



4

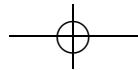
Modality and Normativity: From Hume and Quine to Kant and Sellars

1 The modal revolution

The status and respectability of alethic *modality* was always a point of contention and divergence between naturalism and empiricism.¹ It poses no problems in principle for *naturalism*, since modal vocabulary is an integral part of all the candidate naturalistic base vocabularies. Fundamental physics is, above all, a language of *laws*; the special sciences distinguish between true and false *counterfactual* claims; and ordinary empirical talk is richly *dispositional*. By contrast, modality has been a stumbling block for the *empiricist* tradition ever since Hume forcefully formulated his epistemological, and ultimately semantic, objections to the concepts of law and necessary connection.

Those traditional reservations about the intelligibility of modal notions were underscored, reinforced, and confirmed for twentieth-century versions of empiricism, which had been distinguished, strengthened, and made more precise by the addition of the semantic logicist model of the conceptual articulation of empirical content. Extensional, first-order quantificational languages could express *regularities* and *generalizations* with hitherto undreamed of power and precision. But for philosophers from Russell through Carnap to Quine, that just made it all the more urgent to explain, or explain away, the *lawlikeness* or counterfactual-supporting *necessity* distinctive of at least *some* of those generalizations, which demonstrably

¹ This tension was a principal source of conflict within the Vienna Circle, dividing Neurath and Schlick, for instance, with Carnap trying to mediate.

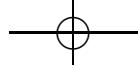


extended beyond what can be captured by the expressive resources of that logical vocabulary.²

This confluence of traditional empiricist with logicist difficulties concerning the content expressed by modal vocabulary had the result that for roughly the first two-thirds of the twentieth century, Anglophone philosophy regarded alethic modal vocabulary with extreme suspicion, if not outright hostility. It ranked, with normative vocabulary, as among the most mysterious and philosophically puzzling forms of discourse, the source of central standing and outstanding philosophical problems, as a prime candidate for the analytic project of semantic clarification in favored terms or, failing that, principled elimination from perspicuous discourse, as Quine famously recommended.

But philosophical attitudes towards modality underwent a remarkable, in many ways unprecedently radical, transformation during the twentieth century. For, starting in the second half of the century and accelerating through the last third, modal vocabulary became the analytic semantist's best friend, and an essential part of the contemporary philosopher's metaconceptual toolkit. I think it is worthwhile reminding ourselves just how surprised and astonished philosophers who lived and moved and had their being in the earlier milieu would have been to discover that, by the end of their century, when questions were raised about the semantics of some vocabulary—for instance, normative, intentional, or even semantic vocabulary itself—not only the dominant strategy, but the very first recourse would be to appeal to *modal* notions such as dispositions, counterfactual dependencies, and nomological relations to explain the questionable conceptual contents. Just how—they would want to know—did what seemed most urgently in need of philosophical explanation and defense suddenly become transformed so as to be unproblematically available to explain other puzzling phenomena? Surely such a major transformation of *explanandum* into *explanans* could not be the result merely of a change of fashion, the onset of amnesia, or the accumulation of fatigue? But if not, what secret did we find out, what new understanding did we achieve, to *justify* this change of philosophical attitude and practice?

² We now know, thanks to Danielle Macbeth's *Frege's Logic*, that Frege's own *Begriffsschrift* notation did not share the expressive impoverishment with respect to modality exhibited by the extensional first-order logic that Russell, and following him, everyone else, drew from it.



Two answers to this question lie ready to hand. First, there was a formal-semantic revolution in modal logic. And second, the Anglophone tradition more or less gave up empiricism in favor of naturalism. I think both those explanations are right, as far as they go, both as a matter of historical fact and in the order of justification. But it is important to understand exactly *which* questions those developments *did* offer responsive answers to, and to which they did *not*.

As to the first point, I think there is a widespread tendency to think that, to paraphrase Alexander Pope:

Modality and Nature's laws lay hid in night,
God said: "Let Kripke be!" and all was light.

But that cannot be right. Kripke's provision of a complete extensional semantic metavocabulary for intensional modal logical vocabulary—and its powerful development, by others such as Montague, Scott, Kaplan, Lewis, and Stalnaker, into a general intensional semantics for non-logical vocabulary—is an adequate response to worries stemming from the *extensional* character of the *logical vocabulary* in which semantics had been conducted. That is, it addresses the difficulties on the *semantic logicist* side of the classical project of analysis that stem from the expressive impoverishment of first-order logical vocabulary. But these formal developments do *not* provide an adequate response to residual *empiricist* worries about the intelligibility of modal concepts. For the extensionality of the semantic metalanguage for modality is bought at the price of making free use of modal primitives: most centrally, the notion of a possible world (as well as that of accessibility relations among such *possibilia*). As Quine emphasized, the modal vocabulary whose use is essential to this semantic approach evidently falls within the circle of terms and concepts to which empiricist suspicions and questions apply. That is, even putting *ontological* issues aside, whether possible worlds are thought of as abstract objects, as concrete particulars spatio-temporally unconnected to our universe, or as *sui generis possibilia*, the *epistemological* question of how we are to understand the possibility of our *knowing* anything about such items (and their accessibility relations), and the question how, if the possibility of such *cognitive* contact is mysterious, the idea of our having the *semantic* contact necessary so much as to *talk* or *think* about them can be made intelligible, are wholly





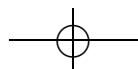
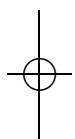
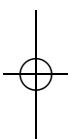
untouched by this formal apparatus, and remain every bit as pressing as before.

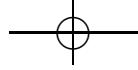
2 The modal Kant–Sellars thesis

How urgent those questions are depends on whether we have grounds to accept criticisms of the empiricist program that undermine the basis for its relegation of modal vocabulary to a suspect, second-class status. I think that the best justification for our new comfort with modal idioms is indeed to be found in the principled rejection of some of the crucial presuppositions of the empiricist critique of the credentials of modal concepts. We can now see that the operative core of both Quine’s and Sellars’s arguments against empiricism consists in objections to its underlying *semantic atomism*.³ Arguing that meaning must at least determine inferential role, and noticing that what follows from or is evidence for or against a claim depends on what other claims are available as auxiliary hypotheses or collateral premises, Quine concludes that the smallest unit of meaning is not a sentence, even in the case of observation sentences, but what he calls a ‘theory’: the whole constellation of all sentences held true, articulated by their inferential relations both to one another and to sentences not held true. Sellars argues that even observational beliefs acquired non-inferentially through perception can be understood as *conceptually* contentful—and hence potentially cognitively significant—only in virtue of their inferential relations to other possible beliefs. He concludes that non-inferential reports, no matter what their subject matter, cannot constitute an autonomous discursive practice: a language-game one could play though one played no other.

It is clear, I take it, how these anti-atomist arguments bear against empiricist foundationalism: the layer-cake picture of a semantically autonomous base of perceptual experience or reports thereof, on which is erected a semantically optional superstructure, in effect, of theories inferentially based on those observations. And insofar as empiricist worries about the status of laws, necessary connections, dispositions, and counterfactual possibilities

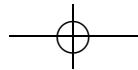
³ In their classic papers of the 1950s, “Two Dogmas of Empiricism” (1953) and “Empiricism and the Philosophy of Mind” (1956).

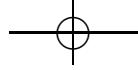




are predicated on the difficulty of justifying the inferences that would add them to the supposedly semantically autonomous base of non-modal reports of actual experiences, Quine's and Sellars's assault on the layer-cake picture, if successful, undercuts those worries by removing the motivation for their ultimately unmeetable constraints on an account of what modal vocabulary expresses. Thought of this way, though, criticism of the semantic presuppositions of the empiricist project does not bear any more directly on its treatment of *modal* vocabulary than it does on its treatment of any other potentially puzzling candidate for empiricist explication: *theoretical* (that is, non-observational, exclusively inferentially applicable) vocabulary, *normative* vocabulary, *probabilistic* vocabulary, and so on.

But there is another, much more intimate and immediate *positive* connection between arguments against semantic atomism and our understanding of what is expressed by the use of modal vocabulary. And it is here that I think we can find the best justification for our current relaxed attitude toward, and even enthusiastic embrace of, modal idioms as suitable tools for serious analytic semantic work. The underlying idea is what I will call the “Kant-Sellars thesis about modality.” Hume found that even his best understanding of actual observable empirical *facts* did not yield an understanding of *rules* relating or otherwise governing them. Those facts did not settle which of the things that *actually* happened *had* to happen (given others), that is, were (at least conditionally) *necessary*, and which of the things that did *not* happen nonetheless were *possible* (not ruled out by laws concerning what did happen). Though initially couched as an *epistemological* question about how one could *know* what rules or laws were in play, Hume's worries run deeper, raising the *semantic* question of what it could so much as *mean* to say that the facts are governed or related by rules or laws. Hume (and, following him, Quine) took it that epistemologically and semantically fastidious philosophers faced a stark choice: either show how to explain modality in non-modal terms or learn to live without it. But that challenge is predicated on the idea of an independently and antecedently intelligible stratum of empirical discourse that is purely descriptive and involves no modal commitments, as a semantically autonomous background and model with which the credentials of modal discourse can then be invidiously compared. One of Kant's most basic ideas, revived by Sellars, is that this idea is mistaken. **The ability to use ordinary empirical descriptive terms such as ‘green’, ‘rigid’, and ‘mass’ already presupposes grasp**





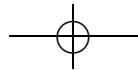
of the kinds of properties and relations made explicit by modal vocabulary. Sellars summed up the claim admirably in the title of one of his early papers: “Concepts as Involving Laws, and Inconceivable Without Them.”⁴

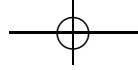
Kant was struck by the fact that the essence of the Newtonian concept of mass is of something that, by law, *force* is both necessary and sufficient to *accelerate*. And he saw that all empirical concepts are like their refined descendants in the mathematized natural sciences in this respect: their application implicitly involves counterfactual-supporting dispositional commitments to what *would* happen *if*. Kant’s claim, put in more contemporary terms, is that an integral part of what one is committed to in applying any determinate concept in empirical circumstances is drawing a distinction between counterfactual differences in circumstances that *would* and those that *would not* affect the truth of the judgment one is making. One has not grasped the concept cat unless one knows that it would still be possible for the cat to be on the mat if the lighting had been slightly different, but not if all life on earth had been extinguished by an asteroid-strike.⁵

In an autobiographical sketch, Sellars dates his break with traditional empiricism to his Oxford days in the 1930s. It was, he says, prompted by concern with the sort of content that ought to be associated with logical, causal, and deontological modalities. Already at that point he had the idea that

⁴ First published in 1948; reprinted at pp. 87–124 in J. Sicha (ed.), *Pure Pragmatics and Possible Worlds: The Early Essays of Wilfrid Sellars* (Ridgeview Publishing Co., 1980), hereafter *PPPW*. This slogan is a good place to start in thinking about Kant’s point, but in fact Sellars’s own view is subtly but importantly different from Kant’s. For Sellars, the laws determining the truth of counterfactuals involving the application of a concept are part of the content of the concept. For Kant, modal concepts make explicit not something implicit in the *content* of determinate concepts, but something implicit in their *empirical use*, in *applying* them to make empirical *judgments*. That is why the pure concepts of the understanding—what he calls ‘categories’, such as possibility and necessity—both are to be understood in terms of the forms of judgment (the table of categories derives from the table of judgments) and express synthetic, rather than analytic, necessities. From Kant’s point of view, a better slogan than Sellars’s would be “The Use of Concepts in Empirical Judgments as Involving Laws and Inconceivable Without Them.”

⁵ It is this observation, unwittingly underscored by Hume (for Kant, the Moses who brought us to within sight of the Promised Land he himself was destined not to enter), that motivates Kant to wheel in his heavy transcendental machinery. For Kant sought to explain the modal commitments implicit in the application of ordinary empirical concepts by placing the modal concepts of law and necessity in the newly postulated realm of *pure* concepts or categories, which must be graspable a priori precisely in the sense that their applicability is presupposed by the applicability of *any* empirical concepts. The concept of vocabularies that are “universally LX,” introduced below, is a successor notion along at least one important dimension.





what was needed was a functional theory of concepts which would make their role in reasoning, rather than supposed origin in experience, their primary feature.⁶

Somewhat more specifically, he sees modal locutions as tools used in the enterprise of

making explicit the rules we have adopted for thought and action. ... I shall be interpreting our judgments to the effect that A causally necessitates B as the expression of a rule governing our use of the terms 'A' and 'B'.⁷

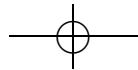
In fact, following Ryle,⁸ he takes modal expressions to function as *inference licenses*, expressing our commitment to the goodness of counterfactually robust inferences from necessitating to necessitated conditions. If and insofar as it could be established that their involvement in such counterfactually robust inferences is essential to the *contents* of ordinary empirical concepts, then what is made explicit by modal vocabulary is implicit in the use of any such concepts. That is the claim I am calling the "Kant-Sellars thesis." On this view, modal vocabulary does not just add to the use of ordinary empirical observational vocabulary a range of expressive power that is *extraneous*—as though one were adding, say, *culinary* to *nautical* vocabulary. Rather, the expressive job distinctive of modal vocabulary is to articulate just the kind of essential semantic connections among empirical concepts that Sellars (and Quine) point to, and whose existence semantic atomism is principally concerned to deny.

As I would like to formulate it, the Kant-Sellars thesis begins with the claim that in using ordinary empirical vocabulary, one already knows how to do everything one needs to know how to do in order to introduce and deploy modal vocabulary. If that is right, then one cannot be in the position the atomist (for instance, empiricist) critic of modality professes to find himself in: having fully understood and mastered the use of *non-modal* vocabulary, but having thereby afforded himself no grip on the use of *modal* vocabulary, and no access to what it expresses. The Humean-Quinean predicament is accordingly diagnosed as resulting from a failure properly to understand the relation between modal vocabulary and what one must *do* in order to *deploy* non-modal, empirical, descriptive vocabulary.

⁶ In H. N. Castañeda (ed.), *Action, Knowledge, and Reality* (Bobbs-Merrill, 1975), 285.

⁷ Sellars, "Language, Rules, and Behavior" (1949), reprinted in *PPPW*, 136, fn. 2.

⁸ Gilbert Ryle, "'If', 'So', and 'Because,'" in Max Black (ed.), *Philosophical Analysis* (Prentice Hall, 1950), 302–18.



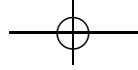
The thought that the expressive role characteristic of alethic modal vocabulary is to make explicit semantic or conceptual connections and commitments that are already implicit in the use of ordinary (apparently) non-modal empirical vocabulary faces at the outset at least two sorts of potentially weighty objection. First, did not Kripke's semantic investigations of modally rigid designators reveal the sort of necessity they articulate as being *metaphysical*, specifically by contrast to the sort of *conceptual* necessity that Quine, for instance, had worried about and rejected? And second, to talk about what is necessary and possible is not to *say* anything about rules for using linguistic expressions, or about what anyone is committed to, since the objective modal claims in question could have been true even if there had never been language users, linguistic expressions, rules, or commitments.

As to the first objection, the philosophical phase of the modal revolution (developing the earlier logical and semantic phases of that revolution) that Kripke precipitated in “Naming and Necessity”⁹ did indeed use the semantic phenomenon of the modal rigidity of some non-descriptive vocabulary to articulate a kind of necessity that is knowable only *a posteriori*. The conclusion that such necessity should not be understood as *conceptual* necessity follows only if one either identifies conceptual content with *descriptive* content (by contrast to the causally-historically acquired content of proper names and demonstratives) or takes it (as Quine, following the tradition, had) that conceptual connections must be knowable *a priori* by those who have mastered those concepts. But both of these are optional commitments, which can and should be rejected by anyone trying to follow out the Kant-Sellars line of thought about modality. McDowell has argued, to my mind convincingly, that the content expressed by demonstrative vocabulary should be understood as thoroughly conceptual (and that Frege already took it to be so).¹⁰ And in *Making It Explicit*, I articulate a broadly inferential notion of the conceptual that incorporates the indirectly inferential roles of substitution and anaphora—including the anaphoric phenomenon that is modal rigidity.¹¹

⁹ In Donald Davidson and Gilbert Harman (eds.), *Semantics of Natural Language*. Dordrecht: Reidel, 1972.

¹⁰ “*De Re* Senses,” in *John McDowell, Meaning, Knowledge, and Reality* (Harvard University Press, 2001).

¹¹ Robert Brandom, *Making It Explicit* (Harvard University Press, 1994), chapters 6, 7 (especially sections III and IV), and 8 (section V).



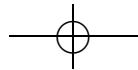
On the other point, Sellars's forthright response to Quine's pragmatic challenge in "Two Dogmas of Empiricism"—to say what it is about the *use* of expressions that distinguishes inferences underwritten by necessary conceptual relations from those underwritten by contingent matter-of-factual ones—is to identify the concept-articulating inferences as those that are counterfactually robust.¹² He cheerfully embraces the consequence that to discover what is contained in the concept *copper* one needs empirically to investigate the laws of nature. (This is a kind of semantic 'externalism' that does not need to take on the dangerous and difficult task of making sense of a notion of the 'internal' with which to contrast.) The issue about conceptual necessities here is not an empirical one: who is right about the conceptual? The Kant-Sellars thesis about modality requires deploying a concept of the conceptual that differs in important ways from the traditional one. As long as such a notion can be intelligibly developed and consistently applied, those differences need only be kept firmly in mind, not counted as fatal flaws.

The response to the second objection (that saying what is necessary or possible is not saying anything about how anyone talks) must be to be clearer about the sort of pragmatically mediated semantic relation the Kant-Sellars thesis takes modal vocabulary to stand in to ordinary, non-modal descriptive vocabulary. The large claim in the vicinity—one that will occupy me not only in this lecture but beyond—is, as Sellars puts it, that "the language of modality is ... a 'transposed' language of norms."¹³ I do not think that Sellars himself ever manages to say clearly just what sort of 'transposition' he has in mind. He appeals to a distinction between what is *said* by the use of some vocabulary, and what is *conveyed* by its use. While admitting that talk of what is necessary does not *say* anything about what language users ought or ought not to do, he nonetheless insists that it "conveys the same information" as "rules to the effect that we may do thus and so, and ought not do this and that, in the way of manipulating expressions in a language."¹⁴ His (only somewhat helpful) example is that when I say "The sky is clear," I have both said something about the weather and conveyed something about my beliefs. The point, I take it, is to distinguish what follows *semantically* from the content of what I have

¹² Sellars, "Is There a Synthetic *A Priori*?" *Philosophical Studies*, 20 (1953), 121–38.

¹³ Sellars, "Inference and Meaning" (1953), reprinted in Kevin Scharp and Robert Brandom (eds.), *In the Space of Reasons* (Harvard University Press, 2007), 21.

¹⁴ *Ibid.*



said from what follows *pragmatically* from my *saying* of it. (Embedding the claims as the antecedents of conditionals will distinguish these two sorts of consequences. “If the sky is clear, then it will not rain,” expresses a good inference, whereas “If the sky is clear, then Brandom believes that the sky is clear,” does not. For only the semantic content, and not the pragmatic force of the utterance, survives such embedding.)

We are in a position to be a little clearer about what Sellars is after with his dark notion of what an utterance ‘conveys’. The view is that what I am *doing* when I say that it is causally necessary that if this piece of copper is heated to 1084 °C, it will melt, is endorsing a certain kind of inference. I am not *saying that* that inference is good; the facts about copper would be as they are even if there were no inferrers or inferrings. When Sellars says “the language of modality is … a ‘transposed’ language of norms,” he is saying in my terms that normative vocabulary codifying rules of inference is a *pragmatic metavocabulary* for modal vocabulary. His ‘transposition’ is just this pragmatically mediated semantic relation between deontic normative and alethic modal vocabulary. The corresponding meaning-use diagram (MUD) is shown in Figure 4.1.

In Lecture 5, I will show in detail how this thought can be exploited to develop a new sort of formal semantics, which yields new insights into the conceptual contents expressed by ordinary empirical-descriptive vocabulary, as well as logical and modal vocabularies.

This claim is merely part of the background of what I have been calling the “Kant-Sellars thesis” about modality, however. That thesis comprises two claims:

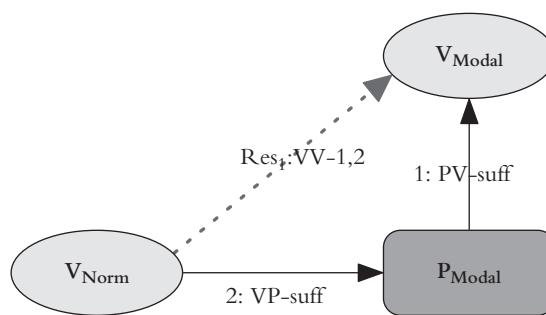


Figure 4.1 “The language of modalities is a ‘transposed’ language of norms”

1. In using ordinary empirical vocabulary, one already knows how to do everything one needs to know how to do in order to introduce and deploy *modal* vocabulary.
2. The expressive role characteristic of alethic modal vocabulary is to *make explicit* semantic, conceptual connections and commitments that are already *implicit* in the use of ordinary empirical vocabulary.

The first says that some practices that are PV-necessary for the use of any empirical vocabulary are PP-sufficient for practices that are PV-sufficient to deploy modal vocabulary. The second says that that modal vocabulary then makes explicit those aspects of practices-or-abilities that are implicit in the use of any empirical vocabulary. In fact, these are ways of saying that modal vocabulary stands to ordinary empirical vocabulary in the complex, pragmatically mediated semantic relation I have already identified (in Lecture 2) as elaborating-explicating: the meaning-use relation called 'LX' for short. The corresponding MUD is shown in Figure 4.2.

Combining these claims yields a MUD asserting relations among modal, normative, and empirical vocabularies (Figure 4.3).

3 Counterfactual robustness and the updating argument

So far, I have only expounded, explicated, and mentioned some of the consequences of the Kant-Sellars thesis about modal vocabulary, but not

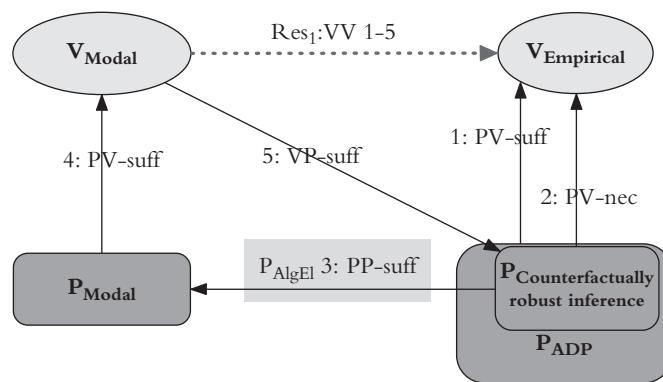


Figure 4.2 The Kant-Sellars thesis: modal vocabulary is elaborated-explicating (LX)

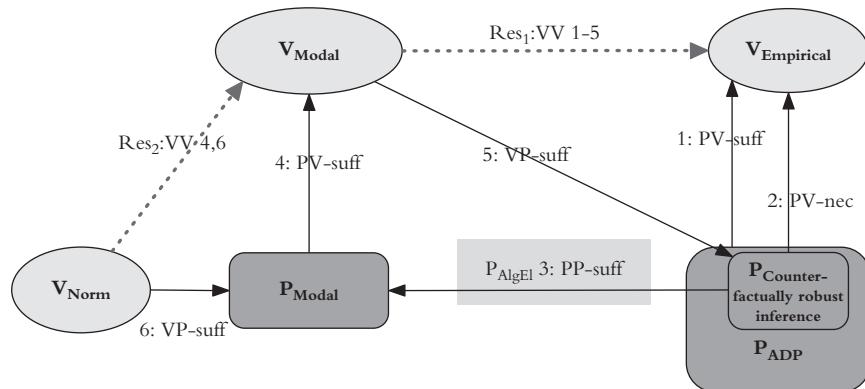
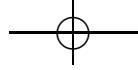


Figure 4.3 Modal, normative, and empirical vocabulary

sought to *argue* for it. What reason is there to think that it is *true*? The analysis of the Kant-Sellars thesis as asserting a complex pragmatically mediated semantic relation between vocabularies that is the resultant of a definite constellation of basic meaning-use relations, as presented in the MUD, tells us exactly what shape such an argument must have. For it tells us just which *basic* meaning-use relations must be established in order to show that the *resultant* one obtains. The key element in this case will be finding some set of practices that can be argued to be at once contained in or exhibited by every autonomous discursive practice, and PP-sufficient for practices PV-sufficient for deploying explicitly modal vocabulary, which is VP-sufficient to specify the original PV-necessary practices-or-abilities. As the labels on the MUDs indicate, for the argument I will mount (picking up on the theme on which my argument last time turned), those practices are *counterfactually robust inferential* practices-or-abilities—more specifically, the practical capacity to associate with materially good inferences *ranges of counterfactual robustness*. If it can be established that deploying any ordinary empirical vocabulary presupposes these practices-or-abilities, and that they in turn suffice to introduce explicit modally qualified conditionals that permit the expression of those practical discriminations, then the universal LX-character of modal vocabulary relative to ordinary empirical vocabulary will have been demonstrated.

I have already claimed that any autonomous discursive practice (ADP) must include practices-or-abilities of distinguishing some inferences as materially good from others that are not. For some bit of vocabulary to



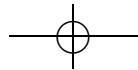
function as a propositionally contentful declarative sentence is for it to be available to serve as the premise and conclusion of such material inferences. Further, it is the expressive job generically characteristic of *conditional* vocabulary to *codify* endorsements of material inferences: to make them explicit in the form of declarative sentences that can themselves serve as the premises and conclusions of inferences. The philosopher most responsible for getting us to think about conditionals in this way is Gilbert Ryle. In his classic essay “‘If’, ‘So’, and ‘Because’,” in which he introduces the idea of hypothetical statements as inference tickets or licenses, he also points out an intimate connection between them and *modal* claims. He says:

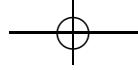
We have another familiar way of wording hypothetical statements. Although the standard textbooks discuss “modal propositions” in a different chapter from that in which they discuss hypotheticals, the differences between modal and hypothetical statements are in fact purely stylistic. There is only one colloquial way of correctly negating the superstitious hypothetical statement “If a person walks under a ladder, he comes to grief before the day is out,” namely, by saying “No, a person may (might, or could) walk under a ladder and not come to grief.” And the only colloquial way of putting a question to which an “if-then” statement is the required affirmative answer is to ask, for example, “Can an Oxford Vice-Chancellor not be (or need he be) a Head of College?” ... [W]e always can reword an “if-then” statement as a statement of the pattern “It cannot be Monday today and not be Tuesday tomorrow.”¹⁵

I think he is right that “It is possible that (p and $\neg q$)” is incompatible with “if p then q ” when the latter is used to codify an ordinary material inference such as the inference from a banana’s being yellow to its being ripe. Endorsing a material inference does involve a commitment of the sort made explicit by the use of modal vocabulary, about what is and is not possible, and what is at least conditionally necessary.

For this reason, the fact that we cannot intelligibly describe someone as deploying a concept unless he makes some distinction between materially good and bad inferences involving it has the consequence that we also cannot understand the practitioner as deploying the concept unless he treats the material inferences he takes to be good as having a certain *range of counterfactual robustness*, that is, as remaining good under various merely

¹⁵ Gilbert Ryle, “‘If’, ‘So’, and ‘Because’,” 313.



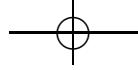


hypothetical circumstances. One grasps the claim “the lioness is hungry” only insofar as one takes it to have various consequences (which *would* be true if it *were* true) and to rule out some others (which *would not* be true if it *were* true). And it is not intelligible that one should endorse as materially good an inference involving it, such as the inference from “the lioness is hungry” to “nearby prey animals visible to and accessible by the lioness are in danger of being eaten,” but be disposed to make no distinction at all between collateral premises that would, and those that would not, if true infirm the inference. One must make *some* distinction such as that the inference would still go through if the lioness were standing two inches to the east of her actual position, the day happened to be a Tuesday, or a small tree ten miles away cast its shadow over a beetle, but not if she were shot with a tranquilizing dart, the temperature instantly plummeted 300 degrees, or a plane crashed, crushing her. The claim is not that one could not fail to assess some, or even all, of *these particular* counterfactuals correctly and still count as grasping the claim that is their premise, but that one could not so qualify if one made *no* such distinctions.

It may initially be tempting to think that the inferences that are counterfactually robust are all and only those underwritten by *laws*. Thus inferences underwritten by the law that all samples of copper melt at 1083.4 °C are counterfactually robust: if this coin (which in fact is silver) *were* made of copper, it *would* melt at 1083.4 °C. Whereas inferences underwritten by the accidental regularity that all the coins in my pocket are copper are not counterfactually robust: if I *were* to put this coin (which in fact is silver) in my pocket, it *would not* be copper. There are indeed real and significant differences between these cases, but I think it is important not to think of them in terms of the difference between inferences that *are* counterfactually robust and inferences that *are not*. The difference is rather one of the character of the particular *ranges* of counterfactual robustness. For the accidental generalization that all the coins in my pocket are copper *does* underwrite counterfactuals such as: “If I *were* to choose a coin at random from my pocket, it *would* be copper.” In fact, *every* claim, whether contingent or not, supports *some* counterfactual inferences, and if one grasped *none* of them one would not qualify as understanding those claims.

I think these considerations suffice to establish that autonomous discursive practices *essentially*, and not just *accidentally*, involve the association of ranges of counterfactual robustness with at least some material inferences. If, as



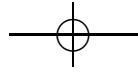


Ryle claims, and as is in any case plausible, *modal* vocabulary specifying what is, at least, conditionally possible and necessary can then be introduced to make explicit those commitments to the, at least, limited counterfactual goodness of material inferences, then we have what is needed for the modal Kant-Sellars thesis. But I think that if we dig deeper, we can learn more. So, rather than leaving things at this point, I want to consider a more detailed line of argument for this, the most potentially controversial element of the complex meaning-use relation that thesis asserts.

For the first premise, I take it to be clear that every autonomous discursive practice must have some vocabulary that can be used *observationally*, in reliably differentially elicited non-inferential reports. This is the core of what I have been referring to as “ordinary empirical vocabulary.” Second, I have already argued that those who engage in any discursive practices must distinguish in practice between materially good and materially bad inferences—where calling them ‘material’ just means that the presence of some non-logical vocabulary is essential to the classification. Recall that this is not to claim that they must have a view about the goodness or badness of every possible candidate material inference; there can be some about which they have no view. And it is not to claim that they always are correct about the goodness of the inferences toward which they do have attitudes. But to count as deploying any vocabulary at all, one must treat some inferences involving it as good and others as bad. Otherwise, one’s utterances are wholly devoid of conceptual content; whatever pragmatic significance they may have, it cannot be thought of as *discursive* significance. Even tokenings that are non-inferentially elicited by environing stimuli—that is, the applications of observational vocabulary mentioned in the first premise—must have inferential *consequences*, if they are not to be cognitively idle.

The third claim is that material inference is in general *non-monotonic*. That is, the inference from p to q may be materially good, even though the inference from $p \& r$ to q is not. Monotonicity of inference is, of course, a familiar feature of inferences within a formal *logical* system, and in mathematical reasoning; and that feature is arguably inherited by fundamental physics. But in the special sciences inferences are almost always *defeasible*, by collateral circumstances that thereby count as ‘special’. Each stage in a physician’s differential diagnosis is like this: the inference from test result, physical finding, or symptom is surrounded by a nimbus

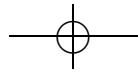


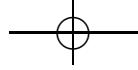


of usually unspoken ‘unless’es. And no one supposes that such probative reasoning can always be turned into dispositive reasoning by making an explicit, exhaustive list of the potential defeasors. Certainly, reasoning in everyday life does not generally admit such completions. If I strike this dry, well-made match, it will light—unless it is done inside a strong magnetic field. But it still will light if, in addition, it is struck inside a Faraday cage—unless there is not enough oxygen. And so on. There need be no definite totality of possible defeasors, specifiable in advance. Even where we have some idea how to enumerate them, unless those provisos are generally left implicit, actually stating the premises so as to draw inferences from them monotonically is impossibly cumbersome in practice.

At this point, one is liable to think of *ceteris paribus* clauses. The careful way to formulate the ordinary inference just mentioned is that if I strike this dry, well-made match, *ceteris paribus*, or other things being equal, it will light. I think that is indeed exactly what we ought to say, and the point I want to make can be made by saying that what such *ceteris paribus* clauses mark is an unavoidable feature of ordinary material inferences. But it is critical to understand what such clauses *do* and do *not* do. They are *not* devices for the wholesale stipulation of the denial of all of the potential defeasors that, even if exhaustively knowable and statable, if denied retail would make the inference unsurveyable. That is, they are not devices that *make* non-monotonic inferences monotonic. The proper term for a Latin phrase whose utterance could do *that* is ‘*magic spell*’. If it is thought of as a wholesale proviso covering all possible defeasors, the effect of adding ‘*ceteris paribus*’ to the statement of the inference that if I strike this dry, well-made match, then it will light, would be to say: “unless for some reason it doesn’t,” or “except in those circumstances when it doesn’t.” That is not producing an inference that is *monotonic*; it is producing one that is *trivial*. The real expressive function of *ceteris paribus* clauses is not *magically* to *remove* the non-monotonicity of material inferences, nor to replace them with other monotonic ones, but rather *explicitly* to *acknowledge* their non-monotonicity: to mark the inference being endorsed as one that has unspecified, but potentially important defeasors.¹⁶

¹⁶ For empirical claims involving theoretical vocabulary, this is obvious. For theoretical vocabulary is, by definition, vocabulary that can *only* correctly be applied as the conclusion of an inference. But the justification even of beliefs acquired *non-inferentially*, through observation typically will involve appealing to the reliability of the observer’s differential responsive dispositions to endorse such claims





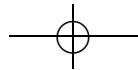
The fourth premise is that at any given time, many, if not most, of a subject's beliefs could only be justified by exhibiting them as the conclusions of material inferences. We might call a believer "epistemically responsible" insofar as she acknowledges a commitment to being able to justify many, if not most, of her beliefs, under suitable circumstances. My fifth premise is that in order to count as a discursive practitioner, one must be at least minimally epistemically responsible. Present purposes will not require that we attempt to quantify what the minimal level of such responsibility is.

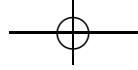
Developing a theme from my previous lecture, we can draw a preliminary conclusion. The five considerations advanced so far together entail that epistemically responsible believers face a potentially intractable *updating problem*. Every change of belief, no matter how small, is *potentially* relevant to the justification of every prior belief. Acquiring a new belief means acquiring what, for any material inference the believer endorses and relies upon for justification, might possibly turn out to be a defeasor. And giving up any belief means giving up not only a premise that might previously have been relied upon in justification, but also a potential counter-defeasor (for instance, a magnetic field's not being a defeasor to the match's lighting if there is a Faraday cage inside the field).

Now, it is not practically feasible explicitly to review all of one's beliefs every time one's beliefs change, in order to check which are and which are not still justifiable. If that were what epistemic responsibility demanded, then it would be a pointless, impossible ideal. Language users who *do* not (because they *cannot*) do *that*, must *practically* distinguish, among all the inferences that rationalize their current beliefs, which of them are update *candidates*, in the light of the current change of belief (let us say, for simplicity, a newly added belief). That is practically to associate with the new belief a set of material inferences of which it is a potential defeasor. The potential defeasors in this way associated with each material inference endorsed in turn define (by complementation) the range of counterfactual robustness practically associated with that inference.¹⁷

under a range of circumstances. The inference from my being a reliable reporter of red things in good light to my responsively elicited claim that something is red being true can be a good material inference. But it is non-monotonic, defeasible by a whole range of collateral circumstances.

¹⁷ Somewhat more carefully put: assuming some length restriction ensuring finiteness of the set of logically non-compound sentences involved, the ability to associate with each sentence a set of





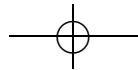
I conclude that in view of the non-monotonicity of material inference, the practical task of updating the rest of one's beliefs when some of them change is tractable in principle only if those who deploy a vocabulary practically discriminate ranges of counterfactual robustness for many of the material inferences they endorse. If that is right, then establishing the modal Kant-Sellars thesis requires further showing how to introduce modal vocabulary on the basis of such counterfactual conditionals, and how to use modal vocabulary to make those counterfactual conditionals explicit. Ryle's remarks suggest a strategy for both: treat "If p were true, q would be true," as equivalent to "It is not possible that p and $\neg q$." In the next lecture I will show how to follow out this strategy in detail, by treating the claim that q follows from p as equivalent to the claim that everything materially incompatible with q is materially incompatible with p —so that to say that "Coda is a dog" entails "Coda is a mammal" is to say that everything incompatible with his being a mammal is incompatible with his being a dog.

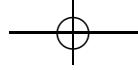
4 The normative Kant-Sellars thesis

Before turning to that project of connecting material inferential relations with an implicitly modal notion of material incompatibility, however, I want to consider an analogue of the Kant-Sellars thesis about *modal* vocabulary that applies instead to *normative* vocabulary.

Kant read Hume's theoretical and practical philosophies as raising variants of a single question. On the side of theoretical reasoning, Hume asks what our warrant is for moving from descriptions of what *in fact* happens to characterizations of what *must* happen, and what could not happen. How, he wants to know, can we rationally justify the move from descriptions of matter-of-factual regularities to formulations of necessary laws? On the side of practical reasoning, Hume asks what our warrant is for moving from descriptions of how things are to prescriptions of how they ought to be. How, he wants to know, can we rationally justify the move from 'is' to 'ought'? In Kant's terminology, these are both species of 'necessity': *practical*

inferences of which it is a potential defeasor can be algorithmically elaborated into (and hence is PP-sufficient for) the ability to associate with each inference a set of potential defeasors, and hence again, the set of non-defeasors.

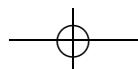




(including moral), and *natural* necessity, respectively. For him, ‘necessary’ (*notwendig*) just means “according to a *rule*.” Hume’s predicament is that he finds that even his best understanding of *facts* does not yield an understanding of *rules* governing and relating those facts, underwriting assessments of which of the things that actually happen (something we can experience) *must* happen (are *naturally* necessary), or *ought* to happen (are *normatively* or *practically* necessary).

As we have seen, on the modal side, Kant’s response is that Hume’s predicament is not a real one. One cannot in fact fully understand the descriptive, empirical employment of ordinary determinate concepts such as cat without at least implicitly understanding also what is made explicit by the modal concepts that articulate laws. Kant mounts a corresponding line of thought on the side of *normative* or *practical* necessity. Normative concepts make explicit commitments that are implicit in any use of concepts, whether theoretically in judgment or practically in acting intentionally—that is, in endorsing practical maxims. *Judgment* and *agency* are implicitly normative phenomena because they consist in the application of concepts, and applying concepts is undertaking commitments and responsibilities whose *content* is articulated by those concepts. (For Kant, specifically *moral* normative vocabulary makes explicit commitments that are already implicit in the practical use of concepts to endorse maxims, ends, and plans.)

I am not going to go into how Sellars builds on this thought, because I will develop it in a somewhat different way. Suffice it to say that in the light of Kant’s parallel responses to Hume’s parallel concerns with the credentials of modal and normative vocabulary—concerns couched in epistemological terms, but at base semantic in character—we can formulate a *normative* Kant-Sellars thesis by analogy to the *modal* one. It is the claim that in order to apply or deploy ordinary, empirical, descriptive vocabulary, including observational vocabulary—and hence, in order to deploy any autonomous vocabulary whatsoever—one must already be able to do everything needed to introduce normative vocabulary. Articulated in terms of meaning-use analysis, it is the claim that there are practices PV-necessary for engaging in any autonomous discursive practice that are PP-sufficient for practices PV-sufficient to deploy normative vocabulary. If, again by analogy to the modal case, we add the claim that normative vocabulary is VP-sufficient to *specify* those aspects of the practices that are PV-necessary for any ADP, we have the full-blown claim that normative vocabulary is elaborated-explicating,



or LX, for all autonomous vocabularies. The by-now familiar MUD for the resultant complex meaning-use relation among vocabularies is set out in Figure 4.4.

How might one argue for the normative Kant-Sellars thesis? I have been working all along with the idea that any autonomous set of practices can be intelligible as deploying a vocabulary—that is, as being *discursive* or *linguistic* practices—only insofar as those practices attribute to some performances the pragmatic significance of *assertions*, and that it is a necessary feature of that pragmatic significance that assertions can serve both as premises and conclusions of *inferences*. The notions of asserting and of inferring are, on this account, essentially and indissolubly linked. This is to say that every autonomous *discursive* practice must include core practices of *giving* and *asking for reasons*. It is playing a suitable role in such a constellation of practices that makes the sign-designs whose production can have in that context the pragmatic significance of being an assertion—something that can both serve as and stand in need of a reason—qualify as *declarative sentences*. And standing in those inferential (justificatory, evidential) relations is a necessary condition of those sentences being intelligible as expressing *propositional contents*.¹⁸

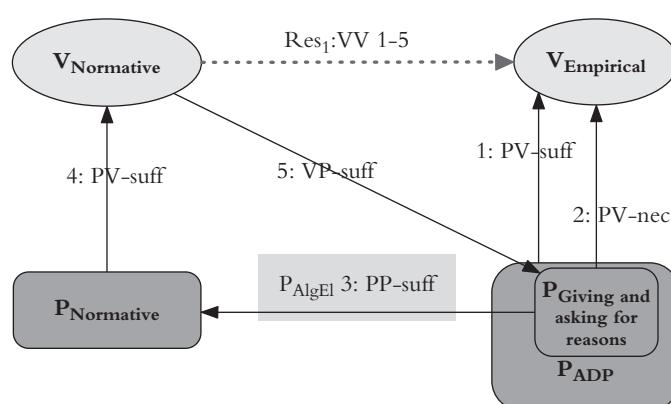
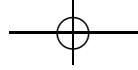


Figure 4.4 Normative Kant-Sellars thesis: normative vocabulary is elaborated-explicating (LX)

¹⁸ For my purposes here I do *not* need to claim that inferential articulation, broadly construed, is *sufficient* to constitute propositional content. I need only the weaker claim that it is a *necessary* feature: that nothing that could *not* play the role of premise and conclusion of an inference could be intelligible as propositionally contentful.

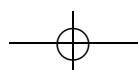
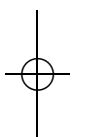
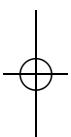


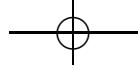
It is these core practices of giving and asking for reasons that I propose as being both PV-necessary for every autonomous discursive practice (as I have just been claiming) and PP-sufficient for, in the sense of algorithmically elaboratable into, practices PV-sufficient for the introduction of normative vocabulary, which can then serve explicitly to specify key features of those practices. In particular, I will argue that no set of practices is recognizable as a game of giving and asking for *reasons* for *assertions* unless it involves implicitly (practically) acknowledging at least two sorts of normative status, *commitments* and *entitlements*, and some general structures relating them.

Suppose we have a set of counters or markers such that producing or playing one has the social significance of making an assertional move in the game. We can call such counters ‘sentences’. Then, for any player at any time, there must be a way of partitioning sentences into two classes, by distinguishing somehow those that he is disposed or otherwise prepared to assert (perhaps when suitably prompted). These counters, which are distinguished by bearing the player’s mark, being on his list, or being kept in his box, constitute his score. By playing a new counter, making an assertion, one alters one’s own score, and perhaps that of others.

Here is my first claim: for such a game or set of toy practices to be recognizable as involving *assertions*, it must be the case that playing one counter, or otherwise adding it to one’s score, can *commit* one to playing others, or adding them to one’s score. If one asserts “The swatch is red,” one *ought* to add to one’s score also “The swatch is colored.” Making the one move *obliges* one to be prepared to make the other as well. This is not to say that all players actually *do* have the dispositions they *ought* to have. One might not act as one is committed or obliged to act; one can break or fail to follow this sort of rule of the game, at least in particular cases, without thereby being expelled from the company of players of the asserting game. Still, I claim, assertional games must have rules of this sort: rules of *consequential commitment*.

Why? Because to be recognizable as *assertional*, a move must not be *idle*, it must make a *difference*, it must have *consequences* for what else it is appropriate to do, according to the rules of the game. Assertions express judgments or beliefs. Putting a sentence on one’s list of judgments, putting it in one’s belief box, must have consequences for how *else* one ought, rationally, to act, judge, and believe. We may be able to construct cases where it is intelligible to attribute beliefs that are consequentially inert and



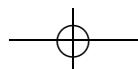
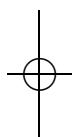
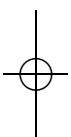


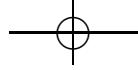
isolated from their fellows: “I just believe that cows look goofy, that’s all. Nothing follows from that, and I am not obliged to act in any particular way on that belief.” But *all* of our beliefs could not intelligibly be understood to be like this. If putting sentences onto my list or into my box *never* has consequences for what else belongs there, then we ought not to understand the list as consisting of my *judgments*, or the box as containing my *beliefs*.

Understanding a claim, the significance of an assertional move, requires understanding at least *some* of its consequences, knowing what *else* (what other moves) one would be committing oneself to by making that claim. A parrot, we can imagine, can produce an utterance perceptually indistinguishable from an assertion of “That’s red.” Our nonetheless not taking it to have asserted that sentence, not to have made a move in that game, *is* our taking it that, unaware as it is of the inferential involvements of the claim that it would be expressing, of what it would be committing itself to were it to make the claim, it has not thereby succeeded in committing itself to anything. Making that assertion is committing oneself to such consequences as that the swatch is colored, that it is not green, and so on.

For this reason we can understand making a claim as taking up a particular sort of normative stance towards an inferentially articulated content. It is *endorsing* it, taking *responsibility* for it, *committing* oneself to it. The difference between treating something as a claiming and treating it just as a brute sounding off, between treating it as making a move in the assertional game and treating it as an idle performance, is just whether one treats it as the undertaking of a commitment that is suitably articulated by its consequential relations to other commitments. These are *rational* relations, whereby undertaking one commitment *rationally* obliges one to undertake others, related to it as its inferential consequences. These relations at least partly articulate the *content* of the commitment or responsibility one undertakes by asserting a sentence. Apart from such relations, there is no such content, hence no assertion.

The next claim I want to make is that practices incorporating a game of giving and asking for reasons must involve acknowledgment of a *second* kind of normative status. I have said that making a move in the assertional game should be understood as acknowledging a certain sort of *commitment*, articulated by consequential inferential relations linking the asserted sentence to other sentences. But players of the game of giving and asking for reasons must also distinguish, among the commitments an





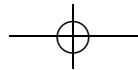
interlocutor undertakes, a distinguished sub-class to which she is *entitled*. *Linguistic rationalism* understands *assertions*, the fundamental sort of speech act, as essentially things that can both serve as and stand in need of *reasons*. Giving reasons for a claim is producing other assertions that *license* or *entitle* one to it, that *justify* it. Asking for reasons for a claim is asking for its *warrant*, for what *entitles* one to that commitment. Such a practice presupposes a distinction between assertional commitments to which one *is* entitled and those to which one is *not* entitled. Reason-giving practices make sense only if there can be an issue as to whether or not practitioners are entitled to their commitments.

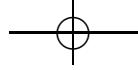
Indeed, I take it that liability to demands for *justification*, that is, demonstration of *entitlement*, is a major dimension of the *responsibility* one undertakes, the *commitment* one makes, in asserting something. In making an assertion one implicitly acknowledges the propriety, at least under some circumstances, of demands for reasons, for justification of the claim one has endorsed, the commitment one has undertaken. Besides the *committive* dimension of assertional practice, there is the *critical* dimension: the aspect of the practice in which the propriety of those commitments is assessed. Apart from this critical dimension, the notion of *reasons* gets no grip.

So the overall claim is that the sense of endorsement that determines the force of assertional speech acts involves, at a minimum, a kind of *commitment* the speaker's *entitlement* to which is always potentially at issue. The assertible contents expressed by declarative sentences whose utterance can have this sort of force must accordingly be inferentially articulated along both normative dimensions. *Downstream*, they must have inferential *consequences*, commitment to which is entailed by commitment to the original content. *Upstream*, they must have inferential *antecedents*, relations to contents that can serve as premises from which entitlement to the original content can be inherited.

5 Conclusion

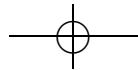
If that is right, then discursive practitioners as such must be able in practice to take or treat each other and themselves as exhibiting *normative* statuses: as being *committed* and *entitled* to contents expressed by the declarative

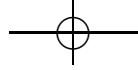




sentences whose free-standing utterance has the pragmatic significance of acknowledging commitments and claiming entitlements. Since, by hypothesis, these practitioners can already make assertions, the introduction of normative vocabulary permitting one explicitly to *say that* someone is committed or entitled to a claim requires only that new vocabulary, “S is committed to *p*,” and “S is entitled to *p*,” be deployed with the *circumstances* of application that one can assert these sentences formed using the new normative vocabulary whenever one would in practice respond to S as having the commitment or entitlement labeled with the sentence *p*, and with the *consequences* of application that whenever one asserts one of these new normative sentences, one must also take or treat S in practice as having the corresponding commitment or entitlement. Introducing vocabulary playing this role requires only the algorithmic elaborative abilities I have called “*response substitution*” (along with the arbitrary formation and permutation of states), together with the sort of basic deontic scorekeeping abilities I have argued one must possess in order to engage in practices of giving and asking for reasons at all. (Compare the instructions for introducing *conditionals* that I offered in Lecture 2.) Further, when used with these circumstances and consequences of application, it is clear that when one of these new normative sentences is asserted, the pragmatic significance of that speech act will be to *say that* someone is committed or entitled to a claim, making propositionally *explicit* a practical attitude—taking or treating someone in practice *as* committed or entitled to a claim—that before the advent of the new vocabulary remained *implicit* in what one *did*.

My overall claim is that both the modal and the normative Kant-Sellars theses are true. In order to be able to talk at all, to make claims and inferences, one must already know how to do everything necessary in principle (in the precise sense of ‘in principle’ given by the notion of algorithmic elaboration) to deploy alethic modal and deontic normative vocabulary. If so, one cannot be stuck in the position Hume took himself to be in: understanding ordinary empirical, descriptive vocabulary, but with that providing no grip on the use of modal and normative vocabulary. The *semantic* relations between what is expressed by the use of empirical descriptive vocabulary, on the one hand, and what is expressed by the use of modal and what (something different) is expressed by normative vocabulary, on the other, are essentially *pragmatically mediated* ones. To





understand the relation between how things merely *are* and how they *must* be or (a different matter) *ought* to be, one must look at what one is *doing* in *saying* how things are, and what is required to *say* what one is thereby doing. Transposing Kant's response to Hume into this pragmatist key requires the metaconceptual resources of meaning-use analysis, which is what enables us to be clear about the pragmatically mediated semantic relations on which that response depends.

Coming to understand both modal and normative vocabulary as standing in the complex resultant pragmatically mediated semantic relation of being LX to—elaborated from and explicating of—practices integral to every autonomous discursive practice will turn out also to be the key to understanding a deep and illuminating feature of the relation of these two vocabularies, not just to vocabulary use in general, but also to each other. It supplies the raw materials for filling out and developing Sellars's suggestive claim that modal vocabulary is a 'transposed' language of norms. In the next lecture I will begin to explore the relations between normative and modal vocabulary that become visible in this way, showing how normative vocabulary can serve both as a pragmatic metavocabulary for modal vocabulary and as the basis for a directly modal formal semantics for ordinary empirical vocabulary that does not appeal in any way to a notion of truth. In the final lecture, that discussion will be brought together with the discussion of modality and normativity from the two lectures that precede it, culminating in an understanding of *discursive intentionality*, the ultimately semantic relations between knowing subjects and their cognitive objects that is expressed by *intentional* vocabulary, in terms of the relations between what is expressed by *normative* and *modal* vocabularies.

