BOOK REVIEW

Stephen Crain & Rosalind Thornton, *Investigations in Universal Grammar: A Guide to Experiments on the Acquisition of Syntax and Semantics*. MIT Press. Cambridge, Massachusetts, 1998.

The principal objective of this book is a practical one: it is intended "to help prepare students of language acquisition to conduct experimental investigations of children's linguistic knowledge" (p. 3). It is not however, merely a manual for setting up acquisition experiments. Crain and Thornton (C & T) maintain, reasonably enough, that experiments cannot be conducted in a theoretical void, and therefore they are quite explicit about their views on language acquisition. It is these views that I will focus on in the following, because I suspect they are more likely to be of interest to the average reader of L & P than the methodological matters that take up the larger part of C & T's discussion. This means, unfortunately, that I will be concentrating on an aspect that is not the authors' chief concern. Still, I believe that my selection is a legitimate one, because C & T spend large tracts of the book developing and defending their theory of language acquisition.

The book divides into three parts of unequal length. The first and longest part presents what C & T call the "Modularity Matching Model", which is their theoretical framework of choice, compares it to competing frameworks, and discusses a number of experimental techniques (such as reaction-time and act-out tasks) that, according to C & T, are less suitable for probing the intricacies of first-language acquisition. Each of the remaining parts is devoted to a method that is argued to be more revealing: part two, the shortest of the three, is about elicited production; and part three, which is nearly as long as the first one, discusses applications of the truth value judgment task.

The Modularity Matching Model is delineated by two fundamental claims. The first, and least controversial, of these is that the language faculty is a separate module within human cognition, as a consequence of which "the construction of syntactic and semantic representations of sentences is not influenced by general cognitive mechanisms — the mechanisms that are used to represent and process real-world knowledge [. . .]"

Linguistics and Philosophy **23:** 523–532, 2000.

(p. 29). The second claim is that “the child’s language-processing system is essentially the same as that of an adult” (p. 30). These two claims pave the way for one of the book’s leitmotifs, which is that well before they reach their fifth year, children have a command of grammar that meets adult standards in most respects. One of the insights I have gained from C & T’s book is that this claim comes a lot closer to being true than I would have expected. The same cannot be said of another theme which recurs again and again, viz. that these feats of language acquisition require a large collection of specific linguistic constraints that are part of Universal Grammar, and therefore innate.

In arguing from linguistic competence to innateness, C & T always employ the well-known argument from the poverty of the stimulus, which in their version goes as follows (p. 19):

(A) All native speakers know some particular aspect of their language, call it property P.

(B) Knowledge of property P could not have been learned on the basis of the primary linguistic data.

So:

(C) Knowledge of property P must be innately specified (i.e., part of Universal Grammar).

This is a precarious argument, not only because it contains a handful of terms which are ill-defined, but also because one of its premises is of negative form: even if it were clear what it means to learn knowledge of a property P on the basis of primary linguistic data, it may be very hard to prove that this is impossible for any given instance of P. Another reason for being wary of this argument is that it is biased towards the view that there is something special about linguistic knowledge, and though it may well be that this is so, we don’t want to prejudge the issue. Of course, the argument from the poverty of the stimulus generalizes to all sorts of knowledge (and beyond), and from time to time it may be useful to remind oneself of this fact, because in the case of language it appears to be more tempting to jump to conclusions than it is in other domains. It is obvious, I take it, that the following argument is no good:

(A) Most people know that there are no three-legged animals.

(B) The knowledge that there are no three-legged animals is acquired in the absence of negative evidence. (Surely none of us have ever observed that there are no three-legged animals, and most of us haven't been told about this, either.)

So:

(C) The knowledge that there are no three-legged animals is innate.

This is a patent howler, but the remarkable thing is that the argument seems to improve if it refers to linguistic knowledge instead of knowledge of the world; several examples of this will be discussed in the following.

Let me begin with what I take to be one of C & T's most spectacular claims, which is that certain aspects of pragmatics are innate. It is a familiar observation that indefinite NPs can, and quantified NPs cannot, function as antecedents for singular pronouns in environments like the following:

(1)a. [A mouse]$_i$ came to Simba's party. He$_i$ wore a hat.

b. ?[No mouse]$_i$ came to Simba's party. He$_i$ wore a hat.

In discourse representation theory and related frameworks, this contrast is explained by assuming that a quantifier like *every* sets up a subordinate context, which is inaccessible to the pronoun, whereas no such context is required by the discourse representation of an indefinite expression. In DRT, (1a, b) are represented by (2a, b), respectively, and the latter is not a "proper" discourse representation, because it contains a free occurrence of x:

(2)a. [x: mouse x, x came to Simba's party, x wore a hat]

b. [: [x: mouse x]⟨ no x⟩[: x came to Simba's party], x wore a hat]

C & T introduce the term "closure constraint" to refer to whatever it is that renders (1b) infelicitous, and although they don't define exactly what the closure contraint is, it may be inferred that they think of it as a syntactic ban on free variables. (Incidentally, for reasons that remain opaque to me C & T prefer to use a rather exotic variety of dynamic semantics instead of DRT or any of its more current relatives; this doesn't make any difference, however.)

Suspecting that the closure constraint is innate, C & T set up an experiment to test this hypothesis. In their experiment, children ranging in age from 3 to 5 were asked to judge if sentences like (1a, b) are true in certain situations. The outcome of the experiment, according to C & T's interpretation, is that a clear majority of their subjects allowed for anaphoric links in cases like (1a), while they disallowed them in cases like (1b). Although I am not entirely convinced that this interpretation is inescapable, I will grant that it is correct, because I want to focus on C & T's contention that their results "support the view that the closure constraint is part of Universal Grammar, that is, of the human biological endowment for language" (p. 284). This conclusion, it seems to me, is precipitate. To begin with, it presupposes a theory of anaphora that is not uncontested. If we treat the pronouns in (1a, b) as E-type anaphors, for example, then the following are approximately equivalent with (1a) and (1b), respectively, and the latter can be dismissed simply because it is contradictory:

(3)a. A mouse came to Simba's party. The mouse that came to Simba's party wore a hat.

b. No mouse came to Simba's party. The mouse that came to Simba's party wore a hat.

Secondly, even if we adopt a DRT framework, there is more than one way of accounting for the contrast between (1a) and (1b). We may or may not follow C & T in supposing that semantic representations with free variables are illicit, but even if we do, it still has to be shown that this is something that cannot be learned. If, on the other hand, we don't want to follow C & T in this, there are indefinitely many ways of interpreting semantic representations containing free variables, and surely some of these will yield readings on which (1b) is either contradictory or odd for other pragmatic reasons.

The problem with C & T's claim is that it is based on an incomplete instantiation of the argument from the poverty of the stimulus. C & T are at pains to show that a certain aspect of linguistic competence manifests itself already at a quite early age. What they fail to establish is that there is no way a child could acquire this knowledge without having been prompted by Mother Nature, so the hard part of the proof is simply left out. This is a recurring pattern in C & T's book.

An intriguing and much-debated puzzle in the acquisition literature is that young children seem to have peculiar difficulties with universal quantifiers, the best-documented of which is that, in a situation like following:

AB AB AB B

there will always be a non-negligible number of children who judge that (4) is false, pointing to the right-most B when asked to explain why.

(4) Every A is paired with a B.

C & T maintain, however, that these findings are an experimental artifact, because earlier studies failed to observe certain pragmatic felicity conditions, which led some children to respond in a way that is out of step with their linguistic competence. In particular, it is crucial, according to C & T, that in a truth value judgment task, both answers be considered:

> In the contexts for yes/no questions, felicitous usage dictates that both the assertion and the negation of a target sentence should be under consideration. (p. 302)

Applied to the artificial example above, this principle demands a scenario in which it *develops* that (4) is true, although it might have turned out the other way, or else children will be puzzled as to why the experimenter wants to know whether or not the sentence is true. C & T report that in an experiment meeting this condition, children performed about as well as adults.

Although I don't wish to dispute that there are all sorts of felicity conditions that constrain speakers' utterances and hearers' interpretations, I am not convinced that the one identified by C & T is among them; and this is a serious worry, because the same condition plays an essential role in several of the experiments presented by C & T. Contrary to what C & T claim, I doubt that a yes/no question is infelicitous unless both the affirmative and the negative answer are "under consideration" in any substantial sense of that term, nor do I see good grounds for maintaining that it is different for kids. When I ask my five-year-old daughter questions like "Am I a dinosaur?", "Do we live in a house?", "Are cars edible?", and so forth, she will produce the right answers without blinking. She may wonder, perhaps, why I should want to ask such questions, but that doesn't prevent her from attaining adult-level performance. So if C & T are right in claiming that previous experiments produced skewed results because they didn't observe their alleged felicity condition, then it seems we may conclude, *pace* C & T, that universal quantification is different for children, after all.

A further complaint about this experiment is that I am at a loss to see what it is supposed to prove. It would have been in line with the rest of the book if C & T had concluded that universal quantification is part of Universal Grammar, but in this case they are less resolute than they are elsewhere: they merely claim to have shown that, as regards universal quantification,

> [. . .] children do not lack grammatical competence. This opens the door for further studies of children's knowledge of universal quantification, and other aspects of quantification. We anticipate that these studies will offer additional results that are in keeping with the precepts of the theory of Universal Grammar. (p. 302)

But what *are* the precepts of Universal Grammar with respect to universal quantification? There are many ways of answering this question. It might be held (i) that the logical concept of universal quantification is innate; or (ii) that the syntactic structures for expressing universal quantification are innate; or (iii) that the syntax-semantics mapping required for the interpretation of universally quantified sentences is innate; or (iv) that several of (i) through (iii) are true; and so forth. But none of these claims seems very promising to me, because I don't see why Universal Grammar should have anything specific to say on the subject of universal quantification (or quantification *tout court*, for that matter). So, as far as I can tell, if C & T's results are in keeping with the precepts of Universal Grammar, it is because Universal Grammar doesn't say anything about quantification. Not that this conclusion comes as surprise, for it is rather obvious that the difficulties children have with universal quantification aren't of a grammatical nature, to begin with. Suppose that the chairman at a meeting asks:

(5) Is everybody present?

Of course, this is not to ask: "Is everybody in this room present?", but rather something like: "Is everybody present who was supposed to come?" On such a construal the universal quantifier is taken to range over a domain of individuals that properly includes the relevant individuals in the immediate context. But this is precisely the type of construal some children favour in the case of (5). It appears, therefore, that these children apply the same kind of pragmatic reasoning as adults do; it is just that children sometimes apply it inappropriately (by adult standards, that is). If this observation is correct, then children's mistakes tell us nothing about their grammars, let alone about Universal Grammar, though they may tell us a lot about the development of pragmatic reasoning.

Much of what has been said in the foregoing carries over to C & T's discussion of the principles B and C of the Chomskyan binding theory, which occupy a prominent place in this book. I will concentrate my attention on the latter principle, but as I will argue later that neither it nor the others can be viewed in isolation, here is the whole set, in one of its more popular versions:

(A) An anaphor must be bound in its governing category.

(B) A pronoun must be free in its governing category.

(C) An R-expression must be free everywhere.

(Let me note in passing that I find this terminology most unfortunate; in particular, I deplore the distinction between anaphors and pronouns, for obvious reasons. But as nobody seems willing to do anything about it, I feel compelled to adopt this jargon, too.) The common opinion in the acquisition literature is that children comply with principle A at a quite early age, but have considerable difficulties mastering principles B and C, which makes it doubtful that these two principles are part of Universal Grammar. C & T challenge this view, just as they challenged the view that children have problems with universal quantification, and they apply more or less the same strategies in both cases. One consistent finding in earlier studies was that children may obtain the same readings for the following sentences, thus violating principle C in the case of (6b):

(6)a. Fred$_i$ thinks he$_i$ is a genius.
b. *He$_i$ thinks Fred$_i$ is a genius.

C & T seek to undermine this conclusion by setting up a series of experiments of their own, designed to show that if things are arranged the right way, children too will reject a statement like (6b), and not because they think it is false or because they fail to obtain the intended (though illicit) reading, but because it is ungrammatical. I recommend C & T's discussion of these experiments to anyone who feels tempted think lightly of experimental research in general and acquisition research in particular, as in these chapters C & T make it abundantly clear how difficult it is to construct a good experiment, demonstrating at the same time that they have the ingenuity and creativity that it takes. But, again, it is when they draw the moral that C & T overplay their hand:

> Presumably, there is nothing in children's experience to tell them that certain sentence/meaning pairs are *not* allowed; therefore, they have no way to learn the structural constraint prohibiting coreference, Principle C. Nevertheless, they appear to know it at an early age. By the logic underlying the poverty-of-the-stimulus argument, we are led to conclude that children's knowledge of Principle C is innately specified. (p. 220)

This conclusion may be challenged on a number of counts. For example, there are well-known reasons for doubting that principle C is correct as it stands, and it has been suggested, plausibly in my view, that to the extent that it is correct it is not a grammatical principle at all, but follows from pragmatic considerations. However, let us grant that not only principle C but all of the binding theory is true, and let us suppose, contrary to fact or not, that the experimental evidence demonstrates that children master these principles already at a very tender age. The point I want to make is that, *even then*, we aren't entitled to conclude that the binding theory is part of Universal Grammar.

I have shifted my focus from principle C to the entire binding theory, because it is easier to see how the principles of the binding theory might be acquired if we consider them collectively. These principles are defined in terms of the notions of governing category and binding. Now let us ask first if a child who already has these two notions, and is aware that they may be implicated in the definition of linguistic constraints, would be in a position to discover these three principles. I submit that the answer to this question is yes, and here, in a nutshell, is how. In view of their obviously restricted distribution, it would seem a good idea to start with what Chomsky has dubbed "anaphors", i.e. reflexive pronouns and their kin. Once this category has been isolated, it is surely not a random guess to conjecture that it is subject to syntactic constraints, nor is it far-fetched to surmise that these constraints, whatever they may be, are defined in terms of relations between a given expression and its syntactic environment. Given, furthermore, that the child is already aware that the notions of governing category and binding are potentially relevant (which holds *ex hypothesi*), the child has now reached a stage at which something like principle A suggests itself rather strongly (and even if it doesn't suggest itself right away, a couple of months should suffice, even if the child's priorities lie elsewhere).

Once the child entertains the hypothesis that principle A, or something like, might hold, it has secured a foothold for tackling other devices for expressing coreference, i.e. the Chomskyan "pronouns" and "R-expressions". For his working hypothesis establishes a relation between anaphoric expressions, on the one hand, and the notions of governing category and binding, on the other. It is natural to suppose, therefore, that the same notions might constrain the interpretation of pronouns and R-expressions. Observing, furthermore, that anaphors and pronouns are in complementary distribution, it is only a small step to hypothesize that principle B, or something like, might hold.

With principles A and B in place, if only as working hypotheses, principle C cannot be far. Let α and β be two coreferential expressions such that α binds β. If α is in β's governing category, then β may be an anaphor but not a pronoun; if not then β may be a pronoun but not an anaphor. This much follows from the principles A and B, and if this much is assumed, it is not such an outlandish idea that β *must* be either an anaphor or a pronoun: R-expressions aren't supposed to be bound, because that is what anaphors and pronouns are "for". This is an inferential leap, to be sure, but it is a natural one, and it brings us, or rather our child, to principle C.

Thus far I have assumed that, before it attends to the vagaries of anaphors, pronouns, and R-expressions, the child already masters the concepts of governing category and binding. Are these part of Universal Grammar, then? I don't think so – or, more accurately: I don't believe that this claim is unavoidable. The crucial ingredient in the notion of binding is c-command, and *that* is one of (at most) a handful of concepts that emerge naturally when one starts building trees. The concept of governing category is more problematic, not the least because there is less agreement about how precisely it is to be defined. But no matter what the outcome of this debate will be, the hope is, surely, that the right concept of governing category will be a natural one, and the more this hope turns out to be justified, the more likely it will be that it can be learned, just as in the case of c-command. Therefore, I don't see how a convincing case can be made for claiming that the concepts of binding and governing category are part of our biological endowment.

My story about how the binding theory is acquired is a plausible one, I hope, but it is no more than that. Indeed, the last few paragraphs are replete with such phrases as "this makes it natural to assume that", "this suggests that", and so on. This is unobjectionable, however. First, it should always be borne in mind that learning language is an inductive enterprise, and, like it or not, there is no such thing as inductive certainty. Secondly, the purpose of my story was to demonstrate the fallacy of C & T's claim that principle C is innate, and in order to *prove* that their argument from the poverty stimulus is flawed, a *plausible* story is all that is needed.

Before concluding this piece I must stress that I have been rather unfair to C & T, because in my remarks I have concentrated on matters that aren't their top priority (which is definitely not to say, however, that these issues are of minor importance to them, for the contrary is true). As I noted at the outset, the main purpose of this book is to demonstrate and explain how experiments in language acquisition should be constructed, and if I haven't given this aspect its due, it is because I take it to be of secondary interest to the readership of this journal. But it is precisely in this respect that, in my opinion, C & T succeed admirably. Through a formidable series of case studies C & T impress on their readers how difficult it is to set up an experiment in such a way that it tests what you want it to test, and nothing else, and one can hardly fail to be awed by the ingenuity, dedication, and patience with which C & T carry through their investigations. In this respect, and it really is the main one, the book is an important achievement.

Department of Philosophy
University of Nijmegen
Postbox 9103, 6500 HD Nijmegen
The Netherlands
bart.geurts@phil.kun.nl

BART GEURTS