

1

Extending the Project of Analysis

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

I The classical project of analysis

I think of analytic philosophy as having at its center a concern with semantic relations between what I will call ‘vocabularies’. Its characteristic form of question is whether, and in what way, one can make sense of the meanings expressed by *one* kind of locution in terms of the meanings expressed by *another* kind of locution. So, for instance, two early paradigmatic projects were to show that everything expressible in the vocabulary of number theory, and again, everything expressible using definite descriptions, is expressible already in the vocabulary of first-order quantificational logic with identity.

2 BETWEEN SAYING AND DOING

The nature of the key kind of semantic relation between vocabularies has been variously characterized during the history of analytic philosophy: as analysis, definition, paraphrase, translation, reduction of different sorts, truth-making, and various kinds of supervenience—to name just a few contenders. In each case, however, it is characteristic of classical analytic philosophy that *logical* vocabulary is accorded a privileged role in specifying these semantic relations. It has always been taken at least to be *licit* to appeal to logical vocabulary in elaborating the relation between *analysandum* and *analysans*—target vocabulary and base vocabulary. I will refer to this aspect of the analytic project as its commitment to ‘*semantic logicism*’.¹

[FN:1]

If we ask *which* were the vocabulary-kinds whose semantic relations it was thought to be important to investigate during this period, at least two *core programs* of classical analytic philosophy show up: *empiricism* and *naturalism*. These venerable modern philosophical traditions in epistemology and ontology respectively were transformed in the twentieth century, first by being transposed into a *semantic* key, and second by the application of the newly available *logical* vocabulary to the self-consciously semantic programs they then became.

As *base* vocabularies, different species of *empiricism* appealed to phenomenal vocabulary, expressing how things appear, or to secondary-quality vocabulary, or, less demanding, to observational vocabulary. Typical *target* vocabularies include objective vocabulary formulating claims about how things actually are (as opposed to how they merely appear), primary-quality vocabulary, theoretical vocabulary, and modal, normative, semantic, and intentional vocabularies. The generic challenge is to show how what is expressed by the use of such target vocabularies can be reconstructed from what is expressed by the base vocabulary, when it is elaborated by the use of logical vocabulary.

As *base* vocabularies, different species of *naturalism* appealed to the vocabulary of fundamental physics, or to the vocabulary of the natural sciences (including the special sciences) more generally, or just to objective descriptive vocabulary, even when not regimented by incorporation into explicit scientific theories. Typical targets include normative, semantic, and intentional vocabularies.

¹ In this usage, the logicism about mathematics characteristic of Frege’s *Grundgesetze* and Russell and Whitehead’s *Principia* is semantic logicism about the relations between mathematical and logical vocabularies.

2 The pragmatist challenge

What I want to call the “classical project of analysis,” then, aims to exhibit the meanings expressed by various target vocabularies as intelligible by means of the logical elaboration of the meanings expressed by base vocabularies thought to be privileged in some important respects—epistemological, ontological, or semantic—relative to those others. This enterprise is visible in its purest form in what I have called the “core programs” of empiricism and naturalism, in their various forms. In my view, the most significant conceptual development in this tradition—the biggest thing that ever happened to it—is the *pragmatist challenge* to it that was mounted during the middle years of the twentieth century. Generically, this movement of thought amounts to a displacement from the center of philosophical attention of the notion of meaning in favor of that of use: in suitably broad senses of those terms, replacing concern with *semantics* by concern with *pragmatics*. The towering figure behind this conceptual sea-change is, of course, Wittgenstein. In characterizing it, however, it will be useful to approach his radical and comprehensive critique by means of some more local, semantically corrosive argumentative appeals to the *practices of deploying* various vocabularies rather than the *meanings* they express.

Wilfrid Sellars (one of my particular heroes) criticizes the empiricist core program of the classical project of analysis on the basis of what one must *do* in order to *use* various vocabularies, and so to count as *saying* or *thinking* various kinds of things. He argues that *none* of the various candidates for empiricist base vocabularies is practically *autonomous*, that is, could be deployed in a language-game one played though one played no other. For instance, no discursive practice can consist entirely of making non-inferential observation reports, for such reliably differentially elicited responses qualify as *conceptually* contentful or *cognitively* significant only insofar as they can serve as *premises* from which it is appropriate to draw *conclusions*, that is, as *reasons* for other judgments. Drawing such conclusions is applying concepts *inferentially*—that is, precisely *not* making non-inferential *observational* use of them.²

FN:2

² This argument occupies roughly the first half of his classic “Empiricism and the Philosophy of Mind” (1956; reprinted by Harvard University Press, 1997). His critique of the phenomenalist version

4 BETWEEN SAYING AND DOING

[FN:3]

Quine offers an even broader pragmatist objection, not only to the empiricist program, but to essential aspects of the whole analytic semantic project, for he attacks the very notion of meaning it presupposes. Quine is what I have elsewhere called a “methodological” pragmatist.³ That is, he takes it that the whole point of a theory of meaning is to explain, codify, or illuminate features of the *use* of linguistic expressions. He, like Dummett, endorses the analogy: *meaning* is to *use* as *theory* is to *observation*. And he argues that postulating meanings associated with bits of vocabulary yields a *bad* theory of discursive practice.

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

Quine thought one could save at least the naturalist program by retreating semantically to the level of reference and truth-conditions. James and Dewey appeal to the same sort of methodological pragmatism in support of more sweeping sorts of semantic revisionism—pursuing programs that Rorty, for instance, argues should be understood as more rejectionist than properly revisionist. And under the banner “Don’t look to the meaning, look to the use,” Wittgenstein further radicalizes the pragmatist critique of semantics. Pointing out, to begin with, that one cannot assume that uses of singular terms have the job of picking out objects, nor that declarative sentences are in the business of stating facts, he goes on to deny, in effect, that such uses even form a privileged center on the basis of which

of empiricism can be found in “Phenomenalism”, in his collection *Science, Perception, and Reality* (Routledge & Kegan Paul, 1963).

³ See “Pragmatics and Pragmatisms,” in James Conant and Ursula M. Zeglen (eds.), *Hilary Putnam: Pragmatism and Realism* (Routledge, 2002), translated as “Pragmatik und Pragmatismus,” in M. Sandbothe, (ed.) *Die Renaissance des Pragmatismus* (Velbrück Wissenschaft, 2000), 29–58.

EXTENDING THE PROJECT OF ANALYSIS 5

one can understand more peripheral ones. (“Language”, he says, “has no downtown.”)

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

More specifically, Wittgenstein uses the image of “family resemblances” to urge that the *kinds* into which linguistic practices and the vocabularies caught up in them are functionally sorted—what belong together in boxes labeled ‘game’, ‘name’, ‘assertion’, ‘observation’, and so on—do not typically admit of specification in terms of underlying principles specifiable in other vocabularies, whether by genus and differentia(e) or any other kind of explicit rule or definition. It is easy to understand this line of thought as

6 BETWEEN SAYING AND DOING

entailing a straightforward denial of the possibility of semantic analysis in the classical sense.

I think that one thought underlying these observations about the unsystematic, unsurveyable variety of kinds of uses of expressions and about the uncodifiable character of those kinds concerns the essentially *dynamic* character of linguistic practice. I think Wittgenstein thinks that an absolutely fundamental discursive phenomenon is the way in which the abilities required to deploy one vocabulary can be practically *extended*, elaborated, or developed so as to constitute the ability to deploy some further vocabulary, or to deploy the old vocabulary in quite different ways. Many of his thought-experiments concern this sort of process of *pragmatic projection* of one practice into another. We are asked to imagine a community that uses proper names only for people, but then extends the practice to include rivers. There is no guarantee that interlocutors can master the extended practice, building on what they can already do. But if they can, then they will have changed the only 'essence' proper-name usage could be taken to have had.⁴ In the old practice it always made sense to ask for the identity of the *mother* and *father* of the named item; in the new practice, that question is often senseless. Again, we are asked to imagine a community that talked about having gold or silver in one's teeth, and extends that practice to talk about having pain in one's teeth. If, as a matter of contingent fact, the practitioners can learn to use the expression 'in' in the new way, building on but adapting the old, they will have fundamentally changed the 'meaning' of 'in'. In the old practice it made sense to ask where the gold was *before* it was in one's tooth; in the new practice asking where the pain was before it was in the tooth can lead only to a distinctively *philosophical* kind of puzzlement.⁵

At every stage, what practical extensions of a given practice are possible for the practitioners can turn on features of their embodiment, lives, environment, and history that are contingent and wholly particular to them. And which of those developments actually took place, and in what

⁴ Cf. Quine's remark in "Two Dogmas of Empiricism": "Meaning is what essence becomes when it is detached from the thing and attached to the word," (1953; reprinted by Harvard University Press, 2006).

⁵ I am indebted for this way of thinking of Wittgenstein's point to Hans Julius Schneider's penetrating discussion in his *Phantasie und Kalkül: Über die Polarität von Handlung und Struktur in der Sprache* (Suhrkamp, 1992).

EXTENDING THE PROJECT OF ANALYSIS 7

order, can turn on any obscure fact. The reason vocabulary-kinds resist specification by rules, principles, definitions, or meanings expressed in other vocabularies is that they are the current time-slices of processes of development of practices that have this dynamic character—and that is why the collection of uses that is the current cumulative and collective result of such developments-by-practical-projection is a motley.⁶ If that is right, then any codification or theoretical systematization of the uses of those vocabulary-kinds by associating with them meanings that determine which uses are correct will, if at all successful, be successful only contingently, locally, and temporarily. Semantics on this view is an inherently Procrustean enterprise, which can proceed only by *theoretically* privileging some aspects of the use of a vocabulary that are not at all *practically* privileged, and spawning philosophical puzzlement about the intelligibility of the rest.⁷ On this conception, the classical project of analysis is a disease that rests on a fundamental, if perennial, misunderstanding—one that can be removed or ameliorated only by heeding the advice to replace concern with *meaning* by concern with *use*. The recommended philosophical attitude to discursive practice is accordingly *descriptive particularism*, *theoretical quietism*, and *semantic pessimism*.

FN:6

FN:7

3 Extending the project of analysis: pragmatically mediated semantic relations

On this account, Wittgenstein is putting in place a picture of discursive meaningfulness or significance that is very different from that on which the classical project of analysis is predicated. In place of *semantics*, we are encouraged to do *pragmatics*—not in the sense of Kaplan and Stalnaker, which is really the semantics of token-reflexive expressions, nor again in

⁶ A patient and detailed investigation of the mechanisms of this phenomenon in basic descriptive and scientific concepts, and an extended argument for its ubiquity can be found in Mark Wilson's exciting and original *Wandering Significance* (Harvard University Press, 2006).

⁷ I would be happy if those who dance with his texts find affinities here with Hegel's insistence that the metaconceptual categories of *Verstand* must be replaced by those of *Vernunft*. It is characteristic of his philosophical ambition that he draws the opposite of Wittgenstein's conclusions from an appreciation of the dynamics of conceptual development and its sensitivity to arbitrary contingent features of the practitioners, devoting himself to elaborating what he insists is the *logic* of such processes and the conceptual contents they shape.

8 BETWEEN SAYING AND DOING

the sense of Grice, which addresses conversational heuristics in terms that presuppose a prior, independent, classical semantics—but ‘pragmatics’ in the sense of the study of the *use* of expressions in virtue of which they are meaningful at all. To the formal, mathematically inspired tradition of Frege, Russell, Carnap, and Tarski, culminating in model-theoretic semantics, is opposed an anthropological, natural-historical, social-practical inquiry aimed both at demystifying our discursive doings and at deflating philosophers’ systematic and theoretical ambitions regarding them. I think that contemporary philosophers of language have tended to draw this opposition in the starkest possible terms, treating these approaches as mutually exclusive, hence as requiring that a choice be made between them, thereby marking out a substantial sociological faultline in the discipline. Those who are moved by the pragmatist picture generally seem to accept the particularist, quietist conclusions Wittgenstein seems to have drawn from it. And those committed to some version of the project of semantic analysis have often felt obliged to deny the significance of pragmatics in this sense, or at the least to dismiss it as irrelevant to properly semantic concerns. In the most extreme cases, the attitude of anti-pragmatist philosophers of language to Wittgenstein’s picture verges on that of the Victorian lady to Darwin’s theory: one hopes that it is not true, and that if it is true, at least that it not become generally known.

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

EXTENDING THE PROJECT OF ANALYSIS 9

If we leave open the possibility that the use of some vocabulary may be illuminated by taking it to express some sort of meaning or content—that is, if we do not from the beginning embrace theoretical semantic nihilism—then the most important positive pragmatist insight will be one complementary to the methodological pragmatism I have already identified. The thought underlying the pragmatist line of thought is that what makes some bit of vocabulary mean what it does is how it is used. What we could call *semantic* pragmatism is the view that the only explanation there could be for how a given *meaning* gets associated with a vocabulary is to be found in the *use* of that vocabulary: the practices by which that meaning is conferred or the abilities whose exercise constitutes deploying a vocabulary with that meaning. To broaden the classical project of analysis in the light of the pragmatists’ insistence on the centrality of pragmatics, we can focus on this fundamental relation between use and meaning, between practices or practical abilities and vocabularies. We must look at what it is to use locutions *as* expressing meanings—that is, at what one must *do* in order to count as *saying* what the vocabulary lets practitioners express. I am going to call this kind of relation “practice-vocabulary sufficiency”—or, usually, “PV-sufficiency” for short. It obtains when engaging in a specified set of practices or exercising a specified set of abilities⁸ is sufficient for someone to count as *deploying* a specified vocabulary.

FN:8

Of course it matters a lot how we think about these content-conferring, vocabulary-deploying practices or abilities. The semantic pragmatist’s claim that use confers meaning (so talk of practices or the exercise of abilities as deploying vocabularies) reverts to triviality if we are allowed to talk about “using the tilde to express negation,” “the ability to mean red by the word ‘red’,” or “the capacity to refer to electrons by the word ‘electron’,” (or, I think, even *intentions* so to refer). And that is to say that the interest of the PV-sufficiency of some set of practices or abilities for the deploying of a vocabulary is quite sensitive to the *vocabulary* in which we *specify* those practices-or-abilities. Talk of practices-or-abilities has a definite sense only insofar as it is relativized to the vocabulary in which those practices-or-abilities are specified. And that means that besides

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

10 BETWEEN SAYING AND DOING

PV-sufficiency, we should consider a second basic meaning–use relation: “vocabulary–practice sufficiency,” or just “*VP-sufficiency*,” is the relation that holds between a vocabulary and a set of practices–or–abilities when that vocabulary is sufficient to *specify* those practices–or–abilities. VP-sufficient vocabularies that *specify* PV-sufficient practices let one *say* what it is one must *do* to count as engaging in those practices or exercising those abilities, and so to deploy a vocabulary to *say* something.

PV-sufficiency and VP-sufficiency are two basic *meaning–use relations* (MURs). In terms of those basic relations, we can define a more complex relation: the relation that holds between vocabulary V' and vocabulary V when V' is VP-sufficient to specify practices–or–abilities P that are PV-sufficient to deploy vocabulary V . This VV -relation is the *composition* of the two basic MURs. When it obtains I will say that V' is a *pragmatic metavocabulary* for V . It allows one to *say* what one must *do* in order to count as *saying* the things expressed by vocabulary V . We can present this relation graphically in a *meaning–use diagram* (MUD), as shown in Figure 1.1.

The conventions of this diagram are:

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

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




Figure 1.1 Meaning–use diagram 1: pragmatic metavocabulary

EXTENDING THE PROJECT OF ANALYSIS II

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

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

12 BETWEEN SAYING AND DOING

FN:9

course it is not expressed in terms of the machinery I have been introducing—is Huw Price’s pragmatic normative naturalism.⁹ He argues, in effect, that although normative vocabulary is not *reducible to* naturalistic vocabulary, it might still be possible to *say* in wholly naturalistic vocabulary what one must *do* in order to be *using* normative vocabulary. If such a claim about the existence of an expressively bootstrapping naturalistic pragmatic metavocabulary for normative vocabulary could be made out, it would evidently be an important chapter in the development of the naturalist core program of the classical project of philosophical analysis. It would be a paradigm of the sort of payoff we could expect from extending that analytic project by including pragmatically mediated semantic relations.

The meaning–use diagram of the pragmatically mediated semantic relation of being a pragmatic metavocabulary illustrates a distinctive kind of *analysis* of that relation. It exhibits that relation as the resultant, by composition, of the two basic meaning–use relations of PV-sufficiency and VP-sufficiency. A complex MUR is analyzed as the product of operations applied to basic MURs. This is *meaning–use analysis*. The same analytic apparatus applies also to more complex pragmatically mediated semantic relations. Consider one of the pragmatist criticisms that Sellars addresses to the empiricist core program of the classical analytic project. It turns on the assertion of the *pragmatic dependence* of one set of vocabulary-deploying practices-or-abilities on another.

FN:10

Because he thinks part of what one is *doing* in saying how things merely appear is withholding a commitment to their actually being that way, and because one cannot be understood as *withholding* a commitment that one cannot *undertake*, Sellars concludes that one cannot have the ability to say or think how things *seem* or *appear* unless one also has the ability to make claims about how things *actually are*. In effect, this Sellarsian pragmatist critique of the phenomenalist form of empiricism consists in the claim that the practices that are PV-sufficient for ‘is’- ϕ talk are PP-necessary for the practices that are PV-sufficient for ‘looks’- ϕ talk.¹⁰ That pragmatic dependence of practices-or-abilities then induces a resultant pragmatically

⁹ See his “Naturalism without Representationalism,” in Mario de Caro and David Macarthur (eds.), *Naturalism in Question* (Harvard University Press, 2004), 71–90.

¹⁰ I discuss this argument in greater detail in the final chapter of *Tales of the Mighty Dead* (Harvard University Press, 2002).

EXTENDING THE PROJECT OF ANALYSIS 13




Figure 1.2 Meaning–use diagram 2: pragmatically mediated semantic presupposition

mediated semantic relation between the vocabularies. The meaning–use diagram for this claim is shown in Figure 1.2.

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

In fact, although Sellars’ argument for the crucial PP-necessity relation of pragmatic dependence of one set of vocabulary-deploying practices—or-abilities on another is different, his argument against the observational version of empiricism—the claim that purely non-inferential, observational uses do not form an autonomous discursive practice, but presuppose inferential uses—has exactly the same form (Figure 1.3).

For these cases, we can say something further about the nature of the pragmatically mediated semantic relation that is analyzed as the resultant MUR in these diagrams. For, instead of jumping directly to this VV resultant MUR, we could have put in the composition of the PP-necessity




Figure 1.3 Meaning–use diagram 3: pragmatically mediated semantic presupposition

14 BETWEEN SAYING AND DOING




Figure 1.4 Meaning-use diagram 4: composition

and second PV-sufficiency relation, yielding a kind of complex pragmatic presupposition (Figure 1.4).

If this diagram were completed by an arrow from $V_{is-\phi}$ to $V_{looks-\phi}$ such that the *same* diagonal resultant arrow could represent *both* the composition of relations 2 and 3 *and* the composition of relation 1 and the newly supplied one, then category theorists would say that the diagram *commutes*. And the arrow that needs to be supplied to make the diagram commute they call the *retraction* of relation 1 through the composition Res_2 (Figure 1.5).

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




Figure 1.5 Meaning-use diagram 5: composition and retraction

4 Automata: Syntactic PV-sufficiency and VP-sufficiency

Now this is all extremely abstract. To make it more definite, we need to fill in (at least) the notions of vocabulary, practice-or-ability, PV-sufficiency, and VP-sufficiency, which are the fundamental elements that articulate what I am calling the “meaning–use analysis” of resultant meaning–use relations—in particular, the pragmatically mediated semantic relations between vocabularies that I am claiming we must acknowledge in order to pursue the classical project of philosophical analysis in the light of what is right about the pragmatist critique of it. We can begin to do that by looking at a special case in which it is possible to be unusually clear and precise about the things and relations that play these metatheoretic roles. This is the case where ‘vocabulary’ takes a purely *syntactic* sense. Of course, the cases we eventually care about—and will be discussing in the remaining lectures—involve vocabularies understood in a sense that includes their *semantic* significance. But besides the advantages of clarity and simplicity, we will find that some important lessons carry over from the syntactic to the semantic case.

The restriction to vocabularies understood in a spare syntactic sense leads to correspondingly restricted notions of what it is to *deploy* such a vocabulary, and what it is to *specify* practices-or-abilities sufficient to deploy one. Suppose we are given an *alphabet*, which is a finite set of primitive sign types—for instance, the letters of the English alphabet. The *universe* generated by that alphabet then consists of all the finite strings that can be formed by concatenating elements drawn from the alphabet. A *vocabulary* over such an alphabet—in the syntactic sense I am now after—is then any subset of the universe of strings that alphabet generates. If the generating alphabet is the English alphabet, then the vocabulary might consist of all English sentences, all possible English texts, or all and only the sentences of *Making It Explicit*.¹¹

FN:11

¹¹ Computational linguists, who worry about vocabularies in this sense, have developed metalanguages for specifying important classes of such vocabularies: the syntactic analogues of semantic metalanguages in the cases we will eventually address. So, for instance, for the alphabet {a,b}, ‘aⁿbⁿ’ characterizes the vocabulary that comprises all strings of some finite number of ‘a’s followed by the *same* number of ‘b’s. ‘a(ba)^{*}b’ characterizes the vocabulary that comprises all strings beginning with an ‘a’, ending with a ‘b’, and having *any* number of repetitions of the sub-string ‘ba’ in between.

16 BETWEEN SAYING AND DOING

FN:12

What can we say about the abilities that count as *deploying* a vocabulary in this spare syntactic sense?¹² The abilities in question are the capacity to *read* and *write* the vocabulary. In this purely syntactic sense, ‘reading’ it means being able practically to *distinguish* within the universe generated by the vocabulary, strings that do, from those that do not, belong to the specified vocabulary. And ‘writing’ it means practically being able to *produce* all and only the strings in the alphabetic universe that do belong to the vocabulary.

FN:13

We assume as primitive abilities the capacities to read and write, in this sense, the alphabet from whose universe the vocabulary is drawn—that is, the capacity to respond differentially to alphabetic tokens according to their type, and to produce tokens of antecedently specified alphabetic types. Then the abilities that are PV-sufficient to deploy some vocabularies can be specified in a particularly simple form. They are *finite-state automata* (FSAs). As an example, suppose we begin with the alphabet {a, h, o, !}. Then we can consider the *laughing Santa vocabulary*, which consists of strings such as ‘hahaha!’, ‘hohoho!’, ‘hahahoho!’ ‘hohoha!’, and so on.¹³ Figure 1.6 is a graphical representation of a *laughing Santa finite-state automaton*, which can read and write the laughing Santa vocabulary. The numbered nodes represent the *states* of the automaton, and the alphabetically labeled arcs represent *state-transitions*. By convention, the starting state is represented by a square (State 1), and the final state by a circle with a thick border (State 4).

As a *reader* of the laughing Santa vocabulary, the task of this automaton is to process a finite string, and determine whether or not it is a licit string of the vocabulary. It processes the string one alphabetic character at a time, beginning in State 1. It recognizes the string if and only if (when and only when) it arrives at its final state, State 4. If the first character of the string is not an ‘h’, it remains stuck in State 1, and rejects the string. If the first




Figure 1.6 The laughing Santa automaton

¹² Here we can safely just talk about *abilities*, without danger of restricting the generality of the analysis.

¹³ In the syntactic metalanguage for specifying vocabularies that I mentioned in the note above, this is the vocabulary (ha/ho)*!

EXTENDING THE PROJECT OF ANALYSIS 17


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

As a *writer* of the laughing Santa vocabulary, the task of the automaton is to produce only licit strings of that vocabulary, by a process that can produce any and all such strings. It begins in its initial state, State 1, and emits an ‘h’ (its only available move), changing to State 2. In this state, it can produce either an ‘a’ or an ‘o’—it selects one at random¹⁴—and goes into State 3. In this state, it can either tack on an exclamation point, and move into its final state, State 4, finishing the process, or emit another ‘h’ and return to State 2 to repeat the process. In any case, whenever it reaches State 4 and halts, the string it has constructed will be a member of the laughing Santa vocabulary.

I hope this brief rehearsal makes it clear how the constellation of nodes and arrows that makes up this directed graph represents the abilities to read and write (recognize and produce arbitrary strings of) the laughing Santa vocabulary.¹⁵ What it represents is abilities that are *PV-sufficient* to

¹⁴ In fact, it can be shown that every vocabulary readable/writeable by a non-deterministic FSA—such as the laughing Santa automaton—is also readable/writeable by a deterministic one.

¹⁵ For practice, or to test one’s grip on the digraph specification of FSAs, consider what vocabulary over the same alphabet that produces the laughing Santa is recognized/produced by this automaton (the “I’ll have what she’s having” automation):



18 BETWEEN SAYING AND DOING

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

It may seem that I am stretching things by calling the digraph form of representation a ‘vocabulary’. It will be useful, as a way of introducing my final point in the vicinity, to consider a different form of pragmatic metavocabulary for the laughing Santa vocabulary. Besides the digraph representation of a finite-state automaton, we can also use a *state-table* representation. The state-table for the laughing Santa automaton (LSA) this is shown in Table 1.1.

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

aHalt3Halh2Halt2oHalt3Halt!HaltHalt4

Table 1.1 State-table for the laughing Santa automaton

Laughing Santa	State 1	State 2	State 3
a	Halt	3	Halt
h	2	Halt	2
o	Halt	3	Halt
!	Halt	Halt	4

EXTENDING THE PROJECT OF ANALYSIS 19




Figure 1.7 Meaning-use diagram 6: specifying the automaton that deploys the laughing Santa vocabulary

This is just a string, drawn from a universe generated by the alphabet of the LSA, together with ‘Halt’ and the designations of the states of that automaton. The strings that specify FSAs that deploy vocabularies defined over the same basic alphabet as the LSA then form a vocabulary in the technical syntactic sense we have been considering. And that means we can ask about the automata that can read and write *those* state-table encoding vocabularies. The meaning-use diagram for this situation is then as shown in Figure 1.7.

5 The Chomsky hierarchy and a syntactic example of pragmatic expressive bootstrapping

Restricting ourselves to a purely syntactic notion of a vocabulary yields a clear sense of ‘pragmatic metavocabulary’: both the digraph and the state-table vocabularies are VP-sufficient to specify practical abilities articulated as a finite-state automaton that is PV-sufficient to deploy—in the sense of recognizing and producing—the laughing Santa vocabulary, as well as many others. (Of course, it does that only against the background of a set of abilities PV-sufficient to deploy *those* vocabularies.) Perhaps surprisingly, it also offers a prime example of *strict pragmatic expressive bootstrapping*. For

20 BETWEEN SAYING AND DOING

in this setting we can *prove* that one vocabulary that is expressively weaker than another can nonetheless serve as an adequate *pragmatic* metavocabulary for that stronger vocabulary. That is, even though one *cannot say* in the weaker vocabulary everything that can be *said* in the stronger one, one can still *say* in the weaker one everything that one needs to be able to *do* in order to deploy the stronger one.

Here the relevant notion of the relative expressive power of vocabularies is also a purely syntactic one. Already in the 1950s, Chomsky offered mathematical characterizations of the different sets of strings of characters that could be generated by different classes of grammars (that is, in my terms, characterized by different kinds of syntactic metavocabularies) and computed by different kinds of automata. The kinds of vocabulary, grammar, and automata line up with one another, and can be arranged in a strict expressive hierarchy: the Chomsky hierarchy. It is summarized in Table 1.2.

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

Table 1.2 The Chomsky hierarchy

Vocabulary	Grammar	Automaton
Regular	$A \rightarrow aB \ A \rightarrow a$	Finite state automaton
Context-free	$A \rightarrow \text{anything}_2$	Push-down automaton
Context-sensitive	$c_1 A c_2 \rightarrow c_1 \text{anything}_2 c_2$	Linear bounded automaton
Recursively enumerable	No restrictions on rules	Turing machine (= 2 Stack PDA)

EXTENDING THE PROJECT OF ANALYSIS 21

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

The next point is that slightly more complex automata *are* capable of deploying vocabularies, such as $a^n b^n$, that are not regular, and hence cannot be read or written by FSAs. As our brief discussion indicated, intuitively the problem that FSAs have with languages like $a^n b^n$ is that they lack *memory*. If we give them a memory, we get a new class of machines: (non-deterministic¹⁶) *push-down automata* (PDAs). In addition to being able to respond differentially to, and produce tokenings of, the alphabetic types, and being able to change state, PDAs can *push* alphabetic values to the top of a *memory-stack*, and *pull* such values from the top of that stack. PDAs can do everything that FSAs can do, but they can also read and write many vocabularies, such as $a^n b^n$, that are *not* regular and so cannot be read and written by FSAs. The vocabularies they can deploy are called “context-free.” All regular vocabularies are context-free, but not vice versa. This proper containment of classes of vocabularies provides a clear sense, suitable to this purely syntactic setting, in which one vocabulary can be thought of as “expressively more powerful” than another: the different kinds of grammar can specify, and the different kinds of automata can compute, ever larger classes of vocabularies. Context-free vocabularies that are not regular require more powerful grammars to specify them, as well as more powerful automata to deploy them. FSAs are special kinds of PDAs, and all

¹⁶ By contrast to FSAs, there need not in general be for every vocabulary computable by a non-deterministic PDA, some deterministic PDA that reads and writes the same vocabulary.

22 BETWEEN SAYING AND DOING

the automata are special kinds of Turing machines. Recursively enumerable vocabularies are not, in general, syntactically reducible to context-sensitive, context-free, or regular ones. And the less capable automata cannot read and write all the vocabularies that can be read and written by Turing machines.

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

The proof of this claim is tedious but not difficult, and the claim itself is not at all controversial—though computational linguists make nothing of it, having theoretical concerns very different from those that lead me to underline this fact. (My introductory textbook leaves the proof as an exercise to the reader.¹⁷) General-purpose computer languages such as Pascal and C++ can specify the algorithms a Turing machine, or any other universal computer, uses to compute any recursively enumerable function, hence to recognize or produce any recursively enumerable vocabulary. And they are invariably context-free languages¹⁸—in no small part just

FN:17

¹⁷ Thomas A. Sudkamp, *Languages and Machines*, 2nd ed (Addison Wesley Longman, 1998), chapter 10.

FN:18

¹⁸ In principle. There are subtleties that arise when we look at the details of actual implementations of particular computer languages, which can keep them from qualifying as strictly context-free.

EXTENDING THE PROJECT OF ANALYSIS 23




Figure 1.8 Meaning-use diagram 7: syntactic pragmatic expressive bootstrapping

because the simplicity of this type of grammar makes it easy to write parsers for them. Yet they suffice to specify the state-table, contents of the tape (or of the dual stacks), and primitive operations of any and every Turing machine. Figure 1.8 shows the MUD characterizing this pragmatically mediated relation between syntactically characterized vocabularies.

I called the fact that context-free vocabularies can be adequate pragmatic metavocabularies for recursively enumerable vocabularies in general ‘surprising’, because of the provable syntactic irreducibility of the one class of vocabularies to the other. But if we step back from the context provided by the Chomsky hierarchy, we can see why the possibility of such pragmatic expressive bootstrapping should not, in the end, be surprising: for all the result really means is that context-free vocabularies let one *say* what it is one must *do* in order to say things they cannot themselves say, because the ability to deploy those context-free vocabularies does not include the abilities those vocabularies let one specify. Thus, for instance, there is no reason that an FSA could not read and write a vocabulary that included commands such as “Push an ‘a’ onto the stack,”—and thus specify the program of a PDA—even though it itself has no stack, and could not *do* what the vocabulary it is deploying specifies. A coach might be able to tell an athlete exactly what to do, and even how to do it, even though the coach cannot himself do what he is telling the athlete to do, does not have the abilities he is specifying. We ought not to boggle at the possibility

24 BETWEEN SAYING AND DOING

of an expressively weaker pragmatic metavocabulary having the capacity to say what one must do in order to deploy an expressively stronger one. We should just look to see where this seems in fact to be possible for vocabularies we care about, and what we can learn from such relations when they do obtain.

6 Looking ahead

Let us recall what motivated this rehearsal of some elements of automaton theory and introductory computational linguistics. I suggested that a way to extend the classical project of semantic analysis so as to take account of the insights of its pragmatist critics is to look analytically at relations between meaning and use. More specifically, I suggested focusing to begin with on two in some sense complementary relations: the one that holds when some set of practices-or-abilities is PV-sufficient to *deploy* a given vocabulary, and the one that holds when some vocabulary is VP-sufficient to *specify* a given set of practices-or-abilities. The composition of these is the simplest *pragmatically mediated semantic relation* between vocabularies: the relation that holds when one vocabulary is a sufficient pragmatic metavocabulary for another. It is a paradigm of the infinite, recursively generable class of complex, pragmatically mediated semantic relations that I propose to lay alongside the other semantic relations between vocabularies that have been investigated by analytic philosophers (for instance, those who address the core programs of empiricism and naturalism): relations such as analyzability, definition, translation, reduction, truth-making, and supervenience. I suggested further that pragmatic metavocabularies might be of particular interest in case they exhibited what I called “expressive bootstrapping”—cases, that is, in which the expressive power of the pragmatic metavocabulary differs markedly from that of the target vocabulary, most strikingly when the metavocabulary is substantially expressively weaker—a phenomenon Tarski has led us *not* to expect for semantic metavocabularies.

We have now seen that all of these notions can be illustrated with particular clarity for the special case of purely syntactically characterized vocabularies. The abilities that are PV-sufficient to deploy those vocabularies, in the sense of the capacity to recognize and produce them, can be thought of as various sorts of automata. There are several well-established,

EXTENDING THE PROJECT OF ANALYSIS 25

different-but-equivalent vocabularies that are known to be VP-sufficient to specify those automata. In this special syntactic case we can accordingly investigate the properties of pragmatic metavocabularies, and when we do, we find a striking instance of *strict expressive bootstrapping* in a pragmatically mediated syntactic relation between vocabularies.

Of course, the cases we really care about involve *semantically* significant vocabularies. Are there any interesting instances of these phenomena in such cases? I have indicated briefly how some of Sellars' pragmatist criticisms of various ways of pursuing the empiricist program can be understood to turn on pragmatically mediated semantic relations. And I mentioned Huw Price's idea that although normative vocabulary is not semantically reducible to naturalistic vocabulary, naturalistic vocabulary might suffice to specify what one must *do*—the practices-or-abilities one must engage in or exercise—in order to deploy normative vocabulary. Here is another example that I want to point to, though I cannot develop the claim here. For roughly the first three-quarters of the twentieth century, philosophers who thought about indexical vocabulary took for granted some version of the doctrine that a tokening *n* of an expression of the type 'now' was *synonymous* with, definable or semantically analyzable as, 'the time of utterance of *n*', and similarly for 'here' and 'the place of utterance of *h*', and so on. During the 1970s, philosophers such as John Perry, David Lewis, and G. E. M. Anscombe, by focusing on the use of indexicals in modal and epistemic contexts, showed decisively that this cannot be right: what is expressed by indexical vocabulary cannot be expressed equivalently by non-indexical vocabulary. This fact seems so obvious to us now that we might be led to wonder what philosophers such as Russell, Carnap, and Reichenbach could have been thinking for all those years. I want to suggest that the genuine phenomenon in the vicinity is a *pragmatically mediated* semantic relation between these vocabularies. Specifically, in spite of the *semantic* irreducibility of indexical to non-indexical vocabulary, it is possible to *say*, entirely in non-indexical terms, what one must *do* in order to be deploying indexical vocabulary correctly: to be saying essentially and irreducibly indexical things. For we can formulate practical rules such as:

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

26 BETWEEN SAYING AND DOING

- If a speaker s at time t and place $\langle x, y, z \rangle$ asserts “ P holds of *me*, *here* and *now*,” the speaker is committed to the property P holding of $\langle x, y, z, t, s \rangle$.

Further (as I show in the appendix to the next lecture, where the necessary concepts have been introduced), those responses can be algorithmically elaborated so as to play the role distinctive of *essential* indexicals. Non-indexical vocabulary can serve as an adequate *pragmatic* metavocabulary for indexical vocabulary. The fact that one nonetheless cannot *say* in non-indexical terms everything that one can *say* with indexical vocabulary just shows that these vocabularies have different expressive powers, so that the pragmatically mediated semantic relation between them is a case of strict pragmatic expressive bootstrapping.

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

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

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

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

EXTENDING THE PROJECT OF ANALYSIS 27

those exercised by a finite-state *automaton*. Indeed, automata are defined by a definite set of *meta-abilities*: abilities to *elaborate* a set of primitive abilities into a set of more complex ones, which can accordingly be pragmatically *analyzed* in terms of or *decomposed* into the other.

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

In my third lecture, I will argue that the notion of the algorithmic decomposability of some practices-or-abilities into others that results suggests in turn a pragmatic generalization of the classical program of artificial intelligence (AI) functionalism—which, though a latecomer in the twentieth century, deserves, I think, to count as a *third* core program of classical semantic analysis. AI functionalism traditionally held itself hostage to a commitment to the purely *symbolic* character of intelligence in the sense of sapience. But broadening our concern from automata as purely syntactic engines to the realm of *transducing* automata puts us in a position to see AI functionalism as properly concerned with the algorithmic decomposability of discursive (that is, vocabulary-deploying) practices-and-abilities. What I will call the “pragmatic” thesis of artificial intelligence is the claim that the ability to engage in some autonomous discursive practice—a language game one could play though one played no other—can be *algorithmically decomposed* into non-discursive abilities. The arguments for and against this pragmatic version of AI-functionalism look quite different from those arrayed on the opposing sides of the debate about the prospects of symbolic AI.

The notion of PP-sufficiency brings into view a slightly more complicated pragmatically mediated semantic relation between vocabularies: that which obtains when practices PV-sufficient for V_1 are PP-sufficient

28 BETWEEN SAYING AND DOING

(in the sense that they can be algorithmically elaborated into) practices PV-sufficient for V_2 . Another basic meaning–use relation of the kind we have been considering is PV-necessity, the converse of PV-sufficiency. It obtains when one cannot deploy a certain vocabulary without engaging in the specified practice, or exercising the specified ability. For example, I have argued elsewhere that nothing could count as engaging in an autonomous discursive practice (hence using a vocabulary one could use though one used no other) that did not include *asserting* and *inferring*. Considering that basic MUR permits the formulation of a complex resultant MUR that is a variant on the prior one: a relation that obtains where practices PV-necessary for V_1 are PP-sufficient for practices-or-abilities PV-sufficient for V_2 .

It can happen, I will argue, that such a V_2 is also VP-sufficient to *specify* the practices-or-abilities that are PV-sufficient to deploy V_1 . A MUD for this is shown in Figure 1.9.

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




Figure 1.9 LX: V_2 is elaborated from and explicative of practices PV-necessary for V_1

EXTENDING THE PROJECT OF ANALYSIS 29

meaning–use relations, and the complex pragmatically mediated semantic relation that results from it.

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

I will begin my treatment of modality, in my fourth lecture, with a consideration of this question, and with a vindication of the role of modal vocabulary that parallels the one I will already have offered for ordinary logical vocabulary: modal vocabulary, too, can be elaborated from and is explicative of, features integral to every autonomous discursive practice—features intimately related to, but distinct from, those made explicit by ordinary logical vocabulary. I will then enter into an extended treatment of the relation between *alethic* and *deontic* (modal and normative) vocabularies. When we look at those vocabularies through the lens of meaning–use analysis, a sequence of startling relations between them emerges.

For a start, I argue that deontic normative vocabulary is *also* universally LX (that it is VP-sufficient to specify practices-or-abilities that are both PV-necessary for deploying any autonomous vocabulary, and PP-sufficient for practices-or-abilities PV-sufficient for deploying the deontic normative vocabulary that explicates them). Although in this regard it belongs in a box with alethic modal vocabulary, the features of autonomous vocabulary use that it explicates are quite different from those explicated by modal vocabulary. I then argue that what lies behind Sellars' dark and pregnant claim that “the language of modality is a transposed language of norms” is the fact that deontic normative vocabulary can serve as a *pragmatic*

30 BETWEEN SAYING AND DOING

metavocabulary for alethic modal vocabulary. In my fifth lecture, I will show how exploiting that relation makes possible a new kind of *directly modal* formal semantics that makes no appeal to *truth*: incompatibility semantics. It in turn gives us a new semantic perspective both on traditional logical vocabulary, and on modal vocabulary. The final lecture will then weave all these strands into a meaning–use analysis of intentionality itself (what is expressed by intentional vocabulary) as a pragmatically mediated semantic relation essentially involving both what is expressed by modal and what is expressed by normative vocabulary.

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

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

Appendix 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 the first lecture, one in which 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. I made this claim already in Lecture 1, but at that point did not yet have on board the theoretical machinery needed to substitute a full argument for that gesture.

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 Anscombe–Perry sort of pragmatic significance *must* have their semantic contents determined in the Kaplan–Stalnaker way.

APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS 57

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 somewhat more surprising claim. The banal 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. (This is what I say in Lecture 1.) The more interesting claim (which is necessary for my overall point) 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,¹⁰ then 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 would be 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

¹⁰ And these really are the dimensions along which the two kinds of vocabulary are distinguished.

58 APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS

FN:11

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.¹¹ To compute the time referred to by a tokening of the temporal indexical type “last Wednesday,” one must know both 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 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 7 of *Making It Explicit*,¹² for instance, the

FN:12

¹¹ 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.

¹² Harvard University Press, 1994.

APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS 59

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 set out in Figure 2.7. (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 .




Figure 2.7 Pragmatically mediated semantic presupposition of anaphoric by indexical and deictic vocabularies

60 APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS

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 $t-1$.

FN:13

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.¹³ 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 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

¹³ G. E. M. Anscombe, “The First Person,” in Samuel Guttenplan (ed.), *Mind and Language* (Clarendon Press, 1975); John Perry, “The Problem of the Essential Indexical,” *Noûs*, 13/1 (March 1979), 3–21; David Lewis, “Attitudes *De Dicto* and *De Se*,” *The Philosophical Review*, 88 (1979), 513–43.

APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS 61

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 envioning 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

62 APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS

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’. These in turn are 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,

APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS 63

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

64 APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS




Figure 2.8 Test–Operate–Test–Exit (TOTE) cycle

the goal, and then followed by another movement calibrated by the results of the prior one, until the goal is reached (Figure 2.8).

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 ‘here and now’ of prior, standing intentions into intentions-in-action, which are practically efficacious, 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-acknowledgements that mediate the cycle of perception and action.¹⁴ 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 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 on to public space. It is natural for us to describe this ability in terms of mappings of indexical on to 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.

¹⁴ 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.

APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS 65

Consider, then, a community that accords some tokenings of non-indexical types the pragmatic significance of immediate acknowledgements 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 recruited and integrated into the states one is already capable of entering in response to the acquisition of

66 APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS

FN:15

a prior intention (a standing practical commitment), by arbitrary state-formation and state-permutation.¹⁵ My claim is that those algorithmic-elaborating abilities are *all* that is needed to 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 acknowledgements of practical commitments (Figure 2.9).

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

¹⁵ 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 8 of *Making It Explicit* (Harvard University Press, 1994).

APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS 67




Figure 2.9 Indexicals as elaborated-explicating vocabulary

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

68 APPENDIX INDEXICAL VOCABULARY: SEMANTICS AND PRAGMATICS

FN:16

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-types¹⁶ 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 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.

¹⁶ 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.