Class 26 - Chomsky's Linguistic Conceptualism
Chomsky, "Language and Problems of Knowledge"

I. Chomsky's I-Language

The last topic on our syllabus concerns the nature of language: What is a language?
We start by reading work by the linguist Noam Chomsky.
Chomsky defends a view of language he variously calls I-language, or generative grammar.
He characterizes I-language in opposition to several competing views.
For platonistically-inclined philosophers (like Frege and Katz), languages are abstract objects.
For early linguists (like Saussure and Bloomfield) and some behavioristically-minded philosophers (like Quine) languages are external, sociological phenomena.
Saussure and Bloomfield were socio-linguists, gathering data on various languages using mainly behavioral evidence.
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 opposed the view that knowledge of language is representational, that terms stand for, or are used to communicate, our private mental representations.
These two broad groups, we can call them platonists and behaviorists, conflict about the nature of language.
But they agree that the nature and properties of language are external to any individual.

In contrast, Chomsky understands I-language to be intensional and internal, an intrinsic property of one's cognitive architecture.
Chomsky is not quite a Lockean conceptualist.
Locke believed that language is a private, psychological phenomenon.
Chomsky believes that language is an evolutionary, neurologically-explicable facet of human cognition.

Like Locke, and unlike Wittgenstein, Chomsky 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 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.

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 who temporarily loses his ability to speak or understand, but who regains those abilities after taking medications.

The Parkinson patient retains his knowledge of the language even though he temporarily loses 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 into 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 as well as earlier theories of language, especially those of Bloomfield and Saussure.
According to the behaviorist, our minds are Lockean blank slates at birth.
Our linguistic abilities are learned by induction over our experiences with language.
Developmental psychologists explored the process of learning language.

The central argument for nativism is called a poverty of the evidence, or poverty of the stimulus, argument (POTS).
We can generate an indefinite set of sentences of a natural language like English.
We can generate this indefinite set from a finite base set of lexical particles, or words.
The lexicon must be finite, since human language-users can learn it.
The fact that we can transform this finite lexicon into an indefinite, perhaps infinite, set of sentences is called the compositionality of the language.
Compositionality is an essential feature of natural human languages, and it helps to distinguish human languages from the languages of more advanced non-human animals, like dolphins or chimpanzees.

In order to generate the indefinite set of sentences of a natural language, we combine lexical particles according to certain rules of formation.
These rules are called a generative grammar.
A generative grammar is a formal system that produces the infinite set, like a set of logical axioms from which we can derive all logical truths.

> 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 (18).

The POTS argument has two distinct portions.
First, children learn too much grammar too quickly for us to account for their grammatical abilities on the basis of behavioral stimulus.
Chomsky considers the following two sentences.

1 I wonder who the men expected to see them.
2 The men expected to see them (8).

1 and 2 each contain the same clause.

In 1, 'them' refers back to the men.
In 2, 'them' does not refer to the men, but to some other people.
If children were learning grammar behaviorally, they would make the reasonable inductive conclusion that 'them' has the same reference in each case.
But, children just do not make that kind of mistake.
Similarly, Chomsky elsewhere argues that children learn without instruction that the structural similarities of 3 and 4 do not entail that they are transformable into 5 and 6, respectively.

3	John is easy to please.
4	John is eager to please.
5	It is easy to please John.
6	It is eager to please John.

Children will make the transformation from 3 to 5, but not from 4 to 6.
If they were learning grammar merely behavioristically, we would expect that they would form sentences like 6 sometimes, requiring instruction to eliminate that formation.
Such instruction is never necessary, leading us to believe that the grammatical rules are built into the brain, in some way, rather than learned.
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 second portion of the POTS argument is that children learn the lexicon (vocabulary) of their first language too quickly to be explained purely behaviorally.
While they learn the specific words behaviorally, these words must hook onto pre-existing concepts.

> 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 (29).

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.

Still, different natural languages differ significantly in grammar and lexicon.
If grammar and lexical concepts are built into our brains when we are born, then since our brains do not know which kind of natural language we will use, there must be a wide range of common features among all different human languages.
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.
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 (12).

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 (25-6).

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.

Prior to Chomsky's work, the dominant approaches to linguistics were behavioristic.
Behaviorists took languages to be identified with people's performance, or ability.
Chomsky distinguishes knowledge and ability.
Knowledge is related to competence; ability is related to performance.
If the study of language were the study of the performance of speakers of the language, linguistics would be extremely messy, as it was according to the behaviorist paradigm.

Performance varies so widely, it would be difficult even to distinguish one language of the various speakers who can all understand each other.
On Chomsky's view, performance errors do not impugn the competence of a speaker, which can be taken as the real locus of the study of language.
We can idealize the object of our study of language by appealing to the competence of native speakers, rather than their actual performance.
It is much more plausible that different speakers share competence, or knowledge, than that they share abilities, or performance, given the variance in what people actually say.

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 an argument why languages should be identified with performance, rather than competence, consider the phenomenon of pronoun binding, as exemplified in the following examples taken from Anne Bezuidenhout.

7	Mary expects to pay for herself.
8	I wonder who Mary expects to pay for herself.
9	Mary expects to pay for her.
10	I wonder who Mary expects to pay for her.

9 and 10 are identical to 7 and 8, except for the substitution of the pronoun 'her' for the pronoun 'herself'.
But, the reference of the pronoun varies.
In 7, the pronoun has to refer to Mary, whereas in 9 it has to refer to someone else.
In 8, the pronoun has to refer to someone other than Mary, whereas in 10 it can refer to either Mary or someone else.
If we take language to be concerned with performance, then the differences among 7-10 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 7-10 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. Linguistic Ontology

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.

There are several ways to refine the notion of an E-language.
One 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, to account for compositionality: 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.
Further, Bloomfield's characterization relies on the concept of a speech community, which is an idealized, homogeneous group of people.
Chomsky argues against Bloomfield that a behaviorist can not really help him/herself to the notions of possible utterances and an idealized speech community.

Another characterization of an E-language, which Chomsky attributes to David Lewis, 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.
Chomsky argues that knowledge of such abstract objects is implausible.
Like the behaviorist, Lewis, by taking language to be external to us, blocks any reasonable account of our being able to know about language.
You might look at Lewis's article "Languages and Language."

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 11 are in the language or not.

11 The child seems sleeping.

It seems like 11 is not part of English, since it is ungrammatical.
But, it also seems like 11 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 (9).

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 prefers to think of language as 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... (10).

By focusing on I-language rather than E-language, Chomsky chooses a conceptualist, or psychologistic, approach to language, rather than an objective one.

An I-language can be construed in different ways.
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.
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, knowledge of a language consists of knowledge of 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 by observing behavior.
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 (23).

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 (25).

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 (11).

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 Locke's 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 (7).

Chomsky's use of 'abstract' is odd.
Ordinarily, we use 'abstract object' as a name for things, like propositions or circles or numbers, that have transcendent, mind- and language-independent properties but which are not available to sense experience.
That's the Fregean third-realm view.
We discover rather than construct such objects.
Chomsky uses 'abstract' to characterize objects that are constructed.

> People may study whatever abstract objects they construct (9).

Chomsky's odd usage of 'abstract' may be the result of his insistence that his theory is supposed to be naturalistically-acceptable.
Chomsky 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 (19).

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 (21).

The plasticity and potential of language is a deep topic and makes for odd alliances.

The purported infinite diversity of language is inconsistent with Chomsky's nativism.

> The principles allow for a limited range of variation. That variation is limited has often been explicitly denied. The leading American linguist Edward Sapir held that languages can vary "without assignable limit", and Martin Joos put forth what he called the "Boasian" view, referring to Franz Boas, one of the founders of modern linguistics: namely, that "languages could differ from each other without limit and in unpredictable ways". Such views echo Nilliam [sic] Dwight Whitney, who greatly influenced Ferdinand de Saussure, and who emphasized "the infinite diversity of human speech". Such views perhaps appeared tenable in some form if one regarded language as a habit system, a network of practical abilities, a complex of dispositions, and the like. In that case, language would be constrained only by whatever general conditions constrain the development of abilities and habits in general, by what are sometimes called "generalized learning mechanisms", if these exist. But this conception does not allow one even to approach the essential features of normal language use... (21).

Chomsky's formal approach to linguistics is incompatible with these earlier, sociological views.
Later, Katz argues that languages are maximally effable, capable of expressing any possible concept; we can say anything that we can think.
It's not clear to me whether maximal effability is the produce of the wide capacity of language of our narrow capacity for thought.

VII. Nativism and Naturalism

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 philosophical tension 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 odd to claim that concepts like 'carburetor' could be innate.
One option 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.
That's Katz's route to which we will return 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 (30-31).

Consider:

12 John persuaded Bill to go to college.
13 Bill decided to go to college.

Chomsky claims that 12 analytically entails 13, 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!