COGNITIVE SCIENCE Vol 24 (1) 2000, pp. 123–149 ISSN 0364-0213

A Local Connectionist Account of Consonant Harmony in Child Language

THOMAS BERG
University of Hamburg

ULRICH SCHADE
University of Bielefeld

This study represents an attempt to sketch a processing model of phonological development in children acquiring their first language. The investigation is framed within a local connectionist network in which activation spreads between levels and inhibition within levels. Three ways are focused on in which an emergent processing system may diverge from a fully developed one—hypoactivation (too little activation due to underdeveloped links), hyperactivation (too much activation due to overloaded links) and impaired self-inhibition. With a view to determining how children's productions relate to these three "error mechanisms," one widespread phonological process, consonant harmony, is subjected to examination. Hypoactivation is found to capture the great majority of harmonic patterns. Hyperactivation and impaired self-inhibition play a more marginal role and are only needed to account for exceptional cases of harmony. The feasibility of the psycholinguistic model is demonstrated with the help of computer simulations that were run with noise on both fully developed and developing networks. The model makes a number of predictions including the claim that difficult sounds are acquired in coda positions first.

I. INTRODUCTION

Despite the many fronts on which the study of the acquisition of phonology has made considerable headway, there is a conspicuous lack of anything approaching an information-processing model that specifies the relationship between the developing retrieval mechanisms in an emergent lexicon and children's linguistic output. The time now seems ripe to take up this challenge as significant advances have been made in the two areas that are of immediate relevance to a processing theory of phonological acquisition. For one

Direct all correspondence to: Thomas Berg, Department of English, University of Hamburg, Von–Melle–Park 6, 20146 Hamburg, Germany; E-Mail: thomas_berg@uni-hamburg.de

thing, a general learning theory is required. For another, it is helpful to draw upon a theory of adult language processing because the fully developed system may serve as a background against which the developing system may be devised.

The past 15 years or so have seen the construction of an explicit psycholinguistic model of language production (e.g., Stemberger, 1985; Dell, 1986; MacKay, 1987; Berg, 1988; Schade, 1992). A defining characteristic of these works is their relatively close adherence to linguistic theory. Specifically, the hierarchical structure of language, the distinction between units and levels as well as that between content and structure are some of the essential linguistic ingredients of the psycholinguistic framework. As far as the phonological component is concerned, it is widely agreed that only a multilevel representation is adequate that consists, minimally, of a syllable, a segment and a feature level. The linguistic units are coded in terms of nodes, namely, simple processing elements whose function is to receive and relay information to their neighbors. The nodes in the network are represented in type-only fashion, that is, there is only one unit at a lower level subserving an arbitrary number of units at the next higher level. For instance, the phonological feature [bilabial] serves all bilabial elements at the segment level.

The exchange of information in the network is ensured by connections between nodes, hence the label "connectionist" for this class of processing models. This exchange occurs largely in parallel to the effect that a higher node not only activates a lower node (feedforward) but is also activated by it (feedback, see Dell, 1985). The production of a linguistic unit is guaranteed by its selection at a particular moment in time from among a number of simultaneously active competitors. After selection, the relevant node undergoes self-inhibition whereby its activation level is momentarily set to zero to prevent its being selected over and over again.

Because connectionist models are basically made up of nothing but nodes and links, there are only two ways in which learning may be conceptualized. A standard assumption in connectionist theory is that learning involves adjusting the links between nodes. Links are hypothesized to vary in strength (or weight) and thus to transport variable amounts of activation. At the onset of language acquisition, children may be presumed to begin with connection weights randomly hovering around zero. Their task would then be to increase the weights to the point that is typical of the full–fledged processing system of competent adult speakers. How do children home in on proper connection weights? We assume that they attempt to produce a certain target word and compare the production result to the perceptual representation of this unit. If they do not discover a mismatch, all weights remain as they are. However, if they do discover a mismatch, weights are changed so as to reduce the likelihood of error in the next production of the same unit (see Stemberger, 1992).

Children's connections may deviate from adults' in only two ways—they may be too weak or too strong so that either too little or too much activation is transmitted. These states will be termed *hypoactivation* and *hyperactivation*, respectively. Furthermore, the mechanism of self-inhibition has to be fine-tuned. Given that the correct functioning of self-inhibition involves the temporary blocking of the just-selected element, any defect of this mechanism is more likely to be of the hypoinhibitory than of the hyperinhibitory kind. We thus add *impaired self-inhibition* as the third possible locus of children's errors.

It is highly likely that children's processing systems are characterized more by hypostates than by hyperstates. Unless it is assumed that children jump immediately from around zero to the optimal value, which would seem unrealistic, learning can be said to necessarily run through hypostates. Hyperstates, in contrast, are not compulsory components of the learning process. They can only be entered after the hypostates have been passed and are reached only if the child does not stop increasing its linkages when the optimal weighting has been accomplished.

In addition to differences in linkage strength, children's networks may either lack nodes or have underdeveloped nodes in the sense that these nodes pass on less information than they would if they worked perfectly. This notion of gradualness is expressed through variable resting levels and serves to explain frequency effects (see Morton, 1970; McClelland & Rumelhart, 1981; Stemberger, 1992). Because we do not focus on frequency effects in this report, we will not recur to variable resting levels, even though they may be needed in a more comprehensive model.

In this paper, we will develop a partial theory of phonological acquisition in which the local connectionist model of adult language processing is combined with the notions of nonoptimal links and linkage adjustment. By focusing our attention on the links, we will concomitantly test the claim that children and adults do not differ much on representational aspects such as the number of levels and their overall organization. In fact, our working hypothesis is that the deployment of the processing machinery is a relatively early process and that a large part of the learning process, at least in the domain of production, can be captured by examining the links in the network.

Of course, it is impossible to investigate the entire range of phenomena that are characteristic of child language. The analysis was therefore restricted to one outstanding phonological process—that of consonant harmony (Vihman, 1978). This process is an especially rewarding area to study because it brings about patterns that diverge in systematic ways from the language of adults. The child thus cannot have picked up these forms and simply retrieved them from memory but must somehow have transformed the ambient language. However, there is no need to assume that the child actively applies a transformational rule in the production process. Stemberger (1992) draws an enlightening analogy between the child's emerging linguistic and walking skills. The toddler does not fall down by using a special rule or strategy. Rather, the unsteadiness of walking evidences the imperfect mastery of motor skills. The same can be said of language. In this perspective, mispronunciations do not evidence rules but are more appropriately interpreted as plans that have been imperfectly carried out. We take it that the children's intention is to sound like their caretakers—given their desire to communicate their ideas successfully and their general striving to conform to adult norms. A plausible interpretation of harmony is then, that it lays bare the limited powers of the learner's processing system, that is, it shows what the processing apparatus is currently capable of and what is beyond its capabilities. These limitations will give testimony of stages that the processing system runs through on its way from imperfect to perfect functioning.

Further reasons to focus upon harmony are its frequency of occurrence and its cross-linguistic nature. This process has been documented for a wide variety of languages

(Vihman, 1978) and can therefore be said to be independent of the structure of any particular language. It is resorted to by most children during the early acquisition period although the extent to which individual children make use of it is variable (Vihman, 1978). It should also be noted that harmony cannot be explained as a low–level articulatory constraint. By definition, it occurs between units at a distance and therefore cannot be reduced to coarticulatory effects. The elements harmony operates on are relatively abstract, corresponding to something like classical phonemes (Berg, 1992). It is this reasoning that motivates the high-level, cognitive analysis that will be presented in the remainder of this article.

II. A CLASSIFICATION SYSTEM OF HARMONY PROCESSES

Consonant harmony is a contextually motivated substitution process whereby a child makes two nonadjacent consonants more similar to each other than they are in the adult model. Most usually, both consonants belong to the same word. In this section, we will describe and classify harmony according to its psycholinguistically most relevant aspects. The first of these is whether harmony leads to identity or nonidentity between the assimilating and the assimilated segment, as exemplified in (1) and (2), respectively. By convention, the adult form is given between slashes, the child's rendition in brackets.

pudding /pudiŋ/ → [pupiŋ] (from Cruttenden 1978) (1)

bunnie /bʌni/ → [bʌmi] (from Cruttenden 1978) (2)

The first example involves *complete* harmony (/p. . .d/ → [p. . . p]) whereas the output in the second is only *partially* harmonized. The /b/ has in all probability conferred its [bilabial] feature upon the alveolar nasal, thereby turning it into [m].

If harmony is conceived of as a reaction to a processing problem, two possible types of motivation have to be reckoned with. In the first instance, the child may not have (fully) mastered a particular segment and as a result of this inability, resorts to harmony. Importantly, it is the segment itself, regardless of the phonological context in which it occurs that assumedly causes trouble. In the second instance, the child may be perfectly capable of producing two given segments in isolation but be incapable of producing these in combination. In other words, the segments can be accurately uttered when they belong to different planning units but they surpass the learner's productive capabilities when they belong to one and the same planning unit. In this case, harmony is used to deal with a sequencing problem. Actually, both sources of difficulty can be identified in children's speech. Refer to (3) and (4).

thing /θiŋ/ → [kiŋ] (from Smith 1973) (3)

duck /dʌk/ → [gʌk] (from Smith 1973) (4)

The word “thing” was attempted at a time when the child was still unable to accurately produce the dental fricative that was customarily replaced by [t] or [d]. By anticipating the velar feature of the nasal (and turning the fricative into a stop), he managed to solve his segment production problem. No such problem is apparent in (4). The alveolar stop had already entered the child’s lexicon. That it was ousted nonetheless suggests that a different problem is here at issue. It seems that the child had difficulty producing the alveolar stop in the presence of a nearby velar stop. Harmony leads to a reduction of the number of disparate place specifications within the same word (e.g., Menn, 1983) and may therefore be argued to alleviate the child’s processing problem.

The next distinction to be introduced is that between *singleton* and *multiple* harmony. The former refers to the fact that a given child may exhibit only one linguistic type of harmony during the acquisition period. For example, recourse may be taken to bilabial harmony but to no other harmony type. Such a case was discussed in Berg (1992). The latter refers to those children who have more than one type of harmony at their disposal. For instance, one and the same child may use both bilabial and lateral harmony at a given stage of development. This is the case with Smith’s (1973) son, Amahl.

Multiple harmony can be subdivided according to whether it may, or may not, apply to the same word. For want of a better term, the former will be baptized *variable*, the latter *invariable* harmony. Going by Smith’s published records, Amahl is an invariable harmonizer, as illustrated in (5–7).

good /gud/ → [gug] (from Smith 1973) (5)

ladder /lædə/ → [dɛdə] (from Smith 1973) (6)

really /riəli/ → [liːliː] (from Smith 1973) (7)

Number (5) is an example of velar harmony, (6) an instance of stop harmony and (7) a case of lateral harmony. It can be taken for granted that these patterns are representative of the child’s output in the sense that the lateral always dominates in the sequence /r. . . l/, the stop in the sequence /l. . . d/ and the velar in the sequence /g. . . d/ (at one point in Amahl’s development). The relationship among the individual harmonies in multiple invariable harmonizers will be detailed in the next section.

Variable harmony depicts a situation in which one word is capable of being harmonized in two different ways. The following example is from Suzanne, a French–speaking girl.

chausser /ʃoːseː/ ↗ a. [ʃoːʃeː] ↘ b. [soːseː] (from Deville 1891) (8)

Number (8) evidences an interaction between the two fricatives /ʃ/ and /s/. In (8a), the alveolar fricative harmonizes the palato-alveolar one whereas the reverse happens in (8b).

Children may differ in the consistency with which they harmonize a given word. Both *consistent* and *inconsistent* users of harmony have been reported in the literature. Although Amahl, as presented in Smith (1973), is a very consistent harmonizer during the larger part of his development, Jacob as presented in Menn (1978) produces harmonized alongside nonharmonized renditions of the same lexical item, as shown in (9).

a. [tita]
teapot /tiːpat/ (from Menn 1978) (9)
b. [tiba]

The example under (9) attests to the optional nature of Jacob's harmony process. The target *teapot* is rendered in harmonized form in (9a) but in nonharmonized fashion in (9b).

The distinction between consistent and inconsistent harmony applies not only to the intraword but also to the interword domain. Although two words may be quite similar in their phonological form, it is possible for the one to be harmonized whereas the other is not at a given point in time. A pertinent example comes again from Amahl. The two words under (10) and (11) were produced in the same temporal stage.

take /teik/ → [k^{h}eik] (from Smith 1973) (10)

tuck /tʌk/ → [t^{h}ʌk] (from Smith 1973) (11)

The only difference between the words in (10–11) lies in the vowels that do not seem to interact with the harmony process.

Finally, we turn to the fact that harmony may interact with other phonological processes such as segment omissions. A child may start out deleting a segment, especially in final positions, whereas accurately pronouncing the initial consonant (12a). When, at a later stage, the child begins to render the erstwhile omitted segment in adult-like fashion, the initial consonant, which used to be correctly uttered, may succumb to the harmonizing force of the newly-acquired sound (12b). Such a case is reported by Menn (1971).

a. (earlier) [buː]
boot /buːt/ (from Menn 1971) (12)
b. (later) [duːt]

The bilabial stop comes out fine provided there is no other consonant to compete with in the same syllable. However, when competition arises, it fails to be generated correctly and is harmonized. This is a kind of *trade-off regression* (Stemberger, 1992) in which the output is improved in one part (here the word-final position) whereas it deteriorates in another (here the word-initial position) (for further types of interaction, see Ohnesorg, 1965; Macken & Ferguson, 1987).

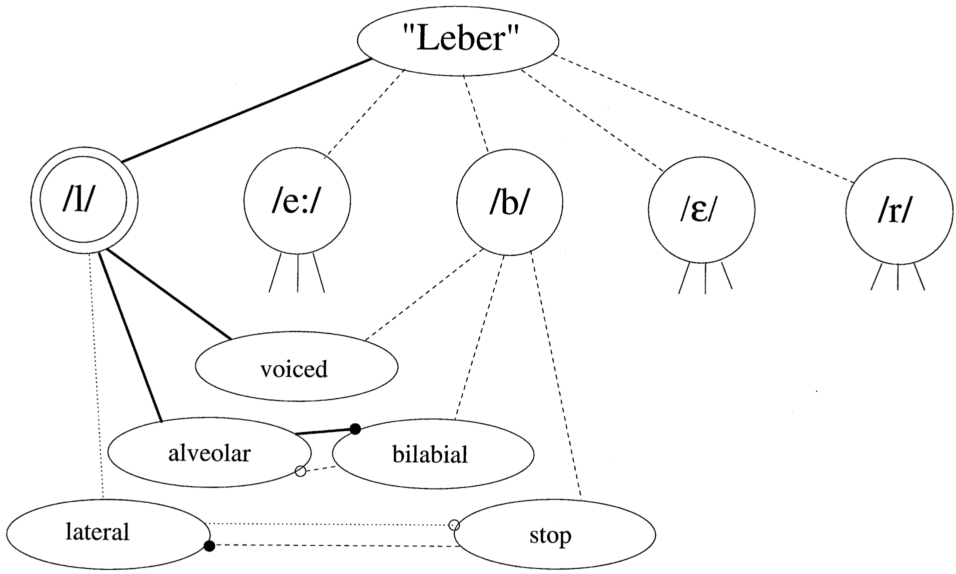

Figure 1. The flow of activation and inhibition at the moment of onset selection in the word *Leber*.

This section has examined the various manifestations of consonant harmony. The analysis offered above lays the groundwork for the processing account to be detailed below. Our major intention will be to explore whether the various harmony types reflect similar or distinct "error mechanisms."

III. A PROCESSING ACCOUNT OF CONSONANT HARMONY

Let us begin with the type of harmony that is instigated by children's inability to pronounce a certain sound and the fact that they produce something in its stead. Basically, this inability is conceivable as either a representational or an access problem. In the former case, a child lacks a specific node for a particular linguistic unit; in the latter, the relevant node exists in the network but cannot be fully activated. As extensively defended in Berg (1992), the segment production problem as illustrated in (13) reflects a difficulty of access. The example is from a German–speaking girl named Melanie.

Leber /leːbɛr/ → [beːbɛr] (from Berg, 1992) (13)

Case (13) highlights Melanie's difficulty with the lateral /l/ that is replaced by the 'easy' segment [b] from the same word. The processes that might underlie the child's harmonization are schematically represented in Figure 1 that captures the moment immediately before the selection of the first segment. Solid lines signal a strong, dashed lines a medium, and dotted lines a weak flow of activation.

Figure 1 illustrates these important principles: While the /l/ is about to be selected, the other segments belonging to *Leber* are simultaneously activated, though to a lesser extent than the lateral. A second claim is that the exchange of information between the lateral consonant and its specification for voicing and place of articulation occurs unhampered. Only the connection between the /l/ and the feature [lateral] is insufficiently developed. The effect is that although the lateral receives some feedback from the feature level, it does not receive enough to be available for production. There is thus a production problem and one for which there is an easy way out in an activation-based processing model. The selection of a linguistic element occurs by inspecting activation levels in the network and simply picking out that unit that has amassed the largest amount of activation (MacKay, 1987). The fact that one particular node is not fully activated does not mean therefore that the production system comes to a standstill. As long as one node is more strongly activated than its competitors, the child will select it, irrespective of whether it is a target or a nontarget unit. The net result is a nonadult-like utterance, a harmonized form.

Stemberger (Stemberger & Stoel–Gammon, 1991; Stoel–Gammon & Stemberger, 1994) accounts for such cases as (13) in terms of underspecification. In underlying representation, /l/ is unspecified for [alveolar] whereas /b/ is specified for [bilabial]. If a given segment is unspecified for a given feature, it is particularly prone to adopting the feature specification from a nearby segment. Therefore, /l/ is harmonized by /b/ rather than vice versa.

As we see it, Stemberger's and our account are not necessarily incompatible. Although ours is a processing account, Stemberger argues that harmony has a representational origin. Let us examine what underspecification means in activational terms. As noted above, /b/ has one feature more than /l/. Given that the activation level of a segment is determined (among other things) by the amount of feedback it receives from the feature level, it follows that /l/ will generally reach a lower activation level than /b/. Thus, underspecification is conducive to hypoactivation.

Bernhardt & Stemberger (1998) provide a further interpretation of harmony from the perspective of Optimality Theory (see also Scott, 1996; Goad, 1996). In this approach, linguistic forms are shaped by satisfying certain output constraints that are of variable importance. For example, nasal harmony may result from satisfying the sequence constraint "NoSequence (−nasal . . . +nasal)." We have a number of reservations about the notion of constraint. To begin with, it is a tautological redescription of the basic observation that a word may not contain a non-nasal and a nasal segment. As such, the notion of constraint has no explanatory power. It remains mysterious how constraints originate and why they occupy certain ranks on the scale of importance. Furthermore, the psychological reality of constraints is less than clear. In particular, the ontological status of negative constraints is questionable. Defining a behavior by a multitude of prohibitive acts is less efficient and more cumbersome than specifying in positive terms what behavior the system is capable of. Note that the number of negative constraints is potentially infinite. Finally, a constraint-based approach presents the learning process in a rather unusual light. It holds that children begin with a great number of constraints that are lost or demoted in the later stages of language acquisition. A much more congenial perspective seems to be to assume that children gradually improve their limited skills as they mature and gain experience with the language.

In the above, segment production difficulties have been ascribed to a hypoactivation problem. Can sequence production difficulties be accounted for in the same way? A few items appear in Melanie's harmony vocabulary in which the consonant being replaced normally poses no hurdle to her. Consider (14).

$$\text{Dom /do:m/} \rightarrow \text{[bo:m] (from Berg, 1992)} \tag{14}$$

This is an instance of partial harmony in which the alveolar stop, which is perfectly mastered by the child, turns into a bilabial one under the influence of the following bilabial nasal. What has to be explained is how a fully producible segment can regularly replace another fully producible segment (or part thereof). Our basic claim is that all that is needed is a difference in the strength of the link between /d/ and [alveolar] and the link between /b/ and [bilabial]. More specifically, the latter must be stronger than the former. There are two ways of looking at this difference. If we take the substituted segment as our point of reference, the connection between /b/ and [bilabial] is characterizable as a hyperactivational one; if, however, we take the substituting sound as our point of reference, the connection between /d/and [alveolar] may be described as a hypoactivational one. These two options are not just notational variants. They contrast in whether the activational difference is located further up or further down the activation scale relative to an average weight value in the network. On the assumption that hypo- and hyperstates have an exceptional status, we may give preference to an account in which the number of such states is kept to a minimum. If there is a small set of segments that are harmonized by all other segments, a hypoactivational problem on the substituted segments is more likely than a hyperactivational problem on the intruding ones because the former requires a lower number of nonoptimal links than the latter. Similarly, if there is a small set of segments that actively harmonize all other segments, the hyperactivation of the intruding units is more likely than the hypoactivation of the substituted elements.

This criterion allows us to arbitrate between the rival accounts. Melanie's harmony was characterized by the intrusion of a single set of segments, the bilabial ones, upon all other elements in her segment inventory. This asymmetry suggests that the hyperactivation of the bilabial consonants provides a superior account than the hypoactivation of all nonbilabial segments.[1]

By way of internal summary, two principles have been argued to be responsible for segment and sequence production problems—hypoactivation (due to insufficient weights on connections) and hyperactivation (due to excessive weights, going by the child's standard). The boundary between segment and sequence repair is a fluid one and the difference between these two types of problem of a gradual nature. In somewhat coarse-grained fashion, we may envision this difference as a sequence of stages. At stage 1, a given segment can be produced under no circumstances whatsoever. At stage 2, this segment can be correctly produced only in the absence of any competitors (e.g., in a CV syllable). At stage 3, it may outweigh weak though not strong competitors, and at stage 4, it can outdo all competing units save hyperactivated ones. This succession of stages mirrors the transition from hypo- to hyperactivation (with normal activation as an intermediate stage). The first stage covers segment repair, the other three sequence repair. Hyperactivation is needed only at the final stage whereas hypoactivation can deal with the first three. Thus, hypoactivation accounts for both segment and (some types of) sequence repair whereas hyperactivation is required for some extreme forms of sequence repair.

The discussion of singleton harmony will be followed by an analysis of multiple harmony. In the case of invariable harmony, there is no obstacle to extending the singleton-harmony account to multiple harmony. Amahl exhibited various harmony types in which one and the same segment is harmonized by another in one context but itself harmonizes another in a different context. Our account of invariable harmony centers around the claim that a rank ordering exists among the individual harmony types. The examples (5) through (7) allow us to establish the following strength hierarchy for Amahl (at one point in his development) (> signifies "stronger than"): /g/ > /d/ > /l/ > /r/. This dominance relation can be straightforwardly captured by adjusting the linkages between the segment and feature levels accordingly. The link between /g/ and [velar] is stronger than that between /d/ and [alveolar]; the link between /d/ and [stop] is stronger than that between /l/ and [lateral] which in turn is stronger than that between /r/ and [approximant]. All that is needed is an appropriate grading of the linkages. This account rules out a "cyclic" harmony pattern whereby the segment S_1 harmonizes S_2, S_2 harmonizes S_3 and S_3 in turn harmonizes S_1. As far as we can tell from Smith's published data, such a cyclic harmony is absent from Amahl's output.

We now proceed to an investigation of variable harmony and inquire whether variable weights can explain the fact that different harmonic strategies may compete for the same word. The analysis will be based upon Suzanne's rendition of the word *chausser* (8) in which a vacillation between palato-alveolar (8a) and alveolar (8b) harmony was observed to occur. It is fairly obvious that many potential explanations immediately fall by the wayside. Hypoactivation on the two links in question would produce a symmetrical effect and therefore does not lead to harmony (which always involves an asymmetry between the harmonizing and the harmonized unit).[2] The same is true of hyperactivation on both connections. Evidently, a change of one link is also not conducive to variable harmony. Because linkage strength is a relatively durable parameter, changing the weights for each new harmonic form is out of the question. Impaired self-inhibition proves equally unsuccessful. It produces a left-to-right bias but no right-to-left bias. Because variable harmony requires both, this error mechanism fails as well.

It has been argued in the foregoing that none of the above error loci taken individually is capable of generating variable harmony. This result prompts the conjecture that variable harmony shows a conspiracy of two separate malfunctions, the one being responsible for the anticipatory (8b) and the other for the perseveratory (8a) effect. Our claim is that while the latter can be naturally explained by impaired self-inhibition, the former most probably evidences a problem with the weights. To understand the nature of this problem, it is fitting to consider Suzanne's anticipatory *invariable* harmonies. If the same processing

TABLE 1
Confusion Matrix of Suzanne's Anticipatory Invariable Harmonies Involving Place-of-Articulation Features

Target Error	Labial	Alveolar	Palato-alveolar	Velar
Labial	—	7	4	2
Alveolar	0	—	2	15
Palato-alveolar	0	0	—	3
Velar	0	0	0	—

principle underlies variable harmony and anticipatory invariable harmony, these two harmony types may be expected to pattern in similar fashion. Table 1 lists all segment substitutions in Deville's corpus that are brought about by anticipatory invariable harmony.

As can be seen from Table 1, Suzanne's anticipatory invariable harmonies exhibit the following rank ordering: [labial] > [alveolar] > [palato-alveolar] > [velar]. Because palato-alveolar segments are weaker than alveolar ones, we propose an underdeveloped connection between /ʃ/ and [palato-alveolar]. This deficit is claimed to be responsible for the anticipatory part of variable harmony as well as for anticipatory invariable harmonies.

This hypothesis entails the following prediction. If hypoactivation accounts for anticipatory variable harmonies, the weaker segment must always occur to the left of the stronger one. If it did not, an anticipation would be prohibited because the leftmost element is too strong to be harmonized. This is in fact the case. Although the number of variable harmonies reported by Deville (1891) is low ($N = 4$), all of them exhibit the weak-before-strong pattern as exemplified in (8).

The low number of variable harmonies suggests that the impairments of the two error mechanisms are differentially severe. To be more precise, the impairment of self-inhibition is most likely to be minimal. If it were stronger, a higher number of variable harmonies would be expected. In contrast, the problem with the weights is larger because it is also responsible for the invariable harmonies that figure so prominently in Suzanne's output.

This complex account reflects the intricacy of variable harmony. We split variable harmony into its component parts and assigned the anticipatory and the perseveratory effect to two independent error mechanisms. These malfunctions explain the existence of variable harmonies but do not by themselves decide which of the two harmony types (8a) or (8b) will win out. We submit that the mechanism that tips the scale in favor of anticipatory or perseveratory harmony is noise that is a general characteristic of adults' and children's processing systems alike (see Stemberger, 1985, 1989).

The account that has been offered of complete harmony can be straightforwardly extended to partial harmony. Because the harmony process usually involves a single subsegmental dimension, the linguistic structure is the major determinant of partial versus complete harmony. Trivially enough, when the interacting segments differ in one feature only, complete harmony must occur. However, when they differ along more than one phonological dimension, partial harmony is most likely. This follows from the claim that only individual connections are underdeveloped whereas other connections of the same segment may be functioning properly.

The aspect to be discussed next is the probability of harmony in harmonic contexts. The previous analysis of compulsory harmony relied heavily upon hypoactivational links. Basically, the same explanation applies to children such as Jacob who vacillate between harmonic and nonharmonic renditions of the same word [see (9)]). Children evincing optional harmony are closer to the adult norm because they show themselves principally capable of pronouncing the relevant words in adult-like fashion (even if they cannot do so consistently). This leads us to posit weakly impaired connections, that is, minimal hypoactivation. However, this mechanism alone does not suffice. To deal with the fact that harmony and no harmony may alternate, noise has to be invoked as the decisive factor.

It bears emphasizing that there is an intimate connection between the gradual elimination of hypoactivation and the recourse to noise. With severe hypoactivation, noise cannot make up for the insufficient transport of information along the pertinent links. As a result, nonharmonic adult-like forms fail to see the light of day. With medium hypoactivation, the insufficiency stays within the range that can be compensated for by noise so that correct and incorrect forms may co-occur. With mild hypoactivation, the correct form is the rule whereas harmonic forms are exceptional.

Up to now, consonant harmony has been treated exclusively as a phonological problem. However, examples (10–11) showed that a harmonic process may apply in one word but not in another even though both words meet the requisite phonological conditions. In this case, the locus of the processing problem has to be shifted upwards from the phonological to the lexical level, that is, to the links connecting segments and words. Specifically, the link between the segment /t/ and the word *tuck* in (11) can be assumed to be fully developed since the onset is produced correctly. By contrast, the link between /t/ and *take* in (10) is claimed to be underdeveloped. As a result, too little activation accumulates on the /t/ that may therefore be outweighed by the most strongly activated nontarget element, which is [k].

The last issue that we have to look at is the interaction between harmony and segment deletion. Let us come back to the shift from a nonharmonic to a harmonic form in the word *boot* (12). In the first stage, the labial is produced fine and the alveolar deleted. This suggests that the [t] suffers from hypoactivation. We assume that this hypoactivation is brought about by a weak link between "boat" and /t/. Because with the appearance of the /t/, the bilabial undergoes harmonization and ends up as [d], it may be assumed that the link from /t/ to [alveolar] is stronger than that from [b] to [bilabial]. We thus posit an underdeveloped link between /b/ and [bilabial] as a second error locus. The change from [bu:] to [du:t] is claimed to be caused by an increased link between "boot" and /t/. This heightened availability of the /t/ may occasion harmony because the /b/ remains insufficiently activated. The preceding discussion is summarized in Table 2.

It can be readily seen that all known manifestations of consonant harmony may be reduced to a minimum number of error mechanisms of which the notion of hypoactivation is clearly the most important, accounting for the great majority of production problems. It thus turns out to be a most powerful tool in the psycholinguistic analysis of phonological development. Hyperactivation and impaired self-inhibition play a rather more marginal role. In the next section, the feasibility of the psycholinguistic model will be examined by means of computer simulations.

TABLE 2
Summary of Mechanisms Assumed to Underlie the Various Harmony Types

Harmony types	Assumed underlying mechanisms
Harmony as segment production problem	Hypoactivation (insufficiently low weights)
Harmony as sequence production problem	Hypoactivation or hyperactivation
Multiple harmony	Graded hypoactivation (several connections affected)
Variable harmony	Hypoactivation + impaired self-inhibition
Compulsory harmony	Severe hypoactivation
Optional harmony	Mild hypoactivation
Inconsistent harmony across similar words	Hypoactivation between word and segment level (one connection affected)
Omission leading up to harmony	Hypoactivation between word and segment level (two connections affected)

IV. THE SIMULATIONS

The ensuing computer simulations were run on an hp-workstation (up to and including an hp 735). They were written in a network description language that is translated into ordinary C by means of a special translation program.

The Simulation Model

The simulation model to be introduced below emulates the psycholinguistic model as closely as possible. It thus follows the tradition of local connectionist networks as developed by Dell (1986) and Stemberger (1985). Although the distributed connectionist approach has greatly advanced our understanding of language acquisition and breakdown (e.g., Seidenberg & McClelland, 1989; Plaut, McClelland, Seidenberg, & Patterson, 1996), the extension of this approach to adult language production has up to now been less than successful. Specifically, distributed models fail to generate contextual slips of the tongue in realistic proportions (Dell & Juliano, 1996). Because child-language harmonies are errors of precisely this type, we decided to work within a localist framework.

The relevant part of the network underlying the simulations consists of a word level, a syllable level, a segment level, and a feature level. The latter is divided into six sublevels. Three sublevels are connected to consonants, that is one for place of articulation, one for manner of articulation and one for voice. The other three sublevels are connected to vowels. They will receive no further attention here because our focus is on consonant harmony. Structural nodes are also included. These nodes are comparable with Dell's (1988) wordshape header nodes and serve the function of serialization (Eikmeyer & Schade, 1991). We—like Stemberger but unlike Dell—use inhibitory links to implement lateral inhibition among the nodes of each (sub)level. The activation function, which calculates the activation values of the nodes during the simulation, is very similar to the one used by McClelland and Rumelhart (1981) (see Schade, 1992 for details).

In connectionist models, errors result from an *incoherent* state in the network at the moment of selection. For example, a feature error arises when the intended segment is the most highly activated node at this level, but the most highly activated node at one of the feature sublevels is not "contained" in this segment. To identify the consonantal output of the system, the activation levels of the following four nodes have to be inspected:

- the node of maximum activation representing a consonant
- the node of maximum activation representing a value of voice
- the node of maximum activation representing a manner of articulation
- the node of maximum activation representing a place of articulation.

The segment that is outputted is the one that can be identified as the triumvirate of the most strongly activated voice, place and manner nodes. In other words, the "material shape" of the segment is determined by the activation pattern at the feature levels. For example, if /d/ is the most highly activated segment node but the most strongly activated nodes at the feature sublevels represent [voiced], [stop] and [bilabial], the segment that is made up of these features (i.e., /b/) is outputted. In this way, a mismatch between the intended and the actual output arises. In the error-free case, the outputs from the segment and feature levels are of course identical.

An incoherence emanates from noise and/or the developmental mechanisms discussed above. Both sources will be investigated in our simulations.

The First Series: Singleton Harmony

The first series aims at simulating singleton harmony. Because this is the simplest type of harmony, it may serve as a point of reference for the other harmony types to be examined later.

Design. The network in this (and most of the following) series consists of 17 word nodes. Lateral inhibition functions to reduce the impact of the lexicon on the production of individual items. Words with no segmental overlap have virtually no effect. We therefore decided to select a set of words that are of maximum phonological similarity.[3] In this way, the effects of the lexicon are magnified. Furthermore, provision was made for the requisite syllable ($N = 17$), segment ($N = 12$), and feature ($N = 6$) nodes as well as the wordshape node CVC. Six combinations of the features correspond to actual segment types. As in Dell (1986), the segment nodes are coded in context-sensitive fashion, that is, there is an onset /g/ and a coda /g/.

The strength of fully developed excitatory links is .2, that of fully developed inhibitory links .5. The production of monosyllabic target words lasts for 30 cycles. At the end of the 20th cycle, the onset is selected, at the end of the 25th, the nucleus, and at the end of the 30th, the coda.

Noise is modeled by adding or subtracting a random amount of activation to or from the activation level of each node in each cycle. Two different noise levels were tested—low noise (SD = .0001) and high noise (SD = .0005). To obtain a broad error distribution, 100 simulations were run for each noise condition.

TABLE 3
Results for the Production of *cap* in the Standard Network (First Two Columns), a Network with Hyperlinks from the Bilabial Segments to the Feature [Bilabial] (the Two Middle Columns), and a Network with Hypolinks from the Nonbilabial Segments to Their Place Features (Last Two Columns), with Two Noise Conditions (Low, High)

	Standard		Hyperlinks		Hypolinks	
	Low	High	Low	High	Low	High
Correct: *cap*	99	62	—	1	—	1
Anticipation: *pap*	—	32	100	99	100	99
Exchange: *pack*	1	6	—	—	—	—
Perseveration: *cack*	—	—	—	—	—	—

The first test was performed on the network with optimal weights. In addition, the network was "lesioned" in two ways. Because we considered both hypo- and hyperactivation in our account of Melanie's harmony, we increased the strength of the link between the bilabial segments and the feature [bilabial] from .2 to .25 and reduced the strength of the link between the nonbilabial stops and their place features from .2 to .15 in separate simulations. Although all words in the network are capable of generating the effects we are interested in, we will focus upon the production of two target words with identical segments but different linear orders, that is, *cap* and *pack*. The antisymmetric makeup of these words serves to elucidate the role of syllable or word position in the harmony process and thereby to distinguish positional effects from nonoptimal connection weights. All in all, 12 experimental conditions (two noise conditions × three network conditions × two target words) were examined.

Results and Discussion. The first notable finding is that in the absence of noise, the network with optimal weights generates all words correctly. The results of the simulations run on noisy networks with nonoptimal weights are summarized in Tables 3 and 4. Table 3 describes the production of the word *cap*, Table 4 that of the target *pack*.

Quite unsurprisingly, the number of correct productions decreases as the noise level rises. Generally, a higher noise level is conducive to a broader distribution of errors. In both the hypo- and hyperactivation conditions in Table 3, low noise always leads to the

TABLE 4
Results for the Production of *pack* in the Standard Network (First Two Columns), a Network with Hyperlinks from the Bilabial Segments to the Feature [Bilabial] (the Two Middle Columns), and a Network with Hypolinks from the Nonbilabial Segments to Their Place Features (Last Two Columns), with Two Noise Conditions (Low, High)

	Standard		Hyperlinks		Hypolinks	
	Low	High	Low	High	Low	High
Correct: *pack*	93	74	100	99	100	99
Anticipation: *cack*	1	19	—	—	—	—
Exchange: *cap*	6	7	—	—	—	—
Perseveration: *pap*	—	—	—	1	—	1

harmonized form *pap*. This is in concert with our interpretation of Melanie's harmonies as regular and consistent production patterns. Somewhat paradoxically, a high noise level may occasionally give rise to an adult-like production that is, in effect, a deviation from the child's typical nonadult-like behavior.

There are vast differences between Tables 3 and 4. Table 3 shows what happens when the stronger segment is in coda position whereas Table 4 shows what happens when it is in onset position. In the former case, harmony is the rule and correct productions are exceptional. In the latter case, it is the other way around. The reason for this difference lies in an asymmetry between onsets and codas. Codas may impact on onsets but not vice versa. This is because onsets undergo self-inhibition after selection and thus cannot interfere with the production of the coda. By implication, even a weak coda can be correctly generated. In contrast, when the selection of the onset is imminent, the coda is concurrently active (due to parallel activation spread) and may interfere with the production of the onset. Our simulation model accordingly predicts a preponderance of anticipatory as opposed to perseveratory harmonies.

It is noteworthy that the hypoactivation and the hyperactivation conditions produce identical results. Both conditions are in principle consonant with Melanie's behavior. To arbitrate between the rival hypotheses, additional simulations were necessary. We examined the target words *tap* and *cat* with low noise. *Cat* is produced correctly when the links between the bilabial segments and the feature [bilabial] are overloaded, but it is subject to anticipatory harmony when the link between /k/ and [velar] is weakened. By contrast, *tap* is harmonized when the link between /p/ and [bilabial] is overly strong but correctly produced when insufficient activation is spread along the links connecting velar segments and the feature [velar]. Thus, unlike hypoactivation, hyperactivation leads to the correct production of words lacking bilabial segments and to harmony in words with bilabial segments. Because this is exactly Melanie's pattern, we are inclined to favor the hyperactivation over the hypoactivation account in Melanie's case.

The Second Series: Invariable Multiple Harmony

We now shift our focus from singleton toward multiple harmony. Our psycholinguistic model sees invariable multiple harmony as a straightforward extension of singleton harmony. In both types, differences in connection weight are presumed to be at the heart of harmony. In the case of singleton harmony, there is only one such difference whereas in the case of multiple harmony, there is more than one. As argued in the preceding section, these differences form a set of dominance relations.

Design. The same network was used as in the first series with one exception. Differentially large differences were worked into the links between the segment level and the place-of-articulation sublevel. The links between the bilabial segments and the feature [bilabial] were optimal (.2), those between the alveolar segments and the feature [alveolar] were set to .175 and those between the velar segments and the feature [velar] to .15. This weighing establishes the following rank order: /b,p/ $>$ /d,t/ $>$ /g,k/. Using this network,

TABLE 5
Results for the Production of the Three Target Words *cap* (First Two Columns), *tap* (Two Middle Columns) and *cat* (Last Two Columns) in a Network with Standard Linkage Strength between the Bilabial Segments and the Feature [Bilabial], Weaker Links between the Alveolar Segments and the Feature [Alveolar], and Even Weaker Links between the Velar Segments and the Feature [Velar], with Two Noise Conditions (Low, High)

	Cap		*Tap*		*Cat*	
	Low	High	Low	High	Low	High
Correct	—	—	—	8	—	1
Anticipation	100	100	100	92	100	97
Exchange	—	—	—	—	—	2
Perseveration	—	—	—	—	—	—

we tested the production of the words *cap*, *tap*, and *cat*. We expected to observe bilabial harmony during the production of *cap* and *tap*, with a more consistent harmony pattern in the case of *cap*, and we expected alveolar harmony during the production of *cat*. As before, we tested the two noise conditions and performed 100 simulation runs for each condition.

Results and Discussion. The results of the second series are given in Table 5. As expected for all three target words, anticipatory harmonies occur more or less regularly. Harmony is most consistent in the case of *cap*. This is due to the large difference in connection weight between /p/ and [bilabial] on the one hand and /k/ and [velar] on the other. In the case of the target *tap*, bilabial harmony also appears, but it is less stable because the /t/ is more strongly connected to [alveolar] than the /k/ to [velar]. The case of *cat* is in-between the other two cases. It is more stable than the *tap* case because the difference in connection weight is .175 to .15 instead of .2 to .175. Relative to the actual strengths, the difference of .025 carries more weight in the *cat* than in the *tap* case. (.025 is one seventh of .175 but one eighth of .2).

The Third Series: Variable Multiple Harmony

We now proceed to variable multiple harmony. We argued above that two opposing forces conspire to give rise to variable harmony. Our aim now is to experimentally demonstrate the viability of this account.

Design. To examine variable harmony in its purest form, we used a network with only one word node, namely, *chausser*. This allows us to show that variable harmony is independent of any particular structure of the lexicon.[4] Seven types of simulations were conducted to study the combined effect of the two assumed forces. In all seven conditions, impaired self-inhibition was modeled by reducing the activation level of selected items to 25% (instead of 0%). The strength of the link between /ʃ/ and [palato-alveolar] was lowered stepwise by .01. Thus, linkage strength was .2 in the first simulation, .19 in the

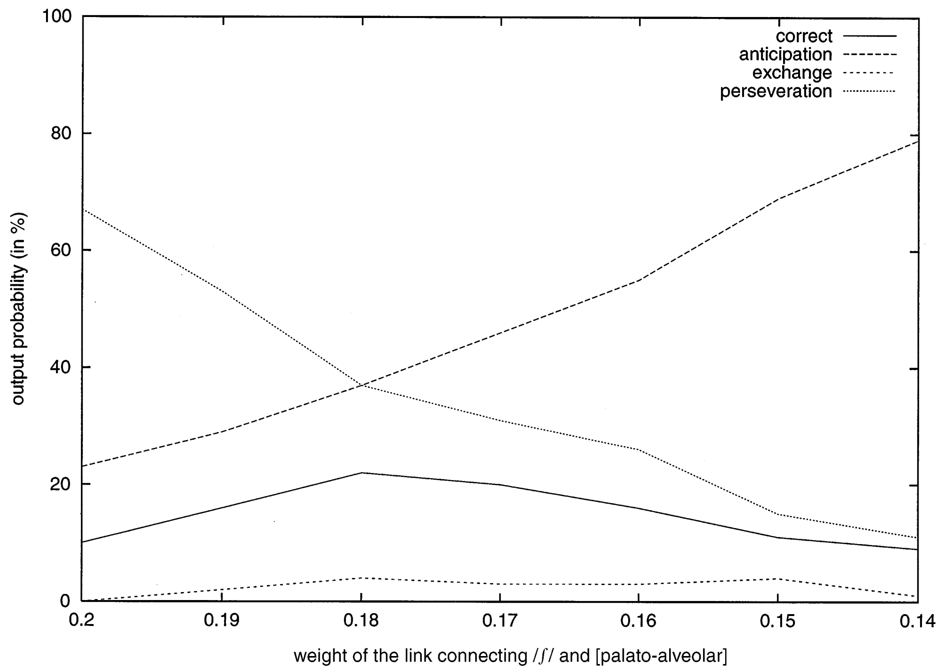

Figure 2. Probability of correct and incorrect output during the production of the word *chausser* in a network with reduced strength of the link connecting /ʃ/ and [palato-alveolar].

second and so forth. The simulations for each network condition were repeated 100 times. The high noise level was chosen.

Results and Discussion. The simulation results are graphically represented in Figure 2. The elevated rate of perseverations as a result of impaired self-inhibition is hardly surprising (67% in the case of optimal links). This percentage descends progressively as the linkage strength between /ʃ/ and [palato-alveolar] is reduced. Anticipatory and perseveratory harmonies are inversely related. The higher the frequency of the one, the lower that of the other. Thus, any proportion of anticipations versus perseverations is possible in a system with compromised links and impaired self-inhibition. Critically, the occurrence of both anticipations and perseverations is feasible in the absence of an appreciable number of exchanges. This is important because exchanges were apparently missing from Suzanne's output. Our conclusion is that genuine variable harmony is brought about by the joint effects of hypoactivation and impaired self-inhibition.

Further Simulations

The following simulations were run with the aim of testing the more detailed aspects of harmony. These simulations will be described in a more condensed manner.

Partial versus Complete Harmony. In the preceding simulations, the critical consonants differed in their place-of-articulation feature only. So, each harmony necessarily was a complete harmony. If, however, the critical consonants differ in more than one feature, partial harmony may also occur. To test this claim, we investigated the production of the target word *tab* using the network with hypoactivated links and low noise from the first series of simulations. We obtained 100 anticipatory harmonies (thereby replicating the result of the first series of simulations) of which only 27 were complete. The other 73 productions were partially harmonized. Only the place feature was harmonized while the onset remained voiceless, to wit: *pab*. Analogous results were obtained with the network used in the second series of simulations.

This result can be interpreted as follows. When complete and partial harmonies have an equal opportunity of occurrence, the latter arise more frequently than the former. This is because it is more difficult to induce complete harmony on words such as *tab*. Either an additional feature change has to take place or the activation pattern at the feature level has to create a substitution error at the segment level (by virtue of feedback). Neither process is very likely. However, the claim that partial harmonies outnumber complete harmonies holds with one reservation that will be discussed in the subsection on consistent versus inconsistent harmony.

Compulsory versus Optional Harmony. We noted above that a large difference in connection weight causes harmony whereas no difference produces adult-like patterns. Networks with small differences take an intermediate position. They exhibit a variable behavior in that, due to the effects of noise, one and the same word may be produced flawlessly on one occasion but harmonized on another. Our simulations show that even low noise may tip the scales in favor of correct versus erroneous productions.

Consistent versus Inconsistent Harmony. According to the psycholinguistic model, inconsistent harmony differs from consistent harmony in respect of the locus of the nonoptimal links in the network. Consistent harmonies stem from differences in the strength of the links that connect segments to their features. By contrast, inconsistent harmonies result from differences in the strength of the links connecting the critical word with its segments. In the latter case, the weaker segment can be harmonized by the stronger one for virtually the same reasons that underlie consistent harmonies. However, there is one notable disparity between consistent and inconsistent harmony. The latter type is always complete. Because it is the substituted element itself that receives insufficient activation, the error must arise at the segment level, not at the feature level. By implication, it is of necessity complete.

The Interaction of Harmony and Deletion. Our explanation of the interaction between harmony and segment deletion relies on our account of inconsistent (i.e., single-word) harmony. On the basis of the network that was used in the first series, we simulated the target *boot* that a child first produced as [bu:] and later as [du:t]. The network was "lesioned" by hypoactivational links between the bilabial segments and the feature [bilabial] ($.2 \rightarrow .15$) as well as by a word-specific hypoactivational link between *boot* and /t/ ($.2 \rightarrow .1$). During the production process at this stage, the alveolar stop never gets enough activation and fails to be produced. Note also that the /t/ does not pass sufficient activation to its features. With time, the connection between *boot* and /t/ improves by learning. When it has reached its optimal strength, the alveolar stop can be selected and forward activation to its features, especially to [alveolar]. Because [alveolar] is more strongly connected to /t/ than [bilabial] to /b/, it can surpass its competitor and induce harmony (i.e., [du:t]).

One aspect of this account bears emphasizing. It has been claimed that in the transition from the first to the second stage, the strength of the link between *boot* and /t/ increases from .1 to .2. This is quite a huge leap. Certainly, it is highly improbable that the child jumps directly from .1 to .2. It is much more probable to assume that such a change is effected bit by bit. Interestingly enough, when this is done, harmonies disappear. We have simulated the gradual transition from .1 to .2 strength of the link connecting *boot* and /t/. With the /b/-[bilabial] link at .15 (hypoactivation) and the *boot*-/b/ link at .2 (optimal strength), the target word is rendered as [bu:] at .1, as [bu:t] from .11 on and as [du:t] from .19 on. This simulation result carries with it an important implication. It allows us to predict that the direct switch from segment omission to harmony as shown in (12) is *not* possible. In-between omissions and harmonies must be a stage during which the target word is correctly produced.

To conclude, our simulations have strengthened the psycholinguistic model sketched above. They have provided detailed insight into the processing mechanisms that characterize children's emergent networks and allowed us to formulate quite specific predictions about their behavior. The next section will highlight several more predictions.

V. PREDICTIONS FROM THE MODEL

The account of consonant harmony as detailed above generates a number of predictions that follow directly from the architecture and processing characteristics of the model. These predictions concern empirical issues that have not been formerly addressed in the literature.

The first set of predictions is about the theoretically most interesting harmony type—variable harmony. We will take up four points. Although invariable harmony is purported to be caused by one underlying mechanism, two mechanisms are claimed to be malfunctioning in variable harmony. Because the simultaneous presence of two error mechanisms is less probable than a single mechanism, variable harmony is predictably less common than invariable harmony. This prediction is in excellent accord with the empirical data. Although invariable harmony is widespread, the variable type is quite rare. Many children never produce it at all (e.g., Smith, 1973), and in those who do, it is restricted to a handful of words (e.g., Deville, 1891).

It was argued in Section 3 that variable harmony should be conceived of as invariable harmony plus impaired self-inhibition. As the anticipatory cases of variable harmony are assumed to be occasioned by the same mechanism that underlies invariable harmony and the perseveratory cases by a different mechanism, we may expect anticipations to pattern differently from perseverations. Specifically, our account makes it logically impossible for anticipations to show the same rank-ordering as perseverations. Let us take the words *shot* and *top* and assume for the sake of the argument that they are harmonically rendered in anticipatory form as [tat] and [pap], respectively. We thus postulate the rank-ordering /p > t > ʃ/ on the basis of a weak link between /t/ and [alveolar] and an even weaker link between /ʃ/ and [palato-alveolar]. If the same target words were rendered in perseveratory fashion as [aʃ] and [tat] and if underdeveloped links were assumed to be the culprit, we would end up with the inverse hierarchy /ʃ > t > p/. Of course, one and the same network cannot accommodate two diametrically opposite rank-orderings. We thus predict that anticipatory variable harmonies should exhibit the rank-ordering that is characteristic of invariable harmony whereas perseveratory variable harmonies should exhibit no such rank-ordering. Assuming that the self-inhibitory mechanism is segment-blind, that is, applies to all segments alike, all segments should be implicated equally often in the perseveratory harmonic process (ignoring frequency effects).

Our model makes predictions not only about general error patterns but also about individual words. Suzanne's variable harmony in cases like *chausser* for instance was argued to evidence a slightly impaired self-inhibition and a weak link between /ʃ/ and [palato-alveolar]. We may wonder how Suzanne would treat words with the inverse order of /ʃ/ and /s/ such as *sachet* [saʃɛ] 'small bag'. The prediction from our model is unequivocal. Because the palato-alveolar fricative is weak and the alveolar fricative likely to be reactivated due to impaired self-inhibition, the overwhelming majority of harmonies should be of the *perseveratory* kind. Unfortunately, we did not find any pertinent examples in Deville's diary to check this.

The next prediction focuses on the relationship between harmonic and nonharmonic forms. Our model of variable harmony predicts not only the absence of exchange errors but also the presence of correct renditions of the target. Whatever the network condition that we simulated, exchanges were extremely uncommon and variable harmonies were always accompanied by a certain percentage of correct productions. We were intrigued by this implicative relationship between variable harmony and adult-like renditions and went into the conditions under which a correct and an incorrect form would be outputted. We ran a simulation to study the interaction of hypoactivation and impaired self-inhibition while varying *both* parameters. The target and only word in the lexicon was *chausser*. The link between /ʃ/ and [palato-alveolar] was progressively diminished from .2 to .1. Impaired self-inhibition was assumed for the places of articulation. It was continually reduced from 100% (optimal self-inhibition) to zero. Continuity was realized by means of nesting of intervals. This method allowed us to compute the critical points at which a change in output occurs. We did not incorporate noise because such a simulation is too complex to be graphically representable in two-dimensional form. Nonetheless, we will resort to noise in our interpretation of the data.

The simulation results are shown in Figure 3. As may be expected, when the deviations from the optimal state are minimal, the target word is produced correctly. When the weight of the link between /ʃ/ and [palato-alveolar] is less than .175, the production result is an anticipatory harmony. When the connection weight is below this critical value and the self-inhibition low (below 66%), perseveratory harmony comes into being. A remark-

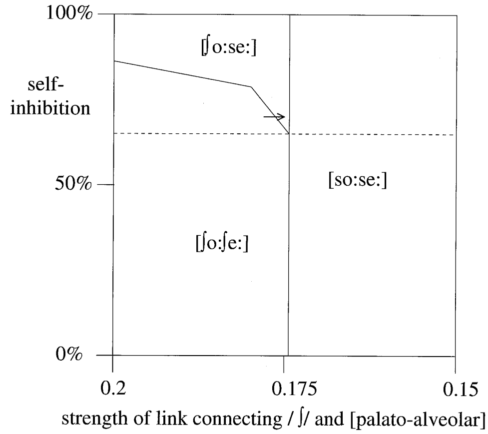

Figure 3. The interaction of the underdeveloped link connecting /ʃ/ and [palato-alveolar] and impaired self-inhibition during the production of the word *chausser* in a network without noise.

able aside is that under certain conditions marked by an arrow in Figure 3, the weakening of the link in question turns the faulty [ʃo:ʃe:] into the correct [ʃo:se:]. This is a consequence of the fact that the two error mechanisms may cancel each other out by neutralizing each other's adverse effects.

The relationship between incorrect and correct productions may now be conceived of as follows. The area of interest appears above the dashed line in Figure 3. The area below depicts an impairment of the self-inhibitory mechanism which is so severe as to be unrealistic. The intriguing implication of Figure 3 is that whatever region the processing system may find itself in, the occurrence of variable harmony implies the error-free production of the target word. This is because the region of correct production is largely between that of perseveratory and anticipatory harmonies. We visualize noise as a momentary shift of the system from one point to another in the diagram. If the system is in the left-hand region (to the left of the arrow) but still capable of generating anticipatory harmonies, it must also, with less extreme noise levels, generate the correct output that is closer to the system's current state. The same holds true if the system is in the right-hand region (to the right of the arrow) and generates perseveratory harmonies. If the system is in the middle, it will evidently be able to produce the target flawlessly.

It comes as no surprise that the pertinent literature, especially Deville (1891), provides no information that would permit us to test this prediction. The production of correct forms is certainly less remarkable than that of deviant forms. This brings us to an important point. As has become apparent throughout this section, research into child phonology suffers from a serious shortage of published empirical data.[5] Theoretical progress will be made only if token-based counts as well as data on correct and incorrect productions are made available. This may necessitate a change in research methodology (see Berg, 1995 for an illustration).

Our model generates predictions not only about harmonic patterns but also about other aspects of phonological acquisition. We will focus here on the syllable or word position in which newcomers to the segment lexicon should show up first. Because, as argued above, processing on the coda is less prone to interference than processing on the onset, we predict that segments in coda positions are easier to produce correctly than those in onset positions. We take it that a new sound is not mastered all at once but its production gradually shifts from a state of vulnerability to a state of robustness. At the initial stage of coping with new sounds, it would be natural to expect segments to appear first in those positions that can be relatively easily accessed, that is, in coda sites.

We hasten to point out two important limitations on this claim. In the first place, it holds only for the period of vulnerability whose length may vary considerably from sound to sound and from child to child. In the second place, it does not hold for the very first segments that are acquired. As is well-known, children's first syllable patterns are almost invariably of the CV type (e.g., Oller & Eilers, 1982; Carreira, 1991). This asymmetry between CV and VC syllables, which requires a separate explanation, does not run counter to our prediction because children begin with only one consonant per syllable. Hence, there is no competition between consonants. Obviously, an absent coda cannot compromise the production of an onset. The processing conflict between onsets and codas arises only when the CVC structure has been mastered. Our expectation is that when new segments are acquired at this stage in the learning process, they will make their first appearance in nonword-initial sites. Succinctly put, whereas early acquisitions may well exhibit a preference for word-initial loci, late acquisitions should favor final sites.

To the best of our knowledge, this hypothesis has never been explicitly addressed in the literature. When positional aspects of sound acquisition have not been ignored, scholars have either focused their attention on one position only (e.g., Ferguson & Farwell, 1975; Ingram, Christensen, Veach, & Webster, 1980) or presented their data in such a way that a direct assessment of the research hypothesis is precluded. In particular, the positional preferences have not been examined as a function of the conflict the acquisition of a new sound creates with respect to already-mastered ones (e.g., Olmsted, 1971; Templin, 1957).

In view of these uncertainties, the only practicable way of testing our prediction seems to be to take the inherent difficulty in the acquisition of segments into account. As noted, the easy-to-acquire and therefore earliest sounds make their first appearance in initial loci. Our hypothesis may now be rephrased to the effect that the more difficult sounds should be attracted more strongly by final positions. Preliminary evidence suggests that this is generally true. Among the stops, the bilabials and alveolars are usually acquired early and the velars late. Significantly, although the former almost invariably appear in initial positions first, the latter gravitate towards final positions (e.g., Chiat, 1983; Dyson, 1986; Schultze–Berndt, 1991; Stoel–Gammon, 1996). Fricatives present a greater learning problem than stops and thus are acquired later. In conformity with the research hypothesis, several studies report that fricatives tend to occur in noninitial positions first (e.g., Ferguson, 1975; Stoel–Gammon, 1985; Chiat, 1989; Edwards, 1996). The same holds good of affricates that are especially difficult to produce (Schultze–Berndt, 1991). As another example one may cite the liquids that are also quite challenging. When the rhotic is mastered at a relatively late stage, it appears first in final sites (Stoel–Gammon, 1985). The same trend was observed for the lateral in a longitudinal study of a German–speaking girl (Elsen, 1991). Finally, consonant clusters pose a particular problem to the language learner. We would accordingly predict that children should acquire CVCC syllables before CCVC syllables. The findings of Lleó and Prinz (1996) are in good agreement with this prediction.

These confirmatory pieces of evidence notwithstanding, it is clear that more detailed and comprehensive studies are required for a fuller evaluation of the "final-first hypothesis." We tentatively conclude that there is a certain pull towards final positions at later stages that does not exist at the earlier stages of phonological development. This shift requires an explanation that our model provides. The access of segments in onset positions is easy as long as simple CV structures are used. However, it becomes more difficult when the more complex CVC structure allows for coda consonants that enter into competition with onset consonants. Inherently difficult segments will be better accommodated by easy than by difficult positions.

VI. CONCLUSION

The principal objective of this study has been to identify the psycholinguistic mechanisms that underlie consonant harmony. Not all possible mechanisms were found to be equally important. For example, impaired self-inhibition was observed to play a minor role. Because of its general effect, it is likely to be fixed before phonological development is in full swing. A possible implication of this is that we might find evidence for impaired self-inhibition at an earlier stage of acquisition. In fact, there is one very early, if not the earliest, phonological process that lends itself well to such an interpretation—reduplication (e.g., Schwartz, Wilcox, Leonard, & Folger, 1980; Fee & Ingram, 1982; Ferguson, 1983). We speculate that reduplication might be evidence of a failure to inhibit selected items so that they are reselected. When this failure affects the syllable node, complete reduplication (e.g., *mama*) occurs; when it affects a segment node, partial reduplication (e.g., [bʌbə] for *ball*) occurs.

Hyperactivation may also have a more important role to play in areas other than consonant harmony. This error mechanism seems to be particularly apt to capture developmental disorders. Indeed, such a claim has recently been put forward by Leonard (1992, 1995) who argues that some characteristics of disordered phonologies can be ascribed to hyper-feedback, that is, hyperactivation from lower to higher levels. He shows that disordered children tend to impose a low number of output templates on adult words with different shapes. For example, the pattern /d/ + V + /s/ was derived from words such as *juice* but enforced on words with a different shape such as *eyes* that was rendered as [das]. Leonard's explanation holds that phonological information is excessively circulated upwards and thus alters the structure of other lexical items. However preliminary it may be, this account is a good example of the deeper understanding that may be gained of normal and disordered phonological development when a connectionist perspective on processing is adopted.

Acknowledgments: Part of the material contained in this paper was presented at the Chateau-d'Oex Conference on "The challenge of natural speech by the computer" in the spring of 1993. The first author is grateful to Eric Keller for having invited him and also to the audience for creating a stimulating atmosphere. We are pleased to express our debt of gratitude to several reviewers, including Gary Dell and Joe Stemberger, as well as the editor Mark Seidenberg, for their valuable comments on earlier versions of this article.

NOTES

1. Note that the term *hyperactivation* is defined within the child's processing system. That is to say, a link may be claimed to be hyperactivated because it is stronger than other links that are strong enough to support correct production in nonharmonic contexts. However, relative to an adult's standard, such a hyperactive link may still be *hypo*active, that is not fully developed.
2. Reducing the weight of the two relevant links has another undesirable effect. As Schwartz, Saffran, Bloch, and Dell (1994) point out, this manipulation increases the rate of exchange errors. We would therefore expect utterances like [so:ʃe:] alongside the harmonic forms. However, such utterances cannot be found in Deville's record. More generally speaking, segmental exchanges are an uncommon aspect of child phonology.
3. These are: *back*, *bad*, *bat*, *pack*, *pad*, *pat*, *dab*, *dag*, *tab*, *tap*, *tag*, *tack*, *gap*, *gad*, *cap*, *cad*, and *cat*.
4. There is no reason to doubt that the relevant effects can be replicated in larger lexicons. This follows from our use of lateral inhibition that reduces within-level interactivity. Indeed, we have simulated the production of *chausser* in the larger lexicon used above and have obtained very similar results.
5. This state of affairs might change with CHILDES, a compilation of primary data on child language (MacWhinney, 1995). However, very few files have been phonetically transcribed. This makes CHILDES a less than optimal data source to exploit for phonological purposes.

REFERENCES

Berg, T. (1988). *Die Abbildung des Sprachproduktionsprozesses in einem Aktivationsflußmodell. Untersuchungen an deutschen und englischen Versprechern*. Tübingen, Germany: Niemeyer.

Berg, T. (1992). Phonological harmony as a processing problem. *Journal of Child Language*, *19*, 225–257.

Berg, T. (1995). Sound change in child language: A study of inter-word variation. *Language and Speech, 38,* 331–363.

Bernhardt, B. H., & Stemberger, J. P. (1998). *Handbook of phonological development. From the perspective of constraint-based nonlinear phonology*. San Diego: Academic Press.

Carreira, M. (1991). The acquisition of Spanish syllable structure. In D. Wanner & D. A. Kibbee (Eds.), *New analyses in Romance languages* (pp. 3–18). Amsterdam: Benjamins.

Chiat, S. (1983). Why "Mikey's" right and "my key's" wrong: The significance of stress and word boundaries in a child's output system. *Cognition, 14,* 275–300.

Chiat, S. (1989). The relation between prosodic structure, syllabification and segmental realization: Evidence from a child with fricative stopping. *Clinical Linguistics and Phonetics, 3,* 223–242.

Cruttenden, A. (1978). Assimilation in child language and elsewhere. *Journal of Child Language, 5,* 373–378.

Dell, G. S. (1985). Positive feedback in hierarchical connectionist models: Applications to language production. *Cognitive Science, 9,* 3–23.

Dell, G. S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review, 93,* 283–321.

Dell, G. S. (1988). The retrieval of phonological forms in production: Tests of predictions from a connectionist model. *Journal of Memory and Language, 27,* 124–142.

Dell, G. S., & Juliano, C. (1996). Computational models of phonological encoding. In T. Dijkstra & K. de Smedt (Eds.), *Computational psycholinguistics* (pp. 328–359). London: Taylor & Francis.

Deville, G. (1891). Notes sur le développement du langage. *Revue de Linguistique et de Philologie Comparée, 24,* 242–257; 300–320.

Dyson, A. T. (1986). Development of velar consonants among normal two-year-olds. *Journal of Speech and Hearing Research, 29,* 493–498.

Edwards, M. L. (1996). Word position effects in the production of fricatives. In B. Bernhardt, J. Gilbert, & D. Ingram (Eds.), *Proceedings of the UBC international conference on phonological acquisition* (pp. 149–158). Somerville, MA: Cascadilla Press.

Eikmeyer, H.-J., & Schade, U. (1991). Sequentialization in connectionist language-production models. *Cognitive Systems, 3,* 128–138.

Elsen, H. (1991). *Erstspracherwerb*. Opladen, Germany: Deutscher Universitätsverlag.

Fee, J., & Ingram, D. (1982). Reduplication as a strategy of phonological development. *Journal of Child Language, 9,* 41–54.

Ferguson, C. A. (1975). Fricatives in child language acquisition. In L. Heilmann (Ed.), *Proceedings of the 11th International Congress of Linguists* (Vol. 2, pp. 647–664). Bologna, Italy: Società Editrice il Mulino.

Ferguson, C. A. (1983). Reduplication in child phonology. *Journal of Child Language*, 10, 239–243.

Ferguson, C. A., & Farwell, C. B. (1975). Words and sounds in early child language acquisition. *Language, 51,* 419–439.

Goad, H. (1996). Consonant harmony in child language: Evidence against coronal underspecification. In B. Bernhardt, J. Gilbert, & D. Ingram (Eds.), *Proceedings of the UBC international conference on phonological acquisition* (pp. 187–200). Somerville, MA: Cascadilla Press.

Ingram, D., Christensen, L., Veach, S., & Webster, B. (1980). The acquisition of word-initial fricatives and affricates in English by children between two and six years. In G. H. Yeni–Komshian, J. F. Kavanagh, & C. A. Ferguson (Eds.), *Child phonology. Vol. 1*: *Production* (pp. 169–192). New York: Academic Press.

Leonard, L. B. (1992). Models of phonological development and children with phonological disorders. In C. A. Ferguson, L. Menn, & C. Stoel–Gammon (Eds.), *Phonological development. Models, research, implications* (pp. 495–507). Timonium, MD: York Press.

Leonard, L. B. (1995). Phonological impairment. In P. Fletcher & B. MacWhinney (Eds.), *The handbook of child language* (pp. 573–602). Oxford, England: Basil Blackwell.

Lleó, C., & Prinz, M. (1996). Consonant clusters in child phonology and the directionality of syllable structure assignment. *Journal of Child Language, 23,* 31–56.

MacKay, D. G. (1987). *The organization of perception and action. A theory for language and other cognitive skills*. New York: Springer.

Macken, M., & Ferguson, C. A. (1987). Phonological universals in language acquisition. In G. Ioup & S. H. Weinberger (Eds.), *Interlanguage phonology* (pp. 3–22). Cambridge, MA: Newbury House.

MacWhinney, B. (1995). *The CHILDES project: Tools for analyzing talk*. Hillsdale, NJ: Lawrence Erlbaum.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception. Part 1: An account of basic findings. *Psychological Review, 88,* 375–407.

Menn, L. (1971). Phonotactic rules in beginning speech: A study in the development of English discourse. *Lingua, 26,* 225–251.

Menn. L. (1978). *Pattern, control, and contrast in beginning speech: A case study in the development of word form and word function*. Bloomington, IN: Indiana University Linguistics Club.

Menn, L. (1983). Development of articulatory, phonetic, and phonological capabilities. In B. Butterworth (Ed.), *Language production* (Vol. 2, pp. 3–50). London: Academic Press.

Morton, J. (1970). A functional model for memory. In D. A. Norman (Ed.), *Models of human memory* (pp. 203–254). New York: Academic Press.

Ohnesorg, K. (1965). Das Phonem und seine Realisierung in der Kindersprache. In E. Zwirner & W. Bethge (Eds.), *Proceedings of the 5th International Congress of Phonetic Sciences* (pp. 441–445). Basel: Karger.

Oller, D. K., & Eilers, R. E. (1982). Similarity in babbling in Spanish- and English-learning babies. *Journal of Child Language*, 9, 565–577.

Olmsted, D. L. (1971). *Out of the mouth of babes*. The Hague: Mouton.

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. E. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review, 103,* 56–115.

Schade, U. (1992). *Konnektionismus: Zur Modellierung der Sprachproduktion*. Opladen, Germany: Westdeutscher Verlag.

Schwartz, M. F., Saffran, E. M., Bloch, D. E., & Dell, G. S. (1994). Disordered speech production in aphasic and normal speakers. *Brain and Language, 47,* 52–88.

Schwartz, R. G., Leonard, L. B., Wilcox, M. J., & Folger, M. K. (1980). Again and again: Reduplication in child phonology. *Journal of Child Language, 7,* 75–87.

Schultze–Berndt, E. (1991). Phonetische Substanz und phonologische Theorie. Eine Fallstudie zum Erstspracherwerb des Deutschen. Working Paper No. 13. Institute of Linguistics, University of Cologne, Germany.

Scott, S. D. (1996). Child language and optimality theory: The case of consonant harmony. In A. Stringfellow, D. Cahana–Amitay, E. Hughes, & A. Zukowski (Eds.), *Proceedings of the 20th Annual Boston University Conference on Language Development* (Vol. II, pp. 682–693). Somerville, MA: Cascadilla Press.

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed developmental model of word recognition and naming. *Psychological Review, 96*, 523–568.

Smith, N. V. (1973). The acquisition of phonology: A case study. Cambridge, UK: Cambridge University Press.

Stemberger, J. P. (1985). An interactive activation model of language production. In A. W. Ellis (Ed.), *Progress in the psychology of language* (Vol.1, pp. 143–186). London: Lawrence Erlbaum.

Stemberger, J. P. (1989). Speech errors in early child language production. *Journal of Memory and Language, 28*, 164–188.

Stemberger, J. P. (1992). A connectionist view of child phonology. Phonological processing without phonological processes. In C. A. Ferguson, L. Menn, & C. Stoel–Gammon (Eds.), *Phonological development. Models, research, implications* (pp. 165–189). Timonium, MD: York Press.

Stemberger, J. P., & Stoel–Gammon, C. (1991). The underspecification of coronals: Evidence from language acquisition and performance errors. In C. Paradis & J. -F. Prunet (Eds.), *The special status of coronals* (pp. 181–199). Orlando, FL: Academic Press.

Stoel–Gammon, C. (1985). Phonetic inventories, 15–24 months: A longitudinal study. *Journal of Speech and Hearing Research, 28,* 505–512.

Stoel–Gammon, C. (1996). On the acquisition of velars in English. In B. Bernhardt, J. Gilbert, & D. Ingram (Eds.), *Proceedings of the UBC International Conference on Phonological Acquisition* (pp. 201–214). Somerville, MA: Cascadilla Press.

Stoel–Gammon, C., & Stemberger, J. P. (1994). Consonant harmony and phonological underspecification in child speech. In M. Yavas (Ed.), *First and second language phonology* (pp. 63–80). San Diego: Singular Publishing.

Templin, M. C. (1957). *Certain language skills in children*. Minneapolis, MN: University of Minnesota Press.

Vihman, M. M. (1978). Consonant harmony: Its scope and function in child language. In L. H. Greenberg, C. A. Ferguson, & E. A. Moravcsik (Eds.), *Universals of human language. Vol. 2: Phonology* (pp. 281–334). Stanford, CA: Stanford University Press.