

On the Acquisition of Principle B

Paul Elbourne

Chien and Wexler (1990) reported that children obeyed Principle B of binding theory when the antecedent was a quantifier but not when the antecedent was referential. This was argued by Grodzinsky and Reinhart (1993) to support Reinhart's (1983) theory according to which Principle B affects only bound pronouns. Since then, other studies have supported the asymmetry between referential and quantifier antecedents. This article, however, argues that previously unremarked experimental factors lessen the force of all these studies, and it points to other relevant experiments that seem to show that children do not obey Principle B at all. It reviews previously offered theories on the acquisition of Principle B that are compatible with the latter view of the facts.

Keywords: acquisition, binding theory, Principle B, Middle English, Rule I

1 Introduction

Grodzinsky and Reinhart (1993) claim that findings by Chien and Wexler (1990) concerning children's knowledge of Principle B support Reinhart's (1983) view of binding theory over the model set out in Chomsky 1981 and related work. In this article, I question that conclusion. I will argue that experimental factors undermine the widely accepted account of the facts set out by Chien and Wexler and others, and I will point to other experiments that do not obtain comparable results.

The relevant results Chien and Wexler (1990) reported were as follows. Consider a scenario involving Mama Bear and Goldilocks, where Mama Bear is not touching Goldilocks but is touching herself. In Chien and Wexler's experiments, 5- to 6-year-olds incorrectly answered "yes" to question (1) 50% of the time.

(1) Is Mama Bear touching her?

By contrast, in a scenario where there are three bears and all three are not touching Goldilocks but are touching themselves, 5- to 6-year-olds incorrectly answered "yes" to (2) only 16% of the time.

(2) Is every bear touching her?

Considered in the light of Chomsky's (1981) binding theory, these findings might seem contradictory. The first fact might be taken to suggest that children do not know Principle B, which in

I am very grateful to Karlos Arregi, Dianne Bradley, Marcel den Dikken, Janet Dean Fodor, Danny Fox, Ted Gibson, Irene Heim, Michal Starke, Rosalind Thornton, Sharon Utakis, Ken Wexler, the anonymous *LJ* reviewers, and the audience at CUNY's Syntax Supper for helpful discussion of this material. Naturally all errors are my own.

adult speech would not allow the test sentence to express the proposition that Mama Bear is touching herself. But the second fact seems to point to a fairly robust knowledge of Principle B. The contrast is evidently between referential and quantificational antecedents for pronouns in Principle B configurations, a phenomenon I will refer to as the *quantificational asymmetry* (QA).

In Reinhart's version of binding theory, Principle B applies only to pronouns interpreted as bound variables. Thus, it will apply to the pronoun in (2), but not necessarily to the pronoun in (1), since the latter could be referential. It looks, then, as if Reinhart's Principle B cuts the empirical pie in exactly the same way as Chien and Wexler's data.

Given that the pronoun in an example like (1) could be referential, what is there in the adult grammar to prevent (3) from meaning the same as (4), according to Reinhart's proposals?

(3) Mama Bear is touching her.

(4) Mama Bear is touching herself.

The answer is that *her* in (3) is prevented from being referential by a special rule called *Rule I*, which is stated in (5).¹

(5) *Rule I*

NP A cannot corefer with NP B if replacing A with C, C a variable A-bound by B, yields an indistinguishable interpretation.

This rule applies to (3) as follows. We consider the result of having *her* simply refer to Mama Bear. That would be the proposition that Mama Bear is touching Mama Bear. We then work out what would happen if we interpreted *her* as a bound variable bound by *Mama Bear*, or rather the λ -abstractor left by moving *Mama Bear* (see Heim 1993, Heim and Kratzer 1998). The result is that we obtain a predicate $[\lambda x. x \text{ is touching } x]$, and this is predicated of Mama Bear. Again, we end up with the proposition that Mama Bear is touching Mama Bear. The two interpretations are indistinguishable, so *her* cannot be interpreted as referential. It must be interpreted as bound, then; but this in turn makes it fall under Principle B, and the sentence is correctly rendered ungrammatical under the interpretation whereby it means the same as (4).

Grodzinsky and Reinhart (1993) propose that the QA in child language comes about as follows. Children know both Principle B and Rule I innately, as parts of Universal Grammar. Their knowledge of Principle B prevents them from interpreting (2) as asking whether every bear is touching herself. But in order to apply Rule I in a case like (1) or (3), they have to keep in mind the representation of it with a referential pronoun and the representation with a bound pronoun; and they must then compare the interpretations which these representations produce. The processing load, according to Grodzinsky and Reinhart, is just too much for young children,

¹ The formulation here is from Grodzinsky and Reinhart 1993:70 and differs slightly from the original version in Reinhart 1983. Heim (1993), however, has shown that restricting the ban effected by the rule to cases of coreference is mistaken and that a broader formulation is necessary. See Reinhart 1997, Fox 2000:109–137, and Elbourne 2002:172–183, for further discussion. This issue is orthogonal to the subject of this article, however, and will not be dealt with here.

and they give up and guess the answer—hence the 50% score.² With a pronoun that has a quantificational antecedent, however, there is obviously no chance of a coreferential reading and no attempt is made to apply Rule I.

One might object at this point that, since Reinhart's binding theory has no Principle C and derives the effects of this principle too from Rule I, children will be predicted to show a similar 50% success rate in Principle C contexts.³ This is not the case, however. Children display a robust knowledge of Principle C (Crain and McKee 1986, Crain 1991). This looks, then, like a serious objection to Grodzinsky and Reinhart's thesis. However, the objection relies on the premise that Reinhart's binding theory has to be accepted or rejected in its entirety, which is not the case. It is possible to suppose that some version of Rule I governs Principle B contexts, and that a separate Principle C exists. In fact, Heim (1993) and Fox (2000) hold positions like this. We need not conclude, then, that children's good performance on Principle C completely undermines Grodzinsky and Reinhart's thesis, although their performance does seem to require it to be revised.

Another variant of Grodzinsky and Reinhart's approach has been proposed by Thornton and Wexler (1999), who also report more experimental results replicating and extending Chien and Wexler's (1990) original data. Thornton and Wexler suggest that it is not the case that a processing overload makes children violate Rule I; rather, the children obey Rule I by construing the relevant sentences in such a way that there is a difference between the interpretation on which the pronoun is referential and the interpretation on which it is bound. This phenomenon is also active in the adult grammar, accounting, for example, for the difference in grammaticality between the following two examples. In most contexts, *She is praising her* could not mean that Mary is praising herself, as we see in (6), for example.

(6) (What's Mary saying?) *She_i is praising her_i.

It is noteworthy, however, that there are certain conditions in which cases like (6) are quite grammatical—for example, when there is a debate about identity. Imagine that Mary is running for political office and that we are wondering whether the speaker on the soapbox over there, whom we cannot see well, is Mary. The dialogue in (7) would then be natural.

(7) A: Is that speaker Mary?

B: She's praising her to the skies. It must be Mary.

² An *LI* reviewer suggests that this account predicts that the 50% success rate will obtain within subjects as well as between subjects. This is not clear to me. If children really do guess in these circumstances, could they not just pick one answer the first time they are confused and then stick with it when they experience similar confusion in subsequent trials? However, the same reviewer also points out that Fritzley and Lee (2003) have claimed that children in the relevant age range are more likely to answer "no" than "yes" when they are confused. If this claim is correct, then Grodzinsky and Reinhart's explanation for the QA is in trouble, since the prediction should actually be that there will be a high rejection rate in cases like (3), contrary to fact. More work is needed on this point, which is potentially important. Note, however, that Thornton and Wexler's (1999) explanation of the QA, explained below, does not rely on the subjects' being confused, and would not be vulnerable to this criticism.

³ I am grateful to an anonymous *LI* reviewer for alerting me to this point, which surprisingly has not been made in the previous literature on this topic.

Person B is convinced (correctly, let us say) that the speaker is Mary, and the grounds for this, as expressed in the first sentence of B's utterance, is that the speaker is praising Mary to the skies. It seems as if *She's praising her to the skies* should be a Principle B violation, then, but instead it is quite grammatical. The crucial fact seems to be that, although both pronouns end up referring to Mary, they get there by different routes, as it were: informally, the first one refers to the speaker over there, who happens to be Mary, and the second refers directly to Mary. Thus, the interpretation of the sentence on which both pronouns are taken to be referential is different from the interpretation that would result if the second one were taken to be bound: if *her* were bound, the sentence would simply mean that the speaker over there is praising herself to the skies, and Mary would not come into it. (The formal details are irrelevant for present purposes; see Heim 1993.) Thornton and Wexler's hypothesis is that children are particularly good at construing referential pronouns in ways that will produce different interpretations for the coreferential and bound variable versions of a sentence. Thus, in answering "yes" to *Is Mama Bear touching her?*, they could be obeying Rule I and Principle B (the latter vacuously) by construing *her* in an imaginative way.⁴

These, then, are the two variants of the approach that explains the QA by means of Reinhart's binding theory. In this article, however, I wish to cast doubt on the robustness of the QA. In section 2, I will examine Chien and Wexler's (1990) experiments and others apparently showing this effect and will argue that it may have been produced by experimental factors. Some of these factors are systematic and apply to more than one study. In section 3, I will describe some comparatively neglected experiments that seem to show that children do not display knowledge of Principle B in any environment, whether or not the pronoun would have to be bound. Some of the latter experiments do not appear to suffer from any relevant methodological flaws. We are left, then, with the distinct possibility that children may not display knowledge of Principle B in their earliest years. This raises a knotty problem for acquisition. Do they really not know it innately? How can they acquire it if they do not? If they do know it innately, what factors interfere with their displaying knowledge of it when they are still comparatively young? In section 4, I will briefly outline the hypotheses I know of that address this problem.

2 Experiments That Show the Quantificational Asymmetry

To my knowledge, there are five sets of experiments that have found the QA for English: subsets of those reported in Chien and Wexler 1990, Thornton 1990, Boster 1991, Avrutin and Thornton 1994, and Thornton and Wexler 1999.⁵ In the interests of clarity, I will first examine the relevant

⁴ It is obviously necessary to ensure that bound variable pronouns are not amenable to the same treatment. I think actually that there is room for debate here, but this is not the topic of this article. See Thornton and Wexler 1999:103–104 for their view.

⁵ There are also two sets that have allegedly found it for other languages (Avrutin and Wexler 1992 for Russian and Philip and Coopmans 1995 for Dutch), but I omit detailed consideration of these in the interests of putting a natural limit on an already unwieldy project. Let me just note briefly that in Avrutin and Wexler's (1992) study, there was no large difference between the percentage of incorrect acceptances of a pronoun with a local referential antecedent (52%) and the percentage of incorrect acceptances of a pronoun with a local DP antecedent involving the word for 'every'

experiments of Thornton and Wexler 1999 (section 2.1), since one of these displays in a particularly clear form the main experimental factor that I will be talking about. I will then examine Chien and Wexler 1990 (section 2.2), Thornton 1990 (section 2.3), Boster 1991 (section 2.4), and Avrutin and Thornton 1994 (section 2.5).

2.1 Thornton and Wexler 1999

2.1.1 Explanation of the Basic Quantificational Asymmetry Data First, then, let us examine Thornton and Wexler's (1999) replication of the basic facts from Chien and Wexler 1990 reviewed above. Thornton and Wexler conducted truth-value judgment tasks in which stories were acted out for 4- to 5-year-olds, and the children were then asked to agree or disagree with a description of some event in the story offered by a puppet. The target sentence testing Principle B with a referential antecedent was (8).

(8) Bert brushed him.

In fact, in the story in question, Bert brushed himself but did not brush any other possible referent for *him*. Nevertheless, 58% of the time children judged this sentence to be true, in violation of the traditional, Chomskyan Principle B (Thornton and Wexler 1999:170).

Now what about the test for Principle B with a quantifier antecedent? The test sentence in this case was (9).

(9) Every reindeer brushed him.

In the story (the same one as used with sentence (8)), it is true that every reindeer brushed himself, but false that every reindeer brushed Bert, the potential referent for *him*. Children judged this sentence incorrectly to be true only 8% of the time, in contrast to the high figure for sentence (8). Is this good evidence that children obey Principle B for bound pronouns but not for referential pronouns, as claimed in the works we are examining?

There are two possible interfering factors worth examining at this point, either one of which would tell us not to take this result at face value. First, it has been observed that children show a preference for deictic readings of pronouns over bound readings (Crain and Thornton 1998: 111, Thornton and Wexler 1999:156). For example, Thornton and Wexler (1999:156) report that children preferred the false Meaning 2 in the experiment described in (10).

(41%), although a larger difference opened up when it came to testing possible binding of pronouns by the word for 'who'. Even on their face, without probing any further, Avrutin and Wexler's results do not seem like good confirmation for the QA; see Utakis 1995:83–89 for further discussion. The same can be said of results reported by Philip and Coopmans (1995): their group of Dutch-speaking 4- to 6-year-olds correctly rejected pronouns with local referential antecedents 36% of the time, and correctly rejected pronouns with local quantifier antecedents 53% of the time. This is a statistically significant difference; but when subjects accept something on 47% of trials, we cannot feel sure that it is not part of their grammar.

- (10) Sentence: Mr. Dog brushed his teeth and every dinosaur did too.
 Characters: Mr. Dog, 3 dinosaurs, Cookie Monster
 Meaning 1, true: Mr. Dog and every dinosaur brushed their own teeth.
 Meaning 2, false: Mr. Dog and every dinosaur brushed Cookie Monster's teeth.

The exact percentage of responses favoring the deictic Meaning 2 is not given, but it is remarkable under the circumstances that there was any noticeable preference for this meaning. This is because brushing teeth can be argued to be inherently reflexive, and thus likely to bias the children in favor of Meaning 1. The verb certainly takes a reflexive clitic in the Romance languages; and the action is generally something that people do to themselves and not to other people. The choice of Meaning 2 here also goes against the observed tendency for children to prefer to answer "yes" if possible in experiments like these. So altogether it looks like there is a powerful case for saying that children have a preference for referential as opposed to bound readings of pronouns. But this allows us to explain the large numbers of children who reject *Every reindeer brushed him* when every reindeer brushed himself: they could in principle construe *him* as bound here. They just choose to respond according to referential interpretations of pronouns because of independent factors. This hypothesis will be called the *Reference Hypothesis*.

What these independent factors could be is another question, and one it is not essential to answer for present purposes, given that there is empirical evidence to support the view that this preference exists. But it is arguable that a bound reading of a pronoun is more complex than a referential one. A referential reading just has the pronoun's index mapped to something in an assignment. The bound reading has that mapping; and, in addition, the thing to which the index is mapped (a variable in this case) is bound. So there is an extra structural relationship.

The second hypothesis is that in tasks like these, children interpret pronouns according to the most salient choice made relevant by the story and the question. To illustrate this idea, we need to know the details of the story used in this experiment. Thornton and Wexler summarize it as follows (1999:142).⁶

Bert and three reindeer friends have a snowball fight, and they all get covered in snow. When they go inside, Bert is shivering, so he asks the reindeer to brush the snow off him. Two of the reindeer (separately) refuse, saying they have too much snow to deal with, and they brush themselves. The third reindeer helps Bert a little bit, but then brushes the snow off himself. Bert thanks the helpful reindeer for starting to brush him. He says he's sorry he can't reciprocate by helping brush the reindeer; he needs to finish brushing all the snow off himself because he's still very cold.

Our literary criticism is hardly fanciful if we say that the whole plot of this story revolves around Bert's efforts to get the snow off himself. The main events are the reindeer's various responses to his requests that they brush him. Bert is the protagonist of the story, and the only character whose potential brushing by the reindeer is a big issue. It is very natural to assume, then, that when the children hear *Every reindeer brushed him*, they will take *him* to refer to Bert, not because

⁶ Unfortunately, the original scripts used are no longer available (Rosalind Thornton, pers. comm.).

the alternative, bound variable interpretation is ungrammatical for them, but because the whole thrust of the narrative propels them this way. This, then, is a second reason why the children in Thornton and Wexler's experiment could reject the bound variable interpretation of (9) even though it is grammatical for them. This hypothesis will be called the *Saliency Hypothesis*.

2.1.2 More Data from Thornton and Wexler 1999 In this section, I will discuss the other data in Thornton and Wexler's (1999) study that seem particularly relevant to children's knowledge of Principle B. I will be especially interested to see if any data distinguish the following three hypotheses:

1. The Asymmetry Hypothesis. The QA reflects a real difference in the way children's grammars treat bound and referential pronouns for the purposes of Principle B. One natural way to elaborate this (but not necessarily the only conceivable way) would be to appeal to Reinhart's (1983) binding theory, as Grodzinsky and Reinhart (1993) do.
2. The Reference Hypothesis. The QA does not reflect a real difference in the way children's grammars treat bound and referential pronouns for the purposes of Principle B; instead, children display an independent and quite general preference for referential readings of pronouns over bound readings, even though bound readings are in principle available. (The preference may be due to processing factors.)
3. The Saliency Hypothesis. The QA does not reflect a real difference in the way children's grammars treat bound and referential pronouns for the purposes of Principle B; instead, it arises because children interpret pronouns in the way made most plausible by the context, and the scenarios used in the relevant experiments make it likely that the pronouns in question will be interpreted as referring to certain prominent characters.

2.1.2.1 The Incredible Hulk brushed his hair . . . Here, the test sentence is (11).

(11) The Incredible Hulk brushed his hair and every troll did too.

The story leading up to this is summarized as follows (Thornton and Wexler 1999:149):

The rock star has a problem: he thinks his hairstyle isn't quite perfect. He asks the Incredible Hulk to brush his hair for him, and the Incredible Hulk agrees. Soon the rock star is worried that his hair is messy again, and this time he asks the Trolls to brush it for him. Although the Trolls consider helping him, they say they can't because they have to get the tangles out of their own long hair.

Children judged this sentence true only 3% of the time (Thornton and Wexler 1999:170).

This result makes perfect sense under any approach that allows children to handle VP-ellipsis in relevant respects as adults do, with the pronoun in the ellipsis site bound if the corresponding pronoun in the antecedent VP is bound, and referential (with the same referent) if the corresponding pronoun is referential. If *his* is taken to be referential, referring to the rock star, the sentence is false because the Trolls did not brush the rock star's hair; if it is taken to be bound by *the Incredible Hulk*, the sentence is false because the Hulk did not brush his own hair.

The various hypotheses under consideration cannot be distinguished by this example, then.

2.1.2.2 *Batman cleaned him . . .* The test sentence is (12).

(12) Batman cleaned him and every turtle did too.

The story is as follows (Thornton and Wexler 1999:152):

The Smurf has been painting his house and dropped some paint on his hat. Unfortunately, he can't reach it, so he asks Batman to help him get clean. Batman says he'd like to help, but he was just spraypainting glitter on his Batmobile and he got paint on his cape; he can't help the Smurf because he has to get the paint cleaned up. So the Smurf asks his turtle friends to help him get clean. However, the turtles have been making decorations and they are covered with glitter too. Two turtles say they can't help the Smurf because they need to get the glitter off their shell; the other turtle helps the Smurf get the paint off him and then gets his own shell clean.

Children judged this sentence true only 14% of the time (Thornton and Wexler 1999:170).

On the Asymmetry Hypothesis, this low figure is expected, because the bound reading (which is true) is excluded by Principle B. On the Reference Hypothesis, the result is expected also, because the preference for a referential construal of the pronoun will make the children take it to refer to the Smurf, or possibly to Batman. The sentence is false under either of these readings. Finally, the result also makes sense according to the Salience Hypothesis: this story is very similar to the story about Bert getting snow brushed off him, so we expect the children to take the pronoun to refer to the Smurf, who is unmistakably the protagonist. This, of course, means that the sentence will be judged false.

So once more the hypotheses under consideration are not distinguished.

2.1.2.3 *Bert lassoed him . . .* The test sentence is (13).

(13) Bert lassoed him and the Tin Man did too.

The story is as follows (Thornton and Wexler 1999:157):

The Indian Chief is teaching Bert and the Tin Man to lasso and he says they need to practice by trying to lasso him. But they are reluctant to lasso the Indian Chief: Bert says the Indian Chief is too far away, and the Tin Man thinks the Indian Chief is too important to lasso. So Bert practices by trying to lasso himself, and the Tin Man uses Bert as a target also.

Children judged this sentence true 43% of the time (Thornton and Wexler 1999:170).

Again there is no trouble for any of the hypotheses. To explain the positive responses, all that is necessary is that *him* refer to Bert, by whatever mechanisms the various hypotheses allow. Then this interpretation is given the elided VP too, and a true sentence results. The negative responses presumably are caused by children taking *him* to refer to the Indian Chief, the other character who was under consideration for being lassoed. All hypotheses seem to need the assumption that children often stick with whatever interpretation first occurs to them.

2.1.2.4 *Gonzo covered him . . .* The test sentence is (14).

(14) Gonzo covered him with sunblock and Snuffy did too.

The story is as follows (Thornton and Wexler 1999:153):

Porky Pig and his friends Gonzo and Snuffy decide to go to the beach. At the beach, they decide they need to put on some sunblock so they won't get sunburned. Porky Pig is very concerned about his fair skin and asks Snuffy to put sunblock on for him. Snuffy says he'd like to help, but he's concerned about his large trunk; he has to get sunblock on it right away. Gonzo is also unable to help because he needs to pay immediate attention to his large nose.

Children judged this sentence true 22% of the time (Thornton and Wexler 1999:170).

Here, the Asymmetry Hypothesis has its first hint of trouble. The 22% figure is fairly high, high enough to need explanation, as Thornton and Wexler observe. And yet the only way the sentence could be true, it seems, would be for the pronouns in both VPs to be taken as bound, in violation of Reinhart's Principle B. In an attempt to defend the Asymmetry Hypothesis, Thornton and Wexler suggest that this sentence be related to another one, namely, (15).

(15) The lizard man lifted him and the ugly guy did too.

In this story, the lizard man lifted the Smurf and the ugly guy lifted Mickey Mouse. There is no way for the sentence to be true, then, in the adult grammar; but children judged it true in Thornton and Wexler's experiment 21% of the time (Thornton and Wexler 1999:170). The authors suggest that the parallelism condition on the interpretation of an elided VP is modular, and that the requirement that the referents of referential pronouns in antecedent VP and elided VP be the same is one that can be detached and overridden by children. This would account for the 21% 'yes' response here. But if this is so, it is tempting to tell a similar story about (14): in 22% of cases, the overt *him* is referential and refers to Gonzo, and the elided *him* is referential and refers to Snuffy. Thus, a true sentence would result.

This account predicts, however, that there should be a correlation between the children who give positive responses to (15) and the children who give positive responses to (14). As Thornton and Wexler show, however, there is no statistically significant correlation between the two groups (Thornton and Wexler 1999:192–193). This is problematic for the Asymmetry Hypothesis.

Now let us consider the other hypotheses. The Reference Hypothesis and the Salience Hypothesis can say that in the 22% 'yes' response, children indeed construe the pronoun as bound, resulting in a sloppy reading and a true sentence. It seems that the Reference Hypothesis might have trouble here, since it has no obvious means to explain the difference between the 8% 'yes' response on (9) and the 22% 'yes' response here. But the Salience Hypothesis does well: the story is like the story with Bert and the reindeer, except that the characters from whom the protagonist seeks help are much better individuated in the current story. The reindeer did not have names, for example, whereas Gonzo and Snuffy do; and Gonzo and Snuffy have their own individual reasons for refusing to help, as opposed to the bland uniformity of the responses of the unhelpful reindeer. Looking at the story here, the Salience Hypothesis would predict that most children will find the test sentence false, since Porky is much the most dramatically salient referent for *him*, but that a significant number of children will take the sentence to be about Gonzo and Snuffy, since these are well-individuated characters. This, of course, is precisely what we find.

2.1.2.5 *Every Cabbage Patch boy . . .* The test sentence is (16).

(16) Every Cabbage Patch boy said Superman likes him.

The story is as follows (Thornton and Wexler 1999:145). (There are three Cabbage Patch boys.)

Each Cabbage Patch boy expresses his certainty that Superman likes him, giving a reason. Two of them say they don't think Superman likes Pokey, but the third disagrees, saying Superman would like to go for a ride on Pokey.

Children judged this sentence true 50% of the time (Thornton and Wexler 1999:170).

The clearest implication of this sentence is that the Salience Hypothesis is superior, once more, to the Reference Hypothesis. Recall that according to the latter hypothesis, the low 8% positive response to *Every reindeer brushed him* is due to an independent and quite general preference for referential readings of pronouns. But if that were the case, we would surely obtain a figure like 8% in response to (16) too. There is no evident reason why the figure should be so much higher here. By contrast, the Salience Hypothesis has a natural story to tell. Bert is obviously the central character in the earlier story, the whole plot centering around his efforts to get the snow brushed off him; thus, it is natural for *him* to be taken as referring to Bert. But in the present story, there is no such central character. In particular, Pokey, the potential referent of *him* if it were taken to be referential, definitely does not dominate the action in the way that Bert did. Thus, the Salience Hypothesis predicts the high figure that Thornton and Wexler found; some children will take *him* to refer to Pokey, but just as many will take it to be bound by *every Cabbage Patch boy*, since there is no reason, given the story, to prefer one reading over the other.

The question for the Asymmetry Hypothesis here is why there are not *more* positive responses, since *him* can be bound by *every boy* consistently with Reinhart's Principle B. Presumably the Asymmetry Hypothesis could appeal to the same sort of pragmatic factors just mentioned: there is no need to take the pronoun to be bound, given the availability of the other reading where it is referential and the fact that there is no compelling dramatic reason to favor one reading over the other. Note, however, that the Asymmetry Hypothesis now needs to rely on the machinery of the Salience Hypothesis in order to stay in good standing.

2.2 *Chien and Wexler 1990*

So far, then, we have reason to prefer the Salience Hypothesis to the Reference Hypothesis and the Asymmetry Hypothesis. It is interesting that when we come to Chien and Wexler's (1990) data, the tables seem to be partially turned, in that the Reference Hypothesis at first sight seems to fare better than the Salience Hypothesis. (The Asymmetry Hypothesis also does well.) I believe this is only a temporary setback for the Salience Hypothesis, however.

The reason why the Salience Hypothesis at first seems hard to apply to Chien and Wexler's data is that the crucial experiment in this study used pictures, not stories. In order to test their response to *Is Mama Bear touching her?*, children were shown a picture that contained only Goldilocks and Mama Bear; Mama Bear was not touching Goldilocks but was touching herself

(Chien and Wexler 1990:273). For *Is every bear touching her?*, children were shown a picture that contained Goldilocks and three bears; the bears were not touching Goldilocks but were touching themselves (Chien and Wexler 1990:273). Since there are no stories to provide a bias toward one interpretation by making one character an obvious protagonist, as in the reindeer and Bert story, it could be argued here that the Salience Hypothesis cannot be successful. On the other hand, the Reference Hypothesis, with its blind application to pronouns in any circumstances, would do the trick.

But since we saw reason to prefer the Salience Hypothesis to its rivals in the previous section, we should take a closer look. To start with, it is evident that the Salience Hypothesis has no difficulties with *Is Mama Bear touching her?* The two characters in the picture are equally prominent in visual terms. So we expect that the children, being bad at casting around for multiple interpretations, will simply select one of the two characters more or less at random as a referent for *her*. (The pronoun could also be bound by *Mama Bear*, of course.) The 50% positive response score is obviously consistent with this.

The difficulty, if difficulty there is, comes with *Is every bear touching her?* Why is the positive response not higher than 16%, given that *her* could in principle be bound and reflexive here, as the Salience Hypothesis demands? Here, it is necessary to look at the picture used—see figure 1. Note, first, that Goldilocks is much more visually prominent than any of the bears. Indeed, casual inspection suggests that she is about three times as tall as each bear! Second, it is actually quite hard to make out the sex of the bears: the only indication seems to be a small ribbon on top of their heads, but that is not obvious and could easily be mistaken for something else or simply not noticed. The bears do not have any facial or other physical features that could be pointers. There is also no indication of the bears' sex in the lead-in or the question (Chien and Wexler 1990:273).

(17) These are the bears; this is Goldilocks. Is every bear touching her?

Contrast the question *Is Mama Bear touching her?*, where it is obvious that both characters are




Figure 1

Goldilocks and the three bears. Picture used in Chien and Wexler's (1990) study to test children's response to *Is every bear touching her?* (From Chien and Wexler 1990. Reprinted by permission.)

female and thus, for a child at the relevant stage, potential referents for *her*. It is quite possible, then, that there is a bias toward a referential interpretation, with *her* referring to Goldilocks in *Is every bear touching her?*, because Goldilocks is the only character in the picture who is obviously female.⁷ I submit, then, that these two factors, size and sex, may have played the role that Bert's dramatic centrality did in the reindeer story.⁸

It seems, then, that the Salience Hypothesis is entirely compatible with Chien and Wexler's (1990) data.

2.3 Thornton 1990

Another experiment that bears directly on the reality of the QA was performed by Thornton (1990) and used binding by a *wh*-operator instead of binding by quantifiers. We are dealing with sentences like those in (18) and (19).

- (18) a. I know who scratched them—Bert and Huckleberry Hound.
 b. I know who washed him—Bert.
- (19) a. That Kiwi and that Kiwi scratched them.
 b. That Smurf and that Smurf washed them.

The claims in sentences like those in (18) were judged true by the children only 8% of the time in cases where the pronoun would have to be locally bound by *who* in order for the claim to be true; in contrast, the children judged sentences like those in (19) to be true 49% of the time in cases where the pronoun would have to be coreferential with the local *c*-commanding subject of type *e* in order to produce a true sentence. These results are obviously very similar to those obtained by Thornton and Wexler (1999) with respect to their story about Bert and the reindeer (8% vs. 58%; section 2.1.1).

Again, however, when we look at the type of scenario given the children, an explanation involving the prominence of certain characters springs to mind. We are once more dealing with truth-value judgment tasks. This is how Thornton describes the format that all the stories followed (1990:179):

In these stories there are two characters who are 'helpless', and are helped by one (or sometimes two) of the 'merciful'. The remaining two characters are 'selfish'. They take care of their own needs, telling the 'helpless' that they cannot provide any help for them. In these stories, the 'helpless' are the prominent characters who take the reference for the pronoun in the test question.

⁷ This sentence was not the only one testing Principle B with quantificational antecedents in Chien and Wexler's study. Similar scenarios were created in which the bears were replaced by Smurfs, mice, and ducks (Chien and Wexler 1990:264). The pictures used in these scenarios are not given in the article, but again it seems that the species used are not easy to differentiate sexually. For experimental confirmation of the use of gender by young children in pronoun resolution, see Arnold et al. 2001.

⁸ After formulating this thesis, I was pleased to discover that Carole Tenny Boster (1991:23) had independently and much earlier come to the same conclusion.

It is noteworthy that, in the last sentence of this passage, the designer of the experiment herself tells us that certain characters seem more prominent than others. This is so, it seems, because the story revolves to a large extent around who will and who will not help the helpless: the merciful characters help them, and the selfish ones explicitly mention the helpless and say they cannot help them. We might expect, then, that the pronoun in the test question will be taken to be referential, referring to the helpless characters, even if the grammar provides for a bound reading.

An example will make it obvious just how natural this distortion would be (Thornton 1990: 180).

This is a story with Big Bird, Snuff[le]upagus, Bert and Huckleberry Hound. In this story they take a walk with RoboCop and Batman just before dark. The problem is that mosquitoes come out at dark. They bit everyone except for RoboCop and Batman because they are wearing special suits. Big Bird and Snuff[le]upagus got the most bites, because they are the biggest. 'Oh, I'm so itchy,' they say. RoboCop and Batman say, 'We'll help you. We don't have any bites.' (RoboCop and Batman scratch Big Bird and Snuff[le]upagus.) Bert says, 'Sorry, I can't help you, Big Bird and Snuffy. I have a lot of bites too, and they're very itchy.' Huckleberry Hound says, 'I can't help you either. If I don't scratch these bites on my leg right now, I'll go crazy.' (Bert and Huckleberry Hound scratch their bites.)

It is evident that this story resembles Thornton and Wexler's tale of Bert and the reindeer: it is a big issue whether certain characters will be given aid, and most of the action undertaken by the other characters consists in either helping or refusing to help the needy ones. The sentence to be evaluated by the children is (20).

(20) I know who scratched them. Bert and Huckleberry Hound.

We can well imagine that if *them* were in principle ambiguous between bound and referential interpretations, the referential reading would end up being favored. At first sight, then, it seems that these data from Thornton 1990 are compatible with the Salience Hypothesis.

There is one complication here, however, that has not previously arisen. In this study, the sentences like those in (19) in which the pronouns do not have to be bound variables were presented after the same stories that were used to test the bound cases. For example, after the story just given, the sentence to be evaluated by the children was (21) (Thornton 1990:181).

(21) Bert and Huckleberry Hound scratched them.

Here, a positive response was expected from children who obey Principle B only for bound pronouns, since Bert and Huckleberry Hound did indeed scratch themselves in the story. For children with the adult grammar, a negative response was expected, since these two characters did not scratch their "helpless" comrades, Big Bird and Snuffleupagus. The problem for the Salience Hypothesis is now the following. Since this hypothesis relies on the prominence of the helpless characters to explain why so many children responded negatively to (20), it seems to predict that the numbers of children responding negatively to (21) will be the same, since exactly the same story has preceded. That is, if Big Bird and Snuffleupagus were prominent enough to be picked up by *them* in (20) with a certain frequency, they should be prominent enough to be

picked up by the pronoun in (21) with the same frequency. But this is manifestly not what happened: children said that sentences like (20) were true 8% of the time, meaning that the helpless characters must have been assigned to the pronoun 92% of the time; but children said that sentences like (21) were true 49% of the time, meaning that, after the very same story, the helpless characters were now assigned to the pronoun only 51% of the time. This looks like a serious blow to the Salience Hypothesis and solid support for the Asymmetry Hypothesis.

Note that this objection to the Salience Hypothesis relies on the environments in which test sentences like (20) and (21) were presented being identical. A closer examination of Thornton's description of her experimental procedure, however, makes it clear that the environments in which these types of sentences were presented were *not* identical. In describing the cases like (21), she says that much of the action in the stories was not relevant to the truth or falsity of these target sentences. And she points out a difference in the environments in which sentences like (20) and (21) were presented (Thornton 1990:181n85):

Sometimes, in a story with a lot of action not strictly relevant to the target, it was more appropriate to say 'I know one thing that happened. Bert and Huckleberry Hound scratched them.'

This difference is crucial. For consider the basic premise of the Salience Hypothesis. This is that the subjects will be concentrating on the main action of the relevant stories, which is by hypothesis that of helping the helpless. Thus, they will interpret pronouns in accordance with the referents made available in consideration of this main action. But now consider the likely effect of the childrens' hearing "I know one thing that happened." I submit that they are reminded that many things happened in the story. Otherwise, why has the experimenter said "one thing that happened" as opposed to, say, "what happened"? And consequently we cannot rule out the hypothesis that they will take into account a wider range of potential interpretations for pronouns than they would have if they had not heard this particular lead-in. In the example we have been discussing, then, they will be more likely to consider Bert and Huckleberry Hound as referents for pronouns than they would otherwise have been. And this, of course, will give exactly the results that were found: more children will accept sentences like (21) than sentences like (20), even if otherwise their acceptance rates for these sentences would have been comparable. Thornton does not tell us how many tokens of test sentences like (21) were preceded by the confounding "I know one thing that happened"; but it seems likely to have been a fairly high proportion, since the stories used would generally have met the condition of containing lots of irrelevant action when the test sentences were like (21).

It seems, then, that the data from Thornton 1990 are after all consistent with the Salience Hypothesis. It has furthermore become evident that the scenarios used in the experiments we have looked at are all open to criticism by the Salience Hypothesis. Thornton's prototype story involving helpless characters in starring roles and selfish characters in supporting roles applies to the stories of Bert and the reindeer (section 2.1.1), the rock star tidying his hair (section 2.1.2.1), the Smurf getting paint off himself (section 2.1.2.2), and Porky Pig getting himself covered with sunblock (section 2.1.2.4), as well as to the experiments discussed here from her own dissertation.

2.4 Boster 1991, Experiment 1

Sentences similar to those just discussed were also tested by Boster (1991, Experiment 1). She conducted a truth-value judgment task involving 10 children ranging in age from 4;6 to 6;0. Each child was presented with 20 sentences that included 4 testing Principle B with referential antecedents, as in (22), and 5 testing Principle B with *who* as an antecedent. Of the latter, 3 had one DP following, as in (23a), and 2 had two DPs following, as in (23b).

(22) Strawberry Shortcake washed her.

(23) a. I know who dressed her—Baby Sally.

b. I know who painted her—Sister Bear and Dorothy.

The stories and test sentences were designed so that in order to respond affirmatively, the children would have to violate Principle B in their interpretation of the test sentence. The results were as follows: the children gave affirmative responses to sentences like (22) 37.5% of the time, but did so in response to sentences like (23a) and (23b) only 4% of the time (Boster 1991:12).

To appreciate the significance of these results, we need, as always, to look closely at the actual scenarios used. Unfortunately, Boster does not give detailed summaries of the stories used to test sentences like (23a) and (23b). But we can perhaps learn something from her description of the general format the stories took (1991:9). For sentences like (22) and (23a), Boster used stories involving two characters (matched for sex), constructed according to the following underlying paradigm (Boster 1991:9):

[T]he first character performs a given action to herself (or himself), the second asks the first to do the same to her, the first one refuses for some plausible reason, and the second one then does the same action to herself.

For sentences like (23b), a third character was added, creating slightly more complex scenarios designed to ensure that there was a plausible context for denying the test sentence.

This means that the story given as background for (23a), for example, contained two female characters, Baby Sally and one other, each of whom dressed herself and neither of whom dressed the other. We can now begin to appreciate how children would be able to interpret *who dressed her* in (23a) as ‘who dressed herself’ and yet would still be likely to answer ‘no,’ as if they had interpreted the pronoun as referring to the other character in the story: *wh*-questions require that their answers be *complete* lists of the entities that performed the specified action (Karttunen 1977: sec. 2.6 and much subsequent literature). The given answer is incomplete, of course, since in the story Baby Sally was not the only character who dressed herself. Boster notes in another connection that the children seemed in general to prefer test sentences that told the whole story, and were apt to respond negatively to sentences that gave accurate but partial descriptions of the story (Boster 1991:15, 25); a fortiori, then, we would certainly expect them to respond negatively to answers to *wh*-questions that were less than complete lists of the relevant characters, since even adults generally demand this kind of maximality.

Boster does not say whether the third character in the background stories for sentences like (23b) also performed the specified action, but it seems reasonable to assume that this was the case, following the model of the other stories. In that event, sentences like (23b) are also unreliable as tests of adherence to Principle B, for the reason just outlined.

Overall, then, the results of this experiment seem to be compatible with the denial of the Asymmetry Hypothesis. This is especially so given the other two experiments that Boster describes in the same paper, which show (one of them using the same subjects as the experiment described here) that the children she tested perform equally well on sentences testing Principle B with a referential antecedent and sentences testing Principle B with an antecedent composed of *every* + NP (see section 3.2).

2.5 *Avrutin and Thornton 1994*

2.5.1 The Experiment More possible evidence for the QA is to be found in work by Avrutin and Thornton (1994), who were investigating Heim, Lasnik, and May's (1991b) hypothesis concerning the distributive and collective interpretations of plural DPs. Heim, Lasnik, and May argue that when plural DPs are interpreted collectively, they just receive an index that refers to the group or plural individual associated with them, as one might expect; but when they are interpreted distributively, they are associated with an indexed distributivity operator that functions as a quantifier, saying that for each member x of the group referred to by the plural DP, x did the action denoted by the VP. Avrutin and Thornton (1994:167) give the following examples:

(24) [The Smurf and the clown]_{*i*} dried Big Bird_{*k*}.

(25) [[The Smurf and the clown]_{*i*} D]_{*j*}] dried Big Bird_{*k*}.

According to Avrutin and Thornton, (24) means that the Smurf and the clown dried Big Bird together, perhaps using a single large towel, whereas (25) means that the Smurf and the clown each dried Big Bird in separate drying events.

Avrutin and Thornton (1994:167) then invite us to consider how this difference between distributive and collective readings might play out in the case of children who allow local coreference in violation of the traditional Principle B, but not local binding. Take a sentence with a plural DP subject and a plural pronoun object, like (26).

(26) The Smurf and the clown dried them.

It is evident that children in the stage we are considering should allow the reflexive collective reading, shown in (27), since the pronoun in this reading is referential.

(27) [The Smurf and the clown]_{*i*} dried them_{*i*}.

But these children should disallow the reflexive distributive reading, shown in (28), since the pronoun in this reading has to be a variable bound by the distributive operator.

(28) [[The Smurf and the clown]_{*i*} D]_{*j*}] dried them_{*j*}.

Thus, say Avrutin and Thornton, it should be possible to test Heim, Lasnik, and May's hypothesis regarding an indexed distributivity operator by seeing whether children get distributive readings in sentences like (26). The prediction is, of course, that they will not.

Avrutin and Thornton tested 33 children between the ages of 3;10 and 4;10 on truth-value judgment tasks in which puppets acted out stories that provided plausible backgrounds for distributive and collective interpretations of plural DPs. For example, in the stories for (26) there were two other characters in addition to the Smurf and the clown who were potential referents for *them*. In one story, the Smurf and the clown refuse to dry these other characters and use a big towel to collectively dry themselves. It was expected that children would take *them* in (26) to refer to the Smurf and the clown and say that this sentence was an accurate description of this scenario, understanding it on the collective reading. In another story, by contrast, the Smurf and the clown separately refuse to dry the other characters and then separately dry themselves. In order for the children to take (26) to be an accurate description of this scenario, they would have to be able to give the subject DP a distributive interpretation, which would involve a bound variable interpretation in violation of Principle B; it is predicted, therefore, that children at the relevant linguistic stage will not take this sentence to be an accurate description of this scenario.

The children were also tested on control sentences to ensure that they were able to obtain both collective and distributive interpretations of DPs. The control sentences worked as follows. Consider a situation in which two turtles have two pet bugs each. Then a puppet might say either (29) or (30).

(29) I know how many bugs they have. Four.

(30) I know how many bugs they have. Two.

(29), of course, will be accepted by children who are able to assign collective interpretations, and (30) by children who are able to assign distributive interpretations.

The results of the experiment were as follows. Seventeen of the children responded like adults, rejecting sentences like (26) every time they were presented in collective contexts. These children were not tested further. The remaining 16 children were the target group. They allowed the pronoun in sentences like (26) to be anaphoric to the subject 93% of the time in collective contexts. In distributive contexts, they allowed the pronoun to be anaphoric to the subject significantly less, 42% of the time.

Avrutin and Thornton (1994:169–170) say further that 4 children among the target group consistently rejected the distributive interpretation of control sentences like (30). These children accepted sentences like (26) in both collective and distributive contexts. Since their rejection of the distributive control sentences made their relevance to the study doubtful, they were excluded from a final analysis of the data, which focused on the remaining 12 children. The responses of this group of children were as follows: they accepted test sentences in collective contexts 93% of the time, and test sentences in distributive contexts only 27% of the time.

Avrutin and Thornton claim that Heim, Lasnik, and May's (1991b) analysis of distributivity receives support from these results, since these authors claim that the distributive but not the

collective interpretation of subject DPs in sentences like (26) involves variable binding; and children who are allegedly at the stage at which they disallow violations of Principle B that involve binding but allow those that do not involve binding indeed show a preference for collective interpretations of sentences like (26).

2.5.2 The Relevance for Principle B One could also interpret these results as confirming the Asymmetry Hypothesis. Suppose we grant that Heim, Lasnik, and May's (1991b) analysis of distributive and collective interpretation is substantially correct, for the reasons given in their article. (See Williams 1991 and Heim, Lasnik, and May 1991a for further discussion.) Then the Asymmetry Hypothesis predicts that children at the relevant stage will accept violations of Principle B in collective contexts, but not in distributive contexts. And this prediction is arguably borne out by Avrutin and Thornton's experiment.⁹

We should be cautious at this point, however. The figure of 42%, which was obtained in the first analysis for the percentage of trials on which the distributive interpretation was accepted, is very high. It would ordinarily be taken to indicate that something is present in the grammar, especially since the scenarios tested in this experiment were ones where the pronoun could plausibly be taken to refer to another set of people anyway, and so presence of something in the grammar would not necessarily entail near-unanimity of 'yes' responses. In other words, Avrutin and Thornton have demonstrated a significant contrast between performance on collective scenarios and performance on distributive scenarios; but they have not shown that the distributive interpretation was not available for the children in question. Their results, in fact, strongly suggest that it was available.

These comments apply even to the reduced figure of 27% obtained by removing 4 children's responses from the analysis. A figure of 27% would normally be thought high enough to need explanation; indeed, as we saw in section 2.1.2, Thornton and Wexler consider a figure of 22% high enough to need explanation.¹⁰ If a figure of 27% 'yes' responses could only be obtained, on the face of it, if a distributive reading were available in the subjects' grammars, we have good reason to believe that a distributive reading was indeed available.

Furthermore, it is not clear that Avrutin and Thornton were justified in removing these 4 children's responses from the data set. These children did not accept distributive readings of DPs in the control sentences, but they seemed to obtain both collective and distributive readings in response to test sentences like (26). It is possible, then, that the children did accept distributive readings and therefore their responses should be kept, and that some factor yet to be discovered led them to reject the distributive control sentences, which were, it should be noted, significantly different in structure from the test sentences; they also involved significantly different types of

⁹ An argument along these lines was suggested by an anonymous reviewer.

¹⁰ See also Crain and Thornton 1998:329, where it is acknowledged that the 27% figure is rather high.

scenarios. It seems, then, that 42% is the figure that really represents the children's acceptance of distributive readings.¹¹

The discussion so far has assumed that the experiment did in fact manage to test the children's responses to structures where the pronoun would have to be locally bound in order to produce a reading that would be true in the distributive scenarios; that is, I have assumed that a structure like (28) would have been forced. However, there is reason to doubt this, as we will see from a closer look at Heim, Lasnik, and May's (1991b) theory and the readings it predicts. The following are the structures and readings predicted by this theory (Heim, Lasnik, and May 1991b:76):

- (31) a. [John and Mary]_i argued that they_i would win \$100.
 b. 'John and Mary together argued that together they would win \$100.'
- (32) a. [John and Mary]_i argued that [they_i D_j] would win \$100.
 b. 'John and Mary argued together that they would each win \$100.'
- (33) a. [[John and Mary]_i D_j] argued that they_i would win \$100.
 b. 'John and Mary each argued that together they would win \$100.'
- (34) a. [[John and Mary]_i D_j] argued that they_j would win \$100.
 b. 'John argued that he would win \$100 and Mary argued that she would win \$100.'
- (35) a. [[John and Mary]_i D_j] argued that [they_i D_k] would win \$100.
 b. 'John argued that he would win \$100 and Mary would win \$100, and Mary argued the same thing.'

These readings do indeed seem to be possible, and to be the only ones possible, for the surface string in question, which is a significant point in favor of Heim, Lasnik, and May's theory.

Avrutin and Thornton (1994) assume that their distributive scenarios will force representations and readings like those in (34), thus ensuring local binding of pronouns (since their examples were monoclausal, like (26)). Against this, however, one can bring the following considerations to bear. It has been convincingly argued that with alternating unaccusatives like *dry* in (26), the transitive alternant has a semantics like this (Dowty 1979, Parsons 1990, Levin and Rappaport Hovav 1995):

- (36) $\lambda x.\lambda y.y$ CAUSE: x BECOME dry

Furthermore, it has been argued that at least some of the semantic elements visible in this paraphrase should actually be realized in the syntax (Marantz 1997, Pytkkänen 2000), with the result that (26) could have a (still simplified) structure along the lines of (37).

- (37) [[[the Smurf and the clown]_i D_j] [CAUSE [[them_i D_k] dry]]]

¹¹ See Utakis 1995:119–124 for important further discussion of this point. In particular, Utakis cites raw data from a 1993 handout of Avrutin and Thornton's, showing that the 4 children whose responses were excluded were not the only ones who uniformly rejected the distributive readings of the control sentences. Of the 12 children who figured in the final analysis, no fewer than 5 also rejected all the distributive readings of the control sentences. It is not clear, to say the least, that excluding the 4 but including the 5 was justified. See Utakis 1995 for further discussion.

The indexing here is analogous to that in (35) rather than (34). The interpretation would be as in (38).

- (38) The Smurf and the clown each caused the following to take place: the Smurf and the clown each became dry.

Is this a plausible description of what happens in the reflexive distributive scenario, in which the Smurf and the clown dry themselves separately? Note that the presence of D_k with *them* does indeed bring about a distributive interpretation, in that we end up with two separate events of becoming dry. The only question arises with the claim that the Smurf and the clown each caused this to happen. Do the Smurf and the clown each have to cause the drying of each of them in order for this to be true? If they did, the reading would not be an accurate description of the reflexive distributive scenario, but in fact the answer seems pretty clearly to be that they do not. It is not specified in the truth conditions that the Smurf and the clown each have to cause both drying events, but simply that they each have to cause (read: play a causal role in bringing about) the result that two separate events of becoming dry took place. Compare a simple sentence like (39), where intuitively all that is necessary is that the Smurf and the clown each have to play a causal role in bringing about the result that Big Bird became dry.

- (39) The Smurf and the clown dried Big Bird.

So (37) could very well produce truth conditions that would accurately describe the reflexive distributive scenario. But note that in this LF structure the pronoun *them* is not bound; it is simply coreferential with *the Smurf and the clown*. So we cannot be certain that Avrutin and Thornton's distributive scenarios forced bound variable readings for the pronouns, as they were intended to.¹²

Overall, then, Avrutin and Thornton's experiment may not have tested children's reactions to bound variable anaphora in violation of Principle B at all. But if we assume it did, we must recognize that the percentage of trials in which this anaphora was found grammatical is rather high, suggesting that it was indeed part of the grammars of the children involved.

2.6 Conclusion Regarding Experiments That Show the Quantificational Asymmetry

We have seen that the experiments that confirm the QA are also compatible with the denial of the Asymmetry Hypothesis. In particular, it seems as if the Saliency Hypothesis may be able to

¹² As an *LI* reviewer notes, one could interpret this point as an explanation of why the percentage of subjects allowing anaphora in the allegedly bound contexts is so high: there is an alternative way of construing the surface string in question that produces the relevant meaning without illegal binding, and many subjects interpreted it this way. The idea is then that my second criticism of this experiment would actually cancel out my first criticism. This point is well taken, but I do not think it allows us to conclude that this experiment provides clear evidence in favor of the QA after all, since it is still not clear what structure was being tested. To provide clear evidence by these means, it would be necessary to test sentences that unambiguously differentiate among the various readings allowed by Heim, Lasnik, and May's system. And one would also have to be sure that Heim, Lasnik, and May's system was correct in all the relevant respects, for children as well as adults, something that is by no means obvious given the complexity of the issues surrounding distributivity; see Winter 2001 for discussion.

explain the data we have examined so far, with a little help from other considerations in the cases of Boster 1991 and Avrutin and Thornton 1994.

This is an appropriate point at which to address a methodological issue relevant to the Salience Hypothesis.¹³ Experimenters using the truth-value judgment task generally assume that, given an ambiguous sentence and a scenario on which one reading is true and the other false, children will prefer to give the speaker the benefit of the doubt and will interpret the sentence as being true (Crain and Thornton 1998:111, 211). So if children claim that a speaker spoke falsely, we can conclude that their grammars did not generate any reading on which the sentence was true. We have seen this assumption in action in the literature we have just reviewed, where overwhelmingly negative responses to sentences like *Every reindeer brushed him* were taken to indicate that no bound reading was available.

But this assumption must be false if the above account of the QA data is correct. For the sentences with quantificational antecedents for pronouns are generally true on the bound reading in these experiments, but are nevertheless judged to be false, even though, according to the Salience Hypothesis, the true bound reading is quite grammatical. So the discourse factors that lead the children to assign certain prominent characters as referents for pronouns must be more powerful than the preference, if it exists, to interpret sentences as being true. I do not think this conclusion is troubling. In fact, recent experimental evidence shows the supposed preference being overruled in a different context (Musolino and Lidz 2003). See Drozd, forthcoming, for related criticisms of the truth-value judgment task methodology.

3 Experiments That Do Not Show the Quantificational Asymmetry

Given the conclusion of the previous section, we predict that if any experiments avoided privileging certain characters in the way described by the Salience Hypothesis, and avoided other serious methodological imperfections, the QA would disappear. This seems already to have happened in experiments reported in Lombardi and Sarma 1989 (section 3.1) and Boster 1991 (her Experiments 2 and 3; section 3.2). It may have happened also in an experiment reported in Utakis 1995 (section 3.3), though this is more debatable.

3.1 Lombardi and Sarma 1989

Lombardi and Sarma (1989) conducted an act-out task with 19 children ranging in age from 4;0 to 6;2. The test sentences were variations on those in (40).

- (40) a. Monkey wants Bert to go in the box and get him a toy.
 b. Monkey wants every puppet to go in the box and get him a toy.
 c. Monkey wants Bert to get him a toy.
 d. Monkey wants every puppet to get him a toy.

¹³ I am grateful to an *LI* reviewer for alerting me to this interesting corollary of the Salience Hypothesis.

Four verbs were used: *give*, *feed*, *get*, *find*. The setup was that one experimenter manipulated a monkey puppet. She said that he was shy and would talk only to her; but he wanted to play a game with the child. Monkey then told this experimenter what he wanted the other puppets to do, and the experimenter told the child and the second experimenter. The child manipulated Bert, when Bert was supposed to do something, and a pair of puppets called Cow and Frog when the test sentence involved “every puppet.” By noting who the children had their puppets give toys to (and so on), it was possible to determine the number and distribution of apparent Principle B violations in the children’s interpretations of the test sentences.

The overall results were as follows: the rate of correct responses for sentences like (40b) and (40d) was 41.8%, and the rate of correct responses for sentences with no quantifier antecedent was 48.3%. (A response was counted as correct if and only if it was compatible with the interpretation of the test sentence according to the adult grammar. Lombardi and Sarma imply, but do not state explicitly, that all the incorrect responses involved incorrect reflexive interpretations—that is, trials where the children had the puppets inappropriately give toys to themselves, and so on.) Thus, no significant asymmetry was found between obligatorily bound interpretations and others. This pattern held both in the overall results just reported and in the results analyzed child by child.¹⁴

Even from the brief account just given, it is evident that this experiment was not completely beyond reproach when it came to controlling for extraneous factors; but the extraneous factors that suggest themselves would if anything bias the experiment in the same way as the previous ones we have looked at, toward rarity of bound variable interpretations in violation of Principle B. First, one might wonder if the sexes of the Cow and Frog puppets were made clear—and indeed, one would ordinarily expect that a puppet called Cow was supposed to be female, which might make it hard to construe *him* as being bound by *every puppet* in the relevant examples. This would tend to bias the experiment away from bound variable interpretations; but as we have seen, they were frequent nonetheless. Perhaps the experimenters made it clear that the two puppets used in the *every puppet* sentences were male; or perhaps the children, even if they knew that cows are generally female and assumed that this one was too, were able to disregard the slight incongruity of a masculine pronoun ranging over both males and females. Of course, the children also might not have known that cows are female.

¹⁴ The data broke down as follows according to the verbs used. For each verb, there were 15 trials in referential contexts, like (40a) and (40c), and 15 trials in binding contexts, like (40b) and (40d). The verbs *find* and *get* were always used in conjunctions, like (40a) and (40b), while the other verbs were used in examples like (40c) and (40d). In the referential contexts, then, the following are the numbers of responses consistent with the adult grammar for each verb, out of 15: *find*, 7; *get*, 7; *give*, 8; *feed*, 5. The following are the corresponding numbers for the binding contexts: *find*, 6; *get*, 3; *give*, 6; *feed*, 7. The low score for *get* in the binding contexts is interesting, since adult English too sometimes allows apparent violations of Principle B with this verb, as in *He got him a gun*. Lombardi and Sarma, however, seem to have counted as correct responses only those consistent with Principle B in the case of *get*, as with all the other verbs. It is unclear that anything should be made of the low figure for *get* in the binding contexts, given that it behaved like the other verbs in the referential contexts.

Second, it might be thought that the most psychologically plausible interpretation of sentences like those in (40) is one in which Monkey wants a toy for himself. But this interpretation would tend to make the adult readings of the sentences favored and the locally bound pronouns rare, again in opposition to the facts, it seems.

This experiment avoided the possible biasing effects of both stories and pictures by using neither. And the results showed roughly equal construal of simple pronouns as locally bound variables and referential pronouns. This is obviously incompatible with the Asymmetry Hypothesis.

3.2 *Boster 1991, Experiments 2 and 3*

Boster's (1991) Experiment 2 consisted of a truth-value judgment task with 24 children ranging in age from 3;3 to 4;9. She used a series of pictures as stimuli. The experiment was explicitly modeled after the tasks in Chien and Wexler 1990, and the aim was to see if it was possible to replicate the QA reported by those researchers. One difference was that plural pronouns were tested with *every* + NP as antecedent, as well as singular pronouns. So the test sentences included the following:

(41) Is every monkey patting him?

(42) Is every bear touching them?

(43) Is Chip brushing him?

Again, the pictures and test sentences were designed so that a "yes" answer would be possible only by violating Principle B.

The results were as follows. The children accepted sentences like (43) in violation of Principle B 37.5% of the time; they accepted sentences like (42) in violation of Principle B 41.67% of the time; and they accepted sentences like (41) in violation of Principle B 34.38% of the time. In other words, there was no significant difference between the rate of Principle B violations with referential antecedents and the rate of Principle B violations with quantificational antecedents. This is obviously a serious problem for the Asymmetry Hypothesis.

Boster also makes some interesting observations about differences between her experimental setup and Chien and Wexler's that may have contributed to the difference in results. She says that all of the characters she used were male, as opposed to Chien and Wexler's characters, some of which were male and some female; thus, this setup avoided confounding factors like the necessity of working out that the three bears in figure 1 are supposed to be female (Boster 1991: 23). Similarly, Boster took care to make all the characters in her pictures of roughly equal size, with the explicit aim of avoiding disparities like that discussed above with reference to figure 1. It seems, in fact, that Boster's Experiment 2 is about as close as one could come to taking an experiment open to criticism by the Salience Hypothesis and replicating it with the disparities in salience removed. The fact that this experiment produced no QA thus supports the Salience Hypothesis, and disconfirms the Asymmetry Hypothesis.

The children who participated in Boster's Experiment 3 were 8 of the children from Experiment 1 (section 2.4), the ones who had accepted coreference in violation of Principle B in sentences with referential antecedents. Experiment 3 was also a truth-value judgment task using stories; but here, only DPs of the form *every* + NP were used. Pronouns could be either singular or plural. Examples are given in (44) and (45).

(44) Every seal tickled him.

(45) Every rabbit brushed them.

The format of the stories was rather similar to that of the stories used in Experiment 1 with the difference that there was a group of three characters (seals, bears, rabbits, or cats) who could potentially be quantified over by the quantificational DP, plus one or two other characters, depending on whether the pronoun was singular or plural. The format of the stories was as follows (Boster 1991:24):

[E]ach of the group of three animals performs a given action to himself, then the second character(s) asks each of the animals in turn to do the same to him (or them), every animal refuses, and the second character (or set of two characters) then does the action to himself (or themselves).

The verbs used were *tickle*, *wash*, *brush*, and *put (a Band-Aid on)*.

The results were as follows. The children accepted sentences with *every* + NP and *him* in violation of Principle B 31.25% of the time, and sentences with *every* + NP and *them* in violation of Principle B 50% of the time. This is comparable to the rate at which these same children accepted sentences with referential antecedents in violation of Principle B in Experiment 1 (50%; Boster 1991:14). Again, the Asymmetry Hypothesis is contradicted.

3.3 *Utakis 1995*

Utakis conducted a truth-value judgment task using pictures as stimuli (Utakis 1995:206, her part 2). Her subjects were 30 children ranging in age from 3;4 to 9;5 (mean age 6;2) and, in addition, 22 adults. Examples of test sentences are given in (46), (47), and (48) (Utakis 1995:295).

(46) Did the bear spray him?

(47) Did every bear spray him?

(48) Did every bear spray them?

Verbs used included *spray*, *brush*, and *pour water on*. The pictures for sentences like (46) showed five animals. For example, one picture (the "match case") showed a bear, a parrot, and three baby tigers; the bear sprayed the parrot and did not spray the baby tigers. There was also a "mismatch case" in which the bear sprayed the baby tigers (as a group) but did not spray the parrot (Utakis 1995:205). As in Boster's (1991) Experiment 2, all the animals were male; the

subjects were told this so the gender of pronouns would not be a confounding factor (Utakis 1995:203).

The results were as follows. In cases using (46) and similar sentences where the bear sprayed himself (and so on), the children accepted the sentence in violation of Principle B 37% of the time, while the adult subjects accepted it 6% of the time (Utakis 1995:242). In cases using (47) and similar sentences where every bear sprayed himself (and so on), children accepted the sentence in violation of Principle B 40% of the time, while adults accepted it 11% of the time (Utakis 1995:244). Similar figures were obtained for sentences like (48) in contexts where an affirmative response entailed a Principle B violation (40% vs. 5%; Utakis 1995:245). The rates of acceptance of Principle B violations with referential and quantificational antecedents were very close, then, a fact seemingly inconsistent with the Asymmetry Hypothesis.

It is, however, possible to call this result into question. The reason is that the relevant pictures for trials of sentences like (47) and (48) involved six animals, in the following configuration (Utakis 1995:286–287): there were three bears, for example, each one spraying himself, and three baby tigers who were not being sprayed. In the case of (47), at least, it appears, as Utakis says herself (1995:287), that there was no possible referent for *him* if it was taken referentially, given that none of the tigers or bears was particularly more prominent than the others. Let us imagine a hypothetical child whose grammar does not allow *him* to be bound in (47), faced with this scenario: since there is no evident way to interpret *him* referentially, the danger is that the child, assuming the question must make sense, will construe the pronoun as bound in violation of her own grammar. It has indeed been observed in conducting this type of experiment that some children seem capable of doing this, since they even sometimes answer the question and then correct its grammar (Boster 1991:14–15). It is essential, then, that the experimental setup allow for ‘plausible denial’ (Boster 1991:9–10) or ‘plausible dissent’ (Crain and Thornton 1998: 237). The present experiment from Utakis 1995 seems not to have done this. (Boster, on the other hand, seems to have paid scrupulous attention to this factor; see Boster 1991:9–10.)

Can we at least say that sentences like (48) provided for plausible dissent, since the pronoun *them* could have referred to the group of three baby tigers? This is certainly the case. Unfortunately, however, these sentences also provide less clear evidence than sentences like (47) that there is local binding in violation of Principle B in the first place. In the absence of thorough investigation of these sentences in child grammar, there is a danger that *them* in (48), for those children who seemed to violate Principle B, is really just a referential pronoun referring to the group of bears. Children could have been moved to say ‘yes’ to (48) upon seeing a picture in which every bear was spraying himself because they thought that this process was part of each bear’s spraying the whole group—that is, that the bears might have gone on to spray the other bears, too. (It would be hard to come up with a picture of each bear spraying the whole group simultaneously.) Alternatively, children might think that spraying some of ‘them’ counts as spraying ‘them,’ just as I can be said to have sprayed a group if I aim at it but actually hit only some of its members. (Similarly, I have touched a table if I just touch one of its legs, that is, the member of a proper subset of its parts; see Link 1983, 1997, for the idea that definite plurals like *them* refer to ‘plural

individuals,” of which the individual members are parts.) In the present state of our knowledge of how children deal with plurality and events, I do not think these possibilities can be excluded.

We must conclude, then, that the relevant experiment reported in Utakis 1995 does not offer compelling evidence against the Asymmetry Hypothesis.

3.4 Conclusion on Experiments That Do Not Show the Quantificational Asymmetry

The findings outlined above can be summarized very briefly. The experiments of Lombardi and Sarma 1989 and Boster 1991 offer strong evidence against the Asymmetry Hypothesis. The relevant experiment of Utakis 1995, despite the similarity of its results to the others, does not.

Since the experiments that do show the QA are in fact compatible with the denial of the Asymmetry Hypothesis (section 2), and most of the experiments that do not show the QA are not compatible with the Asymmetry Hypothesis (section 3), I conclude that the Asymmetry Hypothesis is false.

4 Readdressing the Problem of Acquisition

In this section, then, I will assume as a working hypothesis that the Asymmetry Hypothesis is false—that the experiments that do not show the QA accurately reflect the state of the grammars of children learning English, while those that do show it do so because of experimental artifacts, as discussed in section 2. As noted in section 1, this leaves us with a problem. Do children not know Principle B innately? If they do not, how do they learn it? If they do, why do they not initially show knowledge of it, and how do they learn that it in fact applies in their language? I will merely summarize some speculations that have been advanced to answer these questions and will not attempt to decide among them. To do that, more experiments will probably be needed.

4.1 Boster 1991

Boster (1991:30–36) suggests that children at the relevant stage have two homophonous pronouns pronounced *him*, and so on. One variant would be a Principle B pronoun and the other a Principle A anaphor.

This situation, she speculates, may arise as follows. Binding theory is innate. Children first have just one version of each pronoun, subject to Principle B. It is well known that children frequently analyze the reflexive pronouns of English as possessive + *self*, producing *his self*, *her own self*, and similar forms. What happens when a child first notices that adults say *himself*, not *his self*? Since *self* is analyzed as a separate element, it appears that there is a pronoun *him* in a position where the meaning calls for a Principle A anaphor, followed somehow by *self*. (Boster does not put it this way, but we could imagine that *self* might appear to be adjoined to *him*.) Thus, the child posits a new anaphor *him*, and likewise for the other pronouns. This situation only ceases when the child realizes that *himself* is best analyzed as a single word.

Boster (1991:33) offers anecdotal evidence to support this theory: in her own experiments, the children who accepted Principle B violations consistently used immature forms of reflexives

such as *his self*, whereas the children who rejected Principle B violations consistently used the adult forms.

4.2 Fodor 1992

Fodor (1992) also supposes that forms like *him* are ambiguous at the relevant stage, but according to her they are ambiguous between Principle B pronouns and weak anaphors like Dutch *zich*.

It is well known that Dutch, Norwegian, and other Germanic languages display a three-way distinction among pronouns, weak anaphors (called SE-anaphors in Reinhart and Reuland 1993), and strong anaphors (SELF-anaphors in Reinhart and Reuland's terminology). The distinction between weak and strong anaphors seems roughly to be that weak anaphors are used when identity of arguments is somehow to be expected, in the case of certain "inherently reflexive" predicates like *wash*, *be ashamed*, *enjoy*, and *behave*. Strong anaphors are used elsewhere. In Dutch, for example, it is possible to say (49) and (50), with weak anaphors and inherently reflexive predicates. However, it is not possible to say (51a), with a weak anaphor and a nonreflexive predicate; instead, one has to say (51b), with a strong anaphor (Fodor 1992:20).

(49) Jan_i droogde zich_i af.
Jan dried himself off
'Jan dried himself.'

(50) Jan_i schaamde zich_i.
Jan shamed himself
'Jan was ashamed.'

(51) a. *Jan_i verraste zich_i.
Jan surprised himself
b. Jan_i verraste zichzelf_i.
Jan surprised himself

Interestingly for present purposes, in at least one language, Frisian, the weak anaphor is homophonous with the Principle B pronoun, as shown by the grammaticality of (52) and (53), analogous to (49) and (50), and the ungrammaticality of (54) and (55), analogous to (51a) (Fodor 1992:21).

(52) Hy_i droege him_i of.
he dried himself off
'He dried himself.'

(53) Hy_i skammet him_i.
he shames himself
'He is ashamed.'

(54) *Hy_i beoardielet him_i.
he judges himself
'He judges himself.'

- (55) *Hy_i joech him_i it boek.
 he gave himself the book
 ‘He gave himself the book.’

There is also evidence that *him* can function as a normal, nonanaphoric pronoun in Frisian (Fodor 1992:21).

Fodor suggests, then, that Universal Grammar contains entries for a three-way distinction among pronouns along the lines of Dutch, and that learners of languages like English and Frisian, which do not contain three separate forms, have to collapse or alter this tripartite template in some way (Fodor 1992:21–24). For present purposes, the relevance of this is as follows. Children learning English, faced with only two pronominal forms, could conjecture that the weak anaphor in English is homophonous with the pronoun, as in Frisian (Fodor 1992:24–25). If they did this, the forms pronounced *him* and *her* would not be subject to Principle B, since weak anaphors are not subject to Principle B. Fodor suggests that the error would eventually be corrected when the child hears and correctly analyzes a verb like *behave* or *enjoy* used with a reflexive form like *himself*. Only a small number of verbs (*perjure* is another) are such that their internal argument could not possibly be conjoined with the subject, since verbs like *wash* and *dry* do not fulfill this criterion; so the error might easily persist for a while.

Is it plausible to suppose that learners would continue using this putative weak reflexive as the reflexive object of verbs like *touch*, which are clearly not inherently reflexive? Fodor suggests (1992:25) that they may in fact use it more with verbs like *wash* and *brush* (which feature heavily in the experimental literature on this question) than with verbs like *touch*; it would be necessary to do a large-scale analysis of answers verb by verb in order to find out. But we may also imagine that the category of inherent reflexivity is sufficiently vague and subtle that children take a few years to work out what verb meanings it encompasses, and thus allow their weak anaphors after *touch* because their grasp of how this category is realized in their language is not sufficiently sure. After all, one can touch oneself much more easily than one can wash oneself—is it immediately obvious that the latter notion but not the former should be inherently reflexive?

4.3 Utakis 1995

Utakis (1995:194) suggests that for children of the relevant ages, the forms we call pronouns are in fact indefinites, equivalent in meaning to *someone*. So when, for example, children say ‘yes’ to (56) when presented with a picture in which every bear sprays himself and no one else is sprayed, it is because they interpret the question as (57).

- (56) Did every bear spray him?
 (57) Did every bear spray someone?

This makes the following prediction: children who apparently allow Principle B violations in (56) given the scenario just described should also respond affirmatively to it when the picture shows three bears each spraying a different baby tiger. Utakis in fact tested such scenarios, and carried

out an ANOVA to see if there was a significant correlation between these two groups (Utakis 1995:278–281). The results are unclear, however: there is no significant correlation if the test is carried out on all the potentially relevant children, and a significant result is obtained only if one child whose results are arguably slightly anomalous is excluded altogether (Utakis 1995:280–281).

4.4 *Elbourne 2003*

In a previous version of this work (Elbourne 2003), I suggested that Principle B is subject to a parameter, and that children who appear not to have knowledge of it have not yet learned that their language selects it.

It is uncontroversial that some languages show Principle B effects and some do not: languages evidently without Principle B include Old English (Penning 1875:3–4, Farr 1905:8, Nagucka 1981, Mitchell 1985:112, Keenan 1994), Middle English (Mustanoja 1966:153, Visser 1970:422–433, Keenan 1994), and Maori (Keenan 1975, Fodor 1994). Middle English is particularly interesting in this regard, in that it contained pronouns subject to neither Principle A nor Principle B, yet at the same time contained the newly created *self*-anaphors subject to Principle A (Farr 1905:33, Mustanoja 1966:153, Keenan 1994). In the following examples, Middle English simple third-person singular masculine accusative pronominal forms are referential with no local antecedent (58), referential or bound as a reflexive with a local referential antecedent (59), and bound and reflexive with a local quantifier antecedent (60).

(58) He held euir agane the king,
And hatit *hyme* atour all thing.
(1375, Barbour, *Bruce*, ix.465)
'He remained opposed to the king, and hated him above all things.'

(59) . . . He chaunged his array
And cladde *hym* as a pore laborer.
(Chaucer, *CT*, Knight's Tale, 551)
'He changed his clothing and clad himself as a poor laborer.'

(60) Every wight out at the dore *him* dighte.
(Chaucer, *Troil.*, 948)
'Every man threw himself out of the door.'

And (61) shows a *self*-anaphor in Middle English used in a straightforward reflexive sense.

(61) In this wise *himself* he spilte
With his folhaste and deth he nam.
(1290, Gower, *Conf. Am. I*, 328)
'In this way he destroyed himself with his foolish haste and he attained death.'

Given that children learning Modern English do well in all studies with Principle A anaphors, it seems that adult Middle English closely resembled children's Modern English, as far as pronouns are concerned.

On the basis of these observations, I proposed in Elbourne 2003 that Principle B was subject to a parameter and that the default state for this parameter was “off”,¹⁵ that is, learners would assume that there was no Principle B in their language until evidence led them to conclude otherwise. So Middle English-speaking adults and contemporary children would both be in this initial stage.¹⁶ By claiming that Principle B was innate, this position would account for its near-universality and the fact that children manage to display knowledge of such an abstract principle at all; by claiming that it was subject to a parameter, it would account for the fact that learners initially do not know whether it applies in their language.¹⁷

There is a serious problem with this position, however: namely, it seems to run afoul of the problem of negative evidence. How could learners in this position possibly get to know that Principle B applied in their language? Every sentence they heard, after all, would be compatible with their grammar. It looks like the only way out would be to allow that, over time, the nonoccurrence of predicted constructions (here, Principle B violations) could in fact serve as indirect negative evidence for learners, an idea mentioned briefly in the context of Principle B by Fodor (1994:446). This cannot, I believe, be ruled out, but it is clearly more satisfactory if we can assume that the default state for the Principle B parameter is “on,” and that certain types of clear Principle B violations can flick it to “off,” as Fodor (1994) argues.

4.5 Conclusion on Theories of Acquisition

I have not attempted to decide among the theories just reviewed, or even to offer much in the way of critical commentary. We can conclude simply that there are at least four interesting theories already present in the literature that attempt to explain the relatively late mastery of Principle B consistently with the denial of the Asymmetry Hypothesis. These obviously provide fertile ground for future experiments.

5 Conclusion

This article has investigated the hypothesis that I have referred to as the Asymmetry Hypothesis, which claims that the language faculties of children learning English are such that Principle B applies to bound pronouns but not to referential pronouns. If true, this hypothesis would have important consequences for binding theory, giving particular support to the theory developed in

¹⁵ Strictly speaking, I should rather say that the default state for the Principle B parameter with respect to strong pronouns was “off,” since McKee (1992) and others have shown that there is no delay with regard to Principle B in the acquisition of Romance clitics. It is possible that weak pronouns, too, in the tripartite scheme of Cardinaletti and Starke 1999, are also immune from the delay of Principle B. See Cardinaletti and Starke 1995 for discussion of pronoun classification and acquisition. Thanks to Michal Starke for valuable discussion of this point.

¹⁶ Recall that Boster (1991:33) reports that the majority of the children in her experiments who accepted violations of Principle B also used *him* as a reflexive at some point in their spontaneous speech. So the apparent lack of Principle B is encountered in production as well as in comprehension, as we would expect if there was a mis-set parameter.

¹⁷ Alternatively, one could say that all languages have Principle B but in some languages its effects are somehow hidden, perhaps because the domain in which it applies is very small, or perhaps because no pronominal forms are specified as falling under it. Then an analogue of the proposal in the text would claim that the default state of the language faculty is such that the relevant hiding mechanism is in place, and it takes learners some time to realize that it should be lifted.

Reinhart 1983, as pointed out in Grodzinsky and Reinhart 1993. But a review of the experimental evidence makes it seem likely that the hypothesis is false.

A prominent tool in my argument against the Asymmetry Hypothesis was the Saliency Hypothesis. As remarked in section 2.6, the truth of the Saliency Hypothesis implies the falsity of an assumption behind the truth-value judgment task methodology. This methodology has already been questioned in closely related ways by Musolino and Lidz (2003) and Drozd (forthcoming). The project of submitting the Saliency Hypothesis to direct experimental testing naturally suggests itself. This would presumably require some way of manipulating saliency as a variable.

We should not leave this subject without standing back briefly and asking if we have come up with support for theories like Chomsky's (1981) that treat bound and referential pronouns alike for purposes of Principle B. In fact, the appropriate conclusion seems to be that neither Reinhart's theory nor Chomsky's gains any support if the account of the acquisition facts given here is correct. Children between the ages of 3 and 6 simply appear not to know Principle B at all. If, at the age when knowledge of this principle starts to become firm, it turns out that there is still no asymmetry between bound and referential pronouns, supporters of Reinhart's theory can simply say that children at this age are able to compare representations efficiently and apply Rule I effectively. Thus, it seems that lack of support from acquisition facts is equally distributed at this point.

References

- Arnold, Jennifer, Jared Novick, Sarah Brown-Schmidt, Janet Eisenband, and John Trueswell. 2001. Knowing the difference between girls and boys: The use of gender during on-line pronoun comprehension in young children. In *Proceedings of the 25th annual Boston University Conference on Language Development*, ed. by Anna H.-J. Do, Laura Domínguez, and Aimee Johansen, 59–69. Somerville, Mass.: Cascadilla Press.
- Avrutin, Sergey, and Rosalind Thornton. 1994. Distributivity and binding in child grammar. *Linguistic Inquiry* 25:165–171.
- Avrutin, Sergey, and Kenneth Wexler. 1992. Development of Principle B in Russian: Coindexation at LF and coreference. *Language Acquisition* 2:259–306.
- Boster, Carole Tenny. 1991. Children's failure to obey Principle B: Syntactic problem or lexical error? Ms., University of Connecticut, Storrs.
- Cardinaletti, Anna, and Michal Starke. 1995. The tripartition of pronouns and its acquisition: Principle B puzzles are ambiguity problems. In *NELS 25. Vol. 2, Papers from the workshops on language acquisition and language change*, ed. by Jill N. Beckman, 1–12. Amherst: University of Massachusetts, GLSA.
- Cardinaletti, Anna, and Michal Starke. 1999. The typology of structural deficiency: A case study of the three classes of pronouns. In *Clitics in the languages of Europe*, ed. by Henk van Riemsdijk, 145–233. Berlin: Mouton de Gruyter.
- Chien, Yu-Chin, and Kenneth Wexler. 1990. Children's knowledge of locality conditions in binding as evidence for the modularity of syntax and pragmatics. *Language Acquisition* 1:225–295.
- Chomsky, Noam. 1981. *Lectures on government and binding*. Dordrecht: Foris.
- Crain, Stephen. 1991. Language acquisition in the absence of experience. *Behavioral and Brain Sciences* 14:597–650.
- Crain, Stephen, and Cecile McKee. 1986. The acquisition of structural restrictions on anaphora. In *Proceed-*

- ings of NELS 16*, ed. by Stephen Berman, Jaewoong Choe, and Joyce McDonough, 94–110. Amherst: University of Massachusetts, GLSA.
- Crain, Stephen, and Rosalind Thornton. 1998. *Investigations in Universal Grammar: A guide to experiments on the acquisition of syntax and semantics*. Cambridge, Mass.: MIT Press.
- Dowty, David. 1979. *Word meaning and Montague Grammar*. Dordrecht: Reidel.
- Drozd, Kenneth. Forthcoming. Learnability and linguistic performance. *Journal of Child Language* 31.
- Elbourne, Paul. 2002. Situations and individuals. Doctoral dissertation, MIT, Cambridge, Mass.
- Elbourne, Paul. 2003. Are our children speaking Middle English? Paper presented at CUNY Syntax Supper.
- Farr, James Marion. 1905. *Intensives and reflexives in Anglo-Saxon and Early Middle-English*. Baltimore: J. H. Furst.
- Fodor, Janet Dean. 1992. Designated triggers versus the Subset Principle. Ms., CUNY Graduate Center, New York.
- Fodor, Janet Dean. 1994. How to obey the Subset Principle: Binding and locality. In *Syntactic theory and first language acquisition: Cross-linguistic perspectives*. Vol. 2, *Binding, dependencies, and learnability*, ed. by Barbara Lust, Gabriella Herman, and Jaklin Kornfilt, 429–451. Hillsdale, N.J.: Lawrence Erlbaum.
- Fox, Danny. 2000. *Economy and semantic interpretation*. Cambridge, Mass.: MIT Press.
- Fritzley, V. Heather, and Kang Lee. 2003. Do young children always say yes to yes-no questions? A metadevelopmental study of the affirmation bias. *Child Development* 74:1297–1313.
- Grodzinsky, Yosef, and Tanya Reinhart. 1993. The innateness of binding and coreference. *Linguistic Inquiry* 24:69–101.
- Heim, Irene. 1993. Anaphora and semantic interpretation: A reinterpretation of Reinhart's approach. Sfs-Report-07-93, Seminar für Sprachwissenschaft, University of Tübingen. (Also published in *The interpretive tract*, ed. by Uli Sauerland and Orin Percus, 205–246. MIT Working Papers in Linguistics 25. Cambridge, Mass.: MIT, Department of Linguistics and Philosophy, MITWPL, 1990.)
- Heim, Irene, and Angelika Kratzer. 1998. *Semantics in generative grammar*. Oxford: Blackwell.
- Heim, Irene, Howard Lasnik, and Robert May. 1991a. On "Reciprocal scope." *Linguistic Inquiry* 22: 173–192.
- Heim, Irene, Howard Lasnik, and Robert May. 1991b. Reciprocity and plurality. *Linguistic Inquiry* 22: 63–101.
- Karttunen, Lauri. 1977. Syntax and semantics of questions. *Linguistics and Philosophy* 1:3–44.
- Keenan, Edward. 1975. Logical expressive power and syntactic variation in natural language. In *Formal semantics and natural language*, ed. by Edward Keenan, 406–421. Dordrecht: Kluwer.
- Keenan, Edward. 1994. Creating anaphors: An historical study of the English reflexive pronouns. Ms., UCLA, Los Angeles, Calif.
- Levin, Beth, and Malka Rappaport Hovav. 1995. *Unaccusativity: At the syntax–lexical semantics interface*. Cambridge, Mass.: MIT Press.
- Link, Godehard. 1983. The logical analysis of plurals and mass terms: A lattice-theoretical approach. In *Meaning, use and interpretation of language*, ed. by Rainer Bäuerle, Christoph Schwarze, and Arnim von Stechow, 302–323. Berlin: Walter de Gruyter. (Also published in *Algebraic semantics in language and philosophy*, 11–34. Stanford, Calif.: CSLI Publications, 1998.)
- Link, Godehard. 1997. Ten years of research on plurals—where do we stand? In *Plurality and quantification*, ed. by Fritz Hamm and Erhard Hinrichs, 19–54. Dordrecht: Kluwer. (Also published in *Algebraic semantics in language and philosophy*, 163–187. Stanford, Calif.: CSLI Publications, 1998.)
- Lombardi, Linda, and Jaya Sarma. 1989. Against the bound variable hypothesis of the acquisition of Condition B. Paper presented at the annual meeting of the Linguistic Society of America, Washington, D.C.
- Marantz, Alec. 1997. No escape from syntax: Don't try morphological analysis in the privacy of your own lexicon. In *Proceedings of the 21st Annual Penn Linguistics Colloquium*, ed. by Alexis Dimitriadis,

- Laura Siegel, Clarissa Surek-Clark, and Alexander Williams, 201–225. Pennsylvania Working Papers in Linguistics 4.2. Philadelphia: University of Pennsylvania, Penn Linguistics Club.
- McKee, Cecile. 1992. A comparison of pronouns and anaphors in Italian and English acquisition. *Language Acquisition* 2:21–54.
- Mitchell, Bruce. 1985. *Old English syntax*. Vol. 1. Oxford: Clarendon Press.
- Musolino, Julien, and Jeffrey Lidz. 2003. Why children are not universally successful with quantification. Ms., Indiana University, Bloomington, and Northwestern University, Evanston, Ill.
- Mustanoja, Tauno. 1966. *A Middle English syntax*. Part I. Helsinki: Société Néophilologique.
- Nagucka, Ruta. 1981. A reflexive function of the personal pronoun in Old and Middle English. *Bulletin de la Société Polonaise de Linguistique* 38:33–46.
- Parsons, Terry. 1990. *Events in the semantics of English*. Cambridge, Mass.: MIT Press.
- Penning, G. E. 1875. *A history of the reflexive pronouns in the English language*. Bremen: Hevriř Frese.
- Philip, William, and Peter Coopmans. 1995. The role of the Chain Condition in the acquisition of anaphora. Ms., Onderzoeksinstituut voor Taal en Spraak (OTS), Utrecht University.
- Pylkkänen, Liina. 2000. Representing causatives. In *Proceedings from Semantics and Linguistic Theory X*, ed. by Brendan Jackson and Tanya Matthews, 132–148. Ithaca, N.Y.: Cornell University, CLC Publications.
- Reinhart, Tanya. 1983. *Anaphora and semantic interpretation*. London: Croom Helm.
- Reinhart, Tanya. 1997. Strategies of anaphora resolution. UiL OTS Working Paper TL97-007. Utrecht: Utrecht University, Onderzoeksinstituut voor Taal en Spraak (OTS).
- Reinhart, Tanya, and Eric Reuland. 1993. Reflexivity. *Linguistic Inquiry* 24:657–720.
- Thornton, Rosalind. 1990. Adventures in long-distance moving: The acquisition of complex *wh*-questions. Doctoral dissertation, University of Connecticut, Storrs.
- Thornton, Rosalind, and Kenneth Wexler. 1999. *Principle B, VP-ellipsis, and interpretation in child grammar*. Cambridge, Mass.: MIT Press.
- Utakis, Sharon. 1995. Quantification and definiteness in child grammar. Doctoral dissertation, CUNY, New York.
- Visser, F. T. 1970. *An historical syntax of the English language*. Part I. Leiden: E. J. Brill.
- Williams, Edwin. 1991. Reciprocal scope. *Linguistic Inquiry* 22:159–173.
- Winter, Yoand. 2001. *Flexibility principles in Boolean semantics: The interpretation of coordination, plurality, and scope in natural language*. Cambridge, Mass.: MIT Press.

Universität Potsdam
 Institut für Linguistik/Allgemeine Sprachwissenschaft
 Postfach 601553
 D-14415 Potsdam
 Germany
 elbourne@rz.uni-potsdam.de