

WORKING PAPERS
IN
LINGUISTICS

The notes and articles in this series are progress reports on work being carried on by students and faculty in the Department. Because these papers are not finished products, readers are asked not to cite from them without noting their preliminary nature. The authors welcome any comments and suggestions that readers might offer.

Volume 38(6)

2007
(October)

DEPARTMENT OF LINGUISTICS
UNIVERSITY OF HAWAI'I AT MÂNOA
HONOLULU 96822

An Equal Opportunity/Affirmative Action Institution

DEPARTMENT OF LINGUISTICS FACULTY

2007

Victoria B. Anderson
Byron W. Bender (Emeritus)
Benjamin Bergen
Derek Bickerton (Emeritus)
Robert A. Blust (Chair)
Robert L. Cheng (Adjunct)
Kenneth W. Cook (Adjunct)
Kamil Deen (Co-Graduate Chair)
Patricia J. Donegan
Emanuel J. Drechsel (Adjunct)
Michael L. Forman
George W. Grace (Emeritus)
John H. Haig (Adjunct)
Roderick A. Jacobs (Emeritus)
Paul Lassetre
P. Gregory Lee
Patricia A. Lee
Howard P. McKaughan (Emeritus)
William O'Grady
Yuko Otsuka
Ann Marie Peters (Emeritus, Co-Graduate Chair)
Kenneth L. Rehg
Lawrence A. Reid (Emeritus)
Amy J. Schafer
Albert J. Schütz, (Emeritus, Editor)
Ho Min Sohn (Adjunct)
David L. Stampe
Laurence C. Thompson (Emeritus)
Andrew Wong

THE BIOLOGICAL ENDOWMENT FOR LANGUAGE AND ARGUMENTS FROM THE POVERTY OF THE STIMULUS

HUNTER HATFIELD

There is widespread agreement among linguists of almost all theoretical leanings that our biology constrains the types of languages that can be acquired, but the content of that biological endowment has been the source of continual and fierce debate for many years, a debate usually framed as a choice between nativism and non-nativism. This paper attempts to find a way through this debate by exploring the methodologies available to researchers for answering the biological question. In particular, it assesses whether or not Poverty of the Stimulus Arguments (Poverty Arguments) are particularly useful for this endeavor. It is argued that the nature of the biological endowment for language is significantly underdetermined by the evidence from Poverty Arguments. This is demonstrated both as a philosophical point and with a critique of two recent Poverty Arguments. However, we can make significant progress in the future by looking for causal explanations of language knowledge, and not just structural descriptions of that knowledge.

1. INTRODUCTION. This paper attempts to walk a very difficult line rhetorically. On the one hand, it will discuss significant limitations on Poverty of the Stimulus Arguments (henceforth Poverty Arguments), and yet it tries to do so without entering the raging debate about whether there is innate linguistic knowledge or not. Indeed, it assumes a very general nativist position, namely that linguistic knowledge depends necessarily upon the biological endowment that humans share. This starting point is taken for granted and will remain unargued. The next question then becomes: how does the human biological endowment enable language use and language competence? Let's assume that human biology is critical to understanding possible human languages, and now ask about the details.

When looking at language acquisition research, however, one quickly discovers that what scientists conceive of as a biological approach to language can vary enormously. One vein of research that strives to discover the biological endowment for language is the classic "nativist" position that derives from the work of Chomsky. Chomsky has long characterized the resources that humans bring to the linguistic table as a "language organ" in direct parallel to other organs such as the heart and lungs or the visual system. Chomsky specifically refers to the generative nativist approach as "the biolinguistic point of view" (2002:86). Such terminology is also not restricted to a few comments here and there. Indeed, it is a common thread in this tradition of research: Wexler 1996, which concerns the acquisition of verbal inflection, is explicitly entitled "The development of inflection in a biologically based theory of language acquisition"; Crain, Gualmini, and Pietroski (forthcoming:2) state that the Minimalist Program is "an invitation for nativists to ask just what aspects of language must be attributed to biology"; and Anderson and Lightfoot 2002 attempts to define the language organ.

This tradition of research is of course quite multifaceted, and this paper will make no comprehensive attempt to assess its progress towards the goal of ascertaining the biological endowment for language. Instead, here we will focus on one particular piece of this tradition, which is the belief that Poverty Arguments are a particularly fruitful way of discovering how the biology of humans enables language. Anderson and Lightfoot, for instance, argue that Poverty Arguments should form the basis of a biologically informed model of language (2006:378). The conclusion in this paper is that Poverty Arguments, even when methodologically unassailable, make only moderate progress towards discovering the biological endowment for language, but that it is possible to redirect this research in a more fruitful direction by focusing more directly on the precise causes of linguistic behavior.

To show that this is the case, we will first take a look at the form for Poverty Arguments as exemplified in Lidz, Waxman, and Freedman 2003 and Kiguchi and Thornton 2004. As the purpose of this paper is to discover what methodologically sound Poverty papers can and cannot tell us, we will accept all data from these experiments as showing what they purport to show about language use and experience, and make no attempt to critically assess the papers from that perspective. This will be followed by a discussion of the nature of underdetermination in the philosophy of science. Next, the

paper will present an alternate theory of linguistic knowledge to add meat to the bones of the underdetermination argument. After the paper demonstrates this under-determination, it will return to the broader issues related to discovering the biological endowment and how under-determination can be overcome in future research, either in the generative paradigm or in some alternate approach. The hope is to find a way forward and escape many of the burden-of-proof arguments that have become increasingly common in the debates over linguistic nativism. (See, for instance, Akhtar, Callanan, Pullum, and Scholz 2004, Crain and Pietroski 2001, Cowie 1999, Everett 2006, and Scholz and Pullum 2005, which often devolve into who has to prove what to whom in the face of ignorance, a cardinal sign that research is making limited progress in resolving a debate.)

Poverty Arguments have such a long and contentious history in the literature that it is worth taking the time to state some typical questions that will not be addressed here. This paper will not address whether there is a language organ, whether language is domain-specific, or whether there is a Universal Grammar (UG). The focus is instead on methodology. If we wish to know what aspects of our biology enable language (possibly UG, possibly not; this is what we are trying to discover), is a strong Poverty Argument the right sort of model for discovering the answer? The answer given here is “no,” due to the multiple possible answers that are compatible with the data a Poverty Argument provides. This answer is similar to an argument given by Everett in his recent debate with Anderson and Lightfoot about “the language organ.” Namely, he states “... the burden of proof lies with the proponents of a UG/Language Organ, not with their critics. The reason is simple: in studies of human cognition, there are many general principles that could be explored as alternative explanations for the sources of human language” (2006:386). Everett is also asserting a form of underdetermination, but his focus is slightly different than this essay's. He is arguing for a default position of “non-nativism” and asserting that the nativist must meet all objections before a UG is established. Underdetermination, according to Everett, must be overcome to establish a proposition. This essay skips this debate to a large degree and focuses instead on the precise details of what the biological endowment for language is.

2. POVERTY ARGUMENTS AND UNDER-DETERMINATION IN THE PHILOSOPHY OF SCIENCE.

2.1 THE STRUCTURE OF POVERTY ARGUMENTS AND TWO EXEMPLARS. As Chomsky has emphasized by using the term “Plato’s Problem,” alluding to Plato’s dialogue “Meno,” it has long been recognized that people appear to possess knowledge that is not obviously available from their empirical input. The explorations of what humans bring to their experience has continued ever since, resolving into a bewildering mix of positions from, in the Western tradition, Locke’s empiricism and the Blank Slate approach to Kant's arguments for the structuring of perception with innate human categories to certain forms of Idealism in which innate categories overwhelm all experience. In the linguistic context, the modern debate has become a discussion of precisely how linguistic knowledge is constrained by our biology. Is it constrained by learning mechanisms that are perhaps specific to language (e.g., Newport and Aslin 2004) or general neural architecture not necessarily specific to language (e.g., Elman et al. 1996), or by learning abilities that are shared across many domains (Tomasello 2003), or by a large number of constraints and biases that are not particularly linguistic (MacWhinney 2004), or perhaps by a large number of constraints that are explicitly linguistic in nature (Government and Binding)?

This last tradition has long argued that Poverty Arguments are particularly helpful in establishing that at least some linguistic knowledge can only be accounted for by a language-specific biological endowment. Moreover, Poverty Arguments are supposed to show precisely what (a portion of) that endowment is. All nativists in this tradition believe that some components of language are learned and that nonlinguistic faculties such as attention, memory, auditory processing, etc., are necessary for reaching the adult stage of language performance. However, on top of all that are biological capacities for language that concern only language, and supposedly Poverty Arguments can reveal those items. Since that is the claim, we will now look at the form of a Poverty Argument.

The rough outline of a Poverty Argument has been fairly consistent for several decades now. The first step is (1) to establish that children possess some sort of linguistic knowledge. Second, the researchers attempt (2) to demonstrate that there is no input available to a child that is sufficient for learning that knowledge. The conclusion is then that, since the child possesses knowledge X but the

input cannot teach X, then X must be part of the biological endowment for language. The methods of ascertaining whether or not premises (1) and (2) are true have varied over time.

Both Lidz et al. 2003 and Kiguchi and Thornton 2004 use psycholinguistic methodology to assess premise (1), namely that children possess particular linguistic knowledge. Lidz et al. 2003 assesses knowledge of anaphoric *one* with infants of 18 months old. In particular, they wish to know if the children interpret anaphoric *one* as referring to just the head noun, N^0 in classic X-bar phrase structure, or to the larger phrasal unit, N^1 , which includes the noun with its modifiers exclusive of the determiner. To make this determination, infants were first habituated to an object, such as a yellow bottle, on one side of their location and heard a sentence such as *Look! A yellow bottle!* After this familiarization, in the test phase, infants were then presented with a blue bottle and a yellow bottle. Note that both images are compatible with an interpretation of anaphoric *one* that refers to the head noun; i.e., both images are bottles. However, only the yellow bottle matches the entire N^1 phrase (and image) from the familiarization stage. In the control condition, the children heard a sentence that simply asked what they saw. Infants looked significantly more at the novel image, the blue bottle, than the old one. But in the critical test condition in which infants heard *Now look. Do you see another one?*, they looked more at the old image, the yellow bottle. Lidz et al. conclude that “infants at the earliest stages of syntactic production share with mature speakers of English the syntactic knowledge that the anaphoric element *one* refers to the category N^{1*} ” (2003:B72).

Kiguchi and Thornton (2004) are looking at a far more complicated linguistic structure, Antecedent-Contained Deletion (ACD). ACD is a form of verb ellipsis in which the ellipsis occurs inside a relative clause. An example that we will study in great detail later is *Dora baked him the same food that Cookie Monster did*. These are peculiar structures, because if you try to unpack the ellipsis, you can get caught in an infinite loop, as the elided VP is inside of the VP to which it refers. If *did* refers to the verb phrase *baked him the same food that ...*, you end up with: *Dora baked him the same food that Cookie Monster [baked him the same food that cookie monster [baked him the same food that Cooke Monster [...]]]*. If *did* stands for the entire phrase, then it loops forever. Kiguchi and Thornton review the large body of research in the generative tradition that has been developed to handle these issues, which is beyond the scope of this paper. Here, we will focus on their first experiment, which provided evidence of Binding Principle B. The critical sentences with the ellipsis unpacked without the loop are:

1A. Dora baked him_i the same food that Cookie Monster_{*i/k} baked him_i.

1B. Dora baked him_i the same food that Cookie Monster's_i father_k baked him_i.

Notice in the first sentence that the final *him* does not refer to Cookie Monster. It can only refer to some other party, such as the earlier *him* in the matrix clause. However, in the second sentence, the final *him* can refer to Cookie Monster. Kiguchi and Thornton attribute these differences to Binding Principle B in which a pronoun such as *him* must be free in its governing category. (See Haegeman 1994 for an introduction to binding theory in the Government and Binding tradition.) The problem in the first sentence is that *him* is c-commanded by a co-indexed governor, *Cookie Monster*, which is a violation of Principle B. In the second sentence, since *Cookie Monster* is in a possessive construction with *father*, *him* is not bound by a co-indexed c-commanding governor, which makes the sentence acceptable.

The above paragraph gives the adult acceptance of these two sentences and the generative explanation for those judgments, but Kiguchi and Thornton now must determine whether children make these same grammaticality judgments. Due to the complexity of the structure and the attentional demands of their experimental design, the researchers work with children ages 4 to 5. In a wonderfully designed experiment, the children are given the task of judging whether or not one of the target sentences is true or false based upon a story they have watched, told with puppets. Kiguchi and Thornton state that many children have a bias towards agreeing, and so the stories are designed such that children should only choose “no” if they have the adult interpretation of the sentence. The interpretation that is in violation of Principle B, i.e., in which the final pronoun is bound in its governing category, is true in the story, which means that the children can only rely on their understanding of the critical

sentence to make a “false” judgment. In this experiment, the children overwhelmingly chose interpretations that matched the adult interpretations of the sentences. Kiguchi and Thornton conclude that the children have knowledge of Principle B.

At this point, we have discussed how Lidz et al. and Kiguchi and Thornton establish premise (1) of the Poverty Argument. Now, they need to establish that this knowledge cannot be learned from the input available to the child. Lidz et al. are concerned with sentences that unambiguously rule out the N^0 interpretation of *one* in favor of the N^1 interpretation. They examine the speech to Adam and Nina in the Brown corpus and find 792 uses of anaphoric *one*. However, of all these tokens, only two unambiguously rule out N^0 in favor of N^1 . They conclude that there is insufficient evidence, particularly before 18 months of age, with which children might learn the adult interpretation of anaphoric *one*. Kiguchi and Thornton's arguments for the poverty of the input are two-fold. First, the Binding Principle B violation (*him* cannot be bound in its governing category) occurs only within the elided VP, so strictly speaking there is no explicit input at all, since the pronoun never makes it to phonetic form. They follow this general argument with a look at the speech directed towards the child, Adam, in the Brown corpus and find only a single ACD construction. They assume that such a frequency is insufficient to teach the complexities of ACD structure.

Since Lidz et al. and Kiguchi and Thornton have now demonstrated both premises of our Poverty Argument, they conclude that the knowledge they found is innate. As Kiguchi and Thornton (2004:267) put it, “These data provide an argument from the poverty of the stimulus for a rich innate binding theory in children's grammars.” Lidz et al. (2003:B72) state that “This syntactic knowledge [phrase structure in the form of N^1] must derive from linguistic structure inherent in the learners themselves ...” When we combine these statements with others in the tradition who view this approach as biolinguistics and as revealing the nature of the biological language organ, the natural interpretation of such statements is that these Poverty Arguments are purported to reveal that noun phrase structure and Binding Principle B are part of the biological endowment of humans. The follow-up question is, of course, does this argument work? Have Poverty Arguments revealed components of our human biological endowment?

2.2 UNDER-DETERMINATION. As stated in the Introduction, this paper argues that Poverty Arguments do not go as far as their proponents hoped. This is because the explanations offered for the results of Poverty research are significantly underdetermined by the evidence. If we are interested not simply in establishing that something biological undergirds language knowledge and language use, but instead, precisely how biology constrains language, then we should not be happy with this situation. This rough point is not new. In a review of Crain and Thornton's *Investigations in Universal Grammar*, Geurts (2000:531) argues that as long as there is any plausible alternative to Universal Grammar, UG is not proven. Crain and Pietroski (2001) reply, indirectly, by saying that the Geurts-type argument is just a logical point and that without any substantive alternative, the generative framework is the best choice for conducting productive research today. This is exactly the sort of burden-of-proof debate that this paper is trying to escape. This can be accomplished because not all forms of underdetermination are made equal, so to speak.

One form of underdetermination is essentially classical skepticism, in which it is argued that we can never be certain of our experiences or the generalizations we draw from them. While such skepticism can be useful in certain philosophical discussions, it is not helpful in designing research, and so that version of underdetermination will be put aside. However, underdetermination is indeed a more significant concern in the philosophy of science than in general life. Very few people actually worry whether or not their theory of the existence of chairs is underdetermined by the empirical data before they sit down in one. The items of scientific theory such as gravity, quarks, consciousness, or phrase structure, though, are far more reasonably doubted, and scientific history is littered with forces and objects that seemed wholly justified at one point by the evidence but do not appear to actually exist, ether being an oft-cited example. (The reader is referred to Papineau 1996 for a collection of classic articles arguing for and against underdetermination for scientific theories.)

If underdetermination is a legitimate problem for scientific theories, then to handle it, we need a strong account of what scientific explanation is. (The following discussion is largely based upon Kosso 1992, particularly pp. 52–65.) The classic model of scientific explanation comes from the

work of Carl Hempel and is termed the Covering-Law model. In this model, an event is explained if it is covered by a Law. A form for a successful explanation can be derived from this basic concept: if Law L is true and Conditions C1, C2, C3, ... hold, then Event E is true. Successful scientific explanations do fit this mold. The problem is that many poor explanations fit the model as well. For instance, it is possible to fit the explanation for the gravitational force between Earth and the moon due to each one's mass into the frame. However, one can equally "explain" the mass of the Earth by gravity. However, if that is the case, then something is wrong, because we want to say that the mass of the Earth explains its gravitational pull, but not that the Earth's mass is explained by its gravitation. As Kosso expresses it, the problem is that the Covering-Law model is symmetric between a Law and its effects, while genuine explanation is not.

What we seem to want, in addition to covering, is some notion of causation. The earth has a certain gravitational pull because its mass creates and causes this force. Mass has gravity but gravity doesn't have mass. While an appeal to causation to get rid of spurious correlations, such as gravitation explaining mass, is intriguing, it runs into the immediate problem that we now have to define causation, which has been a source of philosophical trouble since Aristotle. How do we know that mass causes gravitation and not the other way around? Kosso states: "The decision is founded on our theoretical background, that is, on our theories about the mechanics of gravitation. Recognition of a causal explanation, in other words, can only be done in a theoretical setting" (1992:64).

Kosso's point is well-taken, but it is not clear that the theory-bound nature of causal explanation should be particularly worrisome for linguists attempting to define a research program. As a first point, most linguists assume that what they are trying to discover are causal explanations. We might be well aware that for some particular research result we've only established a correlation and not cause, but that does not stop us from keeping causation as the goal. Second, while there are as many linguistic theories as there are linguists, most linguists do share a very general theoretical background. For instance, most linguists believe that linguistic knowledge in humans is somehow instantiated in our body, with the nervous system and brain being particularly relevant. Most linguists accept evolution as the general biological background for human behavior. Most linguists accept that linguistics is in some way a mentalistic endeavor such that we are exploring the psychology of speakers when we study language. This is not to say that all linguists believe that all of these avenues of research are particularly fruitful at the current moment. Chomsky 2002, for instance, argues quite forcefully that the divide between knowledge of the brain and knowledge of language right now is so far apart that it is best to proceed with a focus on pure language. At a later point, when brain research has progressed sufficiently, we can try to connect the two. Chomsky's point is well taken, but even it occurs with the background assumption that there is some connection between the brain and the language faculty. What this all amounts to is that it remains practical for linguists to pursue genuine causal explanations of language as a research program today, since we do share much of the theoretical background that Kosso mentions is required for causal judgments.

With this said, we can now finally state the major claim of this paper: The biological endowment for language is underdetermined by Poverty Arguments in a nontrivial way precisely because the explanations being offered are not causal explanations. Although they are descriptions of *what* someone knows, they continue to ignore *how* they know it, and it is this question of *how* that is the fascinating biological question. How does our biology allow us to make grammatical judgments?

This point will be developed in greater detail in the Discussion section. However, before we can take the point too seriously, it is necessary to show that the critical Poverty experiments are indeed underdetermined in an interesting way. It's one thing to simply assert underdetermination, but it's far less interesting if we can only come up with absurd or trivial alternatives. As Crain and Pietroski argued, underdetermination is only a logical point unless the alternative is a genuine one for guiding scientific research. Therefore, the next section introduces a plausible alternative to knowledge of phrase structure and binding theory that can account for the Poverty data. The section's primary purpose is to serve as a proof of underdetermination, and so readers who are already convinced of underdetermination may find it less useful and can skip to section 4; however, those of us who require more proof should proceed onwards.

3. ALTERNATIVE ACCOUNTS. This section functions as a proof of the assertion that there are nontrivial alternate explanations for the behavioral data discovered in the Poverty Arguments from Lidz et al. 2003 and Kiguchi and Thornton 2004. The framework developed here is rooted in the fundamental notion, since at least Lashley 1951, that producing and understanding language involves mapping between a hierarchical representation of meaning and a linear sequence of symbols. From this perspective, the current framework is in the spirit of the Minimalist Program (Chomsky 1995), in which syntax operates between Logical Form and Phonetic Form, though the details are quite different. The goal is to show that Poverty Arguments are not particularly useful tools in discovering precisely what the biological endowment for language is. The goal is *not* to show either (1) that biology does not affect language, which is trivially false, or (2) that there are no language-specific components to our biological endowment. Knowing whether or not (2) is the case depends crucially upon whether or not we have a methodology that successfully makes progress in discovering our biology. If we have no such methodology, we cannot even start to answer (2) productively. In the rest of this section, we will see that there are actually several different possible explanations for the behavior of the children in the experiments, and that it is unclear whether or not each of those items is learned. As we encounter such forks in the road, small tags of the form (F#, such as F1) will be placed so that we can keep track of the various alternatives and refer to them in the Discussion.

The framework adopted here can be termed Semantic Carpentry, because it is a heavily modified version of O’Grady’s Syntactic Carpentry (2005). In O’Grady’s theory, syntax reflects the operations of an efficiency-driven processor that must satisfy the lexical requirements of the words passed to it as it builds a sentence. “Efficiency-driven” means that the processor must satisfy requirements at the first opportunity, so as to minimize the burden on working memory. This simple feature produces much of the sentence structure that traditional grammars try to explain; however, in the dominant metaphor of the theory, there are no blueprints for what sentences should be like; there are only carpenters who build.

The syntactic carpentry theory explores the relations between a sentence processor and lexical requirements, but it does not explicitly address semantics, other than in passing. This is the critical difference between the semantic carpentry offered here and syntactic carpentry. Semantic carpentry utilizes an efficiency-driven processor to transform a hierarchical semantic structure into a linear string of words, each of which has its own lexical requirements that must be satisfied. In production, the processor moves from semantic structure to a string of words, and, in comprehension, it moves in the reverse order such that a string of words arriving over time is transformed into a semantic structure. Due to the nature of the Poverty experiments under consideration here, this paper will focus almost exclusively upon comprehension. It is assumed that comprehension is the successful construction of the semantic structure of a sentence, and the semantic structure to be employed is the generative approach of Jackendoff 2002. Jackendoff 2002 and O’Grady 2005 are selected simply as plausible accounts in order to function as a proof here. Readers interested in the justification for these theories should consult the original works.

Jackendoff’s (2002) approach employs several layers of structure for any sentence, including discourse and information structure, semantic structure, syntactic structure, and phonetic structure, among others. The critical item for us is the semantic structure as represented in two tiers, the Descriptive Tier (DT) and the Referential Tier (RT). The DT represents “the organization of conceptual functions, arguments, and modifiers” (2002:394). The RT handles implications for what entities and states are implied to exist by the sentence. Each state and object will be represented on the RT, but not all also have existential force, meaning that not all items are asserted to exist in the world. Negations, for instance, do not have existential force. This will be a crucial fact for the later assessment of Kiguchi and Thornton 2004. Very likely the best way to follow this discussion is to go through examples, and so that is where we turn now by looking first at the experiment in Lidz et al. and observing what it is like to hear the linguistic stimuli in this proposed framework.

3.1 LIDZ ET AL. 2003. The experiment begins in the familiarization phase with a child seeing a yellow bottle and hearing the phrase, *Look! A yellow bottle!*. It is unfortunate for the purposes of explication that an imperative and a sentence fragment are the first examples offered to us. Here we will con-

centrate on just the phrase *a yellow bottle* and additionally add a verb, *to be*, as in *there is a yellow bottle*, for the purpose of demonstration. The verb plays no role in the argument.

According to the semantic formalism of Jackendoff 2002, comprehending this sentence will involve constructing a Descriptive Tier and Referential Tier representation of the input. The completed DT looks like example 2.

2. Descriptive Tier

$$\left[\text{STATE BE} \quad \left(\left[\text{OBJ} \quad \begin{array}{l} \text{TYPE BOTTLE} \\ \text{INDEF} \\ \text{[PROP YELLOW]} \end{array} \right] \right) \right]_2$$

This formalism says that the State exists in which there is an Object of Type bottle and this Object has the Property of yellow. Modifiers are simply listed under the Object Type. The words *bottle* and *yellow* in the notation stand for the actual concepts, not the words as such. The RT for this sentence would simply assert the existence of the State and the Object.

3. Referential Tier



The numbers on the RT correspond to indices on the DT, so this notation says that the existence of the State implies the existence of the Objects within that State. Now, let's see how a semantic carpentry processor would construct these tiers in the process of language comprehension.

During familiarization in Lidz et al., the child first hears the article *a*. The parser must satisfy as soon as possible the lexical requirements for the word. An indefinite article takes a noun argument to its right, but so far this requirement cannot be met. As for mapping this word into the DT, we have our first fork in the road. We might attempt to immediately create an Object on the DT but leave the Type undefined. This intermediate DT would look like 2¹.

2¹. DT after hearing the word *a*.

$$\left[\text{STATE BE} \quad \left(\left[\text{OBJ} \quad \begin{array}{l} \text{TYPE} \\ \text{INDEF} \end{array} \right] \right) \right]_2$$

Such an account would be plausible with research that indicates that semantic interpretation is immediate during sentence processing. An alternative, however, is to argue that only Noun Phrases map to Objects. Jackendoff (2002:192) suggests that the mapping between NP and Objects is a prime candidate for Universal Grammar, and, if we accept that, then our parser has not yet built an NP and so cannot map to the Object in semantics.

The next word the child hears is *yellow*, which functions quite similarly to an article and requires a noun to its right. No noun has been heard yet, and so the NP remains unconstructed. If the DT is built immediately without waiting for a noun to head the NP, a property of being yellow would now be created in the DT with the Object remaining undefined. How could a child know to do this? There are several possibilities: (F1-1) Part of the lexical meaning of the word *yellow* is that it is a Property; knowing the meaning of the word includes knowing that it is a Property and not, say, an Object. (F2-1) A statistical regularity is noticed between adjectives and Properties. (F3-1) A mapping between properties and adjectives might be part of UG. (F1-1) and (F2-1) differ in that, with the former, *yellow* is mapped directly to Property, while, with the latter, *yellow* is mapped directly to adjective only, and then adjectives are mapped to Properties. Only (F3-1) requires adding content to UG. Also, none of them are needed if a Property does not get mapped to the DT until the NP is completed, as the steps in handling NPs themselves could subsume those for handling adjectives within the NP.

Finally, the word *bottle* becomes available. Now a noun is present to Resolve the requirements of the previous adjective, which the parser does immediately by copying the index of the noun to the adjectives. It can also now fill in the Object Type in the DT as well as copy the index of the syntactic NP onto the Object. The possible ways in which it might do this are very similar to (F1-1) through (F3-1) above, namely: (F1-2) part of the lexical meaning of *bottle* is that it is an Object; (F2-2) Knowledge of a statistical regularity between nouns and Objects; and (F3-2) NPs are mapped to Objects according to UG. (F3-2) is Jackendoff’s proposal. Once the mapping to an entire Object is done, the properties are immediately tied to the object Type as both are part of the complete Object. In other words, the structure of the DT itself bundles Properties with Object Types. Of course, one might ask how the entire DT structure is acquired: (F2-3) There is a statistical regularity between objects and properties. To put it lightly, objects typically come with properties and so the evidence is ubiquitous. (F3-3) Properties are part of Objects in UG. This last solution is close to an innate Whole Object bias.

It’s worth noting the relationship between O’Grady’s syntactic carpentry and the semantic carpentry here. In the former, lexical requirements are satisfied by resolving dependencies at the first opportunity. Semantic carpentry extends this so that an NP requires being mapped to the first available semantic Object, which is also performed by copying the index from an NP to an object. Similarly, all states, events, and objects on the DT would require mapping to the RT, and so the indices are copied to that tier at the first opportunity as well.

If we take a look back to our goal of explaining how the child knows in the test utterance *Now look. Do you see another one?* that anaphoric *one* refers to a yellow bottle and not just a bottle, we can see some progress already by simply looking at the DT. If *one* is mapped back to the index 1 on the descriptive tier, then it will be mapped to the entire Object with all of its properties and not just to the Type identification. This will result in *one* referring to a yellow bottle and not just a bottle. Jackendoff (2002) happens to examine anaphoric relations, including *one*, and the next paragraph outlines his discussion.

He first looks at anaphoric *it*.

4A. Syntax/Phonology: [[Joan]₁ bought [a car]₂]₃. [[Fred]₄ liked [it]₅]₆.

4B. DT: [BUY (JOAN₁, CAR₂)]₃ [LIKE (FRED₄, [3SingNeut]₅)]₆

4C. RT: 1 2 3 4 5=2 6

(Jackendoff 2002:397)

Note how the indices are copied from tier to tier with the critical item being the identification of indices 5 and 2 on the RT. So the anaphoric properties of *it* involve referential identification; however, this is not how *one* behaves.

5A. Syntax/Phonology: [[Joan]₁ bought [a car]₂]₃. [[Fred]₄ bought [one]₅ too]₆.

5B. DT: [BUY (JOAN₁, CAR₂)]₃ [BUY (FRED₄, [3Sing]₅₌₂)]₆

5C. RT: 1 2 3 4 5 6

(Jackendoff 2002:397)

In this example, the unification of indices 5 and 2 (which are *one* and Joan’s car, respectively) occurs on the descriptive tier, not the referential tier. This is because it is understood that Fred’s car is descriptively the same as Joan’s, but it is a different car. *One* does not mark out the exact same entity, as *it* does, but a separate entity that shares a description with *another*. Jackendoff allows for an interpretation in which *one* is co-indexed on both the DT and the RT, but the addition of *another* in *another one* explicitly rules out the referential interpretation. All this can then be taken to imply that *another one* has a requirement that it be co-indexed with an entity on the DT, but not the RT.

Returning to our 18-month-olds listening to the phrase *see another one* (assuming that this is the critical part of the sentence they hear), the first word they will hear is *see*. Following O’Grady 2005, this verb will require two noun arguments—one to the left and one to the right. The deictic *you* is already available, and so the index of that noun is immediately copied into the verb’s matrix. The rightward argument cannot be satisfied yet; however an Event of Type See, as well as its first object argument, can already be created in the descriptive tier.

Next, *another one* arrives. Because Jackendoff treats *another one* as a single unit, I will as well, but it could easily be changed so that *another* is a modifier of the simple pronoun *one*. The processor can now Resolve the verb's lexical requirements by copying the index of *another one* to complete the verb's matrix. It also must satisfy the semantic requirements of the NP as a unit and the specific noun in play, which is the anaphor *another one*. The NP requires that its index be mapped to the DT as an Object. The anaphor *another one* also then requires that its index be united with a previous Object, which is the earlier yellow bottle, on the DT, not the RT.

How does a child know that anaphoric *one* requires resolution on the DT but not the RT? The only evidence that is required is that *one* be used in a context where it refers to a descriptively similar but referentially separable entity, and this is the default use of *one*. Any case such as *here's an apple for me and one for you* would provide this evidence. This is very different from the type of evidence that Lidz et al. found wanting in their corpus study. They needed to find examples of sentences where the N⁰ interpretation was explicitly ruled out, such as one person holding a red ball and another holding a blue ball. We only need find examples that *one* is a descriptive anaphor, not a referential one. The processor's design and the structure of semantics take care of the rest.

Finally, re-examining our forks along the road, one path, F3, posited specific content to universal grammar (adjectives to Properties, NPs to Objects, or Properties to Objects). Path F1 posited lexical information, while path F2 proposed regularities between word classes. Most importantly, the Lidz et al. experiment provided equally confirming evidence for them all, and the only way to know which is accurate is to rely upon evidence outside of the Poverty argument.

3.2 KIGUCHI AND THORNTON 2004. We turn now to Kiguchi and Thornton's examination of Antecedent-Contained Deletion (ACD) structures, with sentences such as (1A) and (1B) repeated below as (6A) and (6B) for convenience.

6A. Dora baked him_i the same food that Cookie Monster $_{i/k}$ baked him_i .

6B. Dora baked him_i the same food that Cookie Monster's $_i$ father $_k$ baked him_i .

The true / false judgment in the experiment indicated that children, just like adults, find (6A) ungrammatical if the embedded *him* refers to Cookie Monster. If it does not refer to Cookie Monster (also (6A)), or if *Cookie Monster* is in a possessive construction with *father*, then the sentences are acceptable. As this distinction is explained by Principle B in the generative tradition, and since very few ACD structures were discovered in child-directed input from which to learn the complicated rules, Kiguchi and Thornton conclude that binding principles are part of the human biological endowment. This paper will offer a semantic carpentry alternative that shows that it is not necessary to specifically hear ACD constructions to make accurate judgments of grammaticality. As Jackendoff (2002) has already indicated that pronouns like *him* are handled on the RT of semantic structure, our reconstruction of the sentences will focus only on how the embedded sentences (**Cookie Monster_i baked him_i* and *Cookie Monster's_i father baked him_i*) are constructed in semantics.

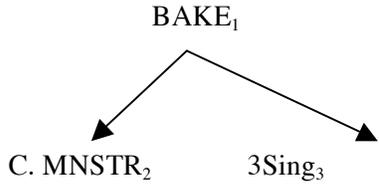
As *him* is the last word in each sentence, we will skip to a completed DT representation.

7A. DT: *Cookie Monster_i baked him_i*

$[_{\text{EVENT}} \text{BAKE} \quad ([_{\text{OBJ}} \text{C. MNSTR} \quad]_2, [_{\text{OBJ}} \text{3SING} \quad]_3)]_1$

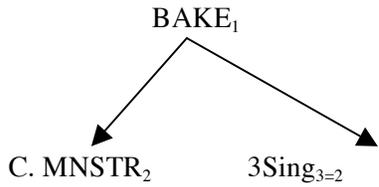
When the parser encounters *him*, it will be able to complete the arguments for *bake*, as well as form a complete NP, which maps into the descriptive tier as an Object with its unique index (in this case index 3). Since *him* requires interpretation on the RT, no further DT actions are required. As items are filled in the DT, mapping to the RT will occur at the first opportunity. This means that as soon as the Event is established on the DT, an index will be created on the RT. The two Objects, Cookie Monster and 3Sing, will then be placed "under" the Event's index, as the existence of the event implies the existence of the entities that participate in the event. The following RT results (subscripted words instead of simple numerals are used here to reduce confusion):

7B. RT: *Cookie Monster_i baked him_i*



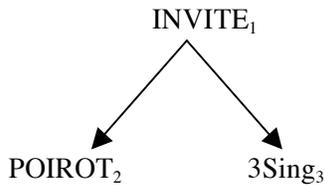
The parser now needs to resolve the referential dependency introduced by *him*, which has been waiting for resolution with the first available Object in the RT. The first available Object is Cookie Monster, index 2, and so the processor will attempt to copy the index from Cookie Monster to 3Sing. This results in 7C:

7C. *RT: *Cookie Monster_i baked him_i*

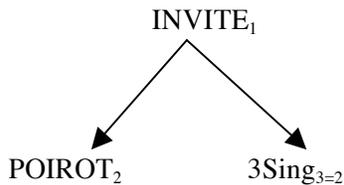


However, here we run into a problem as is indicated by the asterisk. (7C) states that an Event, Bake, implies the existence of a single entity, here Cookie Monster. However, in such situations, language quite uniformly uses a reflexive anaphor. Three examples that demonstrate this are presented in the main text. Appendix A includes a more comprehensive set of examples to establish the plausibility of this suggestion. (8A) – (8C) borrow sentences from Haegeman (1994:215, 224) to reflect the distribution of pronouns and reflexives.

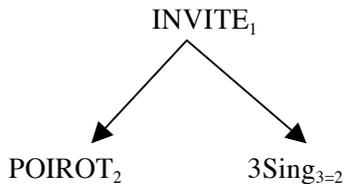
8A. RT: *Poirot_i invited him_k*



8B. *RT: *Poirot_i invited him_i*



8C. RT: *Poirot_i invited himself_i*



In (8C), the anaphor of index 3 is mapped to the closest other Object, index 2. Since this is the classic distribution of reflexives, this is acceptable. (8A) is also allowed, because the pronoun is not the same entity as Poirot and therefore does not encroach on the reflexive's territory. (8B) does attempt to em-

ploy a pronoun in the place that an anaphor is correctly used. Therefore, (8B) is rejected. This allows us to make the Reflexive Generalization.

Reflexive Generalization: If an Event directly implies the existence of a single referential entity, a reflexive must be used. Corollary: A reflexive may not be used when the existence of the entity to which it refers is implied by a different Event than the one that implies the reflexive.

This generalization in semantic structure bears similarities to Binding Principles A and B of generative syntax, of course. This should not be surprising as, whatever its limitations, the binding principles have been one of the most successful components of syntactic research, so it is natural that alternate proposals will share some of the same insights. However, the formulation here is not the same as the binding principles. This generalization relies on very different concepts, such as existential force (again see Appendix A), and means something different. Also, it becomes easier to see how this pattern might be learned. Statistical learning has demonstrated that learners focus on statistical regularities over noisy-but-more-frequent input (Gomez 2002, 2006). Since reflexives are consistently used in this context, learners will extract that information from their input. This would result in a requirement for reflexives in this context and the rejection of other items.

This Reflexive Generalization brings us to another one of our forks in the road. (F4) The reflexive generalization is a regularization of the statistical patterns in the input regarding the environment of reflexives. (F5) The Reflexive Generalization is part of Universal Grammar and part of the biological endowment for language. Either (F4) or (F5) will work for the purposes of the proof underway, as (F5) is a significantly different conception of our biology than the conclusion in Kiguchi and Thornton 2004.

4. DISCUSSION. It is time for a brief review of what we have accomplished so far in this paper. First, it has positioned itself as neutral on the normal nativist vs. non-nativist debate. Indeed, a very broadly nativist position is assumed in that there is some biological endowment for language, and so the question becomes one of details. What precisely is the biological endowment? Before we can productively ask that, however, we must first decide if our methodology is up to the task. Since many linguists argue that Poverty Arguments provide strong support for a language-specific component of the endowment, and, moreover, allow us discover the actual content of that endowment, we have been examining whether or not this is the case. It was then argued that in fact many different biological endowments are compatible with the evidence from the Poverty Argument examples we have. Several different forks in the road were identified, such that the Lidz et al. experiment is simultaneously providing evidence of innate phrase structure (Lidz et al.'s explanation), alternative content of UG, such as innate mapping of noun phrases into objects (F3), regularities between word classes (F2), and the specific semantic content of lexical items (F1). Similarly, Kiguchi and Thornton 2004 is simultaneously providing evidence of Principle B, a UG that contains the Reflexive Generalization (F4), and regularization of the statistical environment of reflexives (F5). Which of these we choose will then depend on the theoretical framework we bring to the question.

Two objections are immediate here. Crain and Pietroski 2001 argues that one must bring a theory to the table when conducting acquisition research, or else one cannot even ask what is being acquired. This point is quite valid if the debate is the normal nativist / non-nativist one. However, this paper is acknowledging that the biological endowment might very well include language-specific content or it might "all" be learned. Instead, now we want to know precisely what that content is. Much nativist research has an underlying tension between, on the one hand, actively choosing to let the biological specifics remain unanswered until a later day, while, on the other, claiming that their approach is particularly biological in a way that other approaches are not. This is presumably because the language parameters themselves are biological in a way that many non-nativists reject. If this is correct, then it means the nativist tradition is already gaining substantial evidence of what the biological endowment for language is, and we are not waiting to fill in those details at a future time. However with the underdetermination problems we are seeing, the Poverty types of experiment have added little to our knowledge of the biological endowment, as which choice we make concerning the explanatory alternatives depends almost entirely on factors external to the Poverty Argument—and yet each Poverty

paper we examined ended with a conclusion that the research presented has provided strong evidence for a certain innate endowment.

The second, related objection is that this paper is asking too much of a Poverty experiment. Such experiments do not occur in theoretical or experimental vacuums. It is entirely appropriate that the justification of a theory depend on far more than one bit of experimental data. This second objection is valid and accepted. In that case, how do linguists decide which is the better explanation and therefore the most likely content of the biological endowment? To make the question specific, how do we decide between the proposals of Semantic Carpentry and those of the Poverty Argument experimenters?

One approach is the one most often used currently. We can look at each theory's simplicity, its universality, its predictive power, its coherence, its corroboration from independent facts, and its coverage—in short, common criteria for assessing scientific theories. These questions can be productive and make some progress. However, as is revealed in many of the nativism debates, this progress can be slow going. One side will hold up a critical fact that binding theory cannot explain, while the other side holds up the limits of its competitor. Science often moves in precisely this way, and such a path should not be dismissed. However, there is another method of adjudicating between theories that might be more productive in the near-term, and that is to return to our notion of causal explanation as discussed in section 2.2.

If we assume that there is a biological endowment for language and we want to discover its content, then we are asking a question of *how*. How does our biology allow us to acquire language? We can see that children possess some knowledge about language. How does their biology allow them to possess that knowledge? Describing that knowledge is a very important scientific endeavor, but describing what the knowledge is is not the same as describing how that knowledge is enabled biologically. If a researcher asserts that grammar is learned due to statistical pattern recognition, that is an important observation, but it does not yet tell us what about our biology allows that recognition to occur. Similarly, if our best description of knowledge about pronoun interpretation is Binding Principle B, that does not tell us much yet about how our biology allows us to know that principle. This can remain the case even as one approaches the limit of knowledge description. An analogy, while always dangerous, can help make this point.

Let's take the example of bird flight. A description of the knowledge and behavior of bird flight might include items such as acceleration rates, velocity at certain altitudes, turning capabilities, take-off and landing patterns, and the like. If one theory predicted a maximum turning speed that the bird then exceeded in our empirical data, this theory might be discarded for one that is more comprehensive of the facts. Moreover, the perfect description would describe all of these abilities throughout the lifespan of the bird—when each ability became available, i.e., its maturation, stable states, and loss. Such a description would be a tremendous feat of science and well worth pursuing. However, this description is telling us what bird flight is like, but it is not yet explaining how the bird does any of it with its physical gifts. The question of how a bird's biology enables flight might include aspects such as bone structure and weight, shape of the wing, motor control behavior in the nervous system, etc. When we learn that, we learn not only what bird flight is like, but also what there is about the bird's biology that allows flight.

Some may object to this example as being too oriented in motor behavior, while language is mental or psychological. To meet such an objection, we can switch to a different example: that of drawing a circle with a stick in the dirt (in order to recall Plato's *Meno* once again). When drawing a circle, the basic task is to create a physical instantiation of a circle's circumference with the radius (or possibly diameter) as the guide to the length of the circumference. This line will start when the stick hits the dirt and will not be complete at least until the beginning is reached again. Clearly, the best formal description of the length of a circle's circumference is the diameter multiplied by pi (πd) or twice the radius multiplied by pi ($2\pi r$). This is an accurate description of the length of the circumference of a circle we are trying to draw. However, it is not clear that this mathematical formula is actually how someone draws. Instead, a likely scenario is that the person drawing knows they must keep moving their pencil, but that the pencil must stay the same distance from the center at all times. If they successfully accomplish this, they will create a circle that matches the formal description of what a circle is. However, they will have done it with no knowledge of that description, or, perhaps more accu-

rately, the correct formal description of the length of the line they must draw was not used as a guide in creating the circle. The *how* of circle drawing (the performance, if you will) is distinguished from the *what* of drawing (the structural description of the circle).

To a large extent, this is where we are with both the generative explanations and the semantic carpentry explanations, which are essentially both explanations of *what*. As the theories improve, they tell us more and more about what language is like, but they still are not informing us of how our biological bodies do it unless we can make an argument that the child actually does this when using language. If we can find ways to ask what about our bodies causes this knowledge, then we will then be asking the biological question.

One intriguing fact is that at least one understanding of the classic competence/ performance distinction often has the effect of explicitly preventing researchers from asking this question. It is common among competence accounts of language, such as Government and Binding theory or Optimality Theory, to put aside objections that the mind cannot actually be doing *that*, whatever that is, when actually processing or producing language because the theory is a competence model, not a performance model. In such a model, in whatever manner the brain or mind goes about accomplishing language use, the end result fits the proposed competence model. Yet, with this understanding of competence and performance, a clear tension arises between the desire for a biological account of language and an account which puts aside questions of how the brain or mind does something. If our goal is to discover competence (i.e., to describe linguistic knowledge) without describing performance (how the mind actually produces and comprehends language), then we are actively refusing to ask the biological question. There is a conflict between biological explanation and a too rigid competence / performance distinction.

This is not the only way to understand the generative tradition's conception of competence or linguistic knowledge, though. Indeed, Devitt (2007) discerns a very different idea of linguistic knowledge in the Chomskyan tradition. Devitt argues that the most natural interpretation of Chomsky is that linguistic knowledge is a set of explicit, though unconscious, propositions about language, which then would commit Chomsky to having linguistic theory not just be a model of the speaker's knowledge but also real-time language processing. Devitt then argues that linguistic theory as psychological how-to is not justified.

What was just stated must not be misunderstood as a criticism of the competence/ performance distinction as a matter of scientific idealization. This is entirely reasonable and the enormous body of research that has developed under this distinction demonstrates its practicality. Describing linguistic knowledge is an important enough research task. If our contemporary understanding of the biological instantiation of competence is not sufficiently developed to help us describe that knowledge, then it can be entirely appropriate for the researcher to ignore performance questions temporarily. (Think how far back we would be in our knowledge of circles if we only knew how people draw circles but not that the circumference was equal to $2\pi r$.) However, the idealization does not come without a sacrifice. If our competence description could be instantiated neurally and psychologically in many different ways, and if we grant that the way people actually perceive, produce, and comprehend language might be distinct from our competence description, both claims that are frequently made by competence researchers when facing criticisms that a grammar is not psychologically plausible, then we are not yet providing a causal explanation of how our biology enables language use.

So what is the way forward? There are at least two possibilities. (1) We accept that it is necessary to idealize away from causal biology, and we continue our study of linguistic competence. If this choice is taken, though, we should stop claiming to be discovering the biological endowment of language. A pure competence model is not yet a biological model of language, as its lack of causal explanation makes it particularly prone to the problems of underdetermination as outlined in this paper. (2) We begin asking questions related to how biology allows us to acquire language. The obvious component of this would be a neurolinguistic approach, as there is no doubt that the biological brain is somehow the primary instantiation of human linguistic knowledge. As stated previously though, Chomsky 2002 argues that our knowledge of the brain is simply not far enough along to productively ask such questions yet. A companion biological account would connect our genetic endowment to language. The problem here is that that account seems inexorably tied up with the neurological one, so if Chomsky is correct about the former, it would likely affect prospects for the latter.

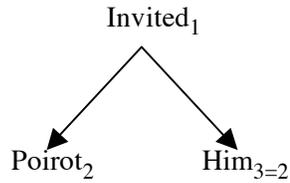
The good news is that, even if we do grant that neurology and genetics are a current dead-end, we are not stuck with ignoring biological causation. If a legitimate evolutionary account of language could be provided, this would also help us understand how biology created language. Assuming that such an account explained what biological changes occurred in our species, it would be another method at assessing what the biological endowment for language is. This author is, however, as wary of the immediate prospects of evolutionary linguistics explaining binding theory as Chomsky is of neurolinguistics doing the same.

If this method fails us, we are still not out of options. The final one, and the one that seems most immediately promising, is simply to take down the competence/performance distinction when asking biological questions (but perhaps not at other times). If performance includes how language is produced, processed, and comprehended, then, if we can understand those topics, we are actively learning how our psychology causes language. An ultimate solution might still need to connect psychology to neurology, but there is no need to take such a reductionist approach or proclaim that we are getting nowhere until such reduction occurs. If we can say not only that anaphor interpretation can be described accurately with binding principles, but also describe how a child actually processes anaphors, then we are making substantive progress towards discovering that biological endowment.

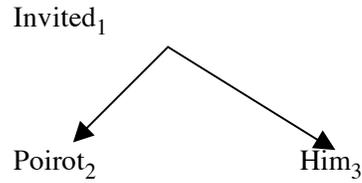
APPENDIX A. THE REFLEXIVE GENERALIZATION

This appendix presents a more comprehensive justification for the Reflexive Generalization. While the notion of referential tiers, indices, and existential assertion are all from Jackendoff 2002, the Generalization, its use here, and justification are innovations of the current paper. The example sentences are all from Haegeman 1994:215, 224 and demonstrate the distribution of *him* and *himself*. For each example, the Referential Tier is constructed, and it is shown that reflexives occur uniformly in situations in which a single Event or Object directly implies the existence of the anaphor and its referent as a single entity. When a pronoun occurs in such a situation, it is considered ungrammatical.

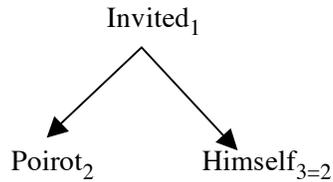
(A1) *Poirot_i invited him_i.



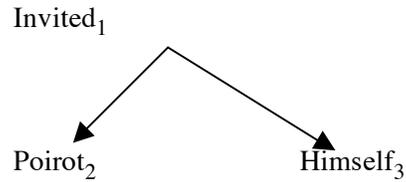
(A2) Poirot_i invited him_k



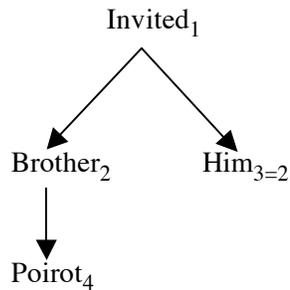
(A3) Poirot_i invited himself_i.



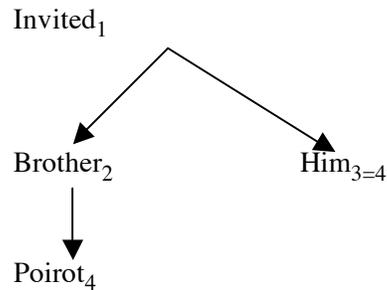
(A4) *Poirot_i invited himself_k



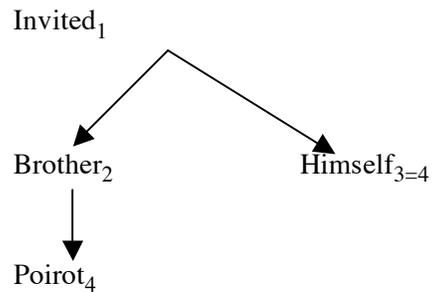
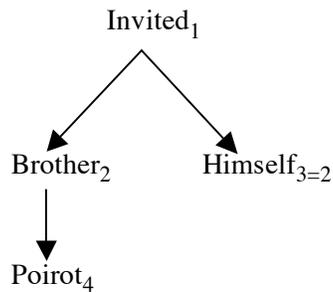
(A5) *Poirot_k's brother_i invited him_i.



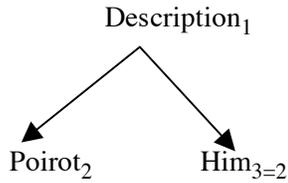
(A6) Poirot_i's brother_k invited him_i



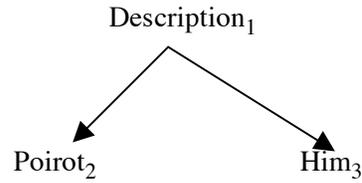
(A7) Poirot_k's brother_i invited himself_i. (A8) *Poirot_i's brother_k invited himself_i



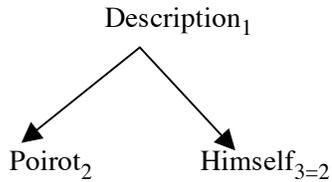
(A9) *Poirot_i's description of him_i



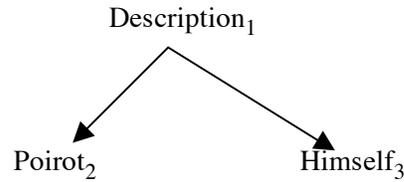
(A10) Poirot_i's description of him_k



(A11) Poirot_i's description of himself_i

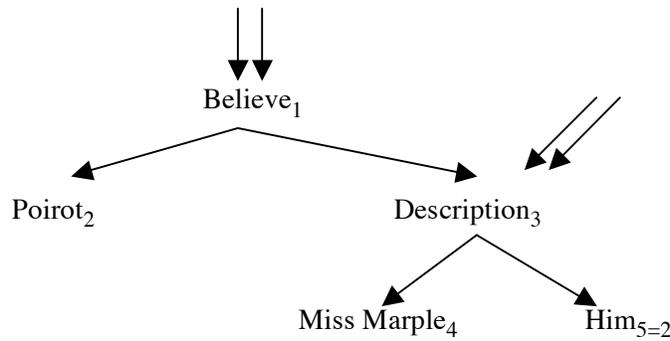


(A12) *Poirot_i's description of himself_k

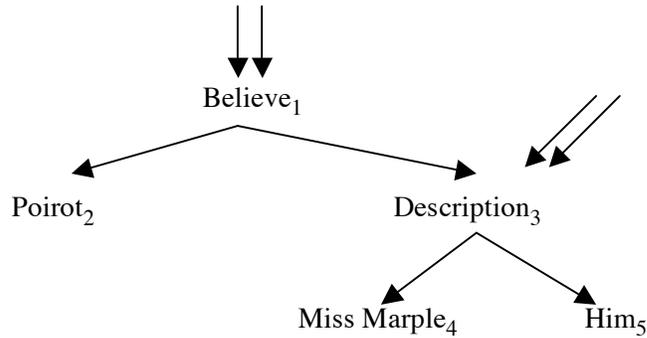


This also works for more complicated sentence structures. To see how, we have to introduce further machinery from Jackendoff 2002, which was only alluded to in section 3. Jackendoff considers a notion of existential force. Simple declaratives state that Events and Objects actually exist. Negation, questions, and possibility, while still represented on the RT, lack existential force. In the case of negation, the negated item is being stated not to exist or not to have occurred. With questions and certain modals, there is no positive assertion of existence. This can be extended to the phrases *Miss Marple's description of him* and *any description of him*. The first phrase implies that a description actually exists and that Miss Marple made the description. The second phrase leaves the matter uncertain. The description may or may not yet have taken place. From this perspective, then, *any description* is existentially transparent, not making an assertion in either direction. In the Jackendoff notation existential force is marked with a double arrow. More complete notations for examples (A1) – (A12) would then have a double arrow marking the top level, and each nominally modified Object (Poirot's Brother) would possess a double arrow as well.

(A13) Poirot_i believes Miss Marple_m's description of him_i.

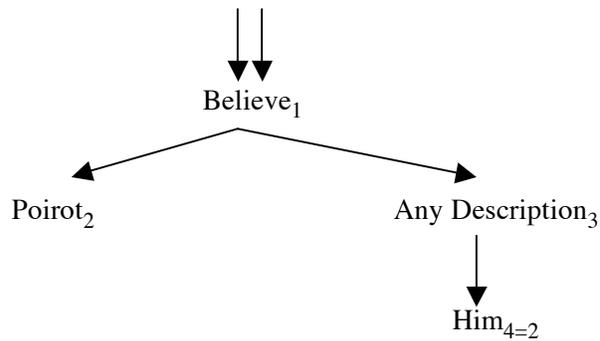


(A14) Poirot_i believes Miss Marple_m's description of him_k.

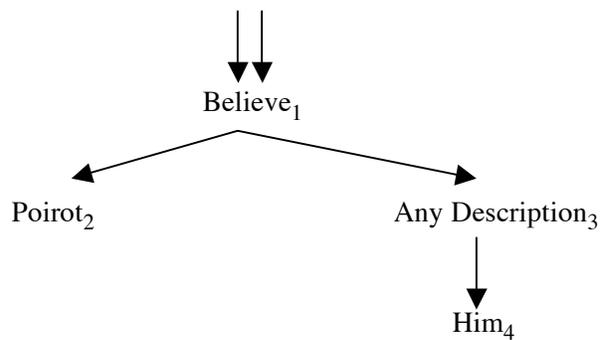


Both sentences are acceptable because Description has existential force and both items which it directly implies the existence of, being Miss Marple and Him, are separate referential objects; i.e., they are not co-indexed.

(A15) *Poirot_i believes any description of him_i.

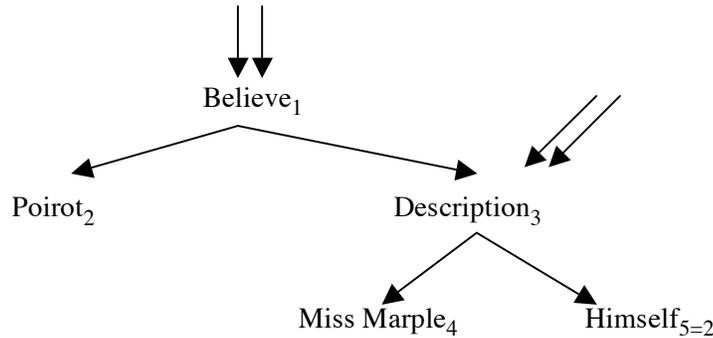


(A16) Poirot_i believes any description of him_k.



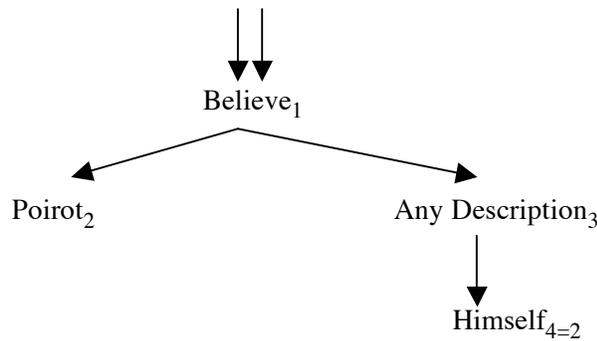
(A15) is unacceptable, because Any Description is existentially transparent, which is indicated by the lack of a double arrow in the notation. Due to this transparency, the Event Believe is directly implying the existence of Poirot and Him as a single entity. However, according to the Reflexive Generalization, this is the domain of reflexives, not independent pronouns.

(A17) *Poirot_i believes Miss Marple_k's description of himself_i.



(A17) is unacceptable because Description has existential force but the Reflexive is being identified with an entity that Description does not imply the existence of. The same Event or Object must imply the existence of the reflexive as well as the entity it is indexed with.

(A18) Poirot_i believes any description of himself_i.



In A18 Any Description is existentially transparent, which then allows the Event Believe to imply the existence of both the reflexive and its co-indexed entity. As stated in the main text, these constraints might be part of Universal Grammar or they might be learned from the regularities of reflexives in the input. Either way, it is a different conception than traditional Binding Principles.

REFERENCES

- AKHTAR, NAMEERA; MAUREEN CALLANAN; GEOFFREY K. PULLUM; and BARBARA SCHOLZ. 2004. Learning antecedents for anaphoric *one*. *Cognition* 93:141–45.
- ANDERSON, STEPHEN, and DAVID LIGHTFOOT. 2002. *The language organ: Linguistics as cognitive psychology*. Cambridge: Cambridge University Press.
- ANDERSON, STEPHEN, and DAVID LIGHTFOOT. 2006. Biology and language: A response to Everett. *Journal of Linguistics* 42:377–83.
- CHOMSKY, NOAM. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.
- CHOMSKY, NOAM. 2002. *On nature and language*. Cambridge: Cambridge University Press.
- COWIE, FIONA. 1999. *What's within? Nativism reconsidered*. New York: Oxford University Press.
- CRAIN, STEPHEN, and PAUL PIETROSKI. 2001. Nature, nurture, and universal grammar. *Linguistics and Philosophy* 24:139–86.

Hatfield: The Biological Endowment for Language and Arguments
from the Poverty of the Stimulus

- CRAIN, STEPHEN; ANDREA GUALMINI; AND PAUL PIETROSKI. Forthcoming. Brass tacks in linguistic theory: Innate grammatical principles. Downloaded September 2, 2007, from <http://www.maccs.mq.edu.au/~scrain/publications.htm>.
- DEVITT, MICHAEL. 2007. *Ignorance of language*. Oxford: Oxford University Press.
- ELMAN, JEFFREY; ELIZABETH BATES; MARK JOHNSON; ANNETTE KARMILOFF-SMITH; DOMENICO PARISI; and KIM PLUNKETT. 1996. *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.
- EVERETT, DANIEL L. 2006. Biology and language: Response to Anderson and Lightfoot. *Journal of Linguistics* 42:385–93.
- GOMEZ, REBECCA L. 2002. Variability and detection of invariant structure. *Psychological Science* 13:431–36.
- GOMEZ, REBECCA L. 2006. Dynamically guided learning. *Attention & performance XXI: Processes of change in brain and cognitive development*, ed. by Yuko Munakata and Mark Johnson. Oxford: Oxford University Press.
- GUERTS, BART. 2000. Review of Stephen Crain and Rosalind Thornton, *Investigations in universal grammar: A guide to experiments on the acquisition of syntax and semantics*. *Language and Philosophy* 23:523–32.
- HAEGEMAN, LILIANE. 1994. *Introduction to government and binding theory*. (Second edition). Oxford: Blackwell Publishing.
- JACKENDOFF, RAY. 2002. *Foundations of language: Brain, meaning, grammar, evolution*. New York: Oxford University Press.
- KIGUCHI, HIROHISA, and ROSALIND THORNTON. 2004. Binding principles and ACD constructions in child grammars. *Syntax* 7:234–71.
- LASHLEY, KARL. 1951. The problem of serial order in behavior. *Cerebral mechanisms in behavior*, ed. by A. Jeffress, 112–46. New York: Wiley.
- LIDZ, JEFFREY; SANDRA WAXMAN; and JENNIFER FREEDMAN. 2003. What infants know about syntax but couldn't have learned: Experimental evidence for syntactic structure at 18 months. *Cognition* 89:B65–73.
- MACWHINNEY, BRIAN. 2004. A multiple process solution to the logical problem of language acquisition. *Journal of Child Language* 31:883–914.
- NEWPORT, ELISSA L., and RICHARD N. ASLIN. 2004. Learning at a distance I: Statistical learning of non-adjacent dependencies. *Cognitive Psychology* 48:127–62.
- O'GRADY, WILLIAM. 2005. *Syntactic carpentry: An emergentist approach to syntax*. Mahwah, NJ: Erlbaum.
- PAPINEAU, DAVID, ed. 1996. *The philosophy of science*. New York: Oxford University Press.
- SCHOLZ, BARBARA C., and GEOFFREY K. PULLUM. 2006. Irrational nativist exuberance. *Contemporary debates in cognitive science*, ed. by Rob Stinton, 59–80. Oxford: Basil Blackwell.
- TOMASELLO, MICHAEL. 2003. *Constructing a language: A usage-based theory of language acquisition*. Cambridge, MA: Harvard University Press.
- WEXLER, KENNETH. 1996. The development of inflection in a biologically based theory of language acquisition. *Towards a genetics of language*, ed. M. L. Rice, 133–44. Mahwah, NJ: Erlbaum.

hunterh @hawaii.edu