

Naturalism—Week 8—Overall Plan:

Introduction:

1. Let me start with a reminder of why we are worrying about Sellars's arguments here. Jackson showed us that besides *scientific* naturalism, there is also such a thing as *descriptive* naturalism—specifically and paradigmatically, about the normative (in his case, the moral—though I pointed out that almost nothing he says depends on its being this sort of normative vocabulary he is discussing). Jackson assumes that naturalists will be descriptivists—about everything. Or perhaps it is just that the “location problem” only arises for the descriptive part of the language. Sellars is not even a descriptivist about *modality*, never mind normativity. As we have seen, this issue can only be sensibly addressed in the context of some way of demarcating descriptive vocabulary, or descriptive uses of language (the activity of describing). (We saw that there are various options for thinking about the relations between these two: does vocabulary count as descriptive if it has *any* descriptive uses? If it can *only* be used descriptively? If its other uses are parasitic on or derivative from its descriptive uses?) One set of criteria of adequacy for such demarcation of the descriptive is set by a Scylla and a corresponding Charybdis:

- On one side, for any paradigmatically descriptive bit of vocabulary ('happy', 'red', 'having a mass greater than 1 gram'), we *can* introduce a special speech act, thought of as distinct from describing, which is performed exactly whenever the term is applied in what we otherwise would have thought was a characterizing or classifying way. This is what I called (following Geach's example), 'macarizing' that bit of descriptive vocabulary. The corrosive effects of rampant macarization then threaten to undercut any demarcation that leaves *any* vocabulary as descriptive.
- On the other side, the strategy Geach suggests for responding to this threat—to look at *embedded*, rather than *free-standing* uses of the vocabulary, to see what contribution its unasserted occurrence makes to the contents of larger compounds (paradigmatically, conditionals in which it appears as an antecedent)—leads to *declarativism* about the descriptive. This is the view that any vocabulary that can be used in declarative sentences is descriptive. This is a workable criterion of demarcation, but it leaves *every* kind of vocabulary philosophers have been worried about on the descriptive side: normative, aesthetic, modal....

So if Sellars's view that alethic *modal* vocabulary is *not* descriptive is not to be either *trivial* (because modal vocabulary is not the only vocabulary that can be macarized: *any* vocabulary can) or *false* (because we are driven to declarativism about the descriptive), then there must be some way of demarcating a class of vocabulary (or uses) as descriptive, which threatens neither emptiness (via promiscuous macarization) nor all-inclusiveness (via declarativism).

2. The discussion will be in two parts:

- a) From labels to descriptions, via a “space of implications”: “It is only because the expressions in terms of which we describe objects, even such basic expressions as words for the perceptible characteristics of

molar objects locate these objects in a space of implications, that they describe at all, rather than merely label.” [§107—final section of essay]

b) Premises *from which* we argue vs. Principles *in accordance with* which we argue:

“Mr. E....conceives of induction as establishing principles *in accordance with which* we reason, rather than major premises *from which* we reason.” [§83—first section of Part IV].

[Note: The connection with induction will be explored next week, as will the distinction—crucial to the articulation of the refined version of Mr. E’s position that Sellars finally endorses—between what is *said* and what is only *conveyed* or “conventionally implied” by what is said.]

Part I:

3. Labeling, nominalism, and the classificatory theory of consciousness:

- a) What I’ll call “*semantic nominalism*” was the universally held semantic theory until Kant, and is still probably dominant. It holds that:
 - i. Proper names are labels, stuck onto or otherwise associated with the objects named. This is the ‘Fido’/Fido theory.
 - ii. Predicates are like names, but they are *general* labels, labels that are stuck onto or otherwise associated with *many* objects (those they are true of), via the *properties* of those objects. They are general in that we stick them onto more than one thing. These labels specify *properties* of labelable objects, or their *kinds*. (Depending on whether they are sortals or not, that is, whether in addition to criteria of application—see below—they have criteria of identity and individuation—as Frege taught us in the *Grundlagen*. Cf. Sellars claim here: “The point is the more radical one that the relation of a thing-kind word to the criteria for belonging to that kind of thing is different in principle from the relation of words for characteristics of things to the criteria for the presence of these characteristics. "Lemon" and "bald" may both be vague, but they are so in radically different ways.” [§46] [But bracket all these considerations.]

Semantic nominalism is the view that the relation between a *name* and its *bearer*, what is a name *of*, construed on the model of *labeling*, is the paradigmatic semantic relation. Predicates name (label) properties, and sentences name possible states of affairs.

- b) Complex labels are *descriptions*. So something can be *described* by pinning on it (associating with it in whatever the way distinctive of *semantic* association is) the labels ‘red’, ‘juicy’, ‘apple’.
- c) Language, accordingly, consists of a bunch of *descriptive* terms, labels. And what one *does* with language is to *describe* things. The result is a picture of language as essentially a system of *classification*. The idea that this is what language is is *descriptivism*.
- d) So semantic nominalism is a principal route to *descriptivism* about language: the view that what language is *for* is to *describe* the world.

- e) Notice that the language-as-labeling view is semantically *atomistic*: Applying the label ‘red’ is *independent* of applying the label ‘apple’. Even though there turn out to be, as a matter of fact, connections between applying the labels ‘ripe’, ‘Macintosh’, and ‘red’, that sort of fact is not an *essential* feature of the semantic connection between the labels and what is labeled. For the relation between one label and what it labels does not depend on the relation between any other label and what it labels.
- f) Thus the logical tradition before Kant started with
 - i. a doctrine of *terms* or *concepts*, particular and general, and advanced from there to
 - ii. a doctrine of *judgements*, thought of as *classifications*: paradigmatically of something *particular* as falling under something *general* (‘universal’), but perhaps also of one class as including another. The tradition then proceeded to
 - iii. a doctrine of *consequences*, understood syllogistically, in terms of the kind of classification involved in the premises and conclusion of the inference.
- g) So *judging* is thought of as *classifying*: applying a *general* (predicative) label to something one also applies a *particular* (term) label to, which is associating the particular and general labels in a distinctive way, marked by a *copula*: “That rose is red.”
- h) The *classificatory theory of consciousness or awareness* says that to be aware of something is to be aware of something *as* something. (This is sometimes—Heidegger picks up this traditional usage in *SZ*—called the “apophantic ‘as’”.) And that means, aware of something *particular as* something *general*. (But it had better not turn out that a chunk of iron, which reliably differentially responds to its environment by rusting under some conditions and not others is *in the relevant sense* classifying those environments as being of one of two general kinds—say, wet and dry.)
- i) There is a critical ambiguity here. The notion of classifying is being appealed to to do two different, and ultimately incompatible jobs.
 - On the one hand, it is supposed to be what one *does* in order to count as *judging*.
 - On the other, it is the way judgeable *contents* (propositions) are built up out of labels of different kinds of things (objects and properties).
- j) We can see the tension between these if (following the Kant-Frege-Geach strategy that should by now be familiar) we look, not at *free-standing*, but at *embedded* uses of sentences, where they occur as *unendorsed components* of more complex sentences—paradigmatically as the antecedents of conditionals. When I say “*If* that rose is red, then it is colored,” have I classified the rose as red? If so, then classifying is *not* judging; for I have *not judged that* the rose is red. I merely entertained the possibility. If not, then how are we to understand the sentence, with its copula (‘is’) associating a general label with a particular one? What more is needed for classification? And if it is not classification going on, what

is happening in that sentence? Further, if the rose is *not* being classified as red in the conditional, and it *is* being so classified in the free-standing assertion “That rose is red,” then it looks as though we are *equivocating* when we move from “*If* that rose is red, then it is colored,” and “That rose is red,” to “That rose is colored.” For the first does *not* contain a classification of the rose as red, and other premise, from which we are going to draw our conclusion by detaching from the conditional, *is* a classification. What we have is an attempt to think of *classification* as an *operation* in two quite different senses:

- One that attaches a certain kind of pragmatic *force* to a content, having the effect of *endorsing* that content, or taking it to be *true*.
- Another that is used to build up sentential semantic *contents* out of subsentential ingredients.

4. From labeling to describing:

- a) *Mere* labels: Consider a tray of disparate objects, each of which is labeled with either a blue or a red dot. They have been *labeled*. Have they been *described*? Evidently not. For what have they been described *as*?
- b) One way of seeing that such mere labels don’t *mean* anything (or at least, that we don’t *understand* them) is that we have no idea how to *go on* labeling things with red and blue dots. If a few more objects are added to the tray, we don’t know which, if either, label is appropriate. The mere labeling of some objects does not establish a *standard*, *norm*, or *practice* we can appeal to in determining how it would be *correct* to continue labeling new objects. (This is *one* of the threads Wittgenstein is pursuing in his discussion of “going on in the same way” in the *PI*.)
- c) [This next bit is something of a detour, given our current goal. But there is enough of a philosophical point here that it is worth tarrying a bit with it.] It actually may be that as a matter of what we could call “conceptual psychology” (by analogy to “moral psychology”), this last claim is not true. Gombrich, in *Art and Illusion* (made much of by Goodman in *Languages of Art*) suggests that mere labels (He’s concerned with “representational art” and “non-representational art.” Goodman is more concerned with pairs such as ‘empiricist’/ ‘rationalist’) may be all too projectible, at least if those doing so share sufficient training and background. He suggests that if take two nonsense labels, ‘ping’ and ‘pong’, and use them to distinguish pairs of well-known cultural figures, people as a matter of empirical fact will largely agree in how to apply those terms to further cases.
 - So: Aristotle, Hegel, and the early Wittgenstein are *ping*. Plato, Kant, and the later Wittgenstein are *pong*. Now what about Quine and Davidson? Which is *ping* and which is *pong*?

- Or painters: Hieronymos Bosch, William Blake, and Salvador Dali are *ping*, while Rembrandt, Constable, and Picasso are *pong*. What about Cezanne and Gauguin?
- Or, writers: Euripedes, Baudelaire, and Pynchon are *ping*, Sophocles, Flaubert, and Joyce are *pong*. What about Henry James and Proust?

The lesson Gombrich and Goodman want us to learn from these exercises is not to think that just because there is a learnable disposition to agree in extending labels to new cases, there is some (interesting?) property being picked out. (I actually think this is not a bad way to find *candidate* interesting properties.) This line of thought gives aid and comfort to those who hope that “all ‘ism’s are ‘wasm’s.”

d) Be that as it may, for our purposes, we may conclude that the *only* labels that have any prospect of counting as *descriptions* are those associated with reliable differential responsive dispositions to apply them to new cases. These are *projectible* labels. They must be associated, explicitly or implicitly, with *standards*, or *norms*, or at least learnable-teachable *practices* that settle when it is and when it is not *correct* to *apply* the label to new cases. At the least, some notion of *mislabeling* must have been put in place, for labels to be even candidate descriptions. Descriptive terms, unlike mere labels in the thinnest sense, must at least come with *circumstances of appropriate application*.

e) So, if we have such RDRDs embodying standards of appropriate application, will *such labels* be descriptions? It is not hard to see that they will not. ‘Gleeb’. Consider possession of an infallible ‘gleebness’ tester. Point the device at something, and it lights up if and only if the thing is *gleeb*. This, by hypothesis, is projectible. It establishes a standard, with respect to which things can be *mislabeled* as ‘gleeb’. But when one has found out that something is *gleeb*, what has one found out? (One can know what is a K without knowing what a K is.) One has not *described* it, but *merely labeled* it. One knows what things are *gleeb*, but has not thereby found out *anything* about them, since one does not, we want to say, yet know what *gleebness is*. Once again we can ask: What is it you are describing things *as* when you label them as ‘gleeb’? The conclusion is that it is *not* enough to have a *description* of something that one have not only a label that has in the past been applied to some things and not to others, but also a *reliably differential responsive disposition* to discriminate things to which the label is and is not correctly applied.

f) What more is needed? Consider a *parrot* who can respond to the visible presence of red things by uttering tokens of “Awrk! That’s red!” And suppose that he does so in just the same circumstances in which we do. He reliably differentially responds to red things by correctly applying a vocalized label. [Kvetch about ‘vocal’ vs. ‘verbal’.] Is he *describing* things as red? The noise he makes is *just* a noise to him, as ‘gleeb’ was to us. For the parrot, that label is not something that *contrasts* with other labels in that it excludes their proper applicability. And the applicability

of that label does not have any further consequences, for instance, making further labels such as ‘colored’ appropriate.

g) Conclusion: For a description, we must have *both circumstances* of appropriate application *and* appropriate *consequences* of application. Dummett argues this (in his 1974 *Frege’s Philosophy of Language*). He excoriates theories of meaning that are “one-sided”, in focusing on one or the other of these to the exclusion of the other. Thus assertibilist, reliabilist, and verificationist theories of meaning treat the meaning of an expression as consisting in its *circumstances* of appropriate application. But this can’t be right. Things can have the *same* circumstances of appropriate application, and *different* consequences of appropriate application. Look at “I *will* write a book on Hegel,” and “I *foresee* that I will write a book on Hegel.” Can regiment a technical use of ‘foresee’, *stipulating* that these two are appropriate (as free-standing assertions) under just the same circumstances. But they behave very differently as antecedents of conditionals:

- If I will write a book on Hegel, then I will write a book on Hegel.

and

- If I foresee that I will write a book on Hegel, then I will write a book on Hegel.

The first is a “stuttering inference”, as sure as can be. The second depends on how good at foreseeing I am, how practically resolute I am, and whether I am hit by a bus.

And concepts can have the *same consequences* of application and *different circumstances*. We know a lot about the more or less forensic *consequences* of application of the term ‘personal identity’—about what *follows* if we say that this is the same person as that. But using different *criteria* to specify the *circumstances* of application—say, bodily continuity vs. memory or psychological continuity, gives us *different* concepts of personal identity, in spite of the agreement in consequences of application. And that difference in concepts means that one is offering a different *description* of someone as “the same person as N.N.” if one uses the different criteria of application.

h) We saw that Sellars says [§107] that what is needed is that the description be put in a “space of implications”. We need *inferential articulation* to have *description*. To be more than a *mere* label, the label must be one that one can offer *reasons* for applying in one case and withholding in another (corresponding to the *circumstances* of appropriate application) and whose applicability can itself offer reasons for the application of *other* characterizations (corresponding to the appropriate *consequences* of application). This is what a *parrot* lacks, who can reliably differentially respond to red things by saying “That’s red” in the same circumstances we do, but who does *not*, just on that basis, count as *describing* anything as red, or *reporting* or *observing* that something is red. To be a *description*, the *label* must be situated in a web of *connections* to *other* labels/descriptions. And those connections are broadly *inferential*: a matter of what is *evidence* or *reason* for or against what, of what *obliges*

one to apply further labels/descriptions, or *precludes* one from doing that. Absent that context, labels are not descriptions. That is what Sellars means by saying:

“It is only because the expressions in terms of which we describe objects, even such basic expressions as words for the perceptible characteristics of molar objects locate these objects in a space of implications, that they describe at all, rather than merely label.”

- i) Since every expression must have *both circumstances of appropriate application and appropriate consequences* of application, each incorporates an *inference*: an inference *from* the obtaining of the *circumstances to* the obtaining of the *consequences*.
- j) And that must be a *counterfactually robust* inference. That is, the commitment involved in using the descriptive expression in question is that if anything *were to* satisfy the circumstances of its application, it *would* satisfy the consequences. For otherwise, the term cannot be applied to *new* cases. For one would need to find out in advance if the inference held in *that* case. But the point is that the circumstances of appropriate application need *not* include checking whether the consequences also obtain.
- k) In general, one cannot count as understanding *any* descriptive expression (or the concept it expresses, what it describes something as being) unless one distinguishes at least *some* of the inferences it is involved in (some of the connections within the “space of implications” it is situated in) as *counterfactually robust*, at least in the sense that they would *remain* good inferences if some further premises were added that do not in fact obtain. Thus one must know such things as that a lion would still be a mammal if the lighting were slightly different, it were a different day of the week, it was transported to a zoo, we clipped its fur....

5. Sellars talks about the “tendency to assimilate all discourse to describing” as principally “responsible for the prevalence in the empiricist tradition of ‘nothing-but-ism’ in its various forms (emotivism, philosophical behaviorism, phenomenism)” [§103] There is also a platonist, reistic version of descriptivism, by contrast to this minimalist empiricist one. It, too, looks for something described by, say, normative vocabulary, and posits metaphysically peculiar items such as *values*, or the property of *to-be-doneness* that an action might be taken to have. His aim is to argue against that tendency for the specific case of what is expressed by *alethic modal* vocabulary, the “language of causal modalities”. The general course of the argument is this:

- a) “It is only because the expressions in terms of which we describe objects, even such basic expressions as words for the perceptible characteristics of molar objects locate these objects in a space of implications, that they describe at all, rather than merely label.”
- b) Some uses of language play the distinctive expressive role of making explicit the relations articulating that space of implications.
- c) Doing that is *not* describing how things are.
- d) The statements that make explicit the inferential relations that articulate the space of implications, in which descriptions as such are embedded, are

subjunctive conditionals, which are formed using the vocabulary of *causal modalities*.

- e) Their characteristic use is not in *description* but in *explanation*.
- f) So the reason why “The descriptive and the explanatory resources of language advance hand in hand...” [§107], why *explanation* is an essential element of any discursive practice that includes *description*, is that descriptions must be located in a space of implications, and explanation of the applicability of some descriptions in terms of the applicability of others proceeds by asserting subjunctive conditionals that make those implication relations explicit.
- g) [We’ll see next time that the issues of offering *inductive evidence* for subjunctive conditionals, and of the *criticism* and hence *process of changing* the contents of our descriptive contents are intimately bound up with these others. But that is a topic we’ll postpone discussing until then.]

6. One issue, then, concerns the claim that the ‘implications’ in which descriptive vocabulary must be caught up in order to count as more than mere labels must be *counterfactually robust* inferences.

7.

- a) Sellars says (in the Introduction to CDCM), that “the framework [note the word] of what objects of a certain kind K *would* do in circumstances C is *basic*.” (Q: In what sense ‘basic’? And why?) We have seen, in effect, that the implications in which genuine descriptive terms (as opposed to mere labels) are involved are *counterfactually robust*. That is, they must extend to *possible* cases. This is just another way of saying that there must be a *norm* or *standard* for the *correct* application of the term in cases that have not actually arisen. We are seeing the general shape of an argument that *modality* (what is expressed by modal vocabulary, such as that used to express subjunctive conditionals—*one* kind of counterfactually robust conditional, as we shall see) is implicated in the *framework* that makes *description* possible. Cf. Sellars’s essay “Concepts as Involving Laws, and Inconceivable Without Them.” And it is this same line of thought that will implicate *explanation* with *description*.
- b) There are *two kinds* of counterfactually robust conditionals, which Sellars distinguishes [ref.] as:
 - i. *Genuine subjunctive conditionals*, and
 - ii. *Subjunctive identicals*.

[Explain this difference:]

Consider the difference between:

- iii. Copper (all samples of copper) melts at 1084° C.. and
- iv. All the coins in my pocket are copper.

The first, (iii), supports the subjunctive conditional:

- v. If this coin (which is in fact a nickel) *were* copper, it *would* melt at 1084° C..

The second, (iv), does *not* support the corresponding subjunctive conditional:

vi. If this coin (which is in fact a nickel) *were* in my pocket, it *would* be copper.

(iv) *does*, however support the claim

vii. If I *were* to choose a coin at random from my pocket, it *would* be copper.

So we cannot distinguish between the “accidental generalization” or “contingent regularity” (iv) and the “lawlike statement” (iii) by saying something like “only laws support counterfactuals”. (Though you do hear things like that a lot.)

Rather, we must distinguish what *kinds* of counterfactuals they support. Sellars does that by introducing the idea of “subjunctive *identicals*”. The idea is that *all* the counterfactuals supported by statements such as (iv) must be derived via *identities* of the objects they concern with some of the members of the class *actually* picked out by the description in the generalization. The generalization (iv) supports counterfactuals concerning objects identical with one of the coins that in fact are in my pocket. It does *not* support counterfactuals concerning what would happen if *other* things *did* fit that description, i.e. *were* in my pocket. So, if any coin at all, in *any* possible world, *were* identical to one of the coins that is *actually* in my pocket, it *would* be copper. For all *those* coins, the ones *actually* in my pocket *are* copper. But the generalization in (iv) tells us *nothing at all* about any coins in other worlds just on the basis of the fact that in *that* world, *their* world, the description “coin in Bob’s pocket” applies to them. Sellars gets at this difference in the significance of the two sorts of claims—genuine subjunctive *conditionals* and mere subjunctive *identicals*—by saying that the former, but not the latter, tell us something about what follows from the applicability of a description “*just on the basis of the applicability of that description*.”

- c) Part of what is at issue here is Sellars’s claim that it is a mistake to think of “lawlike generalities” as just a kind of regularity, like accidental or contingent regularities—only a kind that has some special extra feature. We *could*, he says, have used the expression ‘ $\forall x[Fx \supset Gx]$ ’ to express a relation between *concepts*, as he takes it that ‘ $\forall x[Fx \Box \rightarrow Gx]$ ’ does. (Perhaps Frege did so. See Macbeth’s important recent book, *Frege’s Logic*.) But since we do not use the expression that way, we must resist the temptation to think that that quantified statement about *objects* is the basic phenomenon, to which we must add some special modal ‘oomph’ to get the ‘lawlike’ statement.
- d) **The claim Sellars is going to make is that it is an essential feature of the meaning of *any* genuinely *descriptive* term that its applicability supports at least *some* genuine subjunctive conditionals.**
- e) It is these subjunctive conditionals that are the basis for *explanations* that show why something that *is* true *had* to be true, under the collateral conditions that obtained, how it is at least *conditionally necessary*. We need *not* assume that *all* explanations are of this kind in order to conclude from this claim that the notion of description makes sense only in a *framework* that includes also explanation, (what is expressed by) subjunctive conditionals, and so what is expressed by *modal* vocabulary.

- f) *Understanding* descriptive terms requires knowing at least *some* of the modal significance of their application that is made explicit in subjunctive conditionals.
- g) Looking even further ahead, these considerations are also the background for Sellars's claim that laws and the subjunctive conditionals they support should be understood *not* as providing premises *from which* we reason, but as expressing principles *in accordance with* which we reason.
- h) That is why he will say that modal vocabulary is *not descriptive* vocabulary, and so does *not* have the expressive job of *saying how things are*. It rather makes explicit essential aspects of the *framework in which* it is possible to describe how things are.

8. If Sellars is right about description being intelligible only in the context of a "space of implications," and the various consequences of that claim just sketched, then the fact that we *can* introduce a domain—possible worlds—that we *can* take modal vocabulary to *describe* suggests that there is an issue being *put off* rather than *solved*. For we should ask: "What is the 'space of implications' that articulates the content or meanings of the possible-world-descriptive modal vocabulary?" Implications in the base language are determined by descriptions in the metalanguage of set-theoretic inclusion relations among sets of possible worlds. What settles the "space of implications" for the descriptions of that possible-worlds modal metalanguage?

- a) Is there a *regress* in the offing here?
- b) Or is there rather an *equilibrium* of some reassuring sort being achieved?
- c) Or is it the case that we only need to use *extensional logic* in order to determine the "space of implications" for the modal metalanguage, and so that is where the regress ends in a foundation?

(Cf. the Agrippan trilemma of justification: regress, circle, or foundation.)

Part II:

9. On descriptivism and declarativism:

- a) The conditional $p \Box \rightarrow q$ (cf. "c-box"), a *counterfactually robust*, true subjunctive conditional, gives us a new kind of *sayable*. But it is **not supposed to be a *description of the world***. So Sellars *must* be *rejecting declarativism* about description.
- b) It *does* embed as the antecedent of a conditional.
- c) And we *should* be able to understand its *embedded* use in terms of its *free-standing* use as codifying inferences, in the sense of being a principle in accordance with which we reason—its *ingredient sense* in terms of its *free-standing* sense.
- d) But, think of my T&A stuff on *assertibility* conditions and *truth* conditions. What if we *cannot* specify the *ingredient sense* from the *free-standing* one, just as compounds like conditionals are not assertibility-functional (as *shown* by the 'will'-‘foresee’ example in (1-a-i- ζ) above). If there is *not* a way to *derive* the *ingredient sense* of modal vocabulary from the *free-*

standing sense needed to specify its role in codifying principle *in accordance with which* we reason:

- i. How *is* the ingredient sense determined?
- ii. How, if at all, is that question or issue related to the *descriptiveness* of the vocabulary?
- e) So, for instance, “I *will* Φ ,” thought of as a *prediction*, *is* descriptive. But “I *foresee* that I *will* Φ ,” is *not*—or at least, it describes at most *me*, not the rest of the world. (Cf. in my T&A, ESPs: expression-statement pairs. The ‘*foresee*’ claims are on the expression side.) So the ‘*will*’-‘*foresee*’ case can seem like a *displaced* description, where *two different sorts* of description, one of the world, one of *me*, *systematically* have the *same* circumstances of application in free-standing cases, but *different* ingredient senses. If that is all right, something need *not* be *non-descriptive* in order to fail in the derivability of ingredient sense from free-standing sense.
- f) Q: So, what to do? A: One might think to look to the *consequences* of application. For these distinguish *what* one is describing (what properties one is attributing): I can draw inferences about *me* from the ‘*foresee*’ statement, and about the world (as well as ones about *me*) from the ‘*will*’ statement.
- g) Q: Can one use *this* strategy—looking to the consequences of application—to get ingredient senses for subjunctive conditionals, even when one understands their free-standing use as non-descriptive? For one *can* talk about the *consequences* of endorsing an inference or pattern of inference. And it seems we can talk about the consequences of the inference endorsed, not just the consequences of endorsing it. We would usually distinguish these two by the embedding test. But is that *circular*, and so methodologically forbidden, in the present context? After all, we are asking about what *conditionals* mean here. How can we appeal to them as having unproblematic contents so as to use the embedding test? We still can distinguish the *force* of endorsing an inference from saying *that* I endorse it (though not, perhaps, from saying *that* it is a good one). (Think of the ESP consisting of this pair.)
- h) Q: So is there any *other* way of distinguishing between what follows from the (pattern of) inference endorsed by asserting a subjunctive conditional and what follows from the endorsing of it (as an *act*), *besides* the embedding test? A: Yes.
 - i. Try 1: By abstracting from the effects of *other* commitments. But what about other attributions to *me*—A1 abstract from these *too*, i.e. for any collateral commitments attributed to . . . Just collateral premise that are non-ascriptional, used as premises.
 - ii. Try 2: Just look at inferences I am committed to *before* I say something that makes that explicit—i.e. what *is* made explicit rather than the making explicit of it. So . . . some

inference whether *by this attribution*, apart from *his* collateral commitments, S is attributed to it just by asserting $p \Box \rightarrow q$.

2 cases:

- iii. Can use—*how we use*—the *right* consequences, *without* appealing to embedding test, and
- iv. *If we can, does that let us go from FSS to IS?* A: for *conditionals*, arguing I may not *need* to get *ascription* (hence negation, the same way).
- i) So this is a promising line to keep $p \Box \rightarrow q$ a) a description of the world *or* are (the 2 alternatives) *but* b) ...an IS. If it *can* get both (i) and (ii), then we are *not* being *declarativist* about *descriptivism*.
- j) Q: So what *is* the notion of description? Poss A: Projectible sorting *plus* inferential consequences. “Of the right sort”? Yes: by counterfactually robust inferences made explicit by true subjective conditionals and (so) establishable by induction.

10. Subjunctive conditionals express “*principles in accordance with which* (PAWW) we reason, rather than *premises from which* we reason (PFW). How should we understand this distinction? What kind of a distinction is it?

- a) Any particular PAWW *can* be treated as a PFW. It is a claim, and we can add it to our premises. That, presumably, is the reason we need to be warned not to make this mistake.
- b) Of course, we can’t do that to *all* our PAWWS. That is what Lewis Carroll’s argument in “Achilles and the Tortoise” shows. [Explain.] First level account of LC’s A&T: need *rules* of inference—*in accord with which* we reason—and can’t treat those *in general* as just more *premises from which* we reason.
- c) Nonetheless, we always can put any particular rule of inference into the form of a conditional premise, as long as we are left with at least *some* rules in the form of rules.
- d) There is a line of thought that says that unless we do, our inferences will be *enthymemes*. That is, they will have ‘suppressed’ or ‘missing’ premises. This view depends on *formalism* about the goodness of inferences: the claim that the only sort of good inferences are *formally* good inferences: inferences that are instances of some *formally valid* pattern or form of reasoning. An inference is “*enthymematic*” if it is not formally valid, but seems to be a good inference. The ‘missing’ premises are those that must be supplied—those it is taken that those committed to the propriety of the inference must implicitly be endorsing—to turn the result into a formally valid one. Of course, we can *always* do that (subject to assumptions of finiteness, or at least compactness) in the context of a system with the rule of *modus ponens* or detachment from conditionals, by supplying a suitable conditional.
- e) But there are problems with formalism about the goodness of inferences—a kind of totalitarianism that says *formal* goodness is the *only* kind of

goodness inferences can have. For instance: Massey's "Are There Any Good Arguments that Bad Arguments Are Bad?"

- f) And there is an alternative: allow *primitive material* proprieties of inference. Sellars insists that the inference from "It just rained," to "The streets will be wet," or from "Lightning now," to "Thunder soon," and from "Philadelphia is to the East of Pittsburgh," to "Pittsburgh is to the West of Philadelphia," are *not* enthymemes.
- g) And there is this to be said for acknowledging the existence of such material proprieties of inference: Given such material inferences, one can *define formal validity* of inference. An inference is formally good, good in virtue of its logical form, just in case:
 - i. It is a materially good inference; and
 - ii. It cannot be turned into a materially bad inference by substituting non-logical for non-logical vocabulary.

But there is no converse road, permitting the definition of materially good inferences from formally good ones.

- h) Notice that nothing in this definition turned on the privileged vocabulary held fixed upon substitution being *logical* vocabulary. The same trick can be used to define inferences good in virtue of their *theological*, or *astrological* form.

11. Q: What about going in the other direction? Can we transform PFWs into PAWWs? A: Yes. At least, we can get exactly the same effect as any PFW by looking instead at some *pattern* according to which we reason.

- a) For *any* actual premise *p*, in a field of *material* inference, we can always replace *p* by a principle that, for any inference in which it appears as a *premise*, ...*p*... \therefore *q*, simply add ... \therefore *q* as a primitively good material inference. The collection of all of them will then be an inference principle equivalent inferentially to the presence of *p*. Then instead of premises we take true, we'll have inferential principles we endorse.
- b) Q: But is the collection, possibly infinite, of inferences we get by doing that a 'principle' of inference? A: It is at least a *pattern* of inference.
- c) Q: So what are the rules of this game? If we insist that each inference be an instance of a *formally valid scheme* (cf. "Are There Any Good Arguments That Bad Arguments Are Bad?"), then we can't do this move, because the resulting inference will be *logical enthymemes*. (Another related question: outside the *logical* framework, once we allow *material* inferences, what does it even *mean* to call something an 'enthymeme'? But *if*, as Sellars wants, we allow *primitive material* proprieties of inference, then *this* move is *not* available.
- d) Q: Can we make this premise \rightarrow pattern, PFW \rightarrow PAWW move for *all* the sentences of the language? (We have seen that we can do it for *each*.) A: In some abstract sense, Yes. But really No. (Let's assume all the inferences have a finite number of premises.) Think of the field of materially good inferences as being sentences of the form $p_1 \& p_2 \& \dots \& p_n \therefore q$, and Ramsify on everything but \dots . Then get an expression of *pure*

inferential functional roles of all sentences in the language. Are they *all PAWWs, patterns, (pure patterns)*? But there is pretty clearly no *content* left here: *all* we get are *functional roles* w/res to inferential relations.

- e) Q: What if 'p' expresses a *conclusion*? A: The Ramsification trick is indifferent to that. But the 'PAWW' characterization does *not* seem apt.
- f) Q: Is it an issue that the *pattern* I get by the PFW→PAWW move to patterns *cannot*, at least not smoothly, be codified, formulated, or expressed as a *principle*? A: Since we *can* formulate modal claims, *if* they are PrinAWWs, *then* they *are* different from what we get if we *arbitrarily* apply the mechanical procedure for transforming PFWs into PattAWWs. It is only for the latter kind that the issue: PFW or PrinAWW? arises. So, to restate the problem: in the sub-case, how do we *tell* (not an epistemic issue: what does the difference *consist* in?) whether we have a PFW or a PrinAWW?
- g)

12. So we can't turn *all* principles *according to which* we reason into premises *from which* we reason, or *vice versa*. But as long as we don't try to do it with *everything*, we can trade off between the two categories, according to general, systematic techniques. So what is the force of the claim that some particular sort of expression *belongs* in one or the other category?

- a) Example: My treatment of *normative* vocabulary as codifying *patterns of practical reasoning*. [Explain.]
- b) My account of the role in practical reasoning distinctive of normative vocabulary (including 'desire' and 'preference'). One *can* treat 'ought' of various sorts as PrinAWWs (*not* just PattAWWs) rather than PFWs.
- c) Could replace: "I desire to stay dry" by "if X will (help) keep me dry, then do X": "Bank employees are obliged to wear neckties" by "if X is a bank employee and goes to work, X *should* wear a necktie."
- d) Q: Can I replace 'X is an acid' by 'X is, if sour, then red'? No, because *everything* is *that*. (cf. McD) 'X tastes sour, so it is red.'
 $\forall x[\text{sour}(x) \rightarrow \text{red}(x)]$ is a *presupposition* of the applicability of that concept:
So **acid(t) ≈ sour(t) & $\forall x[\text{sour}(x) \rightarrow \text{red}(x)]$**
- e)

13. "If one is going to explain our thinking in causal matters by using the idea of physical entailment, one must do more than defend the idea that "there are" such entailments; *one must make plausible the idea that these entailments play a role in causal reasoning analogous to the role of 'formal' entailments in less problematic forms of inference.*" [§59]