

Classes 25-6 - Chomsky's Linguistic Conceptualism

I. Chomsky and Wittgenstein

The last topic on our syllabus concerns the nature of language: What is a language?

We start by reading the linguist Noam Chomsky.

I had assigned Chomsky's work intending to start a discrete unit exploring the ontological status of languages.

But there is a substantial connection between our last unit on the PLA and the current one.

Wittgenstein thought that the meanings of the terms of our language are instructions for their use, that our knowledge of language was an ability to use sentences according to conventional rules.

Wittgenstein, and Kripke's Wittgenstein, opposed the view that knowledge of language is representational, that terms stand for, or are used to communicate, our private mental representations. Millikan, too, accepted the claim that any solution to the rule-following puzzle had to ground our knowledge of rules and meanings in some non-representational aspect of human beings.

Thus, Millikan argued that our ability to follow rules, e.g. to choose plus over quus, is grounded in our evolutionary biological purposiveness, not in our conscious awareness of the rules.

Chomsky, in contrast, takes language to be representational.

He argues against the Wittgensteinian Anthony Kenney, who claims that knowledge of language is the possession of an ability.

For the debate between Kenney and Chomsky, which is not central to our work, we should understand a classic philosophical distinction between two kinds of knowledge: knowledge-how and knowledge-that. This distinction, due to another philosopher with sympathies to Wittgenstein, Gilbert Ryle, is between propositional knowledge, like knowing that the Orinoco River is in Venezuela, and practical ability, like knowing how to ride a bicycle.

Wittgenstein, one could argue, tried to show that all knowledge-that was reducible to knowledge-how. More recently, some philosophers have tried to argue in the other direction, that all knowledge is essentially propositional; here's a nice article from the insufferable Jason Stanley on [knowledge-how](#).

Against the Wittgensteinian approach, Chomsky presents two arguments.

First, we can increase our linguistic abilities without increasing our knowledge, as when we improve our public speaking.

Since we can improve our abilities with language without improving our knowledge of language, that knowledge and ability must be distinct.

Second, we can lose our abilities to use language without losing our knowledge of the language.

Chomsky considers a person with Parkinson's disease, Smith, who temporarily loses his ability to speak or understand, but who regains those abilities after taking medications.

Smith retained his knowledge of the language even though he temporarily lost his ability to use it.

It would be difficult to maintain that our knowledge of language is identical to our ability to use it.

In short, Chomsky has argued against Wittgenstein's claim that meaning is use.

We are back to considering language in terms of mental representations.

II. The Chomsky revolution: nativism and UG

Chomsky made (at least) two revolutionary claims about language.

The first claim is epistemological: our knowledge of language is, in part, built-in to our brains.

This epistemological claim is called nativism, for the innate brain structures governing language that Chomsky posits.

Chomsky's nativism was developed in response to both Skinnerian behaviorism and Piagetian developmental psychology.

According to the behaviorist, we are born with a Lockean blank slate, and our linguistic abilities are completely learned.

Developmental psychologists explored the process of learning language.

The central argument for nativism is called a poverty of the stimulus argument (POTS).

Consider the notion of a grammar: a set of rules for generating the indefinite number of sentences of a language, the formal system that produces the infinite set.

We can generate the indefinite stock of sentences from a finite base set of lexical particles.

The lexicon must be finite, since human language-users can learn it.

Grammars are essential for the linguist's analysis of the compositionality of languages.

The central problem of the theory of language is to explain how people can speak and understand new sentences, new in their experience or perhaps in the history of the language (682).

The POTS argument is in part that children learn too much grammar too quickly for us to account for their grammatical abilities on the basis of behavioral stimulus.

For example, consider the following four sentences, taken from earlier Chomsky work.

1. John is easy to please.
2. John is eager to please.
3. It is easy to please John.
4. It is eager to please John.

1 and 2 are identical, grammatically.

If children were learning grammar behaviorally, they would make the reasonable inductive conclusion that since 3 can be inferred (and used) on the basis of 1, 4 should similarly be derivable from 2.

But, children just do not make that kind of mistake.

For more examples, see "What I Know When I Know a Language" by Barry C. Smith in *The Oxford Handbook of Philosophy of Language*, on reserve in the library.

The POTS argument also relies on the claim that children learn the lexicon (vocabulary) of their first language too quickly to be explained purely behaviorally.

It is a very difficult matter to describe the meaning of a word, and such meanings have great intricacy and involve the most remarkable assumptions, even in the case of very simple concepts, such as what counts as a possible "thing." At peak periods of language acquisition, children are "learning" many words a day, meaning that they are in effect learning words on a single exposure. This can only mean that the concepts are already available, with all or much of their intricacy and structure predetermined, and the child's task is to assign labels to concepts, as might be done with very simple evidence (689).

Thus, Chomsky concluded, our abilities to use language must be built into our brains. There might not be a specific language module of the brain, though Broca's area and Wernicke's area are both important for speech and language processing. Chomsky is committed only to a built-in language center at some abstract level of organization.

The dedicated mental organ for learning language, the language faculty, is described or explained, in its initial state, by a very general universal grammar (UG). UG may be transformed into the particular grammars of our particular languages by transformations according to set parameters. The grammars of all particular languages, Chomsky claims, differ only in trivial ways. All languages have essential common features, aside from their differences in lexicon, explicable by biology.

From an angel's point of view, all languages would appear identical, apart from trivialities, their fundamental features determined by facts about human biology (687).

When we acquire a language, we develop the language structure of the brain. Chomsky denies that we learn languages, preferring to say that we grow them.

Acquiring language is less something that a child does than something that happens to the child, like growing arms rather than wings, or undergoing puberty at a certain stage of maturation (680).

The claims about nativism and UG are controversial, but they may be supported or refuted empirically. To establish that there is a universal grammar, we would need to evaluate Chomsky's claims about the triviality of differences among natural languages. We would need a linguistic theory of each language, and of UG, and a description of the parameters and transformations that take UG to those natural languages. Thus, what I have called Chomsky's first revolutionary claim led to an intense and productive research project in linguistics, and to the opening of linguistics departments, in the 1960s and 1970s, in universities around the world.

Establishing nativism is trickier. One way to defend nativism is to show that behaviorism is as explanatorily vacant as Chomsky claims. We would need to show that the stimulus is really that poor. Such a defense will require appeal to Chomsky's second revolutionary claim.

III. The Chomsky revolution: the competence/performance distinction

Chomsky's second revolutionary claim is methodological: there is a distinction between competence and performance in language. People often fail to use their own languages correctly: they use words they do not intend, they fail to finish their sentences, they speak ungrammatically. These are errors in performance. People's performance varies widely, even in their native language. If the study of language were the study of the performance of speakers of the language, linguistics would be extremely messy.

Performance varies so widely, it would be difficult even to distinguish one language of the various speakers who can all understand each other.

Performance errors, though, do not impugn the competence of a speaker, which can be taken as the real locus of the study of language.

That is, we can idealize the object of our study of language by appealing to the competence of native speakers, rather than their actual performance.

We can see the competence/performance distinction at the beginning of our selection, where Chomsky contrasts a formal approach to language with our common, or folk, understanding of what a language is. The common understanding of language involves broader concepts about performance: dialect, interpretation, class structure, and authority.

Explanations of these broader concepts all, it could be argued, require appeals to social conventions.

Chomsky argues that these broader facets of language fail to explain many linguistic phenomena.

For example, consider the phenomenon of pronoun binding, as exemplified in the following examples taken from Anne Bezuidenhout.

5. Mary expects to pay for herself.
6. I wonder who Mary expects to pay for herself.
7. Mary expects to pay for her.
8. I wonder who Mary expects to pay for her.

7 and 8 are identical to 5 and 6, except for the substitution of the pronoun 'her' for the pronoun 'herself'. But, the reference of the pronoun varies.

In 5, the pronoun has to refer to Mary, whereas in 7 it has to refer to someone else.

In 6, the pronoun has to refer to someone other than Mary, whereas in 8 it can refer to either Mary or someone else.

If we take language to be concerned with performance, then the differences among 5-8 should be explicable in terms of some sorts of social conventions.

But, there are no social conventions that dictate the binding of pronouns.

We can choose to switch our conventions: to drive on the other side of the road, or stop using francs and lire and start using euros.

What govern the references in 5-8 are something more like linguistic rules than conventions.

Chomsky's approach to competence, then, takes language to be a narrower object, independent of the social forces on language.

IV. EIEIOntology

Both Chomsky's epistemological and methodological claims leave open the question of the ontology of language.

We might take languages to be abstract objects, independent of us.

Or, we might take languages to be psychological objects, products of our minds.

In Chomsky's terms, we can take language to be extensional (E-language) or intensional (I-language).

An E-language is extensional in that it is a set of objects, perhaps inscription types or meanings.

It is external in the sense that it is not a mental object.

E-languages transcend any particular users, since they are not constructed by us, and are objective.

In contrast, Chomsky thinks that language is intensional, and argues against those who take linguistics to study E-languages.

Chomsky discusses two possible ways to refine the notion of an E-language.

The first is Bloomfield's characterization of language as the totality of utterances that can be made in a speech community.

A speech community is an ideally homogeneous group of language users.

Bloomfield's account of language is essentially behaviorist, relying on a taxonomy of language in use.

The actual uses of language are not sufficient, though, since people can form novel sentences on the basis of their understanding of lexicon and grammar.

Thus, Bloomfieldians had to include possible utterances in their ontology.

Notice that Bloomfield's characterization relies on several notions that are awkward for a behaviorist.

There are possible utterances, utterances which are not actually used.

And, there is a speech community, which is an idealized, homogeneous group of people.

Chomsky's argument against Bloomfield, which will be important for next week's work, is that a behaviorist can not really help him/herself to these notions.

The second characterization of an E-language, from David Lewis for example, is a bit more technical and formal.

Lewis takes a language to be the relevant set of all ordered pairs of sentences or utterances and meanings.

Here we encounter abstract objects in the sets and ordered pairs.

We also can take meanings to be either intensional (as Frege does) or extensional, as the set of possible worlds at which the sentence is true, say.

I will put aside the Lewis proposal, and return to the account of language as an abstract object next week.

You might look at Lewis's article, in Martinich, called "Languages and Language."

(We skipped this difficult piece.)

Chomsky presents several puzzling arguments against E-languages.

He says that if we think of language as E-language, we have difficulty determining whether sentences like 9 are in the language or not.

9. The child seems sleeping.

It seems like 9 is not part of English, since it is ungrammatical.

But, it also seems like 9 is part of English, since English speakers assign meaning to it, whereas people who do not speak English do not.

This argument does seem to present a challenge.

But, it is a challenge for all theories of language to individuate their objects of study; we should admit no entity without identity.

In our article, Chomsky does not make it clear how taking languages to be intensional solves the problem of individuating languages.

What Chomsky calls a second problem is related to Wittgenstein's rule-following puzzle.

If we take language to be a set (of pairs of utterances, say, and meanings), there will be many different ways to generate that set.

It looks as if the choice among these options is arbitrary.

Grammar is the way in which a language is generated from its lexicon.

So, on an extensional theory of language, grammar is conventional, rather than guided by the implicit rules of an I-language.

Empirical research in linguistics, on the structure and parameters of both UG and natural languages,

could support Chomsky's I-languages, if that research generated a rule-guided theory of language. On the other hand, a straight solution to the rule-following problem would deflate this criticism. Whatever accounts for our use of plus rather than quus could account for our choice of a particular grammar, even understood as a set of formal rules for generating sentences of the language.

Lastly, Chomsky argues that taking languages to be extensional sets raises troubling questions about the formal properties of those sets.

Are they context-free, or recursive, or denumerable? All of these choices have been affirmed, and denied, but the point is that the questions are taken seriously, though it is far from clear that the questions are even meaningful. The answers are also thought to have some crucial bearing on questions of parsing and learnability, but quite wrongly, for reasons discussed years ago... The notion of an E-language is an artifact, with no status in an eventual science of language (678).

I fail to see any significant argument against E-languages, here.
But, it might be worthwhile to look at Chomsky's other writings.
Also, you might again see the Smith article in the *Oxford* collection.
For now, I will move ahead to Chomsky's characterizations of I-languages.

V. I-languages, rules, and principles and parameters

Chomsky is considering an I-language, a set of mental representations.

"I" is to suggest "intensional" and "internalized." The I-language is what...grammar purports to describe: a system represented in the mind/brain, ultimately in physical mechanisms that are now largely unknown, and is in this sense *internalized*; a system that is *intensional* in that it may be regarded as a specific function considered in intension - that is, a specific characterization of a function - which assigns a status to a vast range of physical events... (679).

There are different ways that an I-language can be construed.

Chomsky contrasts taking language as a set of rules with taking what he calls a principles-and-parameters approach; only the latter posits UG.

Chomsky defends the principle-and-parameters approach in part on the basis of its fruitfulness for further areas of research.

For example, there are remarkable differences between French and other Romance languages; pp 688-9. These similarities and differences seem traceable to both a common origin and a simple choice of different parameters for a few types of constructions.

More influentially, Chomsky defends the principles-and-parameters approach by arguing that the rules approach fails to explain our language-learning.

On the rules approach, there are a set of rules: context-free rules, lexical rules, transformational rules, phonological rules, and others.

We have to describe how we could come to know these rules.

Chomsky calls this demand for an explanation of how we know the rules Plato's problem.

Again, Chomsky uses a version of the POTS argument; there are just too many possible rules systems for it to be plausible that we learn them.

Furthermore, the rules approach leaves the reason we adopt one language rather than another

unexplained.

Even if appropriate rule systems could be constructed, and even if these systems were found to be restricted in type, we would always want to know why we have these kinds of rules and not others (685).

In contrast, if we take the principles of language (UG) to be built into our minds/brains, as the principles-and-parameters approach suggests, then all we have to learn, in addition to lexical items, are the simple parameters that separate languages.

Chomsky mentions parameters like head directionality, whether a language is head-initial ('Corvette little red') or head-final ('little red Corvette').

English is head-final; other languages (e.g. Japanese) are head-initial.

There are only a few possible structures for languages, all of which, Chomsky claims, have transformations to UG.

There is little doubt that this problem of "poverty of stimulus" is in fact the norm rather than the exception. It must be, then, that the values of parameters are set by the kinds of simple data that are available to the child, and that the rich, complex, and highly articulated system of knowledge that arises, and is shared with others of somewhat different but equally impoverished experience, is determined in its basic features by the principles of the initial state...of the language faculty. Languages may appear to differ, but they are cast in the same mold (687).

Again, UG is a theoretical posit, defended by the value of the theory of language which contains that posit.

A theory of universal grammar, like a particular proposed grammar, is true or false in whatever sense any scientific theory can be true or false (679).

The explananda of the theory in question includes our knowledge of how to speak and understand a language, how we know about the meanings of sentences, and, especially, compositionality.

VI. A return to psychologism?

One worry about taking linguistics to be the study of I-languages is that Chomsky seems to have returned us to psychologism.

Frege proposed a third-realm view of language in order to make talk about language objective.

If we are going to take language to be built into human minds, it seems that we are back in the realm of private, subjective psychology.

In contrast, Chomsky thinks that it is perfectly acceptable to talk about minds, as shorthand for talking about brains.

As I will use the terms, talk about mind is simply talk about the brain at some level of abstraction that we believe to be appropriate for understanding crucial and essential properties of neural systems, on a par with discussion in nineteenth-century chemistry of valence, benzene rings, elements, and the like, abstract entities of some sort that one hoped would be related, ultimately, to the then-unknown physical entities (676). (Note Chomsky's odd use of 'abstract entities'; I think that he really means 'posits'. Compare to his use of 'abstract' on p 678, line 10.)

Chomsky is a naturalist who wants to explain our uses of language in such a way that they are compatible with an understanding of ourselves as physical beings: brains and bodies.

The difference between Chomsky's psychologism (which has come to be known as conceptualism) and the psychologism that Frege opposed is that we have a better understanding of the brain, now.

If the mind is the brain, in some sense, then studying languages as produced by minds can be just as objective as studying languages as abstract objects.

In fact, psychological objects might be taken to be even more objective than abstract objects, since they are subject to the laws of physics.

That is why Chomsky defends the study of cognitive linguistics, of the minds (i.e. brains) of language-speakers.

Chomsky is especially interested in computational theories of mind, for their compatibility with I-languages.

In the past few years it has been shown that a wide range of phenomena from typologically quite different languages can be explained on the assumption that the language faculty of the mind/brain carries out digital computations following very general principles, making use of representations of a precisely determined sort, including empty categories of several kinds. This work then provides evidence, quite strong evidence I believe, for some rather striking and surprising conclusions: that the language faculty, part of the mind/brain, is in crucial part a system of digital computation of a highly restricted character, with simple principles that interact to yield very intricate and complex results (683).

The ability of our brains to perform recursions, in both language and mathematics, is especially striking. If the brain is just a complex digital computer, then a theory of language which relies on a built-in grammar, analogous to a built-in operating system, might be satisfying.

Such a view would undermine, as Chomsky notes, the view of human speech as potentially infinitely diverse.

But, Chomsky doesn't think that view worth saving, anyway: "The conception has been entirely unproductive" (684).

We may return briefly to this topic when we look at Katz's work.

Katz says that languages are maximally effable, capable of expressing any possible concept.

For Chomsky's views on the purported infinite diversity of human speech; see p 684.

VII. Summing up

There is no doubt that Chomsky's conceptualist approach to language, while having matured over the years, has from the start supported an enormously productive research program.

The real tension, philosophically, in Chomsky's approach is between his naturalism about the mind and his apriorism, and realism, about meanings and analyticity.

Putnam criticizes Chomsky for claiming that concepts (for which our learned lexical items stand) are innate.

It does seem that concepts like 'carburetor' could not plausibly be innate.

One option, to which we will return next week, would be to take the meaning of 'carburetor' to be an abstract (Fregean third-real) object.

Then, we would not have to find a place in the brain into which it were built.

But we would have to come up with a theory of how we could come to know about abstract objects.

For Chomsky's views on taking languages to be abstract objects of this sort, see the discussion of P-linguistics on pp 677-8.

Again, we will return to the conception of languages as abstract objects next week.

The natural tension between Chomsky's conceptualism and his nativism is also apparent in Chomsky's arguments against Quine that there are analytic meaning relations.

Quine, Chomsky says, only considered overly-simple examples to deny analyticity.

Philosophers have, I think, been led to this dubious conclusion [that there is no principled distinction between matters of fact and questions of meaning]...by concentrating on an artificially narrow class of examples, in particular, on concepts that have little or no relational structure: such sentences as "cats are animals"... When we turn to more complex categories with an inherent relational structure such as *persuade* or *chase*, or to more complex syntactic constructions, there seems little doubt that analytic connections are readily discerned (691).

Consider:

10. John persuaded Bill to go to college.

11. Bill decided to go to college.

Chomsky claims that 10 analytically entails 11, on the basis of the meaning of 'persuade'.

We can know of the entailment, Chomsky claims, a priori, by reflecting on the meaning of 'persuade'.

Still, Chomsky argues that the question of whether there are a priori, analytic truths is empirical: can we construct a satisfactory theory of language from which these necessary entailments follow?

There's something to contemplate: an empirical justification of a priori knowledge!