

Week Nine: Pragmatics, Metaphysics and Possibility

1. *Pragmatics and Metaphysics*

There's two points left over from last week's seminar still to discuss. The first is whether, as Lewis claims, we are justified in positing an asymmetry in the role of pragmatics. The second is whether this approach is at all justified. We'll look at that before going on to the material scheduled for this week.

The Asymmetry Assumption

Pragmatics is important to metaphysics because of the following methodological assumption. When we consult our intuitions about whether something, call it Fred, has some property, or falls under some concept, say being slithy, our intuitions are primarily intuitions about the potential assertibility of "Fred is slithy", not, as we may have expected, about the slithiness of Fred. So when we intuit that, for example, Al Gore's poor poll numbers are not caused by the absence of mind-control techniques, as opposed to say Bill Bradley's campaign work, our intuition is in the first instance about a sentence. It is the intuition that we could not properly say "Al Gore has poor poll numbers because he hasn't mastered mind-control techniques." Two things follow if this is correct. First, we now have a new and potentially quite powerful way to explain away deviant intuitions. Secondly, philosophy of language becomes relevant to metaphysics in yet another way. So being someone who likes to explain away deviant intuitions, and who spends more time than is healthy doing philosophy of language, I think this is correct.

There is one simple way to explain why a sentence can't be asserted. The sentence is simply false, we are aware of its falsity, and we don't like to utter falsities. What Grice's work is commonly taken to show is that there is at least one other way to explain these intuitions away. Some sentences are true, but it would be unhelpful to assert them, so we don't. After a while, we internalise these rules of proper assertion, not the rules of truthful assertion. Or maybe we just learn the rules of proper assertion, not the rules of truthful assertion. Actually that's closer to what I think.

Last week we went through some ways in which the Gricean maxims could be relevant to metaphysics. As a refresher exercise, how could we use the maxims of quality to make the following position look plausible: There is a sharp line between the tall people and the not-tall people, a precise height that one must be to be tall, and if you are below that you are not tall, but that height is not knowable in principle. As many of you know, this is the principle called epistemicism about vagueness.

Lewis claims, without much argument, that there is an asymmetry in how the pragmatics can be applied. We can use the pragmatics to explain away intuitions of falsity, or of non-truth generally, but we

can't use it to explain away intuitions of truth. As a digression, one other methodological result we get from Grice is that intuitions purporting to distinguish between types of non-truth, such as falsity, lacking a truth value or having a third truth value, are worthless. This was even known to some pre-Griceans, and is well argued for in Dummett's famous paper "Truth."

There are two broad classes of reasons for thinking that Lewis is right in positing such an asymmetry: any false utterance will breach one of the maxims, the one barring false utterances, and while there are untendentious examples of improper true assertions, all purported examples of proper false assertions are tendentious. Let's look at these in reverse order. Here are four classes of sentences which seem to breach Lewis's Principle, with an example of each.

Quantificational

Everyone has heard of David Lewis.

Metalinguistic Negation

Lewis isn't a good philosopher, he is a great philosopher.

Presuppositional

The senator's husband is waiting outside.

Semantic

It isn't true that Al Gore is trailing because of a lack of mind-control techniques.

The first two should be reasonably clear, and we said a little about them last week. The third and fourth might need some explaining. The goal of conversation, it is agreed on all sides of this debate, isn't the rote recitation of truths, it is the imparting of helpful information. Sometimes, it is suggested, that information is more easily imparted by saying something literally false than by saying something true. One class of cases in which this might happen is if your hearer has some false presuppositions. So imagine that a senator comes to a party with someone appearing to fill the functional role of her husband, and I know that you believe this fashion accessory is her husband. If I want to tell you something about that person, the most efficient mechanism of communication may be to refer to him by the phrase "The senator's husband." The resultant sentence will be one I believe to be false, but it may communicate just what it is intended to

communicate, and it may do so while being significantly more brief than any other sentence which does the same job, like “The person you believe to be the senator’s husband is waiting outside.” So if we think brevity and accuracy can be traded-off, this will be a proper false assertion.

Some worries with this example, though these are far from conclusive. First, when we think about the true alternatives to this sentence, we notice that they have some further implication. The implication will usually be that the person is not the senator’s husband, or at least that there are grounds for doubt on this question. So perhaps the reason for uttering the brief sentence is not to implicate these things; perhaps for good reasons. So this case can be subsumed under those cases where it is *morally* permissible to mislead people, as when there’s a killer at the door. But those cases are no counterexamples to Lewis’s principle, so this isn’t either.

Secondly, it might be that in practice definite descriptions can quickly become denoting phrases by newly established conversational conventions. Compare the well-known example, apparently closer to natural language than most examples we use, “The ham sandwich is getting impatient.” Maybe trying to explain away one counterexample to a principle by analogy to another counterexample is a sub-optimal strategy here, so I won’t push this too hard.

The fourth kind I think are very interesting. Certainly sentences like this can be properly asserted. And on one natural interpretation of Lewis, the interpretation on which he keeps semantic compositionality, it is true. In fact we can find examples like this in real life, assuming television is real life. (On one TV show last night a character was accused of confusing TV with real life, and he replied that he did know the difference, but he liked TV more.) This example is from a *Seinfeld* episode. For various reasons, everyone thinks that Jerry is anti-social. So in one scene he sees his maintenance manager, and politely asks for his shower to be fixed. I don’t have the exact transcript, but the conversation goes like this.

Maintenance Guy: So, when you want something done you’re nice to people.

Jerry (whining): That’s not true!

For a long time (well actually about a week, but it felt like an eternity) I thought these were counterexamples to the principle that intuitions of well-informed competent speakers about proper assertion are veridical. Notice that so far we haven’t questioned the folk’s intuitions about when sentences can and cannot be properly asserted. Intuitively, it is only proper to write a reference like “Blah is punctual and neat” if Blah has no other talents, so it is only proper to write a reference like “Blah is punctual and neat” if

Blah has no other talents¹. Since I thought Lewis's principle was obviously correct, and so is semantic compositionality, the only possible way out seemed to be to deny that this inference always went through. I then thought that only speakers who have studied Grice in sufficient detail could count as well-informed for these purposes, which would have a similar effect. And then I learned to love the bomb, the sentence is really true even though the sentence "Al Gore is trailing because of a lack of mind-control techniques" is also true. When you don't believe the semantics for English is compositional you can say the most amazing things! There are no counterexamples to Lewis's principle, just surprising consequences...

There is one other kind of objection to Lewis's move here. (This was suggested to me by Michael Glanzburg, but it was late at night so don't blame either him for the idea or me for the transcription.) Pragmatics is just the wrong part of philosophy to settle questions like the dilemma about causation Lewis discusses. Pragmatics is concerned with what we can say, metaphysics with what there is. What Lewis is doing is no more plausible than the attempt to derive the physical structure of the universe from the grammar of 20th century English. Given what we have said so far, what might we say in response?

2. *Objections to the Gricean Program*

The first barrage of objections turns on the specifics of the program. Almost every specific prediction which has been made by Griceans has turned out to be false. I just want to give you a little flavour of the problems the Griceans face before looking at some more philosophical problems. First, compare what the Gricean would say about answer (1) with the apparent facts about answers (2) and (3).

Q: Did any students get As?

- (1) Some did.
- (2) Yes.
- (3) Some got Bs.

Secondly, note that the order implication which we strongly hear in some sentences we don't hear in others. So compare (4) and (5).

¹ For those who worry about such things, the scope of 'intuitively' in this sentence is narrow, it only covers the first half of the sentence.

-
- (4) Last weekend, Brian had several drinks and drove home.
- (5) Last weekend, Syracuse beat Temple and Penn State lost to Minnesota.

The general form of these objections is quite simple. The Gricean claims that some property or other of the sentence is responsible for generating the implicature. But there are other sentences with the same property with a different implicature, or no implicature at all. And as we see in (3), there are sentences with the same kind of implicature without the salient property.

The second kind of objection is by analogy to a well-know argument in ethics. It is often thought that the semantic content of a sentence, its truth conditional content, is what it brings to the antecedent of a conditional. So “Adultery is wrong” must have some truth conditional content, because it contributes some content in “If adultery is wrong, Bill Clinton will go to hell.” We can also give an argument for this conclusion by noting the role these conditionals play in arguments.

The problem for the Gricean is that the content she calls pragmatic seems to be carried into antecedents of conditionals, and so by this test is semantic content. So compare (6) and (7).

- (6) If Brian had several drinks and drove home, he is a very irresponsible driver.
- (7) If Brian drove home and had several drinks, he is a very anti-social drinker.

These examples also raise a problem for Lewis’s principle. Each of these seems intuitively true, so by Lewis’s principle we should think they are true. But if we assume, as seems consistent with these intuitions, that the two events, Brian’s having several drinks and Brian’s driving home, occurred, and Brian has at most one of the mentioned vices, we have an odd result. Two true conditionals, two true antecedents, at most one true consequent. All very mysterious.

Finally, and I’m not sure how serious this is, note that the psychological story that Grice tells simply cannot be true for all the instances of implicature with which we deal. In every case, Grice assumes that we know the truth-conditions of our sentences, but we choose not to utter them because of the pragmatic rules of assertion. But we don’t know the pragmatic content in most of the interesting cases. So what are we to make of the story he tells in “Logic and Conversation”?

3. *Objections to Jackson*

Now that the Jackson book has arrived at the bookshop, we can return to the original aim of the course. I don't want to keep working through the book quite in the order it is presented. The main reason for this is that I want to look at two the passages using the 'two-dimensional' account of the Kripke-Putnam examples at once. So I'll leave the parts of chapter two where it is discussed until we get to the second half of chapter three. That might be in about half an hour, but it's nice to have a plan. So where we'll go is to look at the objections he discusses at the start of chapter three, go over the 'two-dimensional' stuff, then look at how he applies it. That should take us to at least half way through next week, and possibly to the end of next week. So we're starting with the first 8 to 10 pages of chapter 3.

Objection from Theory Reduction

Jackson's point here is reasonably clear, so I will only make a couple of brief side comments. First, there is a small problem with the example, as temperature is not mean molecular energy, and apparently hasn't been for the last forty years. It is possible, apparently, to have negative temperatures Kelvin, but not negative energies, so the smooth reduction is a bit bumpier than we'd like. (If anyone's interested I will try and track down the reference, or you can just ask a scientist.)

It is also less than obvious that the statements in the little argument on page 59 is right. Obviously it is valid, but it isn't clear that Jackson has properly translated sentences from English into formal language. That is, it isn't clear the premises to which we are really entitled are identities. Orthodoxy distinguishes between two kinds of 'is's, roughly that found in (8) and that found in (9). The first is the 'is' of predication, the second the 'is' of identity.

(8) Aristotle is a philosopher.

(9) Aristotle is the greatest ancient philosopher.

So sometimes it is reasonable to not know the meaning of the word 'is'! Anyway, if the 'is' in 'mean molecular energy is what plays the temperature role' is the 'is' of predication, rather than the 'is' of identity, the argument here will be invalid. From Fa and Fb it doesn't follow that a is b .

Objection from Possible Cases

Again, the main point is fairly clear, so I want to confine my comments here to the side-issues. Jackson seems to equivocate dreadfully on whether analysis can deliver gruesome results. So on the bottom of page

64, we have the following: “I mean that typically we know something useful and non-grue-like, and are giving voice to this knowledge when we classify happenings as examples of grooming behaviour, pain, rational inference, and so on.” Well if that’s meant to be argued for, as opposed to just asserted, he should defend it against the examples Stich gives. For on the bottom of page 61 he looks quite content to have very gruesome analyses count. “The puzzle is that [the?] Roschian view [Stich] describes as opposed to the search for necessary and sufficient conditions is itself a view about the necessary and sufficient conditions for being a bird: as he himself describes it, the view is that being sufficiently similar to the relevant prototype is necessary and sufficient for being a bird.” Well if that is not gruesome, I suppose analyses aren’t gruesome.

There is also an equivocation about what analysis is supposed to do for physicalism. On the middle of page 62 we get the following: “What we require from physicalists who accept the existence of grooming behaviour is enough by way of conceptual analysis to make it plausible that the purely physical account of our world makes true the grooming-behaviour account of our world; and to do that it is *not* required to give necessary and sufficient conditions in physical terms for grooming behaviour.” But on other occasions we are told that the ideal physical story should *a priori* entail the grooming story, or whatever other story you believe refers. If anything like epistemicism is true in vagueness we might have a supervenience thesis which is true but not knowable, so this moderate requirement is all that is needed. (I’m assuming here that *a priori* is elliptical for *a priori* knowable.) Now maybe this is just another refutation of epistemicism, but it seems like a problem. We’ll come back to that when we get to the 2D stuff.

There is a dreadful use-mention confusion on page 64. Quick exercise: spot where it is. A friend of mine once suggested that a point in favour of my current theory of conditionals was that it explained why the following has a true reading: If we called a horse’s tail a leg, horses would have five legs. With friends like that... I don’t think this is a problem, but again it is something to keep in mind when we look at the 2D material.

Codifiability of Rationality

There is an odd little paragraph on page 67 attacking the idea that rationality is uncodifiable. I used to think this was probably true, largely because of worries about the interaction of social and physical sciences, but I’ve been pushed back to being agnostic. Still, as an agnostic I should say something about all the arguments for belief, and here’s what I have to say about this one.

First, it is unclear whether we really are finite in the relevant sense. The issue isn’t whether we are finitely sized, it seems the relevant issue is whether we can take finite or infinitely many different states.

And that seems to depend on some unsolved issues in quantum theory. On the Newtonian picture a thermostat could take infinitely many states, one corresponding to each real number in a particular range. If everything (I mean really everything) does come in quantum packets, then maybe there are only finitely many states. But this isn't the closed question Jackson suggests that it is.

Secondly, not just any codification counts. "X is rational iff X is rational". Ho hum. Or what about "X is rational iff X is rational and a duck or rational and not a duck." Not much better. We have to find non-circular definitions of rationality into other concepts. Interesting question, is there a reduction of our normative concepts to our non-normative concepts. Just because we are finite doesn't guarantee that it is so. Say that we have 100 possible states, and the point of each of our concepts is to distinguish between possible states we could be in. Now also imagine that every non-normative concept is limited in the following way, if state 1 is in, state 2 is in also, and *vice versa*. So no sentence we can write (even in the infinite case) in non-normative language can distinguish between being in state 1, and being in state 2. But some normative concepts can make this distinction. So perhaps we only think "The object of my attention is rational" if we are in state 2, 24 or 67. Then there will be no code, not even an infinitely long code, for rationality in terms of non-normative concepts. But this example didn't presume anything about our being infinite, or having special cognitive access to some normative concepts. It just presumed something about the expressive powers of normative and non-normative languages. Now that presumption may be false, but since it is the presumption which is at issue, so cannot be denied in an argument for codifiability without begging the question at issue.

4. *The Necessary and the A Priori*

We used to think that there was a close connection between being necessary and being *a priori*. Then Kripke came along and cured us of this blindness, and then philosophy was completed. Well, not quite on all fronts. There is still a close link between being necessary and being *a priori*, and it is this close link that Jackson wants to exploit to save conceptual analysis.

Let's first look at some cases to get the hang of how Jackson thinks the dialectic plays out. We'll focus on the boring old 'water' case, because we're a bit familiar with that, and because it is a case that Jackson uses. When we were looking for what water is, we didn't go round looking for H₂O and then run tests on it to determine its chemical composition. Rather, we knew roughly which things around here were water, we just didn't know much about the internal structure of those things. We looked at the examples of things we knew to be water, found out they were H₂O, concluded that water was H₂O. But when we went

hunting for water in other possible worlds, we reacted quite differently. We don't care in that search, it seems, what plays the water role. Rather, we just care whether something is H₂O or not.

This suggests a certain asymmetry between how we investigate our world, the actual world, and how we investigate other possible worlds. Let's make that more vivid. Should we discover tomorrow that the actual world is really like Twin Earth, and it has just been a giant government conspiracy to trick us into thinking water was H₂O, we wouldn't say XYZ was mere twater. Rather, we would say that we've been tricked, and that water is really XYZ. So if Twin Earth is somewhere else, water there is H₂O, the watery stuff is not water, and water is not watery. But if Twin Earth is here, and we've just been tricked into thinking it was merely possible, then water on Twin Earth is XYZ, the watery stuff is water, and water is watery.

Already we have made a huge step towards accommodating the Kripke cases within a very conservative modal framework. Some people thought after Kripke's and Putnam's cases came out that we had large new restrictions on what was possible. What Jackson wants, and what this shows, is that there is no such restriction. The restriction, rather, is on how to describe the possible worlds there are. Just as we thought all along, there are possible worlds in which the stuff which has the external characteristics of water has different internal characteristics. That is, there are worlds in which the stuff which has the characteristics we used to identify water has different internal characteristics. So we can still state the fact that water is H₂O is *a posteriori* in possible worlds terms. There are even worlds in which XYZ looks like water, smells like water, and is called 'water'. These will become important soon. But as a soundbite analysis of the Kripke cases, anything which seems possible still is possible, but there are universal restrictions on how to describe some of these possibilities.

One more step and we can apparently even remove that restriction. Which sentences are true on Twin Earth seems to depend not just on internal facts about Twin Earth, but on whether Twin Earth is actual or not. If it is actual, then water on Twin Earth is XYZ, if it isn't then water on Twin Earth is H₂O. I presume everyone is familiar with the identification of propositions with sets of possible worlds. The proposition *p* is just the set of possible worlds in which *p* is true. (And then all the usual sentence connectives are identified with the usual set-theoretic operations, and life immediately becomes 14% more beautiful.) So the proposition *p* expressed by sentence *S* will be the set of worlds in which *S* is true. But now this should seem ambiguous. Is "Water is watery" true on Twin Earth or not? Yes, if Twin Earth is actual, no if Old Earth is actual. Several possible positions start to present themselves. Jackson and Chalmers identify two.

The **A-proposition** expressed by S is the set of worlds w such that if w is the actual world, S is true at w .

The **C-proposition** expressed by S is the set of worlds w such that given this is the actual world, S is true at w .

S is necessary iff the C-proposition it expresses is the set of all possible worlds. S is *a priori* iff the A-proposition it expresses is the set of all possible worlds. So identifying sentences, like water is watery, will turn out *a priori*. And in that important sense, we don't need to leave the armchair to discover the structure of concepts.

This procedure can be quickly generalised. Say S is true at a pair of worlds $\langle w, v \rangle$ iff S is true at world w given that world v is actual. A sentence is *a priori* iff for all worlds w , S is true at $\langle w, w \rangle$. And letting @ stand for this world, it is necessary iff for all worlds w , S is true at $\langle w, @ \rangle$. Stalnaker suggests a graphic representation of this. Set up an e by e matrix, where e is the number of possible worlds there are. Each cell corresponds to a pair of worlds. The matrix for a sentence has T at a cell iff that sentence is true at the world in the column, given the world in the row is actual. So to get the hang of it, we'll end for today with some questions.

- What will the matrix for a necessary true sentence look like?
- What will the matrix for an *a priori* sentence look like?
- What will the matrix for an analytically true sentence look like?
- Is there a distinction, within the matrix, between being necessary and *a priori*, and being analytic?

Assume there are just three worlds, @ the actual world (with no conspiracy), tw being Twin Earth as usually described, and pw being a world where half the watery stuff is H_2O , and half is XYZ. (On Twin Earth there is some H_2O and it is gooey and sticky.) Draw the matrix for each of the following sentences:

1. All water is H_2O .
2. All water is watery.
3. Water is a gooey, sticky substance.
4. The stuff called 'water' is a gooey, sticky substance.
5. The stuff actually called 'water' is a gooey, sticky substance.

To think about for next week. How should we deal with possible worlds where the meanings of terms is different to here? Is the way you have chosen consistent with Jackson's claim that analyses are *a priori*

true? If so, is it consistent with the idea that the matrix approach can be used to explain meanings, or do we need to know meanings to even apply the matrix method?

The important readings are pages 46-53 of Jackson and the Block and Stalnaker paper which will soon magically appear in the filing cabinet (especially section 10). I will also put the introduction to Stalnaker's recent collected papers in the filing cabinet, though there are only a few pages which are relevant, around 12 to 17. But if you're interested in Stalnaker's larger work, this should be read. Finally, we should at some stage get to Yablo's slightly distinct criticisms of this framework, and that stage may be next week, so that paper will also go in. David Chalmers's web site at <http://ling.ucsc.edu/~chalmers/index.html> has some useful stuff on this question, though you will have to trawl for it a bit. And my paper "Indicatives and Subjunctives" also has a quick (really quick) tour through these fields and a little application to the logic of conditionals, if people are interested in that stuff. It too will magically be filed, soon.