question is still normally understood to concern the *circumstances* of appropriate application of the term “logical vocabulary”. But a question of this kind comes with reasonably definite criteria of adequacy for assessing the correctness of an answer only in the context of a relatively clear specification of the *consequences* of application being associated with the expression at issue.⁷ That is, when we ask whether, say, *modal* vocabulary, or set-theoretic vocabulary, is a kind of *logical* vocabulary, what turns on the answer? What will we have found out about it if the answer turns out to be ‘Yes’? What is the difference that *makes* a difference? The demarcational question can sensibly be addressed only if we address also the (at least co-ordinate, perhaps even prior) question concerning the theoretical, explanatory, argumentative, or constructive role logic or logical vocabulary is being envisaged as playing in some larger philosophical enterprise.

Of course, there are as many possible answers to that question as there are philosophical programs that assign some privileged role to logic. And they are liable to lead to very different answers to the demarcational question. The appeal to logical vocabulary that I want to focus on is that made by the classical project of philosophical analysis. As I characterized that project in my first lecture, it aims to exhibit some target vocabulary as the result of semantically elaborating some base vocabulary—in the core programs I pointed to, some version of those privileged by empiricism or naturalism. I pointed out that *logical* vocabulary typically plays a special role in this enterprise: one is allowed and encouraged to use logical vocabulary in the process of elaborating one vocabulary into another. This is what I called the “semantic logicism” of the classical project of analysis. Insofar as the twentieth century analytic project did assign

⁷ This is a lesson Michael Dummett has taught us. [ref. to the relevant bits of *Frege’s Philosophy of Language*]

such a privileged status to logical vocabulary, it is fair to ask why that privilege could be thought to be legitimate. What are the rules of the analytic game? Why is it all right for the empiricist about theoretical entities to appeal to logical vocabulary, but not to modal or intentional vocabulary in elaborating his analyses? Is it all right for the naturalist to employ modal vocabulary in his elaboration of semantic vocabulary? And if so, is that because it *is* logical vocabulary? Why does logical vocabulary “come for free” in analysis, not itself having to be elaborated from the empiricist’s or naturalist’s basic vocabularies?

This is the larger framework in which I want to place logical vocabulary in order to get a standard for assessing various possible principles of demarcation. I want to ask the general demarcational question with the standard of assessment for answers being set by the demand that it vindicate the semantic logicist commitment of the classical project of analysis. That is, I want a way of picking out vocabulary as distinctively *logical* that explains why it is legitimate to use *logical* vocabulary, but perhaps not other kinds of vocabulary, in showing how the meanings expressed by one vocabulary can be elaborated out of the meanings expressed by another. My idea is that it is precisely whatever is an LX-vocabulary with respect to every autonomous vocabulary that can legitimately be appealed to as an auxiliary elaborating vocabulary in semantic analysis. Semantic logicism is then justified because logical vocabulary is such a universal pragmatically elaborated-explicitating vocabulary.

What conditions must a vocabulary meet in order for it to serve as a suitable auxiliary in the project of semantically analyzing one vocabulary in terms of another, in a suitably broad sense of ‘analyze’? They would seem to be of two different kinds, which pull against one

another so as to create what we might call the “logician’s dilemma” (looking over our shoulders at Hempel’s famous article, “The Theoretician’s Dilemma,”⁸ which makes a cognate point). On the one hand, the auxiliary vocabulary should not add anything distinctive of its own. For any substantive content that it contributes impugns the claim that what can be said in the terms of the target vocabulary is somehow present already in the base vocabulary or at least is intelligible entirely in terms of what that vocabulary expresses. The general thought is that the auxiliary vocabulary can help *articulate* what is expressed in another way by the base vocabulary, can in some sense draw out what is *implicit* in it. And the claim that the version of what is expressed in the target vocabulary that emerges from that process really was already implicit in the base vocabulary, needing only further articulation to become visible as such, is undercut insofar as further new content is imported by the auxiliary vocabulary. Perhaps the semantic content in question is really implicit only in the auxiliary vocabulary, or only in the two when put together. There is no interest to the claim that culinary vocabulary supervenes, for instance, on chemical vocabulary, if it turns out we mean that it does so if we can also help ourselves to the vocabulary of home economics as an auxiliary in securing that relation. I will call this the requirement of “*semantic transparency*” on the suitability of a vocabulary to serve as an auxiliary in the analytic enterprise, that is, as playing the role there traditionally assigned to logical vocabulary.

Something like this line of thought, I think, played a significant role in twentieth-century philosophical approaches to the demarcation of logical vocabulary that appealed to its *formality*. For if logic were purely formal, concerned only with the form of thought and not its content, then logical vocabulary itself would have to be devoid of content, and hence semantically transparent in exactly the way required for it to play the role in semantic analysis that logicism assigns it.

⁸ “The theoretician’s dilemma”, reprinted in Carl G. Hempel *Aspect of Scientific Explanation* [Free Press, 1970].

Tempting as that line of thought is, it has become increasingly difficult to maintain or pursue it.⁹ For how is the contrast between semantic form and content supposed to be drawn, so as to underwrite criteria of demarcation for logical vocabulary? We are no longer in a position to subscribe to the collateral commitments required to appeal to Aristotelian or scholastic hylomorphism, nor to Kant's transcendental faculty-hylomorphism (which are the home language-games of this distinction). And we know that a purely syntactic characterization of form, of Hilbert's sort, cannot do the semantic job we are addressing (even for the target vocabulary of arithmetic). Frege, of course, never thought of logic as purely formal, or of purely logical concepts as contentless.

The closest we come, I think, is the Frege-Bolzano identification of formality with semantic invariance under substitution, of the sort Quine has championed. To say an inference is good or a claim true in virtue of its logical form is to say two things: that it is good or true, and that it remains good or true upon arbitrary substitution of non-logical for non-logical vocabulary. This is indeed a sufficiently crisp concept to be useful for many purposes. But it is important to realize that this notion of form and formality is of no help in the context of the question about semantic logicism we are addressing—and that for two connected reasons. Notice first that it does not *provide*, but rather *presupposes* a criterion of demarcation of logical vocabulary. Unless we can *already* distinguish logical from non-logical vocabulary, we cannot begin to apply the test. And that means, as Frege himself insisted, that the fact that an *inference* is *good* or a *claim true* in virtue of its logical form in this sense is not at all incompatible with its substantive contentfulness. For we can pick any vocabulary we like to privilege substitutionally: an

⁹ For a deep and thought-provoking extended discussion of this issue, see John MacFarlane's *What Does it Mean to Say that Logic is Formal?* [Pittsburgh Ph.D. dissertation, 2000].

inference is good and a claim true in virtue of its *theological* or *geological* form just in case it is good or true and remains so under all substitutions of *non-theological* for *non-theological* vocabulary, or *non-geological* for *non-geological* vocabulary. Theological and geological formality will not just depend upon, but will express an important aspect of the *content* of theological and geological concepts. Because the substitutional notion of formality is promiscuous about the vocabularies it applies to, grinding out a notion of formality for each, it leaves untouched our original question: how to justify the claim that *logical* vocabulary, however demarcated, can play the role of expressive auxiliary in the philosophical project of wholesale semantic analysis of some vocabularies in terms of others, at least insofar as it satisfies the condition of semantic transparency.

The criterion of adequacy complementary to semantic transparency is that a candidate expressive auxiliary vocabulary in such a semantic project must be *analytically efficacious*. That is, using it must *help* in the process of establishing the desired semantic relation between vocabularies—whether that is definability, translatability, reducibility, supervenience, or whatever. After all, it was to be the powerful modern methods of Frege’s *Grundgesetze* and Russell and Whitehead’s *Principia* that within one core program of analysis were to make the difference between *traditional* empiricism and its 20th century successor in *logical* empiricism. For now we had much more powerful ‘glue’ available to stick together and articulate what is expressed by the favored base vocabularies, be they phenomenological, secondary-quality, or observational. And, looking ahead, after the subsequent *modal* logical revolution in the last third of the century, it is the power of modal vocabulary to serve as a mediating expressive auxiliary vocabulary that some philosophers hope will usher in the new age in semantic naturalism.

The logicist's dilemma—or perhaps we should say, challenge—is to explain how logical vocabulary (or any vocabulary) can be at once semantically transparent and analytically efficacious: how it can remain sufficiently modest and unassuming to be eligible while still being in a position to make a substantive contribution to the analytic semantic expressive enterprise. The potential tension between these complementary demands becomes evident when we reflect that even if we did have available a notion of the mere formality of logic that could satisfy us as to the transparency requirement, it would still be very hard to see how to understand its efficacy—how contentless concepts could help us understand contentful ones. I think the attempt to reconcile some version of these two demands has been the motive force behind a fair amount of wriggling in the philosophy of logic over the last hundred years or so. (One might think here about Wittgenstein's early agonizing about the “purity” of logic.)

The claim I want to defend is that being an LX-vocabulary with respect to every autonomous vocabulary is necessary and sufficient for playing the privileged role logical vocabulary is called on to play in the classical project of semantic analysis. I have been saying that showing that requires simultaneously satisfying two principal criteria of adequacy. For the first, the fact that practices sufficient to deploy logical vocabulary can be algorithmically elaborated from practices necessary to deploy any autonomous vocabulary vindicates the *semantic transparency* of logical vocabulary. For that means that anyone who can talk at all, hence can deploy any base vocabulary, can already *do* everything one needs to be able to *do* in order in principle to be able to *say* what logical vocabulary lets one *say*. The capacity to deploy logical vocabulary (or any universally LX-vocabulary) is in this sense always already *implicit* in the capacity to deploy any vocabulary at all that might be chosen to serve as the base vocabulary

of a semantic analysis or explication of any target vocabulary (whether those appropriate to empiricist, naturalist, functionalist, or any other sort of analysis). And the notion of algorithmic elaboration gives a definite sense to the claim that the one set of abilities is *in principle* sufficient for the other. This is the sense in which deploying logical vocabulary requires nothing new on the part of discursive practitioners: anyone who can use any base vocabulary already knows how to do everything needed to deploy any universal LX-vocabulary.

For the other criterion of adequacy, the fact that logical vocabulary makes explicit features of practices PV-necessary to deploy any autonomous vocabulary shows why and how logical vocabulary satisfies the condition of *analytic efficacy*. Here the main point is that the task for which logical vocabulary must prove itself efficacious is an *expressive* task: to show how to *say* in a *different* vocabulary what can already be said using the target vocabulary. Logical vocabulary must make it possible to *say* something one could not *say* without it. Actually introducing this explicating vocabulary and applying it (*using* the expressive resources that turn out to be implicit already in the capacity to deploy other vocabularies) is what one must *do* in order to do that: when things go well, to say in the base vocabulary, together with the logical, everything one could already say using the target vocabulary. Logical vocabulary gives one the capacity to *talk about* the inferential articulation of any other vocabulary—about what follows from what—which is an essential part of what makes that other vocabulary express the semantic content that it does. On this account, semantic logicism is a commitment to the effect that being able to make explicit the relations of material inference and incompatibility that are implicit in the use of target and base vocabularies is a critical element in establishing any relation between them that could count as a semantic *analysis* or *explication* of one in terms of the

other—according to whatever more specific model of that one applies: translation, reduction, etc.. It is to this *expressive* project that the *explicitating* feature of logical vocabulary (the fact that, as a universally LX-vocabulary, it is VP-sufficient to *specify* some practices PV-necessary to deploy any vocabulary) makes its distinctive contribution.

Here is a way of thinking about this vindication of semantic logicism in terms of the fact that (as I have claimed) logical vocabulary stands to any autonomous vocabulary in the complex, pragmatically mediated semantic relation of being both elaborated from and explicative of practices necessary to deploy that vocabulary. The logicist’s dilemma arises if we think in purely *semantic* terms, because the first constraint (transparency) seems to require logical vocabulary to be *contentless*, while the second (efficacy) seems to require it to be *contentful*. Put that way, the conflict between the two demands is going to be hard to resolve. But if we transpose the issue into a pragmatic key, and look at what one must be able to *do* in order to *say* various things (deploy an autonomous vocabulary, deploy logical vocabulary), we see that the resolution requires only the distinction between potentiality and actuality. Transparency arises from the pragmatic fact that one is already in principle *able* to do everything one must do to deploy logical vocabulary, just by being able to talk at all. That this *pragmatic potential* was always already in place is the sense in which “nothing is added” by logical vocabulary. On the other hand, *actualizing* that potential, by introducing and deploying, actually using and applying, logical vocabulary, *does* involving *doing* something new: something one *could in principle* have done before, but had not *actually* done. The dilemma or paradox is resolved by distinguishing between what one *could* (already) do—in a sense made definite by the notion of algorithmic elaboration—and what one *actually does*. At the pragmatic level, nothing more is needed than

the conceptual apparatus Aristotle introduced so long ago (at the dawn of metaphysics) to resolve another apparent dilemma.

The semantic relation between LX-vocabularies and the vocabularies from which they are elaborated is an essentially *pragmatically mediated* one. If that relation really is the one that explains and justifies the utility of logical vocabulary in semantic analysis, then *all* elaboration and explication, including *semantic* elaboration and explication, is implicitly *pragmatically mediated*. For in order to explain the legitimacy of appealing to *logical* vocabulary in the semantic elaboration of one vocabulary into another, we must appeal to the *pragmatic* elaboration of one set of *abilities* into another. If that is right, then supplementing the traditional philosophical analytical concern with relations between the *meanings* expressed by different kinds of vocabulary by worrying also about the relations between those meanings and the *use* of those vocabularies in virtue of which they express those meanings—as I recommended in my first lecture—is not so much *extending* the classical project of analysis as it is *unpacking* it, to reveal a pragmatic structure that turns out already to have been implicit in the semantic project all along. For the conclusion I have been arguing for is that it is *because* some vocabularies are universal pragmatically elaborated and explicating vocabularies that semantic analysis in the twentieth-century logicist sense is a coherent enterprise at all. That strong claim is a central element of a further kind of pragmatism about semantic analysis: what I will call *analytic pragmatism*, whose principal tool is meaning-use analysis.

My primary concern here has been to use the metaconceptual apparatus of meaning-use analysis first to make clear and then to argue for the ubiquity and theoretical centrality of

pragmatically mediated semantic relations, by appeal to the idea of universal elaborated-explicitating vocabularies. Subsequent lectures will consider other philosophically important vocabularies, especially *modal*, *normative*, and *intentional* vocabularies, which I will argue can be understood as having this same general kind of meaning-use structure. In particular, in my fourth lecture I will argue that the addition of *modal* vocabulary to the semantic logicist's toolkit—the hallmark of the modal revolution of the last third of the twentieth century—can and should be vindicated in just the same way that I have suggested for ordinary logical vocabulary—though of course the features of discursive practice from which its use can be elaborated and which it in turn expresses explicitly are different. In my third lecture, building on our discussion of practical algorithmic elaboration, I shall begin an investigation of what I will urge should be considered a third core program of the classical project of philosophical analysis: functionalism about intentionality. Here, too, I will try to show, what is primarily at issue should be understood in terms of pragmatically mediated semantic relations. In particular, I will offer a new way of thinking about the commitments involved in the program of artificial intelligence, in terms of the possibility of algorithmically elaborating a set of nondiscursive practices-or-abilities into one that is PV-sufficient to deploy an autonomous vocabulary.

END

Appendix to Chapter Two: Indexical Vocabulary: Semantics and Pragmatics

In the body of this lecture I introduced the notion of *algorithmic elaboration* as a way of giving definite shape to the notion of PP-sufficiency, offered some reasons to think that it plays a role in *pragmatics* analogous to that played by *logic* in *semantics*, and suggested further that we might exploit that analogy so as to understand the privileged role that logical vocabulary plays in the analytic *semantic* enterprise in terms of the elaboration of practices-or-abilities on the side of *pragmatics*. Here I want to lay out, alongside this intimation of what we might learn about *logical* vocabulary by thinking of it as a kind of universal LX-vocabulary, a slightly more extended elaboration of a case I introduced in my previous lecture of another case where PP-sufficiency relations in the form of algorithmic elaboration underwrite semantic relations that we understand much better when we think of them as pragmatically mediated. That is *indexical* vocabulary, where I claimed we can find an important example of strict pragmatic expressive bootstrapping. For although we can *show* that indexical vocabulary is not semantically reducible without remainder to non-indexical vocabulary, nonetheless purely non-indexical vocabulary *can* serve as an adequate *pragmatic* metavocabulary for indexical vocabulary. That is, one can *say* (that is, describe), in wholly non-indexical terms, everything one needs to *do* in order to *use* indexical vocabulary. Non-indexical vocabulary is VP-sufficient to specify practices-or-abilities PV-sufficient to deploy indexical vocabulary.

Indexicals exhibit two distinctive sorts of discursive behavior. On the *semantic* side, they are token-reflexive expression types: the *content* any tokening expresses depends on the *context*

in which it is produced. And on the *pragmatic* side, their use can have the special pragmatic significance of making explicit the acknowledgment of a practical commitment. Marking these two dimensions of their use by pairs of names of some of those philosophers who have taught us the most about them, I will say that that indexical vocabulary has both a Kaplan-Stalnaker semantics and an Anscombe-Perry pragmatics. (These are not independent. I think we can show that any expression-types whose tokenings characteristically have the latter sort of pragmatic significance must have their contents determined in the Kaplan-Stalnaker way.)

A key semantic point we have come to realize about the relations between indexical and non-indexical vocabularies is that for any proposed translation of some indexical expressions into non-indexical terms, it is possible to describe counterfactual circumstances in which the indexical tokening would, and the non-indexical candidate for semantic equivalence with it would not, have the pragmatic significance to which Anscombe and Perry draw our attention. And that means that the indexical expression will not be intersubstitutable with—and hence not semantically equivalent to—the non-indexical one in some modal contexts. As Perry puts it, indexicals are *essential* for the expression of some thoughts. So indexical vocabulary is not semantically reducible to or analyzable in terms of non-indexical vocabulary. I think that is right and important. But that *semantic* irreducibility is best understood in the context of various other important semantic relations between indexical and non-indexical vocabulary that are *pragmatically* mediated.

In this connection, I want to make an uncontroversial observation, and a more surprising claim. The observation is that the Kaplan-Stalnaker semantics of character and content can be

formulated in an entirely non-indexical vocabulary, and further that that account can be transformed into non-indexically formulated rules for *using* expressions with that semantics. The more interesting claim is that the non-indexically specifiable practices that are PV-sufficient for using expressions that exhibit the Kaplan-Stalnaker semantics can be *algorithmically elaborated* into—and hence are PP-sufficient for—practices of using expressions that have the Anscombe-Perry pragmatic significance of *essential* indexicals. If that is right, anyone who knows how to use non-indexical vocabulary already knows how to do everything necessary in principle to use indexical vocabulary. In spite of the semantic irreducibility of the one to the other, no-one can be in the position of understanding or having mastered the use of non-indexical vocabulary without having all the abilities needed to understand or master the use of indexical vocabulary. So although indexical vocabulary exhibiting the Anscombe-Perry pragmatics cannot be *semantically reduced* to non-indexical vocabulary, a non-indexical *pragmatic metavocabulary* for it is possible. This is a *semantic* example of *pragmatic expressive bootstrapping*. Those are the claims that I want to clarify and at least begin to justify.

Different tokenings of the same indexical expression-type can express different contents: Your utterance of ‘I’ picks out a different individual from mine. But those different contents are associated systematically with features of the tokenings—with, as we say, *indices* of the unrepeatable utterances, such as speaker, and time and place and even the possible world in which they are uttered. The basic idea of the Kaplan-Stalnaker semantics for indexicals is that, however we understand or represent the contents expressed by particular tokenings, the semantic interpretants of indexical expression-*types* should be understood not as such contents, but rather

as *functions* from tokening-indices *to* such contents.¹⁰ To compute the time referred to by a tokening of the temporal indexical type “last Wednesday”, one must know both what on what day that particular, unrepeatable utterance-event was produced and what function is associated with the repeatable type (Kaplan’s ‘character’) “last Wednesday”.

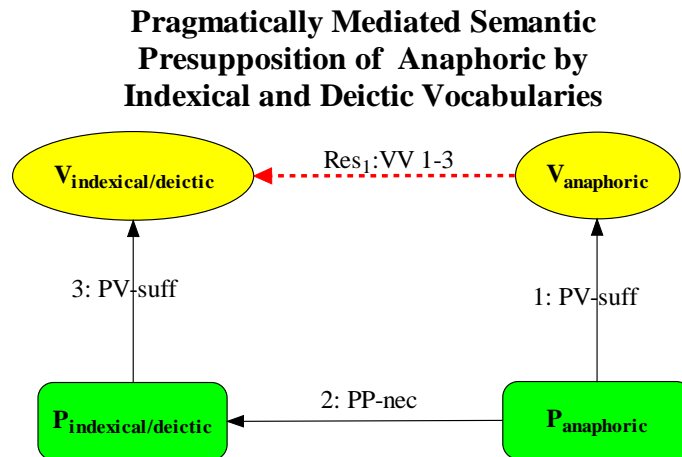
Kaplan thinks that this same semantic apparatus can be applied to demonstratives—that is, that demonstratives are a kind of indexical. To do that, one must think of *demonstrations*, which make some feature of the environment semantically salient, as indices that can be associated with utterances. I think this tempting assimilation elides a very important distinction. What I want to call ‘genuine’ semantic indices are features of utterances that can be read off without knowing anything about what the utterance means. Time, place, speaker, and possible world are properties of tokenings that can be settled and specified before one turns one’s attention to the content expressed by those tokenings. By contrast, virtually *any* feature of an utterance can be used to make something semantically salient. Wittgenstein’s discussion reminds us of how important context is in determining what one is pointing at, even in the paradigmatic case of, as it were, literal demonstration. Prior conversation, for instance, can make any property of an utterance or its environment, whether gestural or not, salience-determining. ‘Demonstration’ as it has to function to serve as an index is not a natural kind, but a characterization of what plays a certain kind of *semantic* role. Deciding what the demonstration is in most cases requires deciding what is demonstrated. It is not, as time, place, and speaker are, something that can be determined independently of that semantic value, and

¹⁰ Here I am abstracting from how Kaplan and Stalnaker combine this traditional way of thinking about indexicals with their functional, possible-worlds approach to content—which is to say, the particulars of their contribution to this tradition.

then appealed to as an input from which the value could then be computed by a character-function.

Kaplan's notion of a semantic index as the input to a character-function captures one way in which features of the use of expressions can contribute to their meaning. Following Stalnaker, he thinks of this phenomenon as exhausting the relevance of pragmatics to semantics. As I have been using the term 'pragmatics', of course, the contribution of the phenomena it studies to semantics is not restricted to token-reflexive expressions such as indexicals and demonstratives. But even within that restricted realm, there is a lot more to the pragmatics of these expressions than the character-relativity of content. As I have argued at length in Chapter Seven of *MIE*, for instance, the use of indexical and demonstrative vocabulary presupposes the use of *anaphoric* vocabulary. An utterance qualifies as cognitively significant and semantically contentful only if it can serve as a premise in inferences. For that reason, securing reference requires making possible non-accidental *co*-reference. For unshared proper names and uniquely identifying definite descriptions, identity of lexical type can guarantee the semantic repeatability of tokenings that supports inferences of the form: *P* of *a*, so *Q* of *a*. What makes it possible for different speakers, at different times and places, and in different situations to draw corresponding conclusions from otherwise unrepeatable indexical and demonstrative utterances is just the possibility of picking up those references anaphorically, by using pronouns, so as to say things like: "If that had been a raccoon you heard in the garden last night, we would see its tracks in the snow there now." Anaphoric uses accordingly come as part of an indissoluble practical package along with indexical and deictic ones, which would otherwise be wholly idle semantically. The underlying relation is one of *pragmatic dependence*: a PP-necessity relation, which induces a

corresponding resultant semantic relation between the vocabularies deployed by the practices-or-abilities that stand in the relation of pragmatic dependence. The meaning-use diagram for this situation is:



(Notice that this resultant MUR is the *retraction* of the PV-sufficiency relation 1 through the other two basic MURs.)

It is clear that one can state Kaplan-Stalnaker rules for associating contents with tokenings of indexical expression-types in wholly non-indexical terms. Temporal and spatial indices can be identified by non-indexical co-ordinates based on any arbitrary origin and using any arbitrary units. Then we can state rules such as:

S) A tokening of ‘It is raining here now,’ uttered at place $\langle x,y,z \rangle$ and time t , is true iff it is raining at $\langle x,y,z \rangle$ at t .

Such rules associate with each indexical character a function from non-indexically specified circumstances of utterance and circumstances of evaluation to truth values. Rules like these show that a VV-sufficient characterization of indexical vocabulary can be provided by an entirely non-indexical semantic metavocabulary. Furthermore, Kaplan and Stalnaker do that by extending the very same functional apparatus used originally to provide an extensional semantic

metavocabulary for intensional vocabulary such as modal operators. And it is straightforward to transform these semantic rules into corresponding pragmatic rules for producing or assessing indexical performances. Assuming suitable units,

P) A tokening of ‘It was raining one mile north of here yesterday,’ uttered at place $\langle x,y,z \rangle$ and time t , is correctly assertible iff it was raining at $\langle x+1,y,z \rangle$ at $t-1$.

But Anscombe and Perry have shown that it does *not* follow that indexicals can be eliminated in favor of or reduced to non-indexical terms.¹¹ Although it is true (with some qualifications that are not important here) that in any utterance u containing the terms ‘I’, ‘now’, or ‘here’, ‘I’ refers to the utterer of u , ‘now’ refers to the time of utterance of u , and ‘here’ refers to the place of utterance of u , those non-indexical characterizations of the extensions of the indexical expressions are not in general equivalent to the indexicals. The difference lies not in their semantic extensions, but in their pragmatic significance, that is, in their *use*. Indexicals play a special role in the behavioral economies of their users. They can have a motivational significance that is not reproduced by non-indexical expressions—even those that are *de jure* co-extensional with the indexical ones. That the meeting is starting *now*, that the bear is going to eat *me*, that the treasure is buried *here* are sayings that can immediately move me to *do* things that I need not be similarly motivated to do by the realization that the meeting starts at noon, the bear is going to eat Bob, and the treasure is buried at $\langle x,y,z \rangle$, even if those are the non-indexical specifications of the very indices that characterize my utterance. For any such non-indexical way of picking out the referents in question, we can come up with counterfactual circumstances in which acquiring the non-indexically expressed belief would not have the same practical

¹¹ [ref.s to G.E.M. Anscombe “The First Person”; John Perry “The Essential Indexical”, and David Lewis “Attitudes *De Dicto* and *De Se*”.]

significance as acquiring the indexically expressed one—and this includes candidates such as “the utterer of *u*”, “the time of *u*,” and “the place of *u*”. Although it takes a somewhat outré scenario, it is even possible for me to acknowledge that the maker of *this very statement* has a certain commitment while not taking it that *I* do, in cases, for instance, where I do not believe that I am making that statement.

The special motivational-epistemological role that Perry points to by focusing on the sorts of practical mistakes or failures to respond appropriately to a situation that are intelligible for any non-indexically formulated beliefs, but not for some indexical ones, and that Anscombe points to by focusing on the kind of failures of identification that are *not* intelligible with some indexically formulated beliefs, but are with any non-indexical ones, is, I think, that some indexicals have *immediately commitment-acknowledging uses*. These are uses that stand in reliable causal-dispositional relations to non-linguistic responses and stimuli. Intentional (which is to say rational) agents must be able to exercise abilities reliably to respond to the acknowledgment of practical commitments whose contents we could make explicit with sentences such as “I shall raise my arm *now*,” and “I shall place the cup *here*,” by doing so—by “straightway acting”, as Aristotle put it. (Such practical responses are ‘immediate’ in the sense of ‘non-inferential’, since they terminate not in a saying, but in a doing.) Correspondingly, being a perceiver in the full-blooded conceptual sense of being a subject of perceptual judgments or beliefs requires the ability to respond to some environing non-linguistic situations such as the visible presence of a white cup by acknowledging doxastic commitments whose contents we could make explicit with sentences such as “There is (or I see) a white cup here now.” The special pragmatic significance of some uses of indexicals to which Anscombe and Perry point is

rooted in the expressive role indexicals such as ‘I’, ‘now’, and ‘here’ play in acknowledging practical commitments that are appropriately responded to immediately as stimuli for the production of non-linguistic states of affairs in exercises of intentional agency, and in acknowledging doxastic commitments that are appropriately elicited immediately as responses to non-linguistic states of affairs in exercises of perceptual judgment.

Of course, it is not merely accidental that it is the *same* expression-types that figure in immediately (non-inferentially) responsively acquired commitment-acknowledgments in perceptual inputs such as “The clock reads 5:05 *now*,” and in immediately (non-inferentially) responded-to commitment-acknowledgments in practical outputs such as “I shall start the lecture *now*.” Intentional agency requires the liability of relatively durable states of prior intention (practical commitment) that would be expressed by sentences such as “I shall start the lecture at 5:05,” to *mature* into causally efficacious events of intention-in-action that would be expressed by sentences such as “I shall start the lecture *now*,” which either immediately bring about or simply consist in the agent’s starting to deliver the lecture. The use of the same locution ‘now’ that appears in the expression of the content of the intention-in-action in making explicit the non-inferentially elicited perceptual judgment “The clock reads 5:05 *now*,” makes explicit the necessary mediating role of that observation in connecting the prior intention with the intention-in-action, and so the action. Vocabulary whose expressive job essentially includes connecting perception with action so as to mediate the maturation of standing commitments into immediately practically efficacious ones will exhibit the sort of immunity to errors of misidentification of time, place, and agent Anscombe and Perry emphasize. For that connection

is forged only when the very subject of a perception at a certain time and place herself acts then and there.

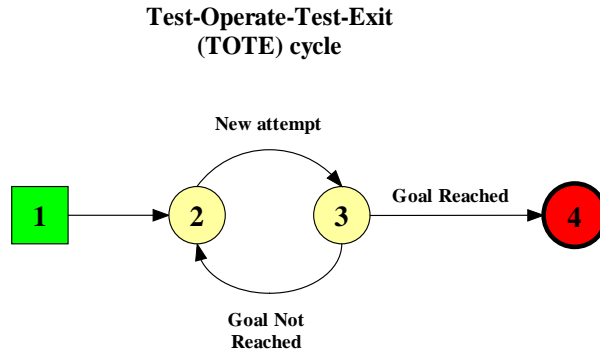
And for this very same reason—as a consequence of playing this same dual observational-practical expressive role—whatever expression-types do make explicit the connections between perception and agency that mediate the maturation of intentions into actions must be characterized by the Kaplan-Stalnaker semantics for expressions like ‘I’, ‘now’, and ‘here’. That is, they must be what might be called ‘context-homogeneous’ indexicals: types whose tokenings have as their extensions at every context of evaluation whatever index characterizes their context of utterance. These contrast, for instance, with context-heterogeneous indexical types, where the time/place/person etc. referred to is not identical to that associated with the utterance, such as ‘yesterday’, ‘two miles north of here’, and ‘my wife’. (They contrast still more with *cross-indexicals* such as time/place/person etc. referred to is not identical to that associated with the utterance, such as ‘yesterday’, ‘two miles north of here’, and ‘my wife’, which are in turn a special case of the even more general category of *cross-sortal* indexicals, such as ‘my mother’s favorite color,’ ‘the dog who lives here,’ and ‘the band that played the festival last year.’ The differences between these kinds of indexicals resides entirely in *how* the character tells us to compute the content from the indices determined by the circumstances of utterance.) Expression-types that mark their tokenings as available for having the pragmatic significance of expressing the immediate acknowledgment of commitments, both doxastic and practical, must function semantically as context-homogeneous Kaplan-Stalnaker indexicals because the maturation of prior intentions into actions requires that the subject, time, and place (and for that matter, world) specified in the contents of the immediately acknowledged

observational and practical commitments, which engage with the content of the standing prior conditional intention, be identical with the subject, time, and place indices of the tokenings which are (when all goes well) events that are immediately causally efficacious in bringing about the intended action under the conditions envisaged. So having the Kaplan-Stalnaker semantics for context-homogeneous indexical expression-types is necessary, though not sufficient, for tokenings of those types to be lexically marked as candidates for having the distinctive Anscombe-Perry pragmatic significance.

Expressing the content of my observation as that it is 5:05 *now* is supposed to mark that in the context of a standing intention to start the lecture at 5:05, that observation should *immediately* be responded to by the practical intention-in-action whose causal efficacy is similarly marked by having its content expressed as “I shall start the lecture *now*.” When ‘now’ is used to express this role in the behavioral economy of the subject, it carries with it immunity to certain kinds of misidentification of the time at which the agent is committed to start the lecture. We might question this theoretical claim on empirical grounds. Is it not intelligible that I should acknowledge that I must start the lecture now, and still not be moved to do so? It is, but this possibility is of a different kind from that to which Anscombe and Perry point. It is a special kind of weakness of the will: *indexical akrasia*. It depends on a certain kind of breakdown in the most basic mechanisms of agency: those that connect acknowledgment of a commitment to do something to doing it, those that permit the maturation of intentions into actions. Such a breakdown is intelligible, but it is a failure of rational agency. Anscombe’s and Perry’s observations concern mistakes that are possible even for fully rational agents, who may fail to start the lecture at 5:05 because they do not realize that it is *then* 5:05, may not respond to threats

to or opportunities for the speaker because they do not realize that *they themselves* are the speaker, and so on.

We can also ask the converse question: Can there be tokenings of *non-indexical* expression-types—that is, ones that do *not* have the Kaplan-Stalnaker semantics characteristic of context-homogeneous indexicals such as ‘I’, ‘now’, and ‘here’—that *do* have the pragmatic significance of immediate acknowledgments of practical and doxastic commitments? It may seem that this question has already been asked, and answered in the negative. But what I argued for above was the claim that if there are expression-*types* that mark their tokenings as candidates for expressing immediate acknowledgment of commitments, they must have the semantics of context-homogeneous indexicals. The current question is whether there must be such types of expression. This is a subtle and difficult question, but I think the right answer is that there *need not* be. Surely, any autonomous discursive practice must make possible the undertaking of standing practical commitments, which must be liable to maturing into intentional doings should the conditions articulated in their contents be observed to be fulfilled. Absent that possibility, practitioners would not be intelligible as agents, and hence not as speakers. Furthermore, it must be possible for actions to be controlled, and not just ballistic (“fire and forget”). That is, it must be possible for agents at least sometimes to adjust what they are doing on the basis of assessments of how successful current attempts are at reaching the desired goal. In automaton-theoretic terms, even so simple an action as reaching for a doorknob must be specified as a Test-Operate-Test-Exit (TOTE) cycle, in which each incremental movement is observed, checked against its approach to the goal, and then followed by another movement calibrated by the results of the prior one, until the goal is reached.



Observation and action both require the immediate acknowledgment of commitments: acknowledgments that can cause non-linguistic performances practically, and be caused by non-linguistic events perceptually. The maturation of prior, standing intentions into intentions-in-action, which are practically efficacious, as we say “here and now” requires that the contents of the commitments acknowledged doxastically in perception and practically in action can sometimes coincide.

But it does *not* follow from this requirement that every discursive practice must include repeatable lexical *types* an essential part of the expressive role of which is to mark their tokenings as fit to express the contents of those immediate commitment-acknowledgments that mediate the cycle of perception and action.¹² It is perfectly intelligible that in some linguistic community agents with a commitment to begin the lecture at 5:05 should be able to respond to clocks by reporting “The clock reads 5:05,” and straightway beginning the lecture. *We* would be right to describe such a person as having a ‘now’ belief, and in expressing the content of the intention-in-action that informs what he does by the use of ‘now’, even if he would not express it that way. For him to be able to act, it is enough that he can respond to the observed fulfillment of the conditions of the prior intention by acting when they are fulfilled, even if he only has

¹² Compare this to the possibility of there being token-token identities (whether ontological or semantic) between vocabularies that are not the result of type-type identities.

available non-indexical specifications of them. Considering a related point may make the situation clearer. Gareth Evans argues convincingly that the capacity to navigate and re-identify physical objects in a spatial environment requires the practical ability to map egocentric space onto public space. It is natural for us to describe this ability in terms of mappings of indexical onto non-indexical specifications of objects and places. But Frisbee-catching dogs and successful predators evidently display the underlying abilities without being able to deploy indexical vocabulary.

Consider, then, a community that accords some tokenings of non-indexical types the pragmatic significance of immediate acknowledgments of commitments, both practical and doxastic, but that lacks expression-types that mark their tokenings as distinctively liable to play that role. Its practitioners have in principle the capacity to use expressions with the Kaplan-Stalnaker semantics of context-homogeneous token-reflexives. For we have seen that one can specify rules for doing so entirely in non-indexical vocabulary. The capacity to use expressions in these ways is intelligible independently of the capacity to use those expressions to mark the immediate acknowledgment of commitments. That is the central point Anscombe makes by describing a community whose members use tokens of the type 'A' token-reflexively, each to refer exclusively to him- or herself, but without according such uses the pragmatic significance characteristic of our use of 'I'. There is no way to *derive* that pragmatic significance just from the semantics of context-homogeneous indexicals. Nonetheless, in virtue of their capacity to accord use tokens of non-indexical types as immediately commitment-acknowledging, practitioners who in addition had learned to use some expression-types according to the rules for context-homogeneous Kaplan-Stalnaker characters-and-contents know how to do everything

necessary in principle to use those latter expressions *as* pragmatically immediately commitment-acknowledging. That is, the abilities they already have can be pragmatically elaborated into the ability to use pragmatically *essential* indexicals: context-homogeneous indexicals with the full-blooded pragmatic significance of immediately acknowledging practical and doxastic commitments.

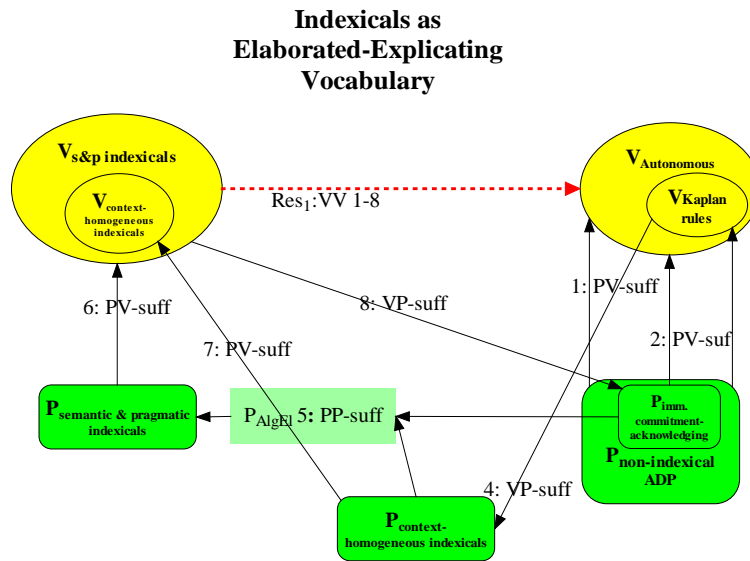
Indeed, those primitive abilities can be *algorithmically* elaborated into the ability to use ‘I’, ‘now’, and ‘here’ with *both* the Kaplan-Stalnaker semantics *and* the Anscombe-Perry pragmatics characteristic of *our* uses of those indexicals. The principal move, for instance, required for Anscombe’s ‘A’ users to become ‘I’ users in the sense she cares about, is an exercise of the algorithmically elaborating ability I earlier called ‘*response substitution*’. Where one would have responded with a non-indexical immediate commitment-acknowledgment in perception, one must be able to respond by producing a different performance, but of a kind one is already capable of: producing a token of a Kaplan-Stalnaker context-homogeneous indexical expression-type. And where before one responded *to* a non-indexically expressed immediate commitment-acknowledgement by acting, now one must produce that response instead to something one could already respond differentially to, namely the use of a Kaplan-Stalnaker context-homogeneous indexical (in each case, one that meshes appropriately with standing practical commitments). These new differential responsive dispositions will then have to be integrable into the states one is already capable of entering in response to the acquisition of a prior intention (a standing practical commitment), by arbitrary state-formation and state-permutation.¹³ My claim is that those algorithmic elaborating abilities are *all* that is needed to

¹³ I have laid out in greater detail how I think this process might work for the particular case of Anscombe’s ‘A’ language in Section V:2 of Chapter Eight of *MIE*.

turn the capacity immediately to acknowledge practical and doxastic commitments involving non-indexicals and the non-indexically specifiable ability to use expressions with context-homogeneous semantic characters into the ability to use indexicals with *both* the semantics and the pragmatic significance characteristic of Anscombe and Perry's *essential* indexicals.

The controversial claim here is that it is intelligible that a genuinely discursive community, which must accord some tokenings the significance of being immediate acknowledgments of practical and doxastic commitments, might nonetheless lack expression-*types* that lexically mark the liability of their tokenings to have that significance. I take it to be clear that *if* there could be such a discursive community, *then* we can see how its practices-and-abilities are in principle sufficient—via algorithmic practical elaboration—for the introduction of expression-*types* with the pragmatics, and *therefore* the semantics, characteristic of context-homogeneous indexicals. If that is right, then such indexical vocabulary stands to the use of non-indexical vocabulary as pragmatically *elaborated* and semantically *explicitating*. That is, it stands to non-indexical vocabulary in the same LX pragmatically mediated semantic relation that conditionals stand in to non-logical sentences related by proprieties of material inference. For though the observation that the clock reads 5:05 might (in the context of the a background commitment to start the lecture at 5:05) implicitly have the pragmatic significance of acknowledging a commitment to start the lecture, expressing it as the observation that it is 5:05 *now* in that context explicitly acknowledges the commitment to start the lecture *now*. For that is, as the mere commitment to start the lecture at 5:05 is not, the form in which intentions-*in-action* are expressed. There can be such intentions without a mode of expression dedicated to marking them out, but when they are available, it is context-homogeneous *indexical* expression-*types* that

make explicit these causally efficacious-and-elicitable acknowledgments of practical commitments.



Here we have a live and significant semantic example, for a vocabulary of independent philosophical interest, of pragmatic expressive bootstrapping. Anscombe and Perry show that the full expressive power of indexical vocabulary cannot be duplicated with non-indexical vocabulary. No non-indexical expression will behave *semantically* the way ‘I’, ‘now’, and ‘here’ do in sentential contexts that include operators that make explicit the situations their pragmatic observations point to: operators such as “It is possible that S rationally believes that... but not that...” But in spite of this *semantic* irreducibility, non-indexical vocabulary is *pragmatically* sufficient for the *use* indexical vocabulary. The Kaplan-Stalnaker semantics permits the formulation in non-indexical vocabulary of rules VP-sufficient to specify how to use expression as context-homogeneous indexicals, and those capacities, together with the ability immediately to acknowledge practical and doxastic commitments whose contents are characterized in non-indexical terms, can be pragmatically elaborated—using only *algorithmic* elaborating abilities

(principally response substitution)—into the capacity to use indexicals in the full-blooded sense. The intermediate appeal to the Kaplan rules as a pragmatic metavocabulary makes this a more complicated case than that of conditional introduction, but what results is recognizably an instance of the form of pragmatically mediated semantic relation between vocabularies I have called ‘LX’.

Why does this matter? What does this meaning-use analysis tell us? It shows us that and how anyone who knows how to use non-indexical vocabulary already knows how to do everything she needs to, in principle (a qualification we can cash out *precisely*, in terms of algorithmic elaborative abilities), to deploy indexical vocabulary. So one could never be in the position of understanding non-indexical vocabulary, but being mystified by indexical vocabulary. Even though what is expressed by the latter cannot be fully expressed by the former, there is nothing about the use of indexicals that must remain hidden from the view of one who professes to know his way around only the non-indexical fragment of a language. And it seems to me that the interest of this result is not hostage to the most controversial claim I made along the way. For even if the claim that there could be an autonomous discursive practice that does not include the use of indexical expression-types¹⁴ is incorrect, I have indicated how the non-indexical *fragment* of an autonomous discursive practice suffices, as a pragmatic metalanguage (the Kaplan rules for using context-homogeneous indexicals) and via algorithmic elaboration (of

¹⁴ Notice, however, that I did not extend that claim to *demonstratives*, which I urge should be put in a different semantic box from indexicals. Including demonstratives in the fragment of a language counted as non-indexical is not cheating in this argumentative context, for two reasons. First, the Anscombe-Perry phenomena show that indexicals are not reducible to demonstratives either. ‘I’ does *not* mean ‘the utterer of this very sentence,’ or anything in the vicinity, as can be seen when we look at the very sort of counterfactuals they consider. Second, I did not appeal to demonstratives in pragmatically reconstructing the use of indexicals from non-indexical ingredients.

non-indexical, implicitly immediately commitment-acknowledging uses into indexical, explicitly immediately commitment-acknowledging ones), for the full-blooded use of indexicals.

It may be that no-one actually boggles at indexical vocabulary in a way that would make this analysis dispositive of an urgent problem rather than just generally enlightening. Indexical skepticism is certainly not widespread. (And it is hard for me to see how this account will help those theologians who really are worried about the apparent inability of a non-spatiotemporal deity to think indexical thoughts). The analytic apparatus I introduced in the first lecture, and developed and applied to a live example in this one, can be thought of as “tooling up” for the discussion of *modal* and *normative* vocabularies in the next couple of lectures: cases where both empiricists and naturalists have had good cause for genuine puzzlement.