Children's forms as derivational steps: external evidence for a new synthetic compound structure

Jessica L. Gamache and Cristina Schmitt
Michigan State University

1 Introduction

Children produce a fair amount of non-adult forms when acquiring synthetic compounds. These "errors" can inform not only how children master compound structures, but also how the syntax of compounds actually works. In this paper we assume a strong version of the Continuity Hypothesis (Pinker 1984), exploring the idea that the non-adult forms that children produce during development reflect the derivational steps of the adult structure. That is, any non-adult form produced by children must be accounted for by the internal structure of synthetic compounds.

To test this hypothesis, two things are needed: (i) child data that is representative of the development of synthetic compounds and (ii) a syntactic structure. This paper is divided accordingly. In the first part of the paper we present the results of two new studies[1], designed to probe into conflicting results from Clark et al. (1986) and Nagpal and Nicoladis (2009). In the first study, using an elicited production task, we replicate the forms discussed in Clark et al. (1986), but fail to replicate the forms in Nagpal and Nicoladis (2009). Essentially, we find children producing *drive-truck* and *driver-truck*, but not *truck-drive*. In the second study, a repetition task, we again fail to replicate the patterns discussed by Nagpal and Nicoladis.

In the second part of this paper we ask, given that these forms (*drive-truck, driver-truck*) are representative of children's development, what structure would account for them? Assuming the strong Continuity Hypothesis, each of the children's forms should reflect a step in the adult derivation. In this section, previous proposals for the structure of synthetic compounds are evaluated against the children's form. Finding that none of the proposals provide adequate coverage of the developmental stages, we propose a new structure to account for the acquisition data, as well as other facts about complex verbal elements in English.

2 Acquisition

2.1 Acquisition Background

Using an elicited production task, Clark et al. (1986) examined the development of synthetic compounds in children ages three to six. Children were asked to name people and machines with the prompt "This is a man who drives trucks. What could we call a man who drives trucks?" The children were told that there were four good

[1]We would like to thank Alan Munn, Drew Boboltz, Kaitlyn Ayres, Emalee Manns and the other members of the MSU Language Acquisition Lab for their help with this project. We'd also like to thank the Eastminster Development Center, Radmoor Montessori, Stepping Stones Montessori, St. Martha's School, and Edgewood Elementary for continually supporting our research.

names for this man, *truck-driver, truck-driving-man, driving-man,* and *driver*. The results showed that it was not until age five that most children were proficient in producing target *truck-driver* forms. Children's non-adult forms followed a number of patterns which decreased with age: (i) object-less forms (*drive-man*), (ii) *–er*-less forms (*drive-truck*), and (iii) forms with *–er* but with sentential ordering (*driver-truck*). Clark et al. also found that children were better at producing target forms when the external argument was animate, and that younger children tended to use both *–er* and a lexical external argument (e.g., *driver-man*). They argued that the slow development of synthetic compounds is due to children's tendency to preserve the canonical sentential word order, as in the forms *drive-truck* and *driver-truck*.

In a recent replication study, Nagpal and Nicoladis (2009) argued that Clark et al.'s stages are not robust enough to account for the data. Instead, Nagpal and Nicoladis found that children produce forms such as *driving-man* in numbers much higher than Clark et al. reported, leading them to suggest that the acquisition of synthetic compounds is delayed by exposure to other forms in the input. These structures, such as *running-man*, are arguably easier for the child since *–ing* is among the first morphemes to be learned (Brown, 1973). Despite having differing results, both sets of researchers attribute their results to interference from the input.

2.2 Shortcomings of Previous Studies

There are two main methodological differences that could cause conflicting results in the previous studies: the prompts and the verb selection. With respect to the prompts, Clark et al. gave four different forms as possible targets (*truck-driver, truck-driving-man, driving-man, driver*), which may have led to confusion for younger children. Nagpal and Nicoladis, on the other hand, gave only one target form (*truck-driver*), but used *–ing* forms in their prompts (e.g., *This man is driving a truck...*), which may have led to the prevalence of *–ing* forms in their results. *–ing* forms, although given as possible target forms in Clark et al.'s study (*truck-driving-man* and *driving-man*), were not nearly as frequent in Clark et al.'s results.

Another possible source of variation in the previous studies involves the selection of verbs used in the studies. Verbs such as *break*, *wave* and *fly*, were included in the previous studies. Perhaps mixing verbs that do not always include an external argument (*break*) and an internal argument (*fly*) may have added some noise to the results.

The next section will introduce two different experimental situations to examine the development of compounds: an elicited production task (Study 1) and a repetition task (Study 2). In these studies we used only one possible target form (*truck-driver*) and controlled for verb types so as to not mix verb types.

3 New acquisition studies

3.1 Study 1: Elicited Production

The first study is a near replication of both Clark et al. (1986) and Nagpal and Nicoladis (2009). The goal of this study is to (i) remove *–ing* from the prompts to avoid priming, (ii) control for verb type, in order to include only mandatorily transitive verbs, and (iii) streamline priming by giving children only one target form,

truck-driver.

3.1.1 Methods

Participants

Sixty-four native English-speaking children in the East Lansing, MI area took part in this study. The children were broken down into three groups based on their ages (Table 1).

	3-year-olds	4-year-olds	5-year-olds
N	16	16	16
Age Range	3;2 - 4;0	4;1 - 5;0	5;1 - 6;3
Mean (*SD*)	3;6 (0;3)	4;7 (0;4)	5;8 (0;4)

Table 1: Study 1 age characteristics by group.

Materials

To elicit synthetic compounds, children were shown colored pictures on a computer screen containing either a person or machine performing a particular action. The nouns and verbs for this study were taken from the MacArthur's CDI (Dale & Fenson 1996). All words chosen were produced by 70% of the children in the inventory at age 2;6. Each child received a total of 12 items, six with a human agent and six with a machine instrument. The purpose of using animate and inanimate external arguments was to explore the effect of animacy in Clark et al.'s results.

The 12 target verb-noun combinations (*chair-builder, baby-holder, dog-finder, boy-feeder, apple-cutter, cat-chaser, car-fixer, bug-thrower, girl-tickler, raisin-pourer, mailman-catcher,* and *bunny-carrier*) were randomized to create a single ordering. Each verb-noun combination was represented either as a person or a machine, with animacy counterbalanced across two lists. The ordering of animacy was pseudorandom: no more than three humans or machines could be adjacent. The two lists were then inverted to balance out ordering effects for a total of four different lists. The children were randomly assigned to one of the lists.

Procedure

Children were tested in preschools in the East Lansing area. Each child was tested individually with the first author and a coder.

In order to elicit the best results possible, the children were told a story that framed the experiment as if they were helping a puppet. In this story, the puppet refused to talk to adults, so he needed the children to help in coming up with names for the people and machines in the pictures. The story facilitated responses, with pilot data showing that children were less willing to give answers directly to the experimenter than they were to the puppet.

The structure of the study was based loosely on the Clark et al. (1986) study. Pilot data showed that giving children four different alternatives in the practice trials (*truck-driver, truck-driving-man, driving-man, driver*) appeared to cause confusion, so a single form (*truck-driver*) was used instead. The study began with the experimenter going through two practice items, giving examples to the children of what a "good name" would be. The two practice examples had one human (a *book-reader*)

and one machine (a *can-opener*). The experiment would say, "This is a man who always reads books. We could call a man who always reads books a *book reader*. Does this sound like a good name to you?" The same prompt was used for the *can-opener* example. The children were then asked to help out with names for the rest of the pictures.

The structure of the prompts was much like that of the two practice examples, "This man always builds chairs. What could we call a man who always builds chairs?" If the children responded with a non-compound answer, they were asked for a second name and that was recorded as their response.

Coding production data

The children's responses were coded schematically as the external argument (S), the verb (V), the internal argument (O), and *–er* or *–ing*. Any response that contained both the external argument (S) and the *–er* affix was coded as a case of doubled morphology (two instances of the external argument). Since there was a single form given to the children in the practice examples (*truck-driver*), only those were coded as "target". Forms such as *driver*, however, are also considered to be adult-like.

3.1.2 Production Results

The overall distribution of responses is shown in Table 2. For this table, instances of doubled morphology are collapsed into the *–er* responses (e.g., *thrower-man* and *thrower* would both be counted as V*er*.) Doubled morphology will be examined separately below.

Form	Example	3-year-olds		4-year-olds		5-year-olds	
OV*er*	*truck-driver*	.23	(45)	.38	(73)	.72	(139)
V*er*O	*driver-truck*	.07	(13)	.07	(13)	.03	(6)
VO	*drive-truck*	.04	(7)	.09	(18)	.04	(8)
OV	*truck-drive*	.01	(1)	.00	(0)	.01	(1)
V*er*	*driver*	.07	(13)	.07	(13)	.01	(2)
VS	*drive-man*	.01	(2)	.01	(1)	.03	(5)
OS	*truck-man*	.07	(14)	.04	(8)	.05	(9)
OV*ing*S	*truck-driving-man*	.01	(1)	.00	(0)	.01	(1)
V*ing*S	*driving-man*	.02	(3)	.02	(3)	.00	(0)
V*ing*O	*driving-truck*	.01	(1)	.00	(0)	.00	(0)
	Total *N*	192		192		192	

Table 2: Distribution of production forms (proportions and raw *N*s) by age group.

Across all age groups, children produced a certain amount of target forms, with the highest proportion (.72) produced by the 5-year-olds. In keeping with Clark et al.'s results, children produced many forms with sentential VO word order (*drive-truck*), some with *–er* morphology (*driver-truck*). Importantly, children basically did not produce OV *truck-drive* word order without *–er* (2 out of 576 total responses). Contra Nagpal and Nicoladis's findings, *–ing* forms are also nearly absent from the children's productions.

From these results, a number of patterns emerge with respect to the acquisition of *–er*. First, the younger children tend to exhibit higher levels of morphology doubling. Around 40% of 3-year-olds *–er* responses also contained a lexical representation of the external argument (e.g., *driver-man*) (Figure 1). Figure 1 shows that the number of *–er* forms produced that are neither doubled nor target (Non-doubled *–er*) is very low across all age groups. This suggests that as soon as children are learning *–er*, they are using it in target forms. In other words, there is not a lag between when children learn *–er* and when they begin to produce *truck-driver*.

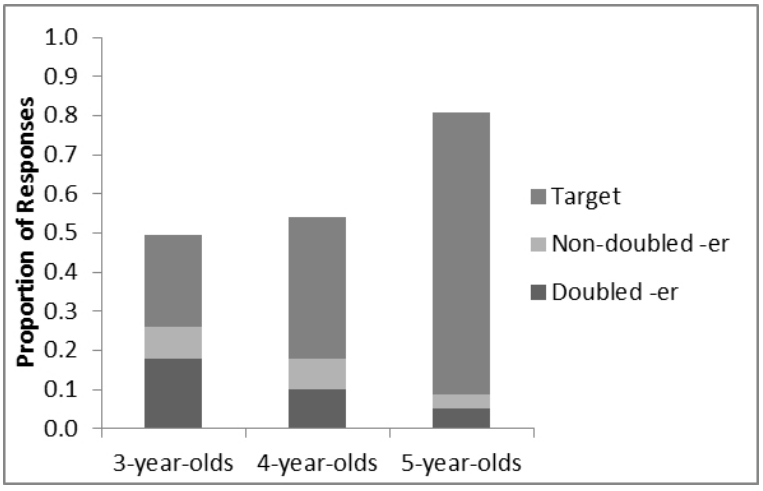

Figure 1: Proportion of forms with doubled *–er*, non-doubled *–er*, and target *–er*.

Along with the tendency to double *–er*, there appears be a tendency to first equate *–er* to only human external arguments. The younger children show a preference for using *–er* with animate external arguments (humans), as opposed to inanimate external arguments (machines) (Figure 2), although even the difference for 3-year-olds does not reach significance ($p = .13$).

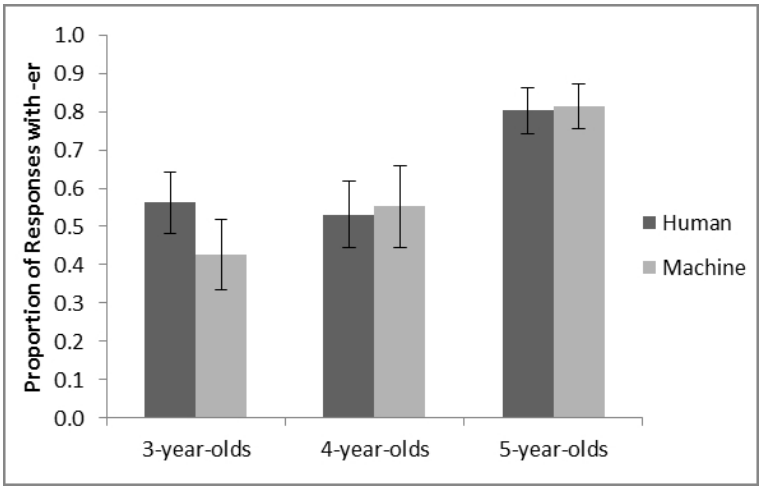

Figure 2: Proportion of *–er* responses by Age Group and Animacy of the external argument.

3.1.3 Discussion

The results replicated the types of forms reported in the Clark et al. study. The younger children produced the lowest proportion of target responses, using other forms such as *driver-truck*, *drive-truck*, and *truck-man*, while the older children predominantly used the target *truck-driver* form. Children almost never produced OV *truck-drive* forms (2 of 576 total responses), showing a strong tendency to maintain sentential word order until *–er* is a part of the form.

Also rare in the children's productions were *–ing* forms, with the total number representing less than 3% of responses (16 of 576 total responses). This suggests that the prompts used by Nagpal and Nicoladis may have led to the high instances of *–ing* in their study, not the interference of other forms during the acquisition process.

The patterns related to the acquisition of *–er* show that the morpheme is acquired relatively late by children, and not without effort. The tendency that younger children showed to equate *–er* more often to a human agent than a machine instrument suggests that children may start with a narrower representation of the morpheme. Moreover, children's tendency to double the *–er* with an overt agent suggests that, at first, *–er* is not capable of fully representing the external argument. However, once the morpheme is acquired, there is little delay in children producing target forms, regardless of their age, shown by the low number of non-target non-doubled forms that include *–er*. Since we know that children can produce noun-noun compounds at age 2 (Clark 1982), the problem is not how to create compounds, but rather mastering the syntax and semantics of *–er*. If *–er* is the missing link, the acquisition of synthetic compounds should proceed instantaneously once *–er* is acquired, and we see that it does in the results presented here.

3.2 Study 2: Repetition

The goal of the second study is to further examine Nagpal and Nicoladis's claim that *–ing* is easier for children than *–er*, possibly leading to the interference of *–ing* forms in the input.

3.2.1 Methods

Participants

38 native English-speaking children in the East Lansing, MI area took part in this study. The children were broken down into three groups based on their ages (Table 3).

	3-year-olds	4-year-olds	5-year-olds
N	13	13	12
Age Range	3;2 - 4;0	4;1 - 5;0	5;1 - 6;3
Mean (*SD*)	3;7 (0;3)	4;6 (0;4)	5;8 (0;4)

Table 3: Study 2 age characteristics by group.

Materials

To elicit responses, children were shown colored pictures on a computer screen containing either a person or machine performing a particular action. All words

included were produced by 70% of the children in the MacArthur's CDI (Dale & Fenson 1996) by age 2;6. Each child received a total of 16 items, eight with a human agent and eight with a machine instrument. Four of the items were well-known synthetic compounds (*can-opener, fire-fighter, bus-driver,* and *coffee-maker*).

The 16 target verb-noun combinations were randomized to create a single ordering. Each verb-noun combination was represented either as a person or a machine, with animacy counterbalanced across two lists. The ordering of the animacy was pseudorandom: no more than three people or machines could be adjacent. The two lists were then inverted to balance out ordering effects for a total of four different list. The children were randomly assigned to one of the lists.

Procedure

Children were tested in preschools in the East Lansing area. Each child was tested individually by the first author and a coder. The type of form given in the prompts in this study was a between-subjects variable, with roughly half of the children receiving *truck-driving-man* forms (n = 20) and half *truck-driver* forms (n = 18).

In a similar manner to the production study, the children were told that the purpose of the study was to help out a puppet with naming things. The children were given a prompt and asked to repeat back just the name for the person or machine. For example, "Because this man always builds chairs, Mary called him a *chair-builder*. What did Mary call him?" should lead to the response of *chair-builder*. If the children responded with a non-compound answer they were asked for a second name and that was recorded as their response.

3.2.2 Repetition Results

Figure 3 shows that children did generally well on this task. Importantly, when children did not produce target forms, they also did not produce *truck-drive* forms.

To examine whether children were doing better on one of the forms, a 2-way (3 Age Group (3-year-olds, 4-year-olds, 5-year-olds) X 2 Repetition Type (*truck-driver* forms, *truck-driving-man* forms)) ANOVA on the proportion of correct responses was conducted. Age Group was found to be significant ($F(2, 32) = 7.542$, p = .002), reflecting the higher scores as children get older. Repetition Type was marginally significant ($F(1, 32) = 3.486$, $p = .071$), reflecting the generally higher scores for those children who received the *truck-driver* forms. The interaction between Age Group and Repetition Type was not significant ($p = .630$). The lack of this interaction shows that the children are not significantly better at one form than another at any given time. Independent samples t-tests for each age group partially confirm this, revealing no significant differences between the two Repetition Types (3 year-olds, $p = .279$; 4-year-olds, $p = .068$; 5-year-olds, $p = .103$).

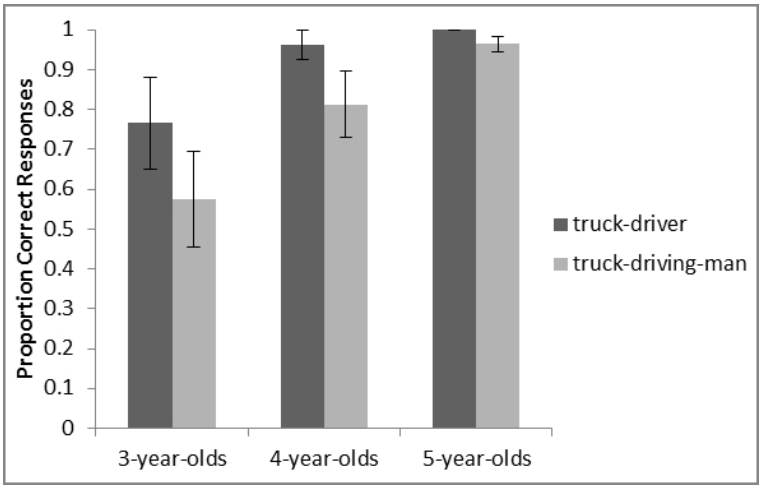

Figure 3: Proportion of Correct Responses by Age Group and Repetition Type.

3.2.3 Discussion

Children in the *truck-driver* forms group performed slightly better than the children in the *truck-driving-man* forms group, regardless of age group. While this is expected based on the number of morphemes being asked to repeat, it may be unexpected based on Nagpal and Nicoladis's claim that *–ing* forms are easier for children and therefore interfere with their acquisition of synthetic compounds. These results further support the idea that the differences in the results of Clark et al. and Nagpal and Nicoladis's studies may have stemmed from methodological differences involving the phrasing of prompts.

3.3 General Discussion

The results of these two studies show that Clark et al.'s stages are representative of children's development of synthetic compounds, and that controlling for morphology in the prompts basically eliminates *–ing* forms. The patterns that show up during the acquisition of the morpheme *–er* (a preference for human agents and the doubling of *–er* with a lexical external argument) suggest that the learning of *–er* is an effortful process that requires children to learn a number of different aspects about the morpheme: (i) *–er* can represent both humans (agents) and machines (instruments) and (ii) *–er* takes the place of the external argument. What the children seem to know implicitly, however, is that *–er* allows the internal argument to move. We see from the lack of OV forms that *–er* must be merged into the structure before the internal argument can move. This rigidity in children's responses is compatible with the idea that children's forms are partial structures. Without the guiding principles of syntax, it would be difficult to explain why children have a near complete absence of some forms (*truck-drive*), but regularly produce some other forms (*driver-truck*).

4 The Structure of Synthetic Compounds

In this section we examine attested proposals for synthetic compounds structures based on three measures: (i) the acquisition forms, (ii) the lack of productive *truck-drive*-type verbs in English, and (iii) the verb-theme relationship exhibited by the verb and internal argument in synthetic compounds. The child forms that need to be accounted for by the syntactic structure of synthetic compounds are *drive-truck* and *driver-truck*. The lack of *truck-drive* forms across all previous studies suggests that this should not be a substructure in the derivation.

The structures in (1) and (2) form the basis of most structures that have been proposed for synthetic compounds.

(1) Selkrik 1982, Lieber 1992

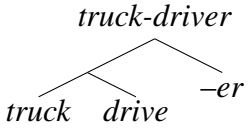

(2) Lieber 1983, Booij 1988, 2009

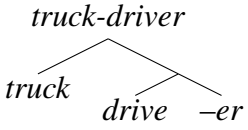

The structure in (1) maintains the sisterhood relation between the verb and its internal argument as it would be in a clausal VP; the verb *drive* takes *truck*, the theme, as its complement. This structure, however, fails to explain why *truck-drive* constructions are largely unattested in English, both as verbs (e.g., *He truck drives for a living.*) and also in child acquisition errors. (1) also cannot account for *drive-truck* and *driver-truck* forms which are attested during development.

The structure in (2) does not maintain the sisterhood relation between the verb and its internal argument. It can account for the fact that *truck-drive* is not attested either as a verb or an error in children's speech. It fails, however, to account for both the *drive-truck* and *driver-truck* forms during acquisition.

Recently Harley (2009) has proposed a structure for synthetic compounds (3) based on Distributed Morphology (Halle & Marantz 1993). In this structure a nominal element is first created by merging *truck* with a nominal head. This nominal element *truck* then merges with the root $\sqrt{DRIVE}$ to create a $\sqrt{}$P, *drive-truck*. Next, *truck* incorporates into *drive*, creating the complex head *truck-drive*, which then incorporates into the nominalizing element *–er*, leading to *truck-driver*.

(3)

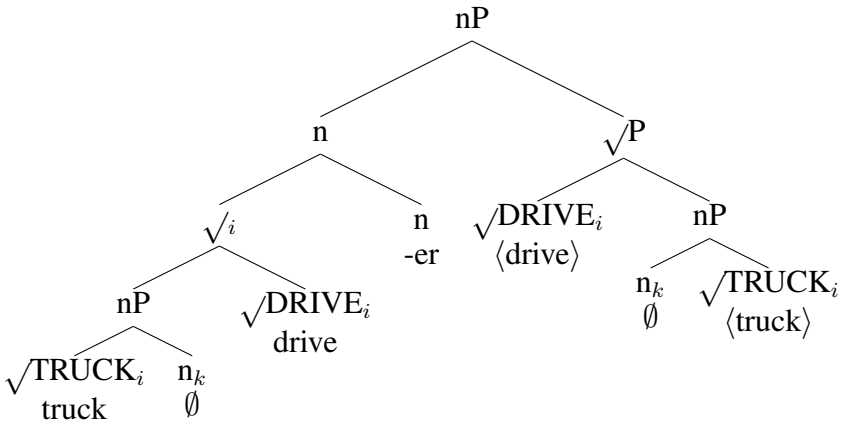

This structure accounts for the verb-theme relationship and the *drive-truck* acquisition stage. It fails, however, to capture the lack of *truck-drive* forms in English, as

this would be predicted as an intermediary step in this structure. Harley's structure also cannot account for the *driver-truck* error found in children's productions.

Given the current proposals, coverage of the empirical facts is summarized in Table 4. As no single structure is capable of accounting for all of the facts, a new structure will be presented in Section 5.

	[[truck drive] er]	[truck [driver]]	Harley (2009)
Acquisition data	no	no	no
**truck-drive* verbs	no	yes	no
Sisterhood relation	yes	no	yes

Table 4: Coverage of previous synthetic compound structures.

5 Proposed Linguistic Analysis

Our structure begins with the merger of the verb and internal argument (4). This first step gives us both the form produced by children (*drive-truck*), as well as the thematic relationship between the verb and internal argument.

(4)

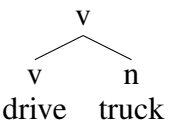

Next, the *–er* morpheme is merged into the structure (5), causing the incorporation of the verb (6), which reflects the *driver-truck* form produced by children. This step in the derivation is where this structure is critically different than Harley's: while her structure incorporates the internal argument into the verb before merging *–er*, the structure proposed here waits to move the internal argument until after *–er* has merged into the structure. This ordering is supported by the lack of OV errors in child language, and also by English's lack of complex OV verbs.

(5)

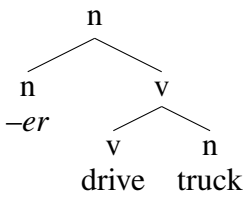

(6)

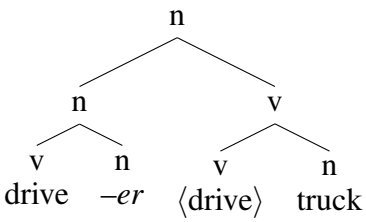

Lastly, with the merger of *–er* completed, the internal argument moves to the specifier of the nominal *driver*, in line with Holmberg's Generalization (Holmberg 1986), which states that head movement, in this case the movement of the verb,

must take place before phrasal movement, the movement of the internal argument. This analysis is also in line with Roeper's analysis of leftward movement in morphology (1999, 2004).

(7)

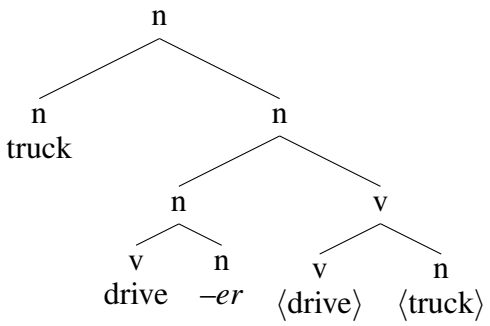

This structure successfully accounts for all empirical facts and begins to explain why children's acquisition of *–er* is so critical: before *–er* is merged, no movement of the internal argument can take place. After *–er* is merged, however, the final stage of the derivation can proceed. Under this proposal there is no room for *truck-drive* forms without *–er*.

6 Conclusion

In this paper we asked, given an accurate picture of the errors that children produce during the acquisition of synthetic compounds, what can we learn about the structure? Two previous studies, Clark et al. (1986) and Nagpal and Nicoladis (2009), reported results of children's productions during the acquisition of synthetic compounds, coming to different conclusions on both the forms produced as well as the source of children's errors. To tease apart these differences, we presented the results of two studies. In the production study, we replicated the forms reported by Clark et al. Namely, children produced *drive-truck* and *driver-truck* forms before producing target *truck-driver* synthetic compound forms, but never produced *truck-drive* forms. Our results failed, however, to produce *–ing* forms in the numbers found in Nagpal and Nicoladis's study, suggesting that the differing results arose form differences in methodologies. In the repetition study, we showed that, contra Nagpal and Nicoladis's claim, *–ing* forms are likely not interfering in the acquisition of synthetic compound forms. Instead, children are likely better at less morphologically complex forms, such as *running-man* vs. *truck-driver*. While the children do generally well in the repetition study, when they do make errors, they are not of the *truck-drive* type.

We examined a number of attested structures for synthetic compounds against children's errors to determine how well they could predict the children's forms (*drive-truck*, *driver-truck*, but not *truck-drive*), as well as the lack of **truck-drive* verbs in English and sisterhood relationship between the verb and internal argument. Finding that the previous proposals did not provide adequate coverage of these, we presented a new structure for synthetic compounds. Our structure reflects not only children's production errors, but also why acquisition proceeds quickly after *–er* is acquired; while children understand compounding early, the syntax and

semantics of *–er* must be acquired before the internal argument of synthetic compounds can move, leading to the target *truck-driver* form. This suggests that it is not sentential word order nor interference from other structures that is delaying acquisition. Rather, the late acquisition of synthetic compounds is due solely to the late acquisition of the morpheme *–er*.

References

Booij, Geert. 1988. The relation between inheritance and argument linking: Deverbal nouns in Dutch. In *Morphology and Modularity*, ed. by Martin Everaert, Arnold Evers, Riny Huybregts, & Mieke Trommelen, 57–74. Dordrecht, Netherlands: Foris.

——. 2009. Compounding and construction morphology. In *The Oxford Handbook of Compounding*, ed. by Rochelle Lieber & Pavol Stekauer, 201–216. New York, NY: Oxford University Press.

Clark, Eve. 1982. The young word-maker: A case study of innovation in the child's lexicon. In *Language Acquisition: The State of the Art*, ed. by Eric Wanner & Lila Gleitman, 390–425. Camrbidge, UK: Cambridge University Press.

——, Barbara Hecht, & Randa Mulford. 1986. Coining complex compounds in English: Affixes and word order in acquisition. *Linguistics* 24.7–29.

Dale, Phillip, & Larry Fenson. 1996. Lexical development norms for young children. *Behavior Research Methods, Instruments, and Computers* 28.125–127.

Halle, Morris, & Alec Marantz. 1993. Distributed morphology and the pieces of inflection. In *The View from Building 20*, ed. by Kenneth Hale & Samuel Keyser, 111–176. Cambridge, MA: MIT Press.

Harley, Heidi. 2009. Compounding in distributed morphology. In *The Oxford Handbook of Compounding*, ed. by Rochelle Lieber & Pavol Stekauer, 129–144. New York, NY: Oxford University Press.

Holmberg, Anders, 1986. *Word Order and Syntactic Features in the Scandinavian Languages and English*. University of Stockholm dissertation.

Lieber, Rochelle. 1983. Argument linking and compounding in English. *Linguistic Inquiry* 14.251–286.

——. 1992. *Deconstructing Morphology*. Chicago/London: The University of Chicago Press.

Nagpal, Jaya, & Elena Nicoladis. 2009. Why are noun-verb-er compounds so difficult for English-speaking children? *Mental Lexicon* 4.276–301.

Pinker, Steven. 1984. *Language Learnability and Language Development*. Cambridge, MA: Harvard University Press.

Roeper, Thomas. 1999. Leftward movement in morphology. In *Morphology and Syntax Cycle Two*, ed. by Vivian Lin, Cornelia Krause, Benjamin Bruening, & Karlos Arregi, volume 34, 35–67. Oxford University Press.

——. 2004. Nominalizations: How a marignal construction reveals primary principles. In *Handbook of Morphology*, ed. by Rochelle Lieber & Pavol Stekauer. Dordrecht, Netherlands: Kluwer.

Selkirk, Elisabeth. 1982. *The Syntax of Words*. Cambridge, MA: MIT Press.