

Class 8 - Reflective Equilibrium in Linguistics

I. The Ethics/Linguistics Analogy

Rawls pursued reflective equilibrium for a theory of justice, balancing intuitions about fairness with more general, more abstract theories.

He compares his method to that used by Chomsky and his followers in linguistics.

A useful comparison here is with the problem of describing the sense of grammaticalness that we have for the sentences of our native language. In this case the aim is to characterize the ability to recognize well-formed sentences by formulating clearly expressed principles which make the same discriminations as the native speaker. This is a difficult undertaking which, although still unfinished, is known to require theoretical constructions that far outrun the ad hoc precepts of our explicit grammatical knowledge (Rawls, *A Theory of Justice*, 47).

For Rawls, the complete theory of justice, in reflective equilibrium, will outrun our initial intuitions about justice.

Indeed, it will conflict with some of those initial intuitions.

We want the theory to yield most of our pre-theoretic intuitions, but we also want the theory to be systematically pleasing, and predictive, and, most importantly, to help us settle the questions we do not know how to answer.

In other words, Rawls applies the same kinds of formal, or structural, constraints on his theory of justice as scientists generally place on their theories.

In my notes on scientific methodology, in the section on underdetermination of theories by evidence, I called these constraints the immanent virtues of a theory.

Rawls and Chomsky both see their theories as adopting the methods of science, i.e. seeking reflective equilibrium.

In the previous set of class notes, I noted that linguists, following Chomsky, distinguish between competence and performance.

Linguistic theory, focused on competence, will not match particular performances (or intuitions about grammaticality) perfectly, just as the theory of justice in reflective equilibrium will not match particular intuitions about justice perfectly.

In this class, we will explore the situation in linguistics a bit further, to cash out the metaphor between Rawls's work and that of Chomsky.

II. The Chomsky Revolution: Epistemology

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 as well as earlier theories of language, especially those of Bloomfield and Saussure.

According to the behaviorist, we are born with a Lockean blank slate and our linguistic abilities are learned by standard stimulus-response conditioning.

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 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.

How does every child know, unerringly, to interpret the clause differently in the two cases? And why does no pedagogic grammar have to draw the learner's attention to such facts... (8)

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.

These are facts that we know, however difficult it may be to determine that our system of knowledge has these consequences. We know these facts without instruction or even direct evidence, surely without correction of error by the speech community (11-12).

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 (Chomsky, “Language and Problems of Knowledge,” 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.

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.

The study of one language may provide crucial evidence concerning the structure of some other language, if we continue to accept the plausible assumption that the capacity to acquire language, the subject matter of UG, is common across the species (37).

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.

The claims about nativism and UG are controversial, but they may be supported or refuted empirically.

These are not topics for speculation or *a priori* reasoning but for empirical inquiry, and it is clear enough how to proceed: namely by facing the questions [(i) What constitutes knowledge of

language? (ii) How is knowledge of language acquired? (iii) How is knowledge of language put to use?] We try to determine what is the system of knowledge that has been attained and what properties must be attributed to the initial state of the mind/brain to account for its attainment (4).

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, the dominant approaches to linguistics were behavioristic.

They 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.

Two people may share exactly the same knowledge of language but differ markedly in their ability to put this knowledge to use. Ability to use language may improve or decline without any change in knowledge. This ability may also be impaired, selectively or in general, with no loss of knowledge, a fact that would become clear if injury leading to impairment recedes and lost ability is recovered (9).

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.

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.

Despite appearances, the problem of accounting for the unbounded character of the E-language and the person's knowledge of language... is not squarely addressed in such approaches... (20).

Chomsky's arguments against E-languages are not particularly clear, either in this chapter, or elsewhere that I have seen.

Nonetheless, 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... ("Language and Problems of Knowledge," 679).

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

I have several concerns about Chomsky's characterization of the objects of linguistics as an I-language, but that characterization is not our main concern, here.

I recommend two articles for further reading:

Jerrold Katz, "[The Unfinished Chomskyan Revolution](#)"

Michael Devitt, "[Linguistics is Not Psychology](#)"

Our concern as always in this course is with method.

V. Chomsky's Methods (Or: The Point of It All)

Once we take language to be the study of competence, rather than performance, behaviorist analyses of languages seem extremely implausible.

The behaviorist studies performance and tries to construct a theory out of actual uses of language.

A completely different methodology would have to govern the study of language as competence.

Thus, Chomsky's revolutionary claims about nativism, UG, and competence lead directly to a methodology in linguistics similar to the one we see in Goodman and Rawls.

In all three cases, we are pushed to constructing abstract, general theories on the basis of our intuitions.

Then, we look to balance intuition and theory in reflective equilibrium.

In Chomsky's case, our goals are the specifications of UG, and the transformation rules which map UG into specific natural languages.

To specify the structure of UG, linguists rely on intuitions of ordinary folk regarding grammaticality.

In actual practice, linguistics as a discipline is characterized by attention to certain kinds of evidence that are, for the moment, readily accessible and informative: largely, the judgments of native speakers. Each such judgment is, in fact, the result of an experiment, one that is poorly designed but rich in the evidence it provides (36).

These intuitions about grammaticality form the starting points of the theory, as our intuitions (or considered judgments) about fairness and justice form the starting points of the theory of justice, and our intuitions about which deductions are acceptable form the starting points of logical theory, and our intuitions about confirmation (and the distinction between green and grue, perhaps) form the starting points of scientific theory.

Just as in the other cases, we can not take these judgments to be unassailable foundations.

In general, informant judgments do not reflect the structure of the language directly; judgments of acceptability, for example, may fail to provide direct evidence as to grammatical status because of the intrusion of numerous other factors. The same is true of other judgments concerning form and meaning... (36).

To be sure, the judgments of native speakers will always provide relevant evidence for the study of language, just as perceptual judgments will always provide relevant evidence for the study of human vision, although one would hope that such evidence will eventually lose its uniquely privileged status (37).

Our final theory will accommodate as many of our considered judgments as possible, though, balancing our desire to comprehend our intuitions with interests in theoretical simplicity, strength, and the other virtues Quine discussed.

In sum, there are two ways in which Chomsky's work in linguistics is relevant to our course. First, his program in linguistics uses and instantiates a more-general scientific methodology of seeking reflective equilibrium.

As we have seen, the method of seeking reflective equilibrium may even be constitutive of rational discourse.

Some approaches to knowledge-gathering ignore or downplay the method of seeking reflective equilibrium.

We have seen, in the failures of foundationalism and in the vacuity of simple coherentism, that rational thought must invoke reflective equilibrium.

Chomsky's program embraces Goodman's method.

Second, reflective equilibrium demands a balance between theory and evidence, but it does not itself determine the nature of that evidence.

In science, we take the evidence to be (at least partially) observational.

Chomsky's revolution in linguistics depends essentially on intuitive evidence.

Its enormous success is largely responsible for much of the philosophical community's interest in intuitions.

Chomsky's success helped make it acceptable to invoke intuitions explicitly in philosophy and paved the way for Rawls's rehabilitation of first-order ethics.

It may also have supported a return to a kind of [armchair philosophy](#) that many philosophers revile.