

Classes 13-15 - Meaning Holism

I. The dogmas

Quine's paper is an attempt to cleanse empiricism from reliance on two presuppositions:

- D1. There is an analytic/synthetic distinction; and
- D2. Reductionism: statements can be translated to terms which refer only to immediate experience.

Quine concludes that the two dogmas are essentially the same.

Reductionism entails that statements are confirmed (or disconfirmed) individually, by the experiences that justify them.

It is the claim we find in Hume, that all ideas must derive from initial impressions, and in Ayer, that the meaning of a statement is the method we use to verify that statement.

Radical reductionism says that every meaningful statement is reducible to a report about immediate experience.

The early reductionists (e.g. Locke and Hume) intolerably focused on reducing terms, rather than statements.

Bentham turned to sentences as the basic units of communication.

The positivists, continued to seek reductions of sentences.

But, the positivists failed to explain how, precisely, we were to build statements out of sense data.

The most valiant attempt actually to carry out the in-principle reductions was Carnap's *Der Logische Aufbau der Welt*, or *The Logical Structure of the World*.

The *Aufbau* included the whole language of pure mathematics in addition to sense data.

Carnap treated spatio-temporal points as quadruples of real numbers, to which qualities then applied.

So the ultimate constituents of the world were taken to be sensory events plus classes, or sets.

Quine notes that there are problems getting from sense data to the appropriate ascriptions.

Further, Carnap leaves the 'is at' relation as a primitive.

The most basic spatio-temporal property is unexplained.

More importantly, even Carnap's more conservative reductionism relies on the claim that statements are confirmed individually, by the experiences we take to justify them or give them meaning.

This isolation of sentences remains with positivism, whatever the status of Carnap's *Aufbau* project.

The notion lingers that to each statement, or each synthetic statement, there is associated a unique range of possible sensory events such that the occurrence of any of them would add to the likelihood of truth of the statement... (2D, 72)

And the isolation of sentences implicit in reductionism is inextricably linked to the belief in an analytic/synthetic distinction.

As long as it is taken to be significant in general to speak of the confirmation and infirmation of a statement, it seems significant to speak also of a limiting kind of statement which is vacuously confirmed, *ipso facto*, come what may; and such a statement is analytic (2D, 72-3).

Ignoring any connections to positivism, there seems to be a distinction between sentences like:

1. Bachelors are unmarried.

and

2. Bachelors are unhappy.

Quine argues that since there is no way to draw a line between analytic and synthetic statements, such statements differ only in degree of empirical confirmation, not in kind.

It is obvious that truth in general depends on both language and extralinguistic fact. The statement "Brutus killed Caesar" would be false if the world had been different in certain ways, but it would also be false if the word "killed" happened rather to have the sense of "begat." Hence, the temptation to suppose in general that the truth of a statement is somehow analyzable into a linguistic component and a factual component. Given this supposition, it next seems reasonable that in some statements the factual component should be null; and these are the analytic statements. But, for all it's a priori reasonableness, a boundary between analytic and synthetic statements simply has not been drawn (2D, 70).

II. Characterizing analyticity

There are different ways to characterize analyticity:

- A1. Truth in all possible worlds
- A2. Can not be false/true
- A3. Denial as self-contradiction
- A4. Conceptual containment
- A5. Truth in virtue of meaning

Quine traces A1 and A2 to Leibniz, who distinguished truths of reason and truths of fact. The notion of possible worlds is metaphorical, and will not suffice as a foundation for analyticity. Quine interprets A1 as A2, which we can see in Hume's distinction between relations of ideas and matters of fact.

We also see A2 in Kant.

Note that truth, which is central to A2, is a metaphysical, rather than semantic, notion.

That is, if we are distinguishing between semantic claims, epistemological claims, and metaphysical claims, A1 and A2 are not useful as explanations of analyticity at all.

They concern whether statements are necessary or contingent.

To understand why Quine would make the serious error of blurring semantic questions with metaphysical ones, one has to recognize that Quine is responding directly to the positivists.

For the positivists, the analytic statements were precisely the same as the necessary ones.

For a statement to be necessary, it would have to be analytic, since as Hume argued, there are no necessary connections of matters of fact.

Conversely, all analytic claims would have to be necessarily true, since experience cannot disconfirm an analytic truth.

The truth of an analytic statement depends only on the meaning of its symbols.

III. Aside on necessary synthetic and contingent analytic statements

One might think that since there are, arguably, necessary a posteriori statements, and contingent a priori ones, the set of necessary statements would be identical to the set of analytic ones.

It is pretty easy to come up with statements that people would think are necessary, but not analytic.

Kant thought that most arithmetic truths were necessary but synthetic.

The theoretical identifications that Kripke argued were known a posteriori (e.g. water is H₂O) also serve.

Finding examples of analytic statements that are only contingently true is more difficult.

Here is an example from Hilary Putnam that might work.

Imagine that all cats turn out to be cleverly disguised robots, sent to Earth from Mars, for some nefarious purposes.

Now, consider:

3. Cats are animals.

You might argue that 3 is analytically true, but false.

It is part of the concept of a cat that it is an animal, but in the actual world, it turns out that 3 is false.

On second thought, there seems to be an equivocation here.

There are (at least) two distinct ways we could describe this situation.

First, we could say that it turns out that cats are robots.

Second, we could say that it turns out that there are no cats.

On the first description, 3 is false, but it is not analytic.

On the second description, we hold that cats are a natural kind, essentially animals.

Thus, we re-gain the analyticity of 3, but we lose its falsity.

In any case, Quine is not concerned with the cleavage of analyticity and necessity.

IV. Returning to the characterizations of analyticity

A3, on which analytic statements are those whose denials are self-contradictory, characterizes Hume's relations of ideas, as well as the Fregean version of analyticity.

Quine notes that A3 is in need of explanation, just as A1 and A2 were.

It is fairly easy to characterize the sentences whose denials are false.

But, self-contradiction is trickier.

The sentence:

4. $P \bullet \sim P$

is clearly a contradiction, given the meanings of the logical terms and that the token 'P' refers to the same thing in both instances in 4.

But:

5. Fred is a married bachelor.

is not of the form of 4.

We need a broader notion of self-contradiction, which is just another way to state the original problem of characterizing analyticity.

So, A3 will not suffice, either.

A4, as we saw in previous classes, is ambiguous between Frege's plant-in-the-seeds analyticity and Kant's beams-in-the-house version.

Quine presents two problems with containment.

One problem is the presumption that all statements are best understood in subject-predicate form.

Russell, in recasting sentences that contain definite descriptions as existential sentences already questioned the identification of grammatical form and logical form.

Consider:

6. All swans are white.

7. $7 > 5$.

6 and other general claims, including most scientific laws, are best understood as universally quantified sentences.

Even the specific sentence 7 may best be understood as a relation between two particular objects, and not as a property ascribed to an individual object.

We can cast sentences such as 6 and 7 in subject-predicate form, but this interpretation is forced.

Another problem with containment is the vagueness of the metaphor.

Quine understands Kant to mean that an analytic statement is true in virtue of meanings, and independently of fact.

Thus, Quine takes as his quarry analyticity in the form A5, which is more broad than the traditional A4. If Quine can argue that there are no statements which are analytic A5, it would cover A3 and A4, as well.

A5 coheres neatly with the positivists' claim that analytic sentences are true independent of facts.

It clearly evokes the picture of a linguistic component and a factual component of a statement.

Quine agrees that truth in general depends on language and extralinguistic fact.

But, he denies that the factual portion ever disappears.

V. Logical truth

We can distinguish two types of analytically true statements.

Compare

8. No unmarried man is married.

9. No bachelor is unmarried.

8 is a logical truth.

It displays its analyticity on its surface.

It is clear, by rules of logic, that the predicate of is contained in the subject.

9 is merely analytic.

Its truth depends on the meaning of 'bachelor'.

The difference between 8 and 9 may be clearer in their logical regimentations:

- 8'. $(x)(\sim\sim Mx \supset Mx)$
- 9'. $(x)(Bx \supset \sim Mx)$

8' is true under any interpretation of the predicates.

9' has false interpretations.

Quine has no complaint about logical truth.

But, the analytic/synthetic distinction can not rest on logical truth.

9 can be turned into a logical truth by substitution of synonyms (i.e. putting 'unmarried man' for 'bachelor').

But, to make such a substitution, we require an explanation of synonymy.

The bulk of "Two Dogmas," then, is an attempt to discredit synonymy.

VI. Synonymy

The argument against the analytic/synthetic distinction rests, in the first half of the paper, on Quine's argument against synonymy.

Quine tries three possibilities to characterize synonymy:

- S1. Logic (meaning postulates)
- S2. Dictionary definition
- S3. Interchangeability (substitutivity) *salva veritate*

S1 fails because it presumes the notion to be explained.

Quine's challenge is not merely to characterize analyticity, say by listing the analytic sentences.

He wants to define it, to express its essence.

Carnap tries to use meaning postulates to characterize synonymy.

If we want to say that 'Fx' and 'Gx' are synonymous predicates, within the theory, we add an axiom, or semantic rule:

- 10. $(x)(Fx \equiv Gx)$

To model a theory, we consider state-descriptions.

State-descriptions ascribe truth values to atomic sentences of a theory.

We can build complex sentences using logical particles and the constraints of the semantic rules.

Carnap's analytic sentences are those that come out true on every state-description.

The semantic rules assure us that all substitutions of synonymous expressions will maintain analyticity.

That is, the synonymy rules preserve substitutivity *salva analyticity*.

Meaning postulates might characterize logical truth.

They do not help to explain analyticity, which was Carnap's intent.

Carnap's procedure involves presenting a list of semantic rules of synonymy, and their consequences.

Such a list would provide a definition of analyticity for a given language, but not a general explanation of analyticity for all languages.

Even if we define the analytic sentences by iteration, we still need to know what property we are ascribing to them.

There are two types of semantic rules.

The first type specify, recursively or otherwise, which sentences are analytic.

This has the problem of not giving us an idea what property we are ascribing.

The second type of semantic rule just says that those statements specified are among the truths.

Then we can define analyticity as 'true according to the semantic rules'.

But then we need to explain 'semantic rules', and why certain semantic rules are picked out as special.

Semantical rules determining the analytic statements of an artificial language are of interest only in so far as we already understand the notion of analyticity; they are of no help in gaining this understanding (2D, 70).

Quine's criticism of Carnap, here, is reminiscent of many of Plato's dialogues.

Socrates repeatedly chides interlocutors for listing examples of some form, rather than characterizing its essence.

See, for examples, the beginnings of *Euthyphro* and *Republic*.

I have put Carnap's paper, "Meaning Postulates," on the course website.

I had considered assigning it, but decided it was too technical, and too ancillary to our central concerns.

Paper topic, anyone?

Saying that analytic statements are true 'by definition', S2, also presupposes, rather than explains, synonymy.

The lexicographer is a sociologist, who reports synonymy, and so can not ground it.

Terms are found to be synonymous, not made so by fiat.

Explication, which adds clarifying information to the definition, relies on other, preexisting synonymies.

Only when a term is created explicitly to act as a synonym (e.g. an abbreviation) do we have a definition that does not rely on synonymy.

For example, scientists might name a planet, or molecule.

But these exceptions are rare.

The characterization of synonymy as substitutivity, S3, comes from linguistics.

Quine presents a few silly examples of failure of substitutivity.

While you can substitute 'unmarried man' for 'bachelor' in 9, you can not substitute (salva veritate or meaning) the one for the other in:

11. When I graduated, I received a bachelor of arts diploma.

Or in:

12. 'Bachelor' has eight letters.

His example of not being able to substitute 'unadorned' for 'plain' in:

13. Burlington is on the shores of Lake Champlain.

is similarly silly.

We can easily appeal to the notion of a word or term to block the above objections.

We should only expect substitution of terms for terms.

'Bachelor of arts' is a complete term, and the 'bachelor' within it is just part of a term.

Similarly, 'plain' in 'Lake Champlain' is not a whole term.

And, we know that the scare quotes provide a context of indirect reference in 12.

We should not expect interchangeability in any of the examples 11-13.

Still, we have not seen whether substitutivity *salva veritate* actually explains synonymy.

Quine pretends to attempt to defend substitutivity *salva veritate* by appealing to modal contexts.

(This is the "hocus-pocus" argument, pp 67-8.)

The idea is that we can isolate synonymous expressions as those whose identities are necessary.

For example:

14. Necessarily, bachelors are unmarried men.

15. Necessarily, anything plain is unadorned.

Claims 14 and 15, Quine argues, are useless as explanations of synonymy.

They are circular, in that they explain one intensional idiom (synonymy/analyticity) in terms of another (modality).

What Quine wants is to explain synonymy without appealing to any other intensional contexts.

He wants a reduction of the intensional.

His preference would be an explanation of analyticity in extensional terms.

In a purely extensional language, we would get interchangeability for all coextensive terms.

But then we would have synonymy for 'creature with a heart' and 'creature with kidneys'.

Also, 'square circle' and 'short tall hairy bald man' would be synonymous, since they both have null extensions.

Interchangeability in an extensional language is different from the version we need to explain cognitive synonymy.

We need a cognitive synonymy that doesn't merely get us the truth of 'no bachelor is married', but its analyticity, as well.

We need something stronger than extensional equivalence.

It may not be possible to reductively analyze meaning at all.

But, without a reduction, the explanations (like those of 14 and 15) seem question-begging.

'Necessity' presupposes analyticity, so we get a kind of circle.

Our argument is not flatly circular, but something like it. It has the form, figuratively speaking, of a closed curve in space (2D, 68).

One response to Quine's accusation of circularity is to distinguish between virtuous and vicious circles.

If we had a set of inter-theoretically linked intensional terms, we could justify the whole group by appealing to their systematic virtues for the intensional idioms themselves.

I'll not pursue this approach, for now, though we will return to it with a brief look at Katz's work at the end of these notes, and a longer look at his work at the end of the meaning section of the syllabus.

VII. The argument against D1, redux

From his arguments against S1-S3, Quine concludes that there is no such thing as analyticity as apart from logical truth.

Synonymy must be explicated either in terms of definition or in terms of interchangeability *salva veritate*. (I omit dictionary definition, S2, which was a non-starter.)

Since both of these options fail, there must be no analytic/synthetic distinction.

QD1. If there is an analytic/synthetic distinction, there must be a good explanation of synonymy.

QD2. The only ways to explain synonymy are by interchangeability *salva veritate* or definition.

QD3. But interchangeability can not explain synonymy.

QD4. And definition (meaning postulates) can not explain synonymy.

QD5. Thus, there is no good explanation of synonymy.

QDC. And thus there is no analytic/synthetic distinction.

I'll mention one objection to QD, and a Quinean response and a half, before proceeding to examine the ramifications of Quine's conclusion.

The strongest response, I think, to QD is to challenge the weak QD2, as we will see that Katz does.

Quine does not consider all possible ways to explain synonymy, only two.

Perhaps a better theory of meaning could generate a good definition of synonymy, one from which analyticity would follow.

Quine's response to Katz's objection would concede that QD2 is not fully defended.

But, he thinks that he has shown enough for us to see that no other ways of explaining synonymy will succeed.

So, Quine would urge us to take QD as enthymemic, an incomplete argument that could easily be finished.

Another way to look at QD is to take it not as an air-tight argument, but as painting a picture.

Quine is describing a world without an analytic/synthetic distinction.

The defense of adopting that world-view would be not a small, tight argument, but, rather, how satisfying the picture is on the whole.

Thus, in order to understand the details, we have to look both at other work of Quine's, and at the ramifications of the failure of the two dogmas.

We proceed to the latter.

VIII. Meanings skepticism

Quine's arguments against analyticity are motivated by Ockhamist concerns about spooky entities.

His argument is twofold.

First, if we are to posit an object, it must have clear identity conditions.

We must be able to determine when two meanings are identical.

Elsewhere, Quine urges that we should admit "no entity without identity".

Quine uses the need for identity conditions, and the related need to be able to individuate objects, in order to deride alethic modal operators: possibility and necessity.

Quine asks, rhetorically, how many possible persons stand in a doorway.

Thus, Quine's response to Frege's problem of substitutivity in opaque contexts is to deny the legitimacy of (some of) those contexts.

The desire for identity conditions for objects in one's ontology is defensible.

A specific demand for identity conditions on meanings is even more defensible.

Meanings, if there are any, play a central role in communication.

A theory of meaning that entails that people can share their (identical) meanings seems a requirement if one wants to explain communication.

Meanings, e.g. Frege's propositions, inhabit a third realm.

So do mathematical objects, like sets.

Quine does not argue that we can avoid all abstract objects.

In fact, he defends the belief in sets, as long as they are construed extensionally.

But, he believes that the empiricist can avoid intensional objects.

The difference between intensional and extensional objects is that we have clear identity conditions for extensional objects.

We know when we have two and when we have one.

Sets are identified with their members; two sets with exactly the same members are the same set.

But, we don't know when two meanings are the same.

There seems little hope of erecting a fruitful science about them. It is not even clear, granted meanings, when we have two and when we have one; it is not clear when linguistic forms should be regarded as *synonymous*, or alike in meaning, and when they should not (2D, 64).

The second part of Quine's Ockhamist argument is that even if we could construct identity conditions for meanings, they would be otiose.

What we need from meanings is an explanation of synonymy and analyticity.

If we can get these without meanings, then we don't need them.

If a standard of synonymy should be arrived at, we may reasonably expect that the appeal to meanings as entities will not have played a very useful part in the enterprise. A felt need for meant entities may derive from an earlier failure to appreciate that meaning and reference are distinct. Once the theory of meaning is sharply separated from the theory of reference, it is a short step to recognizing as the business of the theory of meaning simply the synonymy of linguistic forms and the analyticity of statements; meanings themselves, as obscure intermediary entities, may well be abandoned (ibid).

The spookiness of meanings is related to the spookiness of essences.

Meaning is what essence becomes when it is divorced from the object of reference and wedded to the word (2D, 64).

For Aristotle, objects had essential characteristics, and accidental ones.

Persons had essential characteristics (e.g. rationality) and accidental ones (e.g. two-leggedness).

This difference now supposedly goes into the meaning of a term.

The meaning of 'biped' includes being two-legged but the meaning of 'man' may not.

Kripke's rehabilitation of essences, and necessity, were of course unwelcome to Quine and Quineans.

We will return shortly to Quine's meanings skepticism in "Ontological Relativity," where he calls meanings realism (like that of Frege) the myth of the museum.

IX. Confirmation holism and semantic holism

The final section of Quine's article is called "Empiricism without the Dogmas."

We have been looking at Quine's skepticism about intensional objects: there are no meanings.

Quine does not, though, deny that sentences (and terms, derivatively) have meaning.

He defends meaningfulness without meanings.

Further, he believes that meaning is properly the property of a much larger unit than the word (as Locke and Hume held), or even the sentence (as Frege and the positivists held).

The unit of significance, for Quine, is one's entire theory, the whole of science!

(See Quine's "Five Milestones of Empiricism" for a brief, helpful discussion of these differences.)

Ayer characterized analytic statements as those onto which we hold come what may.

In the absence of an analytic/synthetic distinction, one might think that Quine would argue that there are no sentences that are immune from revision or abandonment.

This is correct, but there is another consequence.

If we see science as a total field of information, then we can hold on to any statements we want, as long as we adjust our body of knowledge and the logical framework along with it.

The claim that any sentence may be held immune from revision because no sentence is absolutely immune, is called confirmation holism.

Confirmation holism follows from the failure of the analytic/synthetic distinction, and the identities of necessity with analyticity and contingency with syntheticity.

It is the claim that all sentences are confirmed or disconfirmed as a whole body.

One can also argue for confirmation holism from weaker premises.

It is just a point of the logic of theory construction.

Quine's stronger argument for confirmation holism also yields the more contentious semantic holism:

16. Semantic holism: the unit of meaning is the entire language.

Most people's beliefs, their entire language, are too loose to serve as a good semantic theory.

Instead, to find meaning, Quine looks at our best theory, the theory which best accounts for all our sense experience.

Quine compares our scientific beliefs to a giant web.

Experience forms the boundary conditions on the web of belief.

Peripheral statements are those most closely tied to sensory experience.

Central statements are those about logic, mathematics, and the self, the guiding principles of science, highly theoretical statements.

Quine also compares the web to a field of force.

Experience forces us to adjust and readjust the whole field, not one sentence at a time, but altogether.

The underdetermination of the field by boundary conditions gives us play among the statements.

We can hold to the truth of any statement, come what may, by merely adjusting other statements.

Some may require large scale adjustments, others minor ones.

In the face of an odd sense experience, we can, drastically, claim hallucination, or revise logic.

Different persons growing up in the same language are like different bushes trimmed and trained to take the shape of identical elephants. The anatomical details of twigs and branches will fulfill the elephantine form differently from bush to bush, but the overall outward results are alike"

(*Word and Object*, 8).


There is much, much more to say about holism.
Semantic holism is another worthwhile paper topic.

X. Ontology and posits

Here is one last note on empiricism without the dogmas.

Along with the loss of reductive justifications of particular sentences, we lose a straightforward method for determining our ontology.

Instead of direct lines from physical objects to sense data to singular terms, we have to determine our ontology by appeal to the whole of science.



Science is a tool, for predicting future experience in the light of past experience.

Physical objects are just convenient posits of a scientific theory. “[C]omparable, epistemologically, to the gods of Homer” (44).

We already accept an ontology of posits for distant objects and very small objects, like electrons.

Quine argues that all our ontology is of that form.

Quine’s method does not denigrate physical objects.

“To call a posit a posit is not to patronize it” (*Word and Object*, 22).

The method of positing is just a result of the failure of reductionism and the turn towards holism.

The difference between questions of the existence of sets, say, or quarks, and questions of the existence of houses is only one of degree, not of type.

Posits are accepted or rejected according to pragmatic considerations of theory construction, as well as their coherence and consistency with our broader theory, the web of belief.

Quine’s procedure for determining our ontic commitments (QP) is as follows.

QP.1: Select a best scientific theory, one which balances simplicity, strength, and fit with sense experience.

QP.2: Regiment that theory in first-order logic with identity.

QP.3: Model the resulting formal theory.

QP.4: Examine the domain of quantification of the theory to see what objects the theory needs to come out as true.

One consequence of QP is that he reconciles mathematical ontology with empiricist epistemology.

Traditionally, empiricists had difficulty explaining how we could have knowledge of the abstract objects of mathematics.

Like Fregean propositions, they inhabit a third realm, and do not impinge on our sense organs.

By turning all ontology into scientific posits, Quine opens the door for any objects which facilitate, in serious ways, the construction and regimentation of scientific theory.

Quine's justification of mathematics is called the indispensability argument, for its claim that mathematics is indispensable to science.

(I wrote my dissertation on the indispensability argument.)

XI. Defending a dogma

Quine declares the analytic/synthetic distinction, not merely useless, inadequately defined, or vague in application.

He argues that it is a philosophical mistake.

H.P. Grice and P.F. Strawson, in their response to Quine, "In Defence of a Dogma," argue that there is a long tradition of distinguishing the analytic from the synthetic.

(Their paper is available in the readings section of the course website.)

In general, if a pair of contrasting expressions are habitually and generally used in application to the same cases... this is a sufficient condition for saying that there are kinds of cases to which the expressions apply; and nothing more is needed for them to mark a distinction (Grice and Strawson, 143).

Instead of examining the actual use that we make of the notion of meaning the same, the philosopher measures it by some perhaps inappropriate standard... and because it falls short of this standard, or seems to do so, denies its reality, declares it illusory. (Grice and Strawson, 147).

Philosophers generally agree on the extensions of the terms 'analytic' and 'synthetic'.

The fact that philosophers make use of the distinction implies that there is a distinction to be made.

Consider the distinction between a three-year-old child being able to explain Russell's theory of types and a three-year-old being an adult.

The first case is impossible, but not in the way that the second case is impossible.

This example draws a distinction which everyone can use, the difference between not understanding what one could mean and not believing some statement that one understands.

The distinction requires clarification.

The question is whether it is philosophically significant.

Use doesn't guarantee philosophical significance, but Grice and Strawson think it does.

Quine might just reformulate his argument by saying that philosophers are misled about the distinction that is made.

Philosophers may be allured by an attractive but mistaken theory.

Grice and Strawson point out that Quine has to dismiss not only analyticity, but cognitive synonymy, and 'meaning the same as', which are each part of the same intensional family.

They are inter-defined.

As I noted in §VI above, Quine wants a definition to be non-circular, in the sense that it doesn't use any other notions from the family circle (in this case, 'self-contradictory', 'necessary', 'synonymous', 'semantical rule', 'definition', etc.)

And, the definition must provide necessary and sufficient conditions for application.

Grice and Strawson argue that these criteria are too strict.

If the availability of such an explanation were required for an expression to make sense, then very few, if any, terms would make sense.

Further, if we get rid of 'means the same as', we are left with only extensionality to mark differences among words.

Extensionalists are forced to insist that it is always absurd to say that predicates x and y in fact apply to the same things, but do not have the same meaning.

The difference between 'creature with a heart' and 'creature with a kidney' seems to be a clear-cut counter-example to extensionalist theories of meaning.

Grice and Strawson argue that we can revise the definition of 'synonymy', in light of Quine's thesis of the interconnectedness of statements of a theory, in order to isolate the meaning of a sentence.

Two statements are synonymous iff any experiences which, on certain assumptions about the truth-values of other statements, confirm or disconfirm one of the pair, also, on the same assumptions, confirm or disconfirm the other to the same degree (Grice and Strawson, 156).

That is, Grice and Strawson think we can hold the rest of the theory constant and achieve a sentence isolation.

This might give us a good definition of statement synonymy, thus saving (perhaps) verificationism. Perhaps the interconnectedness of holism does not work as an argument against analyticity.

XII. Indeterminacy of translation and the myth of the museum

Translation between languages relies on synonymy between terms or sentences of the different languages.

Thus, if there is no analytic/synthetic distinction because there is no synonymy, Grice and Strawson argue, all talk of correct or incorrect translation meaningless.

Consideration of their claim led Quine to consider in more depth the problem of radical translation.

Quine concludes that translation is in fact indeterminate.

Quine's arguments for the indeterminacy of translation are found throughout the work of his middle and later periods.

One excellent place to see them is in the second chapter, "Meaning and Translation," of his most important book, *Word and Object*.

"Ontological Relativity," which I have instead assigned, is a later work.

It includes arguments for indeterminacy, specifically for inscrutability, in the early part of the paper.

We will focus mainly on the first part of the paper, through p 51, though you can skip or skim the material on deferred ostension on pp 39-45.

That a denial of meanings should run into difficulties with translation should come as no surprise, especially when the argument depends on a failure to make sense of synonymy.

Translations are just systems of synonyms.

In "Ontological Relativity," Quine calls meanings realism the myth of the museum.

In the realist's museum, the meanings are objects, whether mental objects or abstract, third-realm objects.

When we translate from one language to another, we switch labels on a meaning, which is independent of any language, and which maintains its determinate properties.

For example, we can switch labels from 'kichwa chake kikubwa' to 'his head is big', both of which express the determinate proposition that his head is big.

Quine makes the connection between the myth of the museum and the indeterminacy of translation clear.

Uncritical semantics is the myth of a museum in which the exhibits are meanings and the words are labels. To switch languages is to change the labels (OR, 27).

The ramifications of Quine's criticism of the myth of the museum thus trace all the way back to the original argument against the analytic/synthetic distinction in "Two Dogmas."

To see the connection explicitly, consider a version of QD from §VII, above:

QM1. If there is an analytic/synthetic distinction, there must be a good explanation of synonymy.

QM2. The only way to explain synonymy is to posit determinate meanings.

QM3. But there are no determinate meanings; the museum of meanings is a myth.

QM4. Thus, there is no good explanation of synonymy.

QMC. And thus there is no analytic/synthetic distinction.

(I do not know whether Quine ever explicitly formulated this argument.)

Both the arguments against the myth of the museum and the arguments for indeterminacy of translation support QM3.

Quine objects generally to mental entities, but says that the main objection here is to the pre-determination of meaning, to the claim that meanings have some kind of nature independently of how we use language.

His primary argument against the myth of the museum is epistemic.

Even if there were meanings, there would be no way to know them, whether we take them to be ideas or abstract objects.

There is no way for us to apprehend meanings, "[Beyond] what may be implicit in [our] dispositions to overt behavior" (OR, 27).

When we translate from one language to another, we do not merely switch labels on internal exhibits, we look for translation manuals which fit all and only the overt behavior of the native.

XIII. Radical translation and behavioral constraints

Quine's thesis of indeterminacy of translation focuses on radical translation, the translation of completely unrelated languages.

In radical translation, the field linguist attempts to translate a completely alien language into her home language.

She starts with no hints, and the presence of no bilinguals.

Quine focuses on radical translation in order to avoid any preconceptions, any pre-made translation equivalences.

We start only with two languages, that of some natives and that of a field linguist, and start to guess what the native sentences mean.

The evidence for a translation is supposed to be all the evidence we ever have for understanding people.

Language, Quine claims, is “[A] social art which we all acquire on the evidence solely of other people’s overt behavior under publicly recognizable circumstances” (26).

The idea here is commonsensical.

We start to learn language mostly by ostension, the paradigm case of which is simple pointing.

By ostension, we learn to put labels on objects.

As our knowledge of language grows, we discover some words that do not ascribe observable traits to observable things.

Learning abstract terms, logical terms, prepositions, plurals, and individuating terms requires a more subtle learning process, which Quine calls deferred ostension.

Still, the tools we have to learn language are exhausted by behavioral evidence.

We will return to this view of language, and language-learning, later in the term, when we look at Wittgenstein’s private language argument.

For now, we will just note that Quine’s behaviorism is not necessarily metaphysical.

He is not primarily denying that there are mental objects or events.

His behaviorism is centrally epistemic: behavioral evidence is all the evidence we have.

If we learn something that can not be traced directly to overt behavior, then we must have learned it indirectly, in some complex way, from behavior.

If there is a fact of the matter about which of two translations of a native sentence is right, or which words are synonyms, then there would have to be behavioral evidence to decide the matter.

In the absence of any observable evidence that could decide which translation is correct, we have to conclude that translation is indeterminate.

XIV. Underdetermination, indeterminacy, and inscrutability

There are three levels of indeterminacy, broadly construed.

- I1. Underdetermination of scientific theory
- I2. Indeterminacy of translation.
- I3. Inscrutability of reference.

Underdetermination of scientific theory is what happens when there is a lack of information to settle a scientific question.

Our own beliefs, too, are underdetermined by empirical evidence.

To re-use an example, we do not know Socrates’ blood type.

Any theory of the blood types of all human beings who ever lived will thus be under-determined by the evidence.

Similarly, we do not know whether and how much of the universe is made of dark energy.

Our current scientific theories are underdetermined by the evidence.

Consider two scientific theories.

From the first, we derive the claim:

ST1: Dark energy makes up 73% of the universe.

From the second, we derive the claim:

ST2: Dark energy makes up 74% of the universe.

At the moment, let us assume, we lack the evidence to decide between the theories which yield ST1 and ST2.

We can call those theories empirically equivalent, for now.

But, this underdetermination is merely an epistemic problem.

We do not conclude that there is no fact of the matter about which theory to choose.

We just do more research.

At some point, we expect, those theories will no longer be empirically equivalent.

That is, underdetermination of scientific theory is resolvable by doing more research

I1 entails no metaphysical conclusions.

I2, in contrast, is a deeper, more troubling phenomenon.

It occurs on, or beneath, the sentence-level.

We already saw Quine's topiary metaphor: entire theories could look the same on the outside, but be constructed quite differently on a piece-by-piece (i.e. sentence-by-sentence) level.

In terms of radical translation, the indeterminacy thesis says that it is possible to have incompatible translation manuals each of which is consistent with all behavioral evidence.

Quine defends indeterminacy in *Word and Object*.

Hartry Field, in his article "Quine and the Correspondence Theory," provides an example from science.

Translating Newtonian theory into relativity theory, we can translate mass as either relativized mass, in which case momentum is mass times velocity, and mass is not invariant.

Or, we can translate it as rest mass, in which case mass is invariant, but the momentum equation doesn't work.

We will not pursue indeterminacy at the sentential level, here.

In "Ontological Relativity," Quine mainly focus on sub-sentential indeterminacy.

Inscrutability of reference, I3, occurs on the sub-sentential level.

Quine's claim is that there is no fact of the matter about the references of my terms.

Quine presents five examples of inscrutability of reference:

IR1: The French ne...rien construction

IR2: Gavagai

IR3: Japanese classifiers

IR4: Concrete general and abstract singular terms

IR5: Gödel numbering and deferred ostension (which we will ignore, here)

In IR1, you do not know whether to translate the French 'rien' into the English 'nothing' or 'anything'.

It depends on how you translate the rest of the construction.

If you take 'rien' as 'anything', then you have to take 'ne' as 'not'.

But, if you take 'rien' as 'nothing', then you have to take 'ne' as empty, or pleonastic, as an essential part of the 'ne...rien' construction.

The museumist can object that Quine is dicing up the pieces of a sentence too small, that larger segments (i.e. the whole 'ne...rien' construction) carry meaning.

But the lesson of IR1 is useful: by adjusting some portions of the translation, we affect others.

More specifically, we can see that adjustments can be made among logical particles.

By playing with the notion of identity, for example, we can map different ways of cutting up the world onto each other.

The phenomenon of adjusting logical particles in order to make distinct interpretations of other terms equivalent is common to all of Quine's examples of inscrutability.

IR2 is an example which runs throughout all of Quine's work.
Consider the field linguist, standing with a native, when a rabbit runs by.
The native says, "Gavagai."
We have various options for translating 'gavagai' into English.

- G1: rabbit
- G2: undetached proper part of a rabbit (urp)
- G3: three-dimensional temporal slice of a four-dimensional rabbit
- G4: instantiation of the universal rabbithood

On G1, we leave our logical particles as they are.
But, we can use G2 if we correspondingly change our translation of the native's individuation terms.
Consider the native's words which we would, consistent with G1, translate into English as 'is the same as'.
If we chose G2 for 'gavagai', we change 'is the same as' to 'is an undetached proper part of'.
The only difference among rabbits and urps and temporal segments of four-dimensional rabbits is the individuation, and individuation cannot be mastered through pure ostension.

The only difference is in how you slice it. And how to slice it is what ostension or simple conditioning, however persistently repeated, cannot teach (OR, 32).

Similar adjustments to logical particles can yield the empirical equivalence of G3, G4, and other possible translations.

For IR3, the example is of a word which comes with a number and an object.
That word either modifies the number, or the object.
Depending on how we use it, the object becomes either a mass term (like water or sepia), or an individuating term (like rabbit).
Either translation is consistent with speech dispositions, just as 'rabbit' and 'urp' are consistent with speech dispositions, as long as we make corresponding changes to the logical and individuating particles of the rest of the language.

For IR4, both 'green' and 'alpha' can be taken either as concrete general terms ('the grass is green') or abstract singular terms ('green is my favorite color').
The only way we can tell them apart is to use our English apparatus of individuation, which itself is indeterminate.

Fuller examples may require complete translations of a complete language, which is a big task.
You can not just translate into different ontology-types term by term; you have to do it all at once.
Proxy functions can map one theory into another, adjusting the logical particles.
But you can get the idea.

One can map the sentences of one's language onto themselves such that behavior (and dispositions to behave) remain the same, and also that the two mappings are clearly different at the sentence and sub-sentential level.

In other words, there can be incompatible manuals for translating one's language which all accord with the totality of speech dispositions.

XV. Simplicity and Chauvinism

Against inscrutability (at IR2), we might argue that ‘rabbit’ is simply simpler than ‘undetached rabbit part’, and that simplicity should be our guide in translation.

The problem with this response is that we discriminate terms by projecting our own attitudes towards grammar, and logical form.

The maxims for determining reference assume the linguist’s own referential apparatus.

When we try to create a translation manual for a radically different language, we will meet the problem of whether to translate into rabbit ontology or urp ontology.

It is linguistically chauvinistic to imagine that simplicity in our language is simplicity over all.

An actual field linguist would of course be sensible enough to equate “gavagai” with “rabbit,” dismissing such perverse alternatives as “undetached rabbit part” and “rabbit stage” out of hand. This sensible choice and others like it would help in turn to determine his subsequent hypotheses as to what native locutions should answer to the English apparatus of individuation, and thus everything would come out all right. The implicit maxim guiding his choice of “rabbit,” and similar choices for other native words, is that an enduring and relatively homogeneous object, moving as a whole against a contrasting background, is a likely reference for a short expression. If he were to become conscious of this maxim, he might celebrate it as one of the linguistic universals, or traits of all languages, and he would have no trouble pointing out its psychological plausibility. But he would be wrong; the maxim is his own imposition, toward settling what is objectively indeterminate (OR, 34).

Quine also considers a people who call pelicans their half-brothers.

When they talk about what we refer to simply as half-brothers, they have to use a longer term, equivalent to ‘half-brother, but not a pelican’.

And they have a short term for our long ‘half brother or pelican’.

So simplicity, for example, will not do as a guide.

XVI. Inscrutability begins at home

Given Quine’s rejection of the myth of the museum, his view of language as a public object, and his exclusive reliance on behavioral evidence limited by speech dispositions, one might start to wonder whether our own uses of language are determinate or not.

It seems as if I know that I mean rabbit, rather than urp, and that I know which translation of ‘gavagai’ is simpler.

But, if we try to determine how I can mean one rather than the other, we need to appeal to my ideas, which seem off-limits.

If we have no internal grounds for determining correct translations, then there seems to be no fact of the matter about what I say.

On deeper reflection, radical translation begins at home (OR, 46).

Consider what Quine calls homophonic translation, when we translate each string of phonemes into itself.

We use homophonic translation when talking with our friends and family, people who use the same language that we do.

We use a principle of charity even in homophonic translation.
And there are some times when we use heterophonic translations, even among friends.
I know some people who start sentences with 'No', even when they agree with me.
Sometimes, I call them out on it: "So, when you said 'no', you meant 'yes'."
Other times, I just do the heterophonic translation quietly, to myself.

If we can adjust another person's words, and we can translate into proper-part talk without affecting behavior, then we lose the ability to understand our neighbor's assertions as correctly referential.
So if he says that he is talking about rabbits, we don't know if he is, or not.

The inscrutability of reference is not the inscrutability of a fact; there is no fact of the matter
(OR, 47)

Note Quine's reductio argument here.

If there is a fact of the matter in our own words, then there should be a fact about our neighbor.
But since we know that there is no fact about our neighbor's terms, then we know there is no fact about our own terms.

If there is determinacy, then there is a fact of the matter about what we mean when we say something: the right translation would thus be grounded.

But if we had that, then we could have analyticity.

If we have analyticity, then we can get synonymy, and then we have a fact of the matter about what we say, which can make a translation correct.

In the other direction, if we have a correct translation, it must be right in virtue of some fact.

That fact would give us synonymy, which could give us back analyticity.

Determinacy of translation is just another member of the intensional family we have to give up.

It is thus meaningless to ask whether 'rabbit' refers to rabbits or urps or time slices, even in one's own words, absolutely.

It makes sense relative to a background theory which we hold fixed.

Once we see that we can make wholesale adjustments to the interpretation of that background theory and still do justice to all speech/ behavioral dispositions, we lose the absolute nature of reference.

XVII. Ontological relativity

We were led to the inscrutability of reference from our considerations of the theory of meaning.

At first, it looked like the problem was one that merely afflicted Fregean abstract objects.

Meaning was always suspect, but reference seemed more solid.

Some of the members of our seminar were wary of meanings, propositions, and senses, anyway.

And, philosophers like Russell had already tried to avoid them.

Now it looks like both meaning and reference are indeterminate.

We were interested in reference because that was how language hooked onto the world.

If we are going to get any philosophical results from examining language, we have to transcend our focus on language.

If Quine is correct that reference is indeterminate, then our ontological commitments seem to disappear into a foggy haze.

The references of our terms depend on an arbitrary choice of the logic of individuation, which can be arbitrarily and variously interpreted in a series of background languages.

We can only interpret a theory, a web of belief, relative to a background theory.

But that background theory is itself liable to various, empirically equivalent interpretations.

It looks as if we are involved in an infinite regress of the background languages.

Quine thus defends a relational theory of reference.

What makes sense is to say not what the objects of a theory are, absolutely speaking, but how one theory of objects is interpretable, or re-interpretable in another (OR, 50).

We can never fully interpret a theory, (i.e. say what the singular terms signify or denote, or what go into the extensions of the general terms) because this would say absolutely what the objects of that theory were.

Quine invites comparisons to questions of inverted spectra, and inverted fields of vision.

The problem of the inverted spectrum, which traces back to Locke, is the question of whether our qualitative experiences of color are the same as other people's experiences of color.

What if every time I saw red, you saw violet; every time I saw yellow, you saw blue?

If I learned to use language the same way that you use it, and there is no possibility of knowing what another person's qualitative experiences are, it seems possible that my color experience is exactly inverted from yours.

Quine argues that ontic commitments, and the referential apparatus, of any one person's theory may be similarly indeterminate.

[The] network of terms and predicates and auxiliary devices is, in relativity jargon, our frame of reference, or coordinate system. Relative to *it* we can and do talk meaningfully and distinctively of rabbits and parts... We contemplate alternative denotations for our familiar terms. We begin to appreciate that a grand and ingenious permutation of these denotations, along with compensatory adjustments in the interpretations of the auxiliary particles, might still accommodate all existing speech dispositions. This was the inscrutability of reference, applied to ourselves; and it made nonsense of reference. Fair enough, reference *is* nonsense except relative to a coordinate system. In this principle of relativity lies the resolution of our quandary (OR, 48).

In practice, we end the regress through pointing (or something else practical).

But in the end, there is no determinacy.

It only makes sense to talk about how to interpret theories in other theories.

Field argues that Quine's thesis is radical because it seems to preclude any correspondence notion of truth.

A correspondence theory is based on correspondence relations between the world and our words.

The correspondence relations are denoting, signifying, and referring.

For example, 'The cat is on the mat' is correspondence-true iff the object denoted by 'the cat' bears the 'is-on' relation to the object denoted by 'the mat'.

These relations are denied by indeterminacy.

Quine suggest that we can hold a background theory (or, a translation manual) fixed, adopting an arbitrary translation manual.

That is, we can forget the indeterminacy of the referential apparatus, and just talk about rabbits, instead of the other options.

Field points out that Quine is too sanguine about fixing a background theory, given the rest of his arguments.

The problem is that we can't make sense of the phrase 'relative to a given translation manual.'

For, in order to explain that notion, we have already to understand what it is to denote (absolutely) relative to a given translation manual.

That is, in order to hold the translation manual constant, we have to have some notion of determinate reference within that scope.

But if indeterminacy is as infectious as Quine says, we can't even get that.

Field actually accuses Quine of being a closet museumist, here.

In the end, Field claims that indeterminacy has no bearing at all on ontological relativity.

Field says we should transcend the individual manuals,

An adequate theory has to look at all of the equivalent translation manuals.

He claims we can construct a correspondence theory with more general correspondence relations, such as partial signification, and partial extensions.

So 'gavagai' partially signifies rabbits and partially signifies undetached rabbit parts.

And these are the partial extensions of 'gavagai'.

Field's paper is excellent, if difficult, and might be worth an ambitious paper.

Lastly, we will look at a different kind of response to Quine.

XVIII. The refutation of indeterminacy

Jerrold Katz, in his response to Quine called "The Refutation of Indeterminacy," argues that Quine's argument against synonymy in "Two Dogmas," and consequently his argument against translation, is missing a piece.

"The Refutation of Indeterminacy" is available in the Readings section of the course website.

Katz points out the connection between the earlier work in "Two Dogmas" and the later work on translation.

Translation is critical because it is the only relation that provides interlinguistic identity conditions that are discriminating enough to individuate the fine-grained propositions of intensionalism (Katz, 229).

Katz interprets Quine's argument as follows:

KIT1. There can be no identity conditions for meanings due to the failures of synonymy.

KIT2. So, meanings can not serve as the common content of sentences and translations.

KIT3. The facts we can ascertain from behavior do not determine unique translations.

KITC. Therefore translation is indeterminate.

Indeterminacy of translation is a skepticism about meanings.

If there were meanings, then KIT3 would be false.

Katz reduces Quine's meanings skepticism to an absolute skepticism about all knowledge. He denies indeterminacy, but admits that it is logically possible. If indeterminacy is merely logically possible, the force of Quine's argument is severely deflated. It is also logically possible that there are no other minds, that the world was created five minutes ago with all its history and memories in place, and that there is no material world. But, we do not really take these possibilities seriously. In other words, Katz argues that we can put aside Quine's skeptical worries about meanings.

Katz argues that Quine erroneously derives an ontological skepticism from epistemological considerations. All Quine has shown is that there is no evidence which will allow us to choose among various analytical hypotheses. There is a difference between the non-existence of meanings and the unknowability of meanings. Quine's argument at best only gets the unknowability of meanings, not their non-existence

Further, Quine's physicalism countenances mathematical objects and theoretical physical particles, for their virtues in systematizing our experience. Thus, it would be inconsistent to argue against meanings and synonymy from behaviorist, verificationist principles. If meanings help us to systematize our experience of language, the epistemic worries Quine discusses dissipate. We can take meanings as theoretical posits. In addition, the physicalist could easily allow meanings if they were reducible to brain states.

More specifically, Katz claims that Quine's argument against synonymy fails, because there are linguistic tools that ground synonymy. Quine's example from linguistics, S3: Interchangeability, begs two questions. First, Quine does not make it clear in which contexts two linguistic forms are supposed to be interchangeable. Second, Quine neglects to correctly specify which property interchangeability must save. Analyticity and synonymy can not be defined, given substitution criteria as the proper method of definition in linguistics, since the context must be either intensional extensional. If we choose an intensional context, we can choose to preserve truth, but then we are circular in the choice of context, because we need synonymy or a related concept to characterize an intensional context. If we choose an extensional context, we can not choose truth as the property to be preserved, because it is not strong enough, as the creature with a heart/creature with a kidney example showed. If we try to get a stronger criterion, we confront circularity again. Instead, Katz argues, we should appeal to semantic theorems just like the syntactic ones we find in Chomskyan generative grammars. Generative grammars include theorems modeled on recursive systems of logic. Theoretical definitions differ from substitution criteria in that they can use the criteria being explained, like recursively axiomatizable definitions do.

The degree of relatedness exhibited among the concepts in the family is thus a measure, not of circularity, but of the systematizing power of the explanation (Katz, "The Refutation of Indeterminacy," 240-1).

Within recursive semantic systems, Katz argues, we can produce theoretical definitions of synonymy.

First, we must construct a formal representation of sense structure, modeled on the constituent structures of syntax, showing how senses of sentences are composed of senses of their constituent parts.

Then, we merely have to find the formal features of those representations that correlate with particular semantic concepts, like analyticity.

In particular, Katz argues that we must pursue what he calls decompositional semantics.

Decompositional semantics represents semantically simple terms (like 'bachelor') as complex, in terms of the senses of the words.

A syntactically simple term can contain senses, and thus do justice to Kant's containment metaphor.

Thus, appealing to semantic theories on the model of Chomskyan generative grammars deflates both the problem of circularity which plagued attempts to define synonymy in terms of substitutivity, as well as the problem that Carnap's meaning postulates had of providing just an un-explanatory list of analytic sentences.

Theoretical definitions in linguistic theory provide a way of defining concepts in the theory of meaning for variable 'S' and 'L' because, in defining a concept at the level of linguistic theory, they define it in terms of features of optimal generative grammars for every natural language (Katz, "The Refutation of Indeterminacy," 243).

Once we establish that there are meanings, the translator can look for a blend of data and methodological considerations in performing radical translation.

They can rely on judgments about senses of expressions, like whether expressions are meaningful, ambiguous, synonymous or redundant.

Linguists can make guesses, following hunches, just like in other sciences.

Further, bilinguals are actually very helpful.

XIX. Summing up

It might be nice to look at Quine's responses to Frege's puzzles.

He solves Frege's problem with opaque contexts by refusing to substitute into opaque contexts.

He likes Russell's solution to the failure of presupposition.

How can he deal with the problem of cognitive content, in the identity puzzle?