1 *Introduction*

How do children take a string of speech sounds, chop it up into discrete units (words), and then assign to that chopped up string of sounds a particular meaning? In many sentences, perhaps most, words that are semantically related to one another are also near to each other in the sentence. For example, a verb and its arguments – the nouns or other phrases that the verb selects – are usually in close proximity to each other (at least, they are generally clausemates): in a simple main clause sentence like *The student read a book*, the verb *read* selects a subject and a direct object noun phrase (NP), and these NPs are positioned right next to the verb that selects them. This is so regardless of the particular language's basic word order or even the degree of rigidity of word order. Many theories of language acquisition exploit this fact to explain (part of) how children begin to tackle the challenge of integrating form and meaning in their language.

But arguments need not be proximal to their selecting predicate, and adjacent or proximal words need not stand in a semantic, selectional relation to one another. This is because human language allows for the semantic relations between words to span long distances – in principle, infinite distances. This book is about how children begin to figure out how to interpret sentences in which the proximity of words belies their semantic relations – how children determine the underlying syntactic structure of sentences in which semantic relations are long-distance, and how knowing the syntactic structure helps children interpret those semantic relations.

I argue that children recruit cues from the conceptual domain, particularly animacy, in solving this puzzle. One very important insight about language acquisition stems from the observation that subjects of basic, canonical sentences are often animate, or more animate than other nouns in the sentence, and that children can exploit this fact to home in on basic sentence structure (Pinker, 1984): find the most animate noun and it will be the subject. My question is how children go beyond these canonical sentences in which subjects are agents or experiencers, and objects are patients or themes, to figure out the

1

2   *Introduction*

structures of more complex configurations. The answer I offer is that because
children expect subjects to be animate in canonical structures, they can exploit
deviations from this expectation, in particular encountering an *inanimate sub-
ject*, to learn that in just these kinds of sentences the underlying structure is
non-canonical and complex. A predicate that allows its subject to be inanimate
does not bear the same type of semantic relationship to its subject as a pred-
icate that requires its subject to be animate; and the non-canonical semantic
relation between subject and predicate translates into a non-canonical syntactic
structure: one in which the subject is derived, or displaced.

The particular type of non-canonical structure I will focus on involves what
I'll call "displacing predicates." These are predicates that fail to select an exter-
nal argument (a semantic subject) – that is, there is no "do-er" or experiencer
of the predicate's action or state. An example of such a predicate is the verb
seem. When we say *John seems to like French fries* it doesn't make sense to say
that John is a "seemer" of anything. Instead, the subject is semantically related
only to the lower predicate, and thus we can say it is "displaced" (or derived) in
the sentence with respect to the locus of its semantic role. The question I seek
to answer is how children figure this out – how they identify just the sentences
of their language in which the subject is in fact displaced, which in turn allows
them to categorize particular predicates as being "displacing" predicates.

One might think that this is such a small corner of the grammar – displacing
predicates are such a tiny piece of what needs to be learned about language, and
they have a rather peripheral feel to them. Surely what matters most in advanc-
ing the study of language and its acquisition is to explain how children acquire
the canonical parts of grammar, the most well-behaved and earliest learned
predicates, so that the exceptional ones can then be accounted for, precisely, as
exceptions. How is the study of these unusual predicates relevant in the larger
scheme of things? There are two related answers to this question. One is that
these very predicates and their non-canonical structures represent one of the
core properties of human language itself: the ability to have non-local depen-
dencies. That words can bear structural relations to other words over an (in
principle) infinite distance is one of the hallmarks of human language. In this
sense, displacing predicates are profoundly non-peripheral.

The second and related answer is that these predicates have formed part
of the argument for generative grammar from the very beginning. Not only
does the learning puzzle addressed in this book involve determining that a
given sentence has a displaced subject, but also the learner must distinguish
those sentences with displaced subjects from superficially identical sentences
whose subject is *not* displaced, but rather is the semantic subject of the main

predicate (e.g. in *John claims to like French fries* John is the "claimer"). This aspect of the question is old and deep, and it forms one of the pillars on which generative grammar was built. What Chomsky (1957) called "constructional homonymity" (*John seems/claims to like French fries*) was exposed as a fatal flaw in a theory of grammar that did not combine both phrase structure rules and transformational rules. Even though these subclasses of verbs can be distinguished by their distribution in other types of sentences (*It seems/*claims that John likes tomatoes*; *What John claims/*seems is to be the strongest*), the fact of their distributional overlap in even one sentence type requires that learners have a means of teasing them apart. It was suggested in Chomsky (1964, 1965) that the challenge presented by constructional homonyms in terms of language acquisition pointed to the need for an account of language learning within the rationalist tradition of epistemology (Chomsky, 1965, p. 25). That is, they bore directly on arguments for innate knowledge of language.

During this era, Carol Chomsky (1969) took up the puzzle these constructional homonyms posed in an empirical study of children's language. She posed the question of how children handle situations in which

> [t]he true grammatical relations which hold among the words in a sentence are not expressed directly in its surface structure.   (Chomsky, 1969, p. 6)

That is, how do children parse a particular word string that is potentially associable with multiple underlying structures? Chomsky focused on sentences of the form in (1).

(1)   The doll is daxy to see.

Without knowing what the adjective *daxy* means, the sentence could have (at least) either of the following interpretations, the first corresponding to the so-called *tough*-construction (2a) and the second to the control adjective construction (2b), as disambiguated by the familiar English adjectives.

(2)   a. The doll is easy to see.
              (= it is easy for someone else to see the doll)

          b. The doll is eager to see.
              (= the doll is eager to see someone else)

The difference between (2a) and (2b) is clearly semantic, but it is also syntactic. Syntax is about not just the ordering of words, but also the logical relations among them: the fact that the relation between *the doll* and *easy* is profoundly different from that between *the doll* and *eager* is linked to a

4   *Introduction*

difference in how the structures of these two sentences are represented. The
nature of these structural differences will be explained in detail in Chapter 2.
For the moment what is significant is that the semantic role of the subject NP is
utterly different in sentences (2a) and (2b), and therefore its syntactic relation-
ship to the main predicate is different. While the main clause subject in (2a) is
understood as the semantic object (patient) of the embedded clause, in (2b) the
main clause subject is interpreted as the semantic subject (agent/experiencer)
of the embedded verb. Assuming that the semantic difference between (2a)
and (2b) corresponds to an underlying syntactic difference between them, in
the terms being used here the subject in (2a) is displaced, but the subject in
(2b) is not. The parallel to the earlier *seem* example is that neither *easy* nor
seem takes an agentive (or experiencer) subject. For both of these predicates,
the subject's semantic ties are to another predicate altogether; in this sense,
both *seem* and *easy* are displacing predicates.

(3)   The girl is daxy to see.
        a. The girl$_i$ is easy [PRO$_{arb}$ to see t_i.] (*tough*-adjective)
        b. The girl$_i$ is eager [PRO$_i$ to see.] (control adjective)

(4)   Mary gorped to be strong.
        a. Mary$_i$ seemed [t_i to be strong.] (raising verb)
        b. Mary$_i$ claimed [PRO$_i$ to be strong.] (control verb)

The semantic difference between the (a) and (b) pairs in (4) is a little subtler
than that in (3). In both (4a) and (4b) *Mary* is the semantic subject of the lower
predicate *to be strong*. The difference has to do with its semantic relation to
the main predicate, *seem* vs. *claim*: as noted above, there really is no semantic
relationship between *Mary* and *seem*, but there is between *Mary* and *claim* (she
is the "claimer"). Again, this semantic asymmetry corresponds to a syntactic
one: the subject is displaced in (4a) but not (4b).

So the problem for language learners is to figure out that the subject of *seem*
or *easy* is not the semantic subject of these predicates, but rather bears a long-
distance semantic relationship to another part of the sentence, even though a
construal of the strings in (3) and (4) involving a local semantic relationship is
possible given the constructional homonyms with *claim* and *eager*.[1]

---

[1] The same surface ambiguity arises in raising-to-object (also called Exceptional Case Marking)
     and object control, as in *Sue wanted/asked Gordon [to cut the grass]*. Since the main focus here
     is on constructions with derived *subjects* I will not have a lot to say about these constructions,
     but they will be discussed briefly in Sections 2.1.3 and 5.3.1.

This was precisely the question that Carol Chomsky posed. But it was not the question she answered in her empirical work. Rather, her experiments addressed the question of *what* children know about predicates like *easy* and *when* they know it. In fact, nearly all of the literature on children's acquisition of *tough*-adjectives and raising verbs has focused on this aspect of the problem, and so in the decades since Chomsky's seminal work, the deeper question of *how* children disentangle the respective constructions has not been tackled directly.

The purpose of this book is to tackle that *how* question. The answer I propose is that hearing an inanimate subject in a sentence like (3) or (4) provides a cue that the subject is displaced, and therefore that the main predicate of the sentence is a displacing predicate. This cue is informative in these cases because an *inanimate* subject is possible with the structure that involves displacement, but not with the other structure:

(5)   a.   The rock is easy to lift.

        b. # The rock is eager to lift/fall.

(6)   a.   The rock seems to be heavy.

        b. # The rock claims to be heavy.

Lest readers be concerned that I have missed a more obvious answer to this puzzle, namely displacing predicates' ability to occur with expletive subjects (*it*, *there*), I should state that I do think predicates' occurrence with expletives is a valuable cue in this learning process, and I have reasons for focusing on inanimate referential subjects instead. These reasons are laid out in detail in Section 2.5 below.

The main focus in this book will be on the two constructions in (3/5) and (4/6), those involving *tough*-adjectives and raising-to-subject verbs. However, there are other constructions that involve subject displacement, such as the passive, and there are other (non-passive) predicates that can be classified as displacing predicates, such as unaccusative verbs. Unaccusative verbs are a type of intransitive verb which, unlike unergative intransitives, select only an internal argument and no external argument. Thus, the subject of an unaccusative verb has been displaced from an underlying object position. But given a surface string containing only a subject and a verb, it is not immediately obvious whether the underlying structure involves an external argument (as in (7b)) or an internal one (as in (7a)).

6   *Introduction*

(7)   John pilked.
           a. John$_i$ arrived t_i. (unaccusative)
           b. John danced. (unergative)

The asymmetry in (7a–b) is even more subtle than in (4a–b), and fairly unintuitive for English speakers. We will see in Section 2.3 that many languages exhibit more obvious distinctions between unaccusative and unergative verbs, and the distinction between these types of verbs is well supported cross-linguistically. The spirit of the distinction is that in (7a) John is the theme of the verb (in a sense, John "undergoes" the arriving event; he does not have an agentive role), while in (7b) John is the agent of *dance*: he "enacts" the dancing event. Thus, while *arrive* and *dance* are both intransitive verbs, the underlying relationship between the verb and its lone argument is different in each case. Once again, these distinctions map onto structural differences that the language learner must be able to identify in order to be said to have adult-like competence in her language. And similar to the first two constructions, the string in (7) can be associated with the displacing structure in (7a) if the subject of the sentence is inanimate.

(8)   a.   The package arrived.
           b.   # The package danced.

My proposal is primarily about how children solve the mapping and categorization problems: I take the view that predicates, with some important exceptions, are fundamentally either displacing or non-displacing. So encountering an inanimate subject tells the child that the sentence involves a structure with a displaced subject, which in turn tells the child that the main predicate is a displacing predicate. Although I do not try to explain how children figure out exactly what these abstract predicates mean, I suspect that the categorization of a predicate as displacing (or non-displacing) in turn provides a clue to the set of possible meanings the predicate might have. That is, displacing predicates will be largely limited to auxiliary-like meanings – meanings having to do with modality, happenstance, appearance, ease/difficulty, possibility, and likelihood (and non-volitional events, in the case of unaccusative verbs). Non-displacing predicates, on the other hand, will have a volitional, intentional, or emotive aspect to their meaning.

As we have seen, sentence strings like (3), (4), and (7) are associable with multiple syntactic structures if the subject is animate and the predicate's meaning is not known. I will refer to these sentence strings as "opaque"

sentences rather than use the term "ambiguous." The reason is that the structural indeterminacy of these sentences is different from the more typical type of structural ambiguity presented by, for example, Prepositional Phrase (PP) attachment.

(9)   a. Put the frog on the napkin in the box.
        b. I saw the man with binoculars.

The string in (9a) is locally ambiguous at the first PP [on the napkin] because this phrase could indicate either a description of where the frog is, or it could indicate the location where the frog should be put (and the string is disambiguated by the second PP). The processing of this type of construction has been explored extensively in both children and adults (Trueswell *et al.*, 1999, i.a.). But the decision about where to attach the first PP does not have an effect on the lexical meaning of the main predicate *put* – *put* means the same thing, whether the PP is attached to the NP or the VP. Similarly in the globally ambiguous example in (9b), the meaning of *see* does not depend on which structure one applies to this string. If a learner encountered an unknown verb in this string (*I gorped the man with binoculars*) the meaning of *gorp* would not necessarily depend on whether the PP was attached to the NP or the VP node. (And, correspondingly, knowing the meaning of *gorp* would not help resolve the attachment puzzle, and so the sentence is truly ambiguous.) In the kinds of constructions under consideration here, on the other hand, the meaning of the main predicate is fundamentally different according to whether the subject is displaced or not. Not only are the verbs *seem* and *claim* different verbs (and this extends to the other pairs of predicates we've seen: *easy/eager*, *arrive/dance*), but if we encounter a novel predicate in a string like (4) the meaning of this predicate *will* depend on how the string is parsed.

On the surface, this might appear to make the learning problem easier. If you know the meaning of the predicate (*seem*, *claim*, etc.) you can choose the right structure: if you know that the main verb means 'seem' then you know the subject is displaced, and if you know the main verb means 'claim' then you know the subject is not displaced. Thus, the sentence *Mary seems to be strong* is not ambiguous – once you know the lexical properties of *seem* the underlying structure of the sentence follows. However, this does not solve the learning problem for children, for two reasons. First, most of the verbs and adjectives that participate in these structures have abstract lexical meanings that are not straightforwardly discernable directly from observation of the non-linguistic world (*eager* and *easy* are both different from *red* in this respect; and *seem* and *claim* are different from *eat* similarly). Secondly, a wealth of empirical studies,

forming the literature surrounding the Syntactic Bootstrapping hypothesis, tell us that children learn the meanings of predicates, in particular abstract ones, in large part *via* the sentence structure they occur in rather than the other way around (Gleitman, 1990, and considerable work following this).

Thus, learners need to rely on the underlying structure of sentences like (4) in order to figure out whether the main verb means something like 'seem' or something like 'claim,' but how do they first figure out the underlying structure? It is in this sense that the sentence strings in (3), (4), and (7) are syntactically opaque. I define syntactic opacity as follows.

(10)   A string is *syntactically opaque* if the underlying syntactic structure that generates the string cannot be determined unequivocally on the basis of the string and knowing only the grammatical categories of the words, without at least some lexical semantic knowledge of the main predicate.

Strictly speaking, all strings are opaque in this sense, until the lexical semantics of the main predicate is known. Even a string like that in (11) could be associated with various structures including, but not limited to, those in (11a–c).

(11)   NP V NP
        a. $[NP_{subj} [VP [V_{trans} NP_{obj}]]]$
        b. $[NP_{subj} [VP [V_{intrans}] NP_{loc}]]$
        c. $[NP_{subj} [VP [V_{ditrans} NP_{obj}] \emptyset_{ind.obj}]]$

However, much previous work on children's learning of verb argument structure has revealed that children are prone to making assumptions about these strings: a verb with one NP is assumed to be intransitive, a verb with two NPs transitive, and a verb with three NPs ditransitive (see Gleitman *et al.* 2005 for a good overview of this literature; though see Tomasello and Brooks (1998); Tomasello (2000) for a different view). This is precisely because, as noted at the beginning of this introduction, proximal words are typically semantically related – and related in particular ways.

But the constructions under consideration here are special, and especially opaque, because the adjacent NPs are not semantically related to the predicate in the usual way. And so the assumptions learners might make about the underlying structures of sentences like (11) will not apply straightforwardly to constructions with displacing predicates.

In addition to *tough*-constructions, raising verbs, and unaccusatives I will discuss the passive, which, unlike some of the other constructions

considered here, has been studied fairly extensively in the acquisition literature
(Slobin, 1966; Maratsos *et al.*, 1979; Borer and Wexler, 1987; Crain *et al.*,
1987; Lempert, 1989; Fox and Grodzinsky, 1998, i.a.). I include the passive
for two reasons. One is that passives involve a displaced subject: the syntactic
subject is understood as the semantic object, or patient of the verb's action. The
second is that in English certain passives are ambiguous between a verbal and
an adjectival reading, a fact which Borer and Wexler exploited in their account
of children's interpretation of the passive. Thus, a short passive, as in (12b), is
structurally ambiguous.

(12)   a. John was kicked by Sam. (verbal passive)

         b. The door was closed. (verbal or adjectival)

Nevertheless, important differences between the passive construction and
the others considered here will explain why the solution I propose for raising
verbs, *tough*-constructions, and unaccusatives actually does not extend to the
passive. Most pointedly, like in the example of PP-attachment above, the mean-
ing of the main predicate does not change radically depending on whether the
sentence has a passive or an active voice. *Kick* and *be kicked by* both denote
a kicking event. Thus, discovering the structure of a passive sentence requires
understanding that the subject is displaced, but it does not involve the task of
categorizing the main predicate as an inherently displacing predicate – that is,
the passive verb should not be assumed to have an auxiliary-like semantics.
While the passive will be discussed in Chapters 2 and 5 this construction will
not occupy a focal point in the overall discussion.

All of the constructions outlined above that involve a displaced subject,
with the exception of unaccusative verbs, have been argued to be difficult
for children to acquire (though see Babyonyshev *et al.* (2001) for claims that
young children represent unaccusatives as unergatives). Chomsky (1969) and
Cromer (1970), among others, argued this for *tough*-constructions, Hirsch and
Wexler (2007) have argued the same for raising-to-subject verbs, and Borer
and Wexler (1987), among many others, argued this for passives. To the extent
that children *can* correctly interpret such structures, for example non-reversible
passives, they are said to do so by relying on "real world knowledge" rather
than syntax. For example, Slobin (1966) showed that children responded more
quickly to non-reversible passives (*The flower is being watered by the girl*) than
reversible passives (*The cat is being chased by the dog*), presumably because
either dogs or cats can chase or be chased, but flowers do not water girls. What
I argue in this book is that children *do* use "real world knowledge," not as a

means of avoiding complex syntax, but rather as a means of *discovering* the complex syntax itself. The evidence comes from experiments of novel word learning, in which children use subject animacy to draw inferences about the argument structure properties of novel predicates.

In this book I revisit Carol Chomsky's original question through the lens of the advances that have been made in the fields of linguistic theory, psychology, language development, and computational modeling, with the goal of integrating cross-linguistic constraints and preferences on argument structure mapping into a theory of how structures involving displaced subjects are acquired. Chapter 2 provides an overview of some formal accounts of the opaque constructions under consideration here, including a description of how these constructions are analyzed under Minimalist approaches. The vocabulary of the Minimalist Program gives us a unified way of talking about raising, unaccusative, and passive constructions: these are predicates whose vP is considered "defective," allowing an NP argument to move out of their "weak" phase into the main clause. (The spirit of this unification is no different from previous incarnations of the theory, but the language of it is different.) One mechanism that has been proposed within Minimalism to account for passives and subject raising, namely "smuggling," has also been invoked in an account of *tough*-constructions (Hicks, 2009). My thesis is not contingent on any particular syntactic framework or formalism, but *some* kind of formalism is required in order to see why the acquisition question I'm addressing is a matter of acquiring syntax.

Chapter 3 then looks at how animacy is realized in grammar along a number of dimensions: how animacy is grammaticalized in various languages, how it relates to thematic roles, and how, in turn, thematic roles relate to argument structure. The emphasis in this chapter is on typological patterns; that is, how animacy surfaces in the world's languages, and how displacing predicates work in different languages as well. Across genetically diverse languages we can observe two rather clear and consistent patterns. One is that languages tend to organize animacy distinctions between more animate and less animate entities according to a hierarchy, according to which humans are the "most animate," followed by non-human animals, followed by inanimates. Though there is diversity in the number of distinctions made in the hierarchy, and where dividing lines are drawn, the hierarchy itself is robust: we do not find languages, for example, which treat humans and inanimates alike to the exclusion of animals.

The second consistent pattern is that languages prefer non-displaced (i.e. canonical) subjects to be animate but rather liberally allow displaced