LANGUAGE LEARNING AND DEVELOPMENT, *2*(3), 195–218

On the Semantics of the Subset Principle

Julien Musolino
Indiana University

Virtually all generative accounts of language development assume that syntactic acquisition is guided by a learnability constraint called the Subset Principle (SP). In essence, SP forces learners to initially select the value of a parameter that generates the smallest possible language. Recent work on the acquisition of semantics suggests that there also are semantic subset problems; therefore, one needs a Semantic Subset Principle (SSP) to solve these problems. This article shows that (a) to the extent that semantic subset problems exist, the SSP is not the correct solution to them and (b) that semantic subset problems most likely do not exist in the first place. These conclusions have important implications for theories of the acquisition of semantic knowledge and for the study of language acquisition more generally. First, they provide a basis for delimiting the class of problems that theories of the acquisition of semantics ought to be responsible for. Second, these conclusions underscore the fact that data from empirical investigations of child language can be brought to bear on the formulation of linguistic theory.

A fundamental tenet of modern linguistic theory is that language acquisition cannot be explained unless one assumes that human beings come to the task already equipped with an understanding of the basic principles of language. In this regard, virtually all generative accounts of language development assume that the acquisition of syntax is guided by a built-in learnability constraint called the Subset Principle (SP; Berwick, 1985; Hyams, 1992; Pinker, 1995; Wexler & Manzini, 1987; among many others). In parameter-setting models of language acquisition, SP forces learners to be conservative and initially choose the value of a parameter that generates the smallest possible language. This guarantees that learners will be able to revise their initial hypothesis on encountering relevant *positive* evidence. Without SP, the argument goes, learners who have hypothesized too large a language

Correspondence should be addressed to Julien Musolino, Department of Speech and Hearing Sciences, Indiana University, 200 South Jordan Avenue, Bloomington, IN 47405–7002. E-mail: musolino@indiana.edu

would need *negative* evidence to revise their initial hypothesis. However, negative evidence is believed to not be available—hence the need for SP.[1]

Recent work on the acquisition of semantics suggests that there are also semantic subset problems; therefore, one needs a Semantic Subset Principle (SSP) to solve these problems (Crain, 1992, 1993; Crain, Ni, & Conway, 1994; Crain & Philip, 1993; Crain & Thornton, 1998). If true, these conclusions have important implications for the study of the acquisition of meaning. To be sure, the SSP makes specific predictions pertaining to the acquisition of a broad range of semantic phenomena including quantification, negation, disjunction, scope, and entailment relations, to mention just a few. In this article, I provide a detailed examination of the logical status and the empirical consequences of the SSP. I conclude that despite the ingenious insight on which it is built, and its many promising applications, the SSP should be abandoned as a learnability constraint on the acquisition of semantics. Specifically, I show that (a) to the extent that semantic subset problems exist, the SSP is not the correct solution to them; and (b) that semantic subset problems most likely do not exist in the first place. More generally, I take these conclusions to have important implications for theories of the acquisition of semantic knowledge. First, they provide a basis for delimiting the class of problems that such theories ought to be responsible for. Second, these conclusions underscore the fact that data from empirical investigations of child language can be brought to bear on the formulation of linguistic theory.

The discussion is organized as follows: The first section introduces the SSP and the existing empirical evidence supporting it. Section 2 lays out the theoretical concepts and vocabulary that are later used in my assessment of the SSP. Section 3 provides a detailed evaluation of the SSP and shows that (a) as currently formulated and implemented, the SSP is flawed; (b) under any reasonable reformulation of the SSP there is solid empirical evidence against it coming from a growing line of work on the acquisition of quantificational phenomena (Empirical Predictions of the SSP); and (c) that due to the compositional nature of semantics and the relation between syntax and semantics, it is quite likely that there are no semantic subset problems in the first place and, hence, no logical *raison d'être* for the SSP (The Logical Status of the SSP). In doing so, I provide a concrete demonstration that a key difference between English and Chinese, which has been taken as motivating

[1]*Positive* evidence is evidence from the input that certain structures or interpretations are available in one's language. For example, a sentence like, "Not every student can afford a new car," would provide positive evidence that "not" can combine with "every" in English. By contrast, *negative* evidence is evidence that certain structures or interpretations are *not* possible in one's language. For example, in English, "not" cannot combine with "some," as illustrated by the deviant nature of a sentence like, "Not some students can afford a new car." Therefore, the fact that a sentence like, "Not some students can afford a new car," is not well-formed would constitute negative evidence.

the need for a constraint like the SSP, can in fact be acquired without recourse to the SSP. Finally, the fourth section offers some concluding remarks.

THE SSP

Crain and Thornton (1998) define semantic subset problems as follows:

> Sometimes, more than one interpretation of a sentence is made available by Universal Grammar. To further complicate matters, these alternatives may form a subset/superset relation; that is, the circumstances that make the sentence true on one interpretation may be a proper subset of the circumstances that make it true on another interpretation. In such cases, a semantic subset problem arises if the target language includes the subset reading, but not the superset reading. To avoid semantic subset problems, the interpretive options for sentences must be ordered in the LAD [Language Acquisition Device] by a principle instructing learners to initially choose the representation that is true in the smallest set of circumstances. This is called the *Semantic Subset Principle* (Crain, 1992, 1993; Crain et al., 1994; Crain & Philip, 1993). (pp. 117–118)

In turn, they formulate the SSP as shown below:

> Suppose that the interpretive component of Universal Grammar makes two interpretations, A and B, available for a sentence, S. If so, then see if S is true in a narrower range of circumstances on interpretation A than on interpretation B. If so, then A will be hypothesized before B in the course of language development. (p. 118)

To test the predictions of the SSP, Crain et al. (1994) investigated children's interpretation of ambiguous sentences like (1) in which the focus operator *only* can either associate with the object noun phrase (NP; a house; 1b) or the entire verb phrase (VP; painting a house; 1a):

(1) The dinosaur is only painting a house.
  a. The only thing the dinosaur is doing is painting a house (VP).
  b. The only thing the dinosaur is painting is a house (NP).

Notice that the two readings of (1) are arranged in a subset–superset configuration because (1a) entails (1b). In other words, if it is true that the only thing that the dinosaur is doing is painting a house, (1a), then it must also be true that the only thing that the dinosaur is painting is a house, (1b), but not *vice versa.* That is, if it is true that the only thing that the dinosaur is painting is a house, (1b), it does not necessarily follow that the only thing that the dinosaur is *doing* is painting a house, (1a). To be sure, the dinosaur could be painting a house while eating an apple or flying a kite, which would make (1b) true but not (1a). The SSP, therefore, predicts that children should initially hypothesize that sentences like (1) are unambiguous and that they can only be interpreted as in (1a). To quote Crain et al.,

> The circumstances corresponding to the alternative readings of [(1)] are in a subset–superset relationship. Therefore, the semantic subset principle compels children to initially hypothesize the (a) reading … . At a later point in development, the (b) reading will become available, in response to evidence from the input. (p. 456)

Crain et al. tested a group of preschoolers (age, *M* = 4;9) using a picture-verification task. One of the pictures used by Crain et al. showed an elephant and a dinosaur. The elephant was holding a balloon and painting a car, and the dinosaur was flying a kite and painting a chair and a house.

Crain et al. (1994) report that children often rejected (1) as a description of the picture mentioned above on the grounds that the dinosaur was also painting a chair and flying a kite.[2] In other words, children rejected (1) by invoking the fact that reading (1a) was false. Crain et al. reasoned that (1a) must, therefore, be the only reading available to children; thus, they took this result as evidence for the operation of the SSP. In a second, related experiment, Crain et al. investigated the way children (and adults) interpret ambiguous sentences like (2):

(2) The big elephant is the only one playing a guitar.
  a. The big elephant is the only thing playing a guitar.
  b. The big elephant is the only elephant playing a guitar.

Here, too, the alternative interpretations of (2) are arranged in a subset–superset relation such that (2a) entails (2b). In other words, if it is true that the only thing that is playing a guitar is the big elephant, then it follows that the only elephant that is playing a guitar is the big elephant. Consequently, the SSP predicts that children should initially hypothesize that (2a) is the only possible reading for sentences like (2). Crain et al. report that 8 of the 12 preschoolers they tested (*M* age = 4;8) consistently rejected sentences like (2) by invoking the fact that interpretation (2a) was false in the relevant context. As before, Crain et al. reasoned that (2a) must, therefore, have been the only interpretation available to children, as predicted by the SSP.

---

[2]In fact, Crain et al. (1994) used sentences like, "Only the dinosaur is painting a house," to test children's interpretation of "only." These authors did so based on the finding that some children always interpret "only" as if it was occurring in preverbal position. In other words, according to Crain et al., so-called "verb phrase-oriented" children interpret sentences like, "Only the dinosaur is painting a house," as though they meant, "The dinosaur is only painting a house."

THEORETICAL BACKGROUND

This section is intended for readers who are not familiar with generative syntactic theory. My goal here is to introduce the theoretical concepts and vocabulary that are later used in my discussion of the SSP. Before I start, I would like to point out that the ideas I am about to describe represent standard notions that very few, if any, generative linguists would take issue with. The point I wish to make, ultimately, is one about psychology; namely, that the SSP is not psychologically plausible.[3] However, because the SSP is a piece of generative theory, a systematic evaluation of this principle requires an understanding of certain aspects of the general framework within which it is couched. Specifically, the notions I need to introduce are (a) syntactic structure; (b) c-command, which is an abstract relation that is defined over hierarchical structure; (c) the notion of scope, which is also defined over hierarchical structure and in terms of c-command; and (d) the concept of displacement or syntactic movement.

One of the central and most famous conclusions of Chomsky's (1957) foundational work in syntax is that linguistic representations are hierarchical and that the rules of grammar make reference to this hierarchal organization. Therefore, a sentence like (3) is not just a string of words but, rather, can be represented as being hierarchically organized as shown in (4a and 4b), which are notational variants of one another.

(3) The man danced with the woman.

(4)

a. [S [NP The man] [VP danced [PP with [NP the woman]]]].

b.

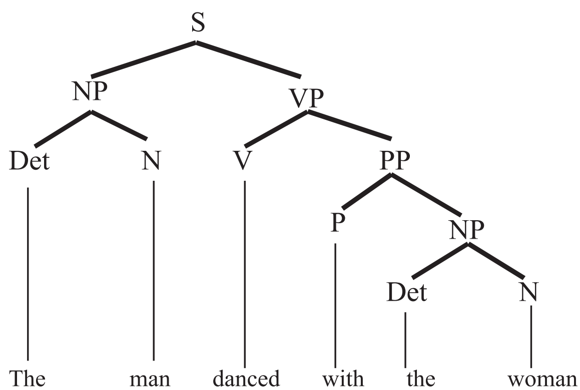

[3]The claim about psychological implausibility I am making is specific to the Semantic Subset Principle and does not extend to other aspects of generative theory.

In addition, syntacticians have noticed that a broad range of seemingly disparate linguistic phenomena can receive a principled explanation if one assumes the notion of c-command. C-command, in turn, is defined as follows:

(5) x c-command y iff.
 a. x ≠ y.
 b. Neither x dominates y nor y dominates x.
 c. The first branching node that dominates x also dominates y.

A useful rule of thumb to calculate c-command without using the formal definition in (5) is to start with the element whose c-command domain one wants to calculate, go up in the tree structure to the first branching node, and then go down. Everything on the way down from the branching node is contained within the c-command domain of the element in question. Therefore, in our example in (3), the NP (the man) c-commands the VP, the prepositional phrase (PP), and the NP (the woman).

Hierarchical structure and c-command, in turn, each play a crucial role in the notion of scope. The notion of scope can be illustrated using a simple mathematical analogy. Consider these mathematical expressions: $2 \times (3 + 5)$ and $(2 \times 3) + 5$. The scope of $2 \times$ (the number 2 followed by the multiplication sign) can be thought of as its domain of application. Therefore, in $2 \times (3 + 5)$, $(3 + 5)$ falls within the scope of $2 \times$. In contrast, in $(2 \times 3) + 5$, 3 falls within the scope of $2 \times$, whereas 5 falls outside of its scope. Finally, notice that different scope relations give rise to different results once the expressions are computed. We can now consider the notion of scope as it applies to language by considering a sentence like (6):

(6) Everybody didn't smile.

Notice that (6) is ambiguous. On one reading, it can be paraphrased as meaning that not everybody smiled (i.e., some people did, but others did not). On this reading, the phrase *everybody* is interpreted within the scope—or domain of application—of negation, giving rise to what I shall refer to as the "not every" reading (abbreviated as "not > every"). Alternatively, (6) can be paraphrased as meaning that everybody is such that they did not smile. In other words, nobody smiled. In this case, the phrase *everybody* is interpreted outside the scope of negation, giving rise to what I shall refer to as the "none" reading (abbreviated "every > not"). Therefore, the ambiguity displayed in (6) can be thought of as a scope ambiguity between the phrase *everybody* and negation. Just as different scope relations in our mathematical analogy gave rise to different results, different scope relations between *everybody* and negation give rise to different interpretations for (6). It should now also be clear from our mathematical analogy that the notion of scope must be defined in hierarchical terms rather than in terms of linear order. To be sure, the expressions $2 \times (3 + 5)$ and $(2 \times 3) + 5$ are identical from the point of view of linear order. Crucially, however, their structures are different as indicated by the parentheses. The same holds of language, and this is where the notion of c-command introduced earlier becomes relevant. Specifically, we can define scope in terms of c-command, thereby defining scope over hierarchical structures because c-command is a relation that is itself defined over such structures. Therefore, we get the scope principle in (7):

(7) *Scope principle:* An expression α is interpreted as taking scope over an expression β iff α c-commands β.

The situation described above now creates a puzzle: How can one sentence, for example (6), in which the structural position of the elements seems fixed, nevertheless be scopally ambiguous? This is where the notion of displacement, or syntactic movement, needs to be introduced. Consider an example like (8) in which the wh-phrase *what* is interpreted as corresponding to the object of the verb "eat" and yet occurs in the initial position of the sentence instead of where objects typically occur in English, namely adjacent to the verb, as shown in (9):

(8) What did John eat __?
(9) John ate an apple.

This illustrates the phenomenon of displacement; namely, the fact that an element can occur in a position that is different from the one in which it is interpreted, giving the illusion of syntactic movement. In the cases of (8) and (9), we are dealing with *overt* displacement in the sense that one can "see" (or "hear") that *what* is displaced. As many other linguistic phenomena demonstrate, however, there is also a *covert* version of displacement. In this case, we must assume that displacement has occurred because of interpretive differences in the absence of phonological consequences. One example of covert displacement in action is our sentence in (6). Notice that in the structural representation corresponding to (6), given here in (10), *everybody* occurs outside the c-command domain of negation. Therefore, by our scope principle, this configuration should correspond to the reading on which *everybody* is interpreted outside the scope of negation; namely, the none or every > not reading (nobody smiled). To get the "not every" reading from the structure in (10),[4] the phrase *everybody* and negation must be

[4]The notation I am adopting here is standard X-bar theoretic notation (e.g., see Haegeman, 1994). According to X-bar theory, all phrases have the same abstract structure and involve three hierarchical levels: a zero level (X^0—read X zero), an intermediate level (X'—read X bar), and a maximal level (XP or X"—read X double bar). To give an example, a verb (V; a zero-level category) would combine with its complement to form an intermediate level category (V'). This intermediate level category would then combine with a specifier to form a fully fledged verb phrase. The X-bar format can also be applied to clausal categories such as a sentence, which is replaced by an inflectional phrase (IP). Following the format described above, IP also contains three hierarchical levels: a maximal level, IP; an intermediate level, I'; and a zero level, I^0.

covertly rearranged such that negation ends up c-commanding *everybody,* thereby yielding the "not everybody" interpretation:

10. [IP [NP Everybody [I' didn't [VP smile]]]].

To show precisely how the not > every reading of (10) is derived via covert displacement, I need to introduce three more assumptions that are both independently motivated and standard (e.g., see Chomsky, 1995; Haegeman, 1994; Hornstein, 1995). The first concerns a level of syntactic representation called *logical form* (LF); the second, the technical implementation of displacement; and the third, the syntactic position from which subjects originate. These assumptions are spelled out below:

(11) LF is the level of linguistic representation at which all grammatical structure relevant to semantic interpretation is provided.
(12) Syntactic movement is copying and deleting (i.e., a full copy of a moved constituent is left at the launching site). Moreover, at LF, all copies but one must be deleted.
(13) Subjects originate from a position internal to VP.

Given these assumptions, let us now see how the relevant interpretive facts regarding sentences like, "Everybody didn't smile," can be derived. The LF structure of "Everybody didn't smile" is given in (14):

(14) [IP [NP Everybody [I' didn't [VP everybody smile]]]].
a. [IP [NP ~~Everybody~~ [I' didn't [VP everybody smile]]]].
b. [IP [NP Everybody [I' didn't [VP ~~everybody~~ smile]]]].

Because subjects start in a VP internal position and movement involves copying, we end up with a copy of *everybody* in VP and one in the inflectional phrase (IP). Recall that at LF, all copies but one must be deleted. This gives us two options: Either the VP or the IP copy may be deleted. Deleted copies in (14) are crossed out. If the IP copy is deleted, notice that the remaining copy (the VP copy) occurs in the c-command domain of negation. Given our scope principle, *everybody* must in this case be interpreted in the scope of negation, yielding the not > every interpretation. On the other hand, if the VP copy is deleted, then the remaining copy (the IP copy) occurs outside the c-command domain of negation and, thus, outside its scope. It follows that in this case, *everybody* will be interpreted outside the scope of negation, thereby yielding the every > not interpretation. Therefore, we have derived the fact that sentences like, "Everybody didn't smile," can either receive a "none" or a "not all" reading in English. At this point, it is important to emphasize the fact that the assumptions described above are independently motivated and not just a set of *ad hoc* assumptions that I am making to account for the possible meaning of one sentence of English.

The last part that I need to introduce is the logical notation that I use to describe scopally ambiguous sentences. Following standard practice in basic logical notation, I use $\forall$ as the symbol for the universal quantifier (corresponding to linguistic expressions like *all* or *every*), $\exists$ as the symbol for the existential quantifier (corresponding to expressions like *some*), $\neg$ for negation, and $\wedge$ for the conjunction "and." Equipped with these symbols, we can now formally represent the different interpretations of sentences like (15) in a compact and precise fashion:

(15) Everyone didn't smile.
a. $\forall x$ [person (x) $\rightarrow \neg$ smile (x)].
b. $\neg\forall x$ [person (x) $\rightarrow$ smile (x)].

The logical representation in (15a) says that for every x, if x is a person, then it is not the case that x smiled. In other words, nobody smiled. This corresponds to the every > not reading on which *everybody* is interpreted outside the scope of negation. The representation in (15b) says that it is not the case that for every x, if x is a person, x smiled. This corresponds to the not > every reading on which *everybody* is interpreted within the scope of negation. Equipped with an understanding of the core theoretical assumptions discussed above, we are now in a position to evaluate the SSP.

EVALUATING THE SSP

Let us begin with the evidence presented by Crain et al. (1994). Notice that the logic of the argument made by these authors is based solely on an examination of the meaning of *only* in English. Missing from Crain et al.'s reasoning is any mention of the existence of a language that, unlike English, would allow only reading (1a) for sentences like (1). If such a language does indeed exist, then learners exposed to that language and initially hypothesizing both (1a) and (1b) as possible interpretations for sentences like (1) would indeed face a semantic subset problem in the absence of negative evidence informing them that (1b) is not an option in their language (however, see "The Logical Status of the SSP" section in this article for a different view). In the absence of such a language, Crain et al.'s argument is untenable because it leads to the following paradox: A child learning English would need to notice that sentences like (1) are ambiguous in her language (otherwise, the SSP would have no reason to apply). Having also noticed that the two readings of (1) are in a subset–superset relation, the child could now follow the SSP and initially hypothesize that only one of the two readings is available—in other words, that sentences like (1) are not ambiguous. Later, on encountering (positive) evidence regarding the availability of a second reading for sentences like (1)—that is, (1b)—the child would revise her initial hypothesis and come to the conclusion that sentences like (1) are ambiguous after all. The problem with this scenario should now be clear: One cannot have a learning principle whose application is contingent on *a priori* possession of the knowledge to be acquired. At this point, one may object that the learner need not engage in the (implicit) computations described above. Perhaps Universal Grammar, in conjunction with the SSP, leaves the learner no other choice but to initially assume that sentences like (1) only have reading (1a). However, then again, in the absence of a language that allows only reading (1a) for sentences like (1)—that is, if other languages behave like English—a child initially hypothesizing readings (1a) and (1b) for sentences like (1) would never be wrong, so there would be no need for the SSP in the first place.[5] As readers can verify for themselves, the same argument applies to Crain et al.'s second example that involves ambiguous sentences like, "The big elephant is the only one playing a guitar."

Empirical Predictions of the SSP

As mentioned earlier, the SSP makes predictions that go beyond the interpretation of sentences with *only.* To test these predictions, I focus here on the semantic interactions between negation and quantified NPs (e.g., every horse, some girls). I believe that there are three good reasons for doing so. First, sentences containing negation and quantified NPs are often ambiguous (i.e., they have more than 1 reading). Second, the alternative readings arising from quantifier–negation interactions often create subset–superset configurations (i.e., 1 reading asymmetrically entails the other). Finally, such interactions have been the focus of a growing body of developmental work, so a fair amount of empirical evidence is now available that can be brought to bear on the predictions of the SSP. Consider the examples in (16) and (17).

(16) Every student can't afford a new car.
  a. $\forall x$ [student (x) $\rightarrow \neg$ can afford a new car (x)].
  b. $\neg \forall x$ [student (x) $\rightarrow$ can afford a new car (x)].
(17) Some students can't afford a new car.

[5]It is worth pointing out that regardless of whether a language in which sentences like (1) can only be interpreted as in (1a) exists, it has recently been shown that the results reported by Crain et al. (1994) regarding young children's interpretation of *only*—and interpreted as evidence for the operation of the Semantic Subset Principle—are more likely to be due to methodological limitations than to a grammatical difference between preschoolers and adults (Paterson, Liversedge, Rowland, & Filik, 2003).

a. $\exists x$ [student (x) $\wedge \neg$ can afford a new car (x)].
b. $\neg \exists x$ [student (x) $\wedge$ can afford a new car (x)].

Recall that sentences like (16) are ambiguous. On one interpretation, (16) can be paraphrased as, "Every student is such that she or he cannot afford a new car." In other words, no student can afford a new car. In this case, the universal quantifier is interpreted outside the scope of negation (every > not). Alternatively, (16) can be paraphrased as, "Not every student can afford a new car." In this case, every student is interpreted within the scope of negation (not > every). Notice now that the every > not (i.e., none) reading entails the not > every (i.e., not all) reading. In other words, if it is the case that no student can afford a new car, it follows that not every student can afford one; but not vice versa. Consider (17) now which contains an existentially quantified subject (i.e., some students) and negation. Unlike (16), (17) is not perceived to be ambiguous. The most natural interpretation of (17) is one on which "some students" is interpreted outside the scope of negation (some > not; i.e., some students are such that they cannot afford a new car). If one were to interpret some students within the scope of negation (not > some), the resulting interpretation could be paraphrased as, "It is not the case that there exists some students that can afford a new car." In other words, no student can afford a new car. Notice that in this case also, the not > some reading entails the some > not reading; that is, if it is true that no student can afford a new car, it must also be true that there are some students that cannot afford a new car, but not vice versa.

Before turning to the predictions of the SSP regarding children's interpretation of sentences containing universally–existentially quantified NPs and negation, we must first settle the question of how to interpret the SSP. Crain and Thornton's (1998) definition of the SSP is repeated below:

> Suppose that the interpretive component of Universal Grammar makes two interpretations, A and B, available for a sentence, S. If so, then see if S is true in a narrower range of circumstances on interpretation A than on interpretation B. If so, then A will be hypothesized before B in the course of language development. (p. 118)

Notice that according to this definition, S must be a sentence of a particular language because Universal Grammar (UG) does not contain a repository of sentences. However, as I showed earlier, this definition is problematic. To repeat the argument, the current definition entails that one must already know whether a given sentence is ambiguous in a particular language (e.g., sentences containing *only* in English) to be able to decide whether the SSP should apply. I can see two ways to fix the SSP so as to preserve its logic and avoid the problem faced by its current formulation. One would be to no longer make reference to sentences and give the SSP an intensional definition instead of an extensional one.[6] Let us call this Option 1. The other option would be to preserve the current extensional definition but to add a clause to the SSP stipulating that it should only apply if there exists a language in which S can only receive interpretation A (the Subset one).[7] Let us call this Option 2. Below are two revised versions of the SSP along the lines of Options 1 and 2:

> Option 1: Suppose that UG allows two categories α and β to interact in such as way as to give rise to two interpretations, A and B when these categories are syntactically combined. Suppose further that A asymmetrically entails B. If so, then a learner should initially hypothesize that when syntactically combined, α and β yield only one interpretation, namely A.
>
> Option 2: Suppose that UG is such that for two languages L1 and L2 a sentence S has meaning M1 in L1 and meaning M1 and M2 in L2. Suppose further that M1 entails M2. If so, a learner should initially hypothesize that S can only be assigned meaning M1.

Let us explore the empirical implications of Option 1 first. To do that, let us consider a particular instantiation of Option 1 in which the categories in question are scope bearing expressions:

---

[6]The intensional–extensional distinction corresponds to the well-known Chomskyan distinction between E-language (external or extensional language) and I-language (internal or intensional language). To use a simple mathematical analogy, the E/I-language distinction corresponds to the intensional–extensional definition of sets. Consider, for example, the set S of multiples of 7. One could extensionally define that set by providing a list of its members (i.e., S: {7, 14, 21, 28 … }). Alternatively, one could provide the formula used to generate the members of S, thereby intentionally defining S (i.e., {x: x = 7y, y an integer}). Transposing the analogy to language, (extensional) English would correspond to the sentences of that language, whereas (intensional) English would correspond to the generative procedure (also known as a generative grammar) used to create the sentences of that language. For an excellent exposition of these notions, see M. Baker (2001). Intentionally defining the Semantic Subset Principle (SSP) would allow us to circumvent the problem of having to deal with sentences of particular languages—because sentences are extensional objects—while preserving the logic of the SSP.

[7]This seems to be the interpretation that Crain et al. (1994) have in mind judging from the following statement:

> The principle orders children's semantic hypotheses in advance, as follows: Default hypotheses are ones that will not subsequently need to be revised (i.e., they are realized universally), and additional (language-particular) hypotheses are added on the basis of positive evidence from the input. (p. 454)

However, as discussed above, these desiderata are not reflected in the way Crain et al. define or implement the Semantic Subset Principle.

> Suppose that UG is such that it allows the phenomenon of scope ambiguity. Specifically, given two scope-bearing expressions, α and β, either α can take scope over β (abbreviated $\alpha\beta$), or β can take scope over α (abbreviated $\beta\alpha$). Further suppose further that $\alpha\beta$ asymmetrically entails $\beta\alpha$. If so, then a learner should initially hypothesize that when syntactically combined, α and β yield only one scopal interpretation: $\alpha\beta$.

To fit our current needs, let us replace α and β in the definition above by $\forall$, the symbol for the universal quantifier; $\exists$, the one for the existential quantifier; and $\neg$, for negation. Therefore, given the pair $\{\forall,\neg\}$, we have two possible scope relations: $\forall\neg$, where the universal quantifier takes scope over negation (corresponding to every > not; i.e., none) and $\neg\forall$, where negation takes scope over the universal quantifier (corresponding to not>every). Moreover, we have seen that $\forall\neg$ entails $\neg\forall$. Given the pair $\{\exists,\neg\}$, possible scope relations are: $\exists\neg$ (some > not) and $\neg\exists$ (not > some; i.e., none). As discussed earlier, $\neg\exists$ entails $\exists\neg$.

According to the intensional version of the SSP then, learners of any natural language should always initially hypothesize that when $\forall$ and $\neg$ combine, the resulting sentence can be assigned only a $\forall\neg$ (none) reading. Similarly, when $\exists$ and $\neg$ combine, the resulting sentence can be assigned only a $\neg\exists$ (none) reading. Notice that this formulation of the SSP circumvents the problem faced by the current, extensional definition. On the intensional definition, one does not need to know whether a sentence S in a given language is ambiguous to determine whether the SSP should apply. Rather, the SSP is such that it dictates that sentences containing $\forall$ and $\neg$ will initially be unambiguous, regardless of the target language. If it turns out that in the target language in question sentences containing $\forall$ and $\neg$ are indeed always unambiguous, then nothing further needs to happen. If, on the other hand, the target language is one like English in which a sentence like, "Every horse didn't jump over the fence," for example, is in fact ambiguous, then a learner who has initially hypothesized that such a sentence can only receive a $\forall\neg$ (none) reading will sooner or later be presented with positive evidence that that sentence can also receive a $\neg\forall$ (not all) reading. What makes the intensional SSP interesting when applied to the acquisition of scope is that it predicts that children should initially compute scope on the basis of entailment relations rather than syntactic position. In other words, the intensional SSP predicts that children should initially assign sentences like, "Every horse didn't jump over the fence" and "The farmer didn't feed every horse," a "none" reading only (i.e., "None of the horses jumped over the fence," and "The farmer didn't feed any of the horses").

To summarize, the intensional SSP predicts that children should initially be restricted to the "none" interpretation of sentences containing universals–existentials and negation, regardless of the syntactic position of the quantified NPs. This follows from the fact that entailment relations between two logical operators are not affected by their syntactic position. Therefore, in the case of (16) and

(17) in English, repeated below as (18) and (19), the SSP predicts that children will initially interpret both sentences to mean that no student can afford a new car (18a and 19b):

(18) Every student can't afford a new car.
a. $\forall$x [student (x) $\rightarrow\neg$ afford a new car (x)].
b. $\neg\forall$x [student (x) $\rightarrow$ afford a new car (x)].
(19) Some students can't afford a new car.[8]
a. $\exists$x [student (x) $\wedge\neg$ afford a new car (x)].
b. $\neg\exists$x [student (x) $\wedge$ afford a new car (x)].

As mentioned earlier, children's interpretation of sentences containing negation and quantified NPs has been the focus of a recent series of experimental investigations. What emerges from this line of work is the observation that preschoolers do, in fact, systematically differ from adults in the way they interpret sentences containing negation and quantified NPs. Crucially, however, young children display a strong preference for the interpretation of such sentences that corresponds to the surface syntactic position of the elements involved. This is known as the Observation of Isomorphism or the Isomorphism Effect (Lidz & Musolino, 2002; Musolino, 1998; Musolino, Crain, & Thornton, 2000; Musolino & Gualmini, 2004; Musolino & Lidz, 2003; Musolino & Lidz, in press).[9] Therefore, this result falsifies the main prediction of the intensional SSP; namely, that children's interpretation of sentences containing negation and quantified NPs should initially be restricted by entailment relations rather than the syntactic position of the elements involved. Table 1 illustrates the predictions of the intensional SSP against actual

[8]The logical representation in (19a) says that there exists one or several *x*s, x being a student, and that it is not the case that x can afford a new car. The logical representation in (19b) says that there does not exist any *x*s, x being a student, and x being able to afford a new car.

[9]The effect I am referring to here has now been replicated several times, by different investigators, and in different languages (English, see references above; French, see Noveck, 2005; Korean, see Han, Lidz, & Musolino, 2006). In all these experiments, the method used was the Truth Value Judgment Task (TVJT; Crain & Thornton, 1998). In the TVJT, children watch short stories that are acted out in front of them using various toys and props. The child's task is to assess the statements of a puppet who describes what happened in the stories. To give an example, imagine a story in which three horses decide to jump over a fence, and only two of them manage to do so. At the end, the puppet would describe this situation by saying that, "Every horse didn't jump over the fence." Notice now that if one interprets this sentence on the "not > every" reading (which is true in the context described above), then one would answer that the puppet is right. On the other hand, if the "none" reading was accessed, then one would say that the puppet was wrong because it is not the case that none of the horses jumped over the fence. The basic finding here, labeled *isomorphism,* is that, whereas adults almost always answer that the puppet is right (because they are accessing the not > every reading), preschoolers almost always answer that the puppet is wrong (because they are accessing the "none" reading).

TABLE 1
Predictions of the Intensional SSP

Sentence Type	*SSP*	*Actual Findings*[a]
(20) Every horse didn't jump over the fence	$\forall\neg$	$\forall\neg$
(21) The Smurf didn't buy every orange	$\forall\neg$	$\neg\forall$
(22) Some girls won't ride on the merry-go-round	$\neg\exists$	$\exists\neg$
(23) The detective didn't find some guys	$\neg\exists$	$\neg\exists$

Note. SSP = Semantic Subset Principle. The results reported for sentences like (20) are from Musolino, Crain, and Thornton (2000) and Musolino and Lidz (in press); those for sentences like (21) from Musolino et al. (2000); those for sentences like (22) from Musolino (1998, 2001); and finally, those for sentences like (23) from Musolino et al. (2000).

[a]Children's interpretations.

experimental findings obtained from preschoolers for sentences containing universally and existentially quantified NPs in subject and object position. Following standard terminology, $\forall$ is used to represent the universal quantifier; $\exists$ represents the existential quantifier; and $\neg$ represents negation. In turn, $\forall\neg/\exists\neg$ indicates the reading of a sentence on which the universal–existential quantifier takes scope over negation and $\neg\forall/\neg\exists$ the reading on which negation takes scope over the universal–existential quantifier.

Notice that although children's interpretations can be systematically predicted by the surface syntactic position of the elements involved, the SSP makes correct predictions only in those cases in which the interpretation predicted by entailment relations happens to coincide with the one predicted by surface syntactic position. Finally, it is worth pointing out that the ages of the children who participated in the experiments reported here are similar to those of the children who participated in Crain et al.'s (1994) study. All these studies involved preschoolers in the 4 to 5 age range.

Intentionally reformulating the SSP allowed us to solve the problems faced by the extensional definition. However, we now find that the SSP thus reformulated makes incorrect empirical predictions. The problematic cases for the intensional SSP on the scenario outlined above are sentences like (21) and (22) in which a universally quantified NP occurs in the object position, and an existentially quantified NP occurs in the subject position. Let us now consider Option 2. One way to solve the problem faced by the intensional definition would be to assume that UG makes sentences like (21) and (22) universally unambiguous. Suppose then that the definition of the SSP could be changed so as to prevent it from applying to sentences like (21) and (22) while at the same time allowing it to apply to sentences like (20) and (23). If so, the SSP would no longer make incorrect predictions for sentences like (21) and (22). So far, so good. In the case of sentences like (20) and (23), the SSP would apply only insofar as languages can be found in which these sentences are ambiguous and other languages can be found in which they are not. Consider sentences like (20), for example, which are ambiguous in English. As it turns out, sentences like (20) are not ambiguous in Chinese. They can be assigned only a "none" reading. In other words, the universally quantified subject must be interpreted outside the scope of negation, as shown in (24):

(24) *Mei-pi ma dou mei tiao-guo langan.*
"Every horse didn't jump over the fence."
a. $\forall x$ [horse (x) $\rightarrow\neg$ jump over the fence (x)].

Given these premises, the following and familiar story can now be told (and has indeed been told by Musolino et al., 2000). Imagine that a learner exposed to Chinese were to incorrectly conjecture that sentences like (24) can receive both a "none" and a "not all" reading. If so, and because the "none" reading entails the "not all" reading, no amount of positive evidence would suffice to compel learners to revise their initial hypothesis. In other words, learners would always hear sentences like (24) used on an intended "none" reading—the only option in the adult grammar. Given that when the "none" reading is true, the "not all" reading is also true, the learner's incorrect hypothesis that sentences like (24) can receive a "not all" reading would never be falsified. Learners would thus fail to converge on the target grammar. However, because this never happens, something must be compelling learners to never make that potentially lethal conjecture in the first place. In other words, there must exist a mechanism that forces learners to be conservative and initially hypothesize that in the language they are exposed to, a sentence like (24) can only receive a "none" interpretation. If it so happens that the language in question is Chinese, then nothing further needs to happen. If, on the other hand, the language is English, then sooner or later, learners will encounter positive evidence informing them that sentences like (24) can also receive a "not all" interpretation. That is to say, sooner or later, learners will hear a sentence like (24) used on an intended "not all" reading. Despite its rigorous logic, the argument just presented faces two serious problems. One is empirical and is discussed below. The other is conceptual and is discussed in the next section.

The empirical problem faced by the SSP on the scenario just outlined has to do with the nature of the Isomorphism Effect. Recall that the central claim underlying the SSP is that children initially *lack* superset readings. In other words, the SSP predicts that children's grammars will initially be different from those of adults. However, it has recently been shown that the Isomorphism Effect reflects differences in the way preschoolers and adults deal with ambiguity resolution (for related evidence, see Trueswell, Sekerina, Hilland, & Logrip, 1999) rather than a grammatical difference between the two populations. Specifically, it has been shown that (a) certain contextual manipulations can lead children to behave more like adults (i.e., to access nonisomorphic interpretations at a much higher rate; Gualmini, in press; Musolino & Lidz, in press); (b) in sentences that require the grammatical operations associated with nonisomorphic interpretations, but lacking a scope ambiguity, children behave identically to adults (Lidz et al., 2004; Syrett & Lidz, 2004); and (c) in certain cases, the Isomorphism Effect can also be induced in adults (Musolino & Lidz, 2003). Therefore, in the case of sentences like, "Every horse didn't jump over the fence," the fact that children display a marked preference for the every > not reading, as shown by Musolino et al. (2000), does not mean that they lack the not > every reading (see Musolino & Lidz, in press). Therefore, the account given by Musolino et al., according to which young children's interpretation of sentences like, "Every horse didn't jump over the fence," can be explained by the SSP (as given in Option 2), must be abandoned (see Musolino & Lidz, in press).

The Logical Status of the SSP

So far, we have seen that to the extent that semantic subset problems exist, the SSP is not the correct solution to them because, under either Options 1 or 2, it makes the wrong empirical predictions. The question that I would like to address now is whether semantic subset problems exist in the first place. According to Crain and Thornton (1998), semantic subset problems arise just in case (a) the interpretive component of UG makes a sentence ambiguous, (b) the alternative readings of that sentence are arranged in a subset–superset relation (i.e., one reading asymmetrically entails the other), and (c) a language exists in which the sentence in question can only be assigned the subset reading. According to this definition, the case of sentences like, "Every horse didn't jump over the fence," in English and Chinese represents a textbook example. Recall that such sentences are (a) ambiguous in English, (b) that the "none" reading entails the "not all" reading, and (c) that Chinese allows only the subset reading (i.e., the "none" reading). Notice now that the logic of the SSP rests on one crucial assumption; namely, that if there were no SSP, there would be no other way for learners to systematically arrive at the correct conclusion regarding the meaning of sentences like, "Every horse didn't jump over the fence," in English and Chinese. This assumption, as I demonstrate below, turns out to be incorrect.

To show that the *raison d'être* of the SSP is not logically necessary, I need to show that children can learn that sentences like, "Every horse didn't jump over the fence," are ambiguous in English but not in Chinese without guidance from the SSP. This, I think, can be done rather straightforwardly. Let me first sketch the main idea and then provide the detailed solution. Given the compositional nature of semantics—the fact that the meaning of a sentence is a function of the meaning of its parts and the way they are arranged (Frege, 1892/1980)—it would seem that a child who knows the meaning of the words that make up a sentence like, "Every horse didn't jump over the fence," (in Chinese or in English) as well as the basic syntax of his or her language (Chinese or English) ought to be able to *deduce* the meaning of such a sentence. Let me now demonstrate how this can be done under the standard and independently motivated theoretical assumptions introduced in the Theoretical Background section. I also assume, with Aoun and Li (1993) and Hornstein (1995), the following, independently motivated, parametric difference between English and Chinese:

(i) Subjects start in a VP internal position in English and then subsequently raise to IP; whereas in Chinese, subjects do not start in a VP internal position. Rather, they are base generated in IP.[10]

Given these assumptions, let us now see how the relevant interpretive facts regarding sentences like, "Every horse didn't jump over the fence," can be derived. Let us go over the derivation of such sentences in English one more time. The LF structure of a sentence like, "Every horse didn't jump over the fence," is given in (25), and its notational equivalent is given in (25').

(25) [IP [NP Every horse [I' didn't [VP every horse jump over the fence]]]].
a. [IP [NP ~~Every horse~~ [I' didn't [VP every horse jump over the fence]]]].
b. [IP [NP Every horse [I' didn't [VP ~~every horse~~ jump over the fence]]]].
(25')

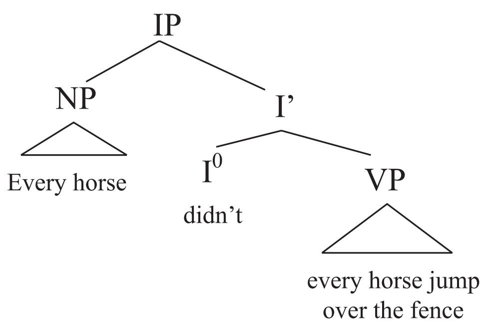

Because subjects start in a VP internal position and movement involves copying, we end up with a copy of *every horse* in VP and one in IP. Recall that at LF, all copies but one must be deleted. This gives us two options: Either the VP or the IP

[10]What is meant here by *base generated* is simply the fact that subjects either begin their syntactic life as part of the verb phrase and are then "moved" to a higher position (c-commanding position; as in English) or that they start their syntactic life directly in that higher position (i.e., inflectional phrase) and do not undergo movement (as in Chinese).

copy may be deleted. Deleted copies in (25) are crossed out. If the IP copy is deleted, notice that the remaining copy (the VP copy) occurs in the c-command domain of negation.[11] Given our scope principle, *every horse* must, in this case, be interpreted in the scope of negation; yielding the not > every interpretation. On the other hand, if the VP copy is deleted, then the remaining copy (the IP copy) occurs outside the c-command domain of negation and, thus, outside its scope. In this case, it follows that *every horse* will be interpreted outside the scope of negation, thereby yielding the every > not interpretation. We have thus derived the fact that sentences like, "Every horse didn't jump over the fence," can either receive a "none" or a "not all" reading in English.

Let us now turn to Chinese. Recall that in this language, subjects are base generated in IP. Therefore, at LF the Chinese equivalent of "Every horse didn't jump over the fence" has the structure in (26). Because *every horse* is based generated in IP and not copied there from a VP internal position, (26) contains only one copy of the subject. It follows that Chinese allows only one scope reading; namely, one on which *every horse* is interpreted outside the scope of negation. Therefore, in Chinese, sentences like, "Every horse didn't jump over the fence," can only be interpreted as meaning that none of the horses jumped over the fence (every > not).

(26) [IP [NP Every horse [I' didn't [VP jump over the fence]]]].

To summarize, given a set of independently motivated assumptions regarding the syntax of natural languages, it is possible to deduce that sentences like, "Every horse didn't jump over the fence," are ambiguous in English but not in Chinese without recourse to the SSP. Put another way, a child exposed to Chinese who has learned that subjects are base generated in IP in his or her language—a fact that children clearly must learn independently of the argument being made here—would be able to (tacitly) deduce that sentences like, "Every horse didn't jump over the fence," are not ambiguous in his or her language. Similarly, a child exposed to English who has learned that subjects originate in a VP internal position in his or her language—also a fact that must be learned independently of the argument being made here—would be able to (tacitly) deduce that sentences like, "Every horse didn't jump over the fence," allow both a "none" and a "not all" reading in his or her language. To the extent that the argument just presented can be generalized to other phenomena that fall under the rubric of semantic subset problems—and I see no reason why it could not—then semantic

[11]This can be seen by looking at the phrase marker in (24') and applying the rule of thumb given in the Theoretical Background section; that is, start with negation in I^0, go up to the first branching node (I'), and then go down. The verb phrase—and everything contained therein—thus falls within the c-command domain of negation.

subset problems do not exist. To be more precise, problems defined as semantic subset problems that can only be solved by the SSP do not exist. Put another way, knowledge of word meaning, in conjunction with knowledge of the principle of compositionality and knowledge of the basic syntax of one's language, is enough to guarantee that the meaning of the sentences of that language is learnable. Let me emphasize the fact that this conclusion follows regardless of whether my analysis of the difference between Chinese and English is correct. I just used this example to show that it was possible, at least in principle, to come up with a plausible account of how certain facts can be learned without recourse to the SSP. The more important point here is that it is the compositional nature of semantics—and the relation between syntax and semantics—that renders the SSP implausible.

To further illustrate this point, consider a classic example from the linguistic literature illustrating the role that syntax can play to ensure semantic learnability. The argument has to do with one-substitution in English. Consider an NP like, "an expensive car," and two possible syntactic representations for it; a flat structure representation (FSR; the one on the left) and a layered structure representation (LSR; the one on the right):

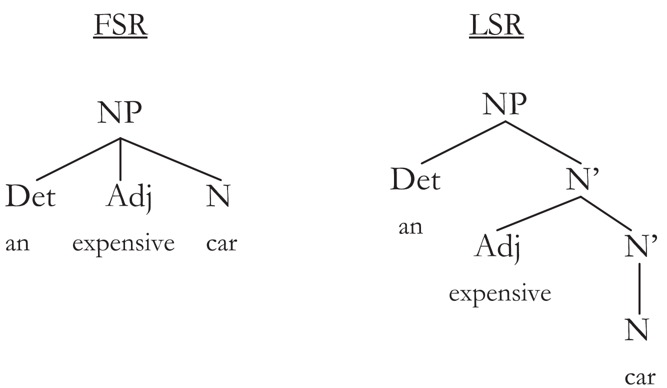

At first sight, LSR and FSR appear to be mere notational variants. However, as linguists have shown, the consequences of adopting one over the other has deep implications for language learnability (C. L. Baker, 1979; Hornstein & Lightfoot, 1981; for recent experimental evidence, see also Lidz, Waxman, & Freedman, 2003). Notice that children who adopt LSR can substitute *one* for either *car* or *expensive car* because both phrases form constituents (i.e., N',[12] on this analysis). By contrast, children who adopt FSR can only substitute *one* for *car* because *expensive car*

[12]Recall that according to X-bar theory, all phrases—including NPs—contain an intermediate or "bar" level, in this case N' (read N-bar), which is recursive. In other words, several N's can be found within the same structure, as shown in the phrase marker corresponding to LSR.

does not form a constituent on this analysis. Consider now the possible meanings of a sentence like (27), which is ambiguous and can either be paraphrased as (27a) or (27b); that is, *one* in (27) can either substitute for *expensive car* (27a) or *car* (27b). Moreover, the preferred reading of (27) for adult speakers of English is (27a). The fact that *one* can indeed substitute for *car* is illustrated in (28):

(27) Julien bought an expensive car and Jeff bought one too.
  a. Julien bought an expensive car and Jeff bought an expensive car too.
  b. Julien bought an expensive car and Jeff bought a car too.
(28) Julien bought an expensive car and Jeff bought a cheap one (one = car).

Notice now that for a child who has adopted FSR, a sentence like (27) can only have the reading in (27b), and so it is unambiguous because on that analysis *one* can only substitute for *car.* Now here is the catch: The extension of the N' *expensive car* is a subset of the extension of the N *car.* In other words, whenever you have *an expensive car* you also have a *car,* but not vice versa. Therefore, for a child who has adopted FSR and who thinks that (27) can only mean (27b), it would be difficult to learn that (27) can also mean (27a). This is so because every time the child would hear a sentence like (27) intended to be interpreted as (27a), this interpretation would be compatible with the child's hypothesis, (27b). The child would, therefore, have no need to revise his or her initial hypothesis. The solution to this potential learnability problem, of course, is to adopt the correct syntactic analysis, namely LSR. To be sure, a child who has adopted LSR would automatically know that (27) can have the readings in (27a & 27b). The important conclusion here is that a potential semantic subset problem can be avoided without recourse to a principle like the SSP. In other words, get your syntax right and the correct semantics will follow.

Let us take stock. I have now shown that (a) on its current formulation, the SSP is untenable; (b) when reformulated, the SSP makes the wrong empirical predictions; and (c) arguments for the logical necessity of the SSP contain serious loopholes. In light of these problems, it seems reasonable to abandon the SSP as a learnability constraint on the acquisition of semantics.

CONCLUDING REMARKS

In this article, I have argued that there are good reasons—both conceptual and empirical—to abandon the SSP as a learnability constraint on the acquisition of semantics. More importantly, I have suggested that semantic subset problems probably do not exist in the first place. I believe that these conclusions are important because they have broader implications for the study of language acquisition. First, they provide a basis for delimiting the class of problems that theories of the acquisition of semantics ought to be responsible for. Second, these conclusions underscore the fact that data from empirical investigations of child language can be brought to bear on the formulation of linguistic theory. Specifically, available experimental evidence bearing on the predictions of the SSP (e.g., results from developmental studies of the interaction between negation and quantified NPs) falsify these predictions. This is an important conclusion because one of the central goals of modern linguistic theory is to construct models of the language faculty that aim at *explanatory adequacy* (i.e., models that are responsive to the demands imposed by the need to explain language acquisition; Chomsky, 1965; Hornstein & Lightfoot, 1981).

Now some important disclaimers. First, I would like to make it clear that although I had a lot to say about the SSP, I have nothing to say about its syntactic cousin, the original SP (Berwick, 1985). In other words, the arguments presented in this article are not intended to be arguments against the SP, whatever form it may take. Rather, they are specific arguments against the SSP as defined and implemented by Crain and colleagues (Crain et al., 1994; Crain & Thornton, 1998). Second, although I have argued that there are probably no semantic subset problems and thus no learnability problems of the kind identified by Crain and Thornton (1998), this does not mean that I believe that there are no learnability problems at all.

Finally, having explained in much detail my reasons for not believing in semantic subset problems or the SSP, I would now like to consider what one would need to show to build a case for a learnability constraint like the SSP. A possible argument would be a demonstration that for some sentences of English (or any other language), knowledge of the meaning of the words that make up these sentences coupled with knowledge of the rules of syntax for that language is not sufficient to deduce the meanings of the sentences in question. Another possibility would be to show that some semantic properties of natural languages are parametrized in such a way that different settings of the relevant parameter would yield the kind of subset–superset configurations discussed by Crain and colleagues (Crain et al., 1994; Crain & Thornton, 1998). Finally, if either of the conditions described above could be met, one would also need psycholinguistic evidence that, in such cases, young children do in fact behave in accordance with the predictions of the SSP. Until this can be done, neither children nor linguists need anything like the SSP.

REFERENCES

Aoun, J., & Li, A. (1993). *Syntax of scope.* Cambridge, MA: MIT Press.

Baker, C. L. (1979). Syntactic theory and the projection problem. *Linguistic Inquiry, 10,* 533–582.

Baker, M. (2001). *The atoms of language: The mind's hidden rules of grammar.* New York: Basic Books.

Berwick, R. (1985). *The acquisition of syntactic knowledge.* Cambridge, MA: MIT Press.

Chomsky, N. (1957). *Syntactic structures*. The Hague, The Netherlands: Mouton.

Chomsky, N. (1965). *Aspects of the theory of syntax.* Cambridge, MA: MIT Press.

Chomsky, N. (1995). *The minimalist program*. Cambridge, MA: MIT Press.

Crain, S. (1992). *The Semantic Subset Principle in the acquisition of quantification.* Paper presented at the workshop on the Acquisition of Wh-Extraction and Related Work on Quantification, University of Massachusetts, Amherst.

Crain, S. (1993). *Semantic subsetutions.* Invited paper presented at the Center for Cognitive Science Conference: Early Cognition and the Transition to Language, University of Texas, Austin.

Crain, S., Ni, W., & Conway, L. (1994). Learning, parsing and modularity. In C. Clifton, L. Frazier, & K. Rayner (Eds.), *Perspectives on sentence processing* (pp. 443–467). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Crain, S., & Philip, W. (1993). *Global semantic dependencies in child language.* Paper presented at the 16th GLOW Colloquium, Lund, Sweden.

Crain, S., & Thornton, R. (1998). *Investigations in universal grammar: A guide to research on the acquisition of syntax and semantics.* Cambridge, MA: MIT Press.

Frege, G. (1980). On sense and meaning. In P. Geach & M. Black (Eds.) *Translation from the philosophical writings of Gottlob Frege* (3rd ed.). Oxford, England: Blackwell. (Reprinted in A. P. Martinich, *The philosophy of language*. Oxford, England: Oxford University Press. Original work published 1892)

Gualmini, A. (2004). Some knowledge children do not lack. *Linguistics, 42,* 957–982.

Haegeman, L. (1994). *Introduction to government and binding theory.* Oxford, England: Blackwell.

Han, C. H., Lidz, J., & Musolino, J. (2006). *Verb-raising and grammar competition in Korean: Evidence from negation and quantifier scope*. Manuscript submitted for publication.

Hornstein, N. (1995) *Logical form: From GB to minimalism.* Cambridge, MA: Blackwell.

Hornstein, N., & Lightfoot, D. (1981). Introduction. In N. Hornstein & D. Lightfoot (Eds.), *Explanations in linguistics: The logical problem of language acquisition* (pp. 9–31). London: Longman.

Hyams, N. (1992). A reanalysis of null subjects in child language. In J. Weissenborn et al. (Eds.), *Theoretical issues in language acquisition: Continuity and change in development* (pp. 249–267). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Lidz, J., McMahon, E., Syrett, K., Viau, J., Anggoro, F., Peterson-Hicks, J., et al. (2004). Quantifier raising in 4-year-olds. In *Proceedings of 28th Boston University Conference on Language Development, 1* (pp. 340–349). Somerville, MA: Cascadilla.

Lidz, J., & Musolino, J. (2002). Children's command of quantification. *Cognition, 84*(2), 113–154.

Lidz, J., Waxman, S., & Freedman, J. (2003). What infants know about syntax but couldn't have learned: Experimental evidence for syntactic structure at 18 months. *Cognition, 89,* B65–B73.

Musolino, J. (1998). *Universal grammar and the acquisition of semantic knowledge: An experimental investigation of quantifier–negation interactions in English.* Unpublished doctoral dissertation, University of Maryland at College Park.

Musolino, J., Crain, S., & Thornton, R. (2000). Navigating negative quantificational space. *Linguistics, 38*(1), 1–32.

Musolino, J., & Gualmini, A. (2004). The role of partitivity in child language. *Language Acquisition, 12,* 97–107.

Musolino, J., & Lidz, J. (2003). The scope of isomorphism: Turning adults into children. *Language Acquisition, 11*(4), 277–291.

Musolino, J., & Lidz, J. (in press). Why children aren't universally successful with quantification. *Linguistics.*

Noveck, I. (2001). When children are more logical than adults: Experimental investigations of scalar implicature. *Cognition, 78,* 165–188.

Noveck, I. (2005). *A role for pragmatics in negation processing: An investigation with autistic participants.* Manuscript submitted for publication.

Paterson, K., Liversedge, S., Rowland, C., & Filik, R. (2003). Children's comprehension of sentences with focus particles. *Cognition, 89,* 263–294.

Pinker, S. (1995). Language acquisition. In L. Gleitman & M. Liberman (Eds.), *An invitation to cognitive science: Language* (pp. 135–182). Cambridge, MA: MIT press.

Syrett, K., & Lidz, J. (2004, November). *Children want to access every interpretation adults do.* Poster presented at the 29th Boston University Conference on Language Development, MA.

Trueswell, J. C., Sekerina, I., Hilland, N. M., & Logrip, M. L. (1999). The kindergarten-path effect: Studying on-line sentence processing in young children. *Cognition, 73,* 89–134.

Wexler, K., & Manzini, R. (1987). Parameters and learnability in binding theory. In T. Roeper & E. Williams (Eds.), *Parameters and linguistic theory* (pp. 44–76). Dordrecht, The Netherlands: Reidel.