

LANGUAGE PRODUCTION

Volume 2

Development, Writing and Other Language Processes

Edited by

B. Butterworth
*Department of Psychology,
University College London,
England*

1983



ACADEMIC PRESS

A Subsidiary of Harcourt Brace Jovanovich, Publishers

London New York

Paris San Diego San Francisco

São Paulo Sydney Tokyo Toronto

Development of Articulatory, Phonetic, and Phonological Capabilities

L. Menn

Boston University School of Medicine

I. Introduction

A. Background

A rather large body of information about the early stages of the acquisition of phonology of English and some other languages (Spanish, Mandarin, Thai) has become available over the last decade or so, and the theory of acquisition of phonology has not only grown, but has changed its nature considerably. Since about 1974, we have moved away from a model in which phonological development was considered to resemble the differentiation of an embryo. In its place we have evolved a notion of the young child as a creature of some intelligence who is trying to solve a problem: the problem of sounding like her companions when communicating with them. This shift of model took place as more diary and small-group studies were published, and in the context of Slobin's similar approach to the acquisition of morphology and syntax (Slobin, 1966; 1973).

In recent years, the study of child phonology has also become distinctly more psychological in the explanatory concepts that it employs. This is largely because the richer data base has made it possible to see a considerable range of individual differences among children. Faced with such diversity, we have had to look below the surface for an underlying unity; and in doing so, we have begun to invoke notions of processing and storage of information in addition to the linguistic notions of articulatory control and phonemic contrast.

In this chapter, I will review the strategies that children are presently believed to use in acquiring phonology, and I will give an account of the psycholinguistic model of early phonology which I think is presently the

most adequate. For more extensive discussion, in addition to the references which will be cited, the reader should see the many important papers collected in Yeni-Komshian *et al.* (1980).

B. Plan of Exposition

We will begin this summary of the acquisition of phonology by looking at the transition from babble to speech in Section II. This is necessary so that we can understand the problems of defining which early vocalizations a theory of child phonology attempts to account for.

Then we shall undertake the construction of a model of child phonology that will allow us to deal separately with three different kinds of information that the child is acquiring: (1) knowledge of how words sound, (2) knowledge of how to pronounce them, and (3) knowledge of allomorphy or abstract phonology, manifested as the relationships among words or morphs that sound somewhat different but are the same in meaning.

In Section III.A, we shall discuss the child's perceptual knowledge of the sounds of language; in Section III.B, we will pause to discuss problems with the notion of "phoneme" in early stages of phonological development. In Section III.C, we will turn to the traditional subject matter of child phonology: the ways in which children pronounce the words of the adult language (Section III.C.2). We will see, however, that modifying the pronunciation of a word is only one type of reaction to the complexities of adult phonotactics; the other type of reaction is the avoidance of particular sound patterns (Section III.C.1). These behaviours can be unified, at least in the early stages of acquisition, by the formal descriptive device of saying that children obey phonological output constraints (Section III.C.3), and data from variety of children are presented which support this description.

In Section IV, we consider how a child may go about inventing rules to derive her output forms from the forms spoken by adults; Section IV.A extends the idea of "ease of articulation" by including skill already acquired at a point in time as a factor affecting the "ease" of a new sound. In Section IV.B, the notion of "naturalness" in child phonology is discussed, and we consider how it can be related to our account of rule-creation. In Section IV.C, we note the non-natural rules that children can also create, including those which appear to stem from a dim awareness of the fact that allomorphy exists in the adult language. Section IV.D describes data on rule origin and growth, and Section IV.E concerns rule overgeneralization.

Section V presents the notion that early phonological development should be viewed as the development of skill in the ability to program and execute complex motor sequences. It begins by noting the theoretical importance of the two irregular phenomena that are most difficult for conventional approaches to deal with: overgeneralizations of rules of child phonology (Section V.A) and phonological idioms (Section V.B). Then the regular

pattern of the arrangement of children's early words into families of canonical forms is recalled in Section V.C. Section V.D attempts to account for these three fundamental types of data within a unified two-stage model of articulatory motor programming.

In Section V.E, we see how this model can be fitted into the overall picture of child phonology that was set up in Section III.A; we flex its explanatory muscles in dealing with rules and rule changes, and we consider some of its conceptual limitations in subsection V.E.3 entitled "Caution: The Limitations of the Programming Metaphor".

In Section VI, we deal with some difficult logical and methodological topics: the relation of imitated to "spontaneous" productions and the nature of children's metalinguistic ability to focus on pronunciation as a task (Section VI.A). These are related to the perennial problem of why some sounds may appear in babble but not in speech (Section VI.B). (Section VI may be considered logically prior to most of the rest of the chapter).

Section VII, "The Acquisition of Allophones and Allomorphs", turns to the other major branch of developmental phonology, and gives a brief outline of this topic, especially as it relates to questions of psychological reality. The reader is referred to MacWhinney's (1978) monograph on this topic, since including a full account of it would double the length of the chapter.

Finally, Section VIII, lists the major findings of the past decade of research in developmental phonology, recalls the motor programming model for the beginning stages of acquisition of phonology which was proposed in Section V, and briefly contrasts the working assumptions of the current approach with those of the preceding "Jakobsonian" era.

Note: Some longitudinal studies will be cited repeatedly in this chapter. For convenience, unless otherwise noted, references to "Hildegard" are from Moskowitz 1970b (originally, of course, from Leopold 1939), to "Daniel" from Menn (1971), to "Si" from Macken (1979); to "Jacob" from Menn (1976a,b); and to "Amahl" from Smith (1973).

II. The Transition from Babbling to Speech

We can usually assume that phonology deals with sound patterns of words, but even in adult languages we must decide whether the phonology that we write should attempt to include certain marginal items, for example onomatopoeic representations of animal cries and noises. In studying early child phonology, this problem, marginal in dealing with most adult languages, becomes central. There is no ready-made solution to it; in this section, I will just attempt to show the nature of the difficulty.

There seem to be three definable types of utterances found during the transition period that we call "the onset of speech": sound-play, proto-words, and modulated babble. Modulated babble refers to the use of strings of sounds which appear to carry meaning only by their intonation contour.

This is also called "jargon" and it can be very eloquent and effective vocal communication. Since our concern is with the development of articulation, we will not discuss modulated babble further here; the reader is referred to von Raffler-Engel (1973) and Menn (1976a,b).

Sound-play, which may include word-practice, is not communicative behaviour; in other words, when we classify an utterance as sound-play, we do so because there is no indication of any association between recurrent context and recurrent sound-play patterns. One can of course say that sound-play is expressive of a cheerful mood, but in that weak sense, any evidence of mood is communicative. Joint sound-play is another matter; it is certainly communicative action, but it seems to be absent or rare in adult-child pairs in our culture when the child appears mature enough to be on the threshold of speech, although it is certainly found with young infants (see Sterne *et al.*, 1975; Snow 1977).

Proto-words are articulated meaningful utterances; some of them are directed to others (one can tell because the child gets annoyed if no-one responds), and some are solo performances. These are our objects of study, for only here can we be certain that the child is trying to say a word – that is, trying to match a desired perceived target. And again, we judge that they are meaningful because of a recurrent association between sound and situation, (although obviously if what appears to be a clear token of an adult word is uttered just once in a context for which it is strikingly appropriate, it is usually included as a meaningful utterance).

A child may have all of these utterance types for a period of several months. Some utterances, furthermore, may contain elements that belong to more than one class: for example, a child may start playing sound-games with a "real" word (Weir 1962; Menn 1976a), or he may address one with an utterance that has a real word or two embedded in modulated babble (Jones, 1967). And of course, in practice, some utterances are hard to classify, since classification depends in part on surmising the child's intent.

The important point here is that clear cases can easily be found, and that a child may have one, two, or all three of these utterance types for a period of many months. The "silent period", despite the emphasis given to it in the older literature, is a rare phenomenon.

There is a fourth type of utterance that we should mention. Some children's early attack on language proceeds by global approximation to long phrases rather than by attempts at single words or short phrases. Their early efforts at speech are characterized by variable and often "loose" articulation which is extremely hard to transcribe; Ann Peters (1977) dubbed these children "mush-mouth kids". In this chapter, we shall consider only children who take the more segmental word-by-word approach to phonology; the reader who is interested in the "global" approach should see the Peters article and also Branigan (1979).

Proto-words now need to be defined more carefully. They are vocables (articulated utterances) which recur in definable contexts. One might fear that this notion of recurrent definable contexts would be very difficult to use,

but it generally is not, because a one-year-old's activities tend to fall into identifiable behavioural routines, some solo and some partnered. These include favourite manipulations on objects (putting things into things), games (peekaboo), directing an adult's attention (pointing), obtaining things (requesting/demanding), offering things, greetings, farewells, and so on. Halliday (1975) describes such pairs of vocalization and behavioural routine in elegant detail; see also Menn (1976a), Clumeck (1977). The meaning of a proto-word is originally very limited, and is best characterized as "what you say when you do X". Proto-words may thus usefully be considered as one type of vocal signal; they are not yet symbols, because each of them is bound to the performance of its routine; it cannot be used freely in new contexts. At some point, however, first singly and then more rapidly, some of the proto-words start to be used in more situations, and thus they begin to acquire the symbolic autonomy of the true word. For example, a "woof-woof" vocable may be initially used only when a child is pointing to a picture of a dog; then it may be generalized rapidly to pointing to real and toy dogs, and yet it may take months to become usable in requesting a toy dog. Incidentally, proto-words do not have to have adult words as models (Halliday, 1975), and some without adult models may even make the transition to becoming true symbols (Menn, 1976a, Menn and Haselkorn, 1977).

Proto-words are, by definition, the first units for which a child is trying to produce a particular articulated sound-pattern for communication (always excepting the whole-phrase efforts of the "mush-mouth kids"). If we wish to make generalizations about the child's first phones, or to evaluate the applicability of terms such as "phoneme" to the onset of speech, we must look into the period when proto-words are first being produced. Sometimes what we see is a handful of nicely-defined CV(CV) shapes, as tradition would have it: [papa], [mama], [dada]. A good example is given in Ferguson, Weeks, and Peizer (1973). But more often, apparently, the early picture shows quite a mixture of forms: some vowelless items, perhaps, such as [m:], or Hildegard's [j:]; some "traditional" CV(CV) shapes and/or some (C)VC and VCV shapes; perhaps an isolated word with a consonant cluster (Hildegard, again); and some wildly fluctuating forms that seem to originate from rather complex adult target words (e.g. Jacob's renditions of "thank-you", which showed an endless variation including [deigʌ], [geigu], [gigo], [gi:do], [dejo], [dido], [dætʌ], [vitʌ]).

Summarizing this section: the transition from babbling to speech is typically gradual, and may involve any combination of four types of utterances: sound-play, modulated babble (using meaningful or possibly meaningful intonation contour), whole-phrase efforts, and proto-words. Proto-words are meaningful utterances with phonetically definable targets; however, the phonetic definition may be quite loose by the standards of adult phonetic target-matching and the meaning may be very limited and situation-bound. We will take child phonology as beginning with proto-words, and, in Section III.B we will examine the problem of applying adult-based phonological concepts to these "first words".

III. Constructing a Model of Early Phonological Knowledge

In this section, we will undertake the description of some aspects of "early phonological knowledge". This includes what children, in the first months of speaking, seem to know about the sounds of words in adult language (perceptual knowledge), about the relations among those sounds (phonological knowledge, including knowledge of segmentation and phonemic contrasts), and about how to pronounce words.

The most striking fact about early child words has always been how simplified most of them are compared to their adult models. What has made child phonology an object for study has been three realizations about these "simplified forms": that there are generally systematic relations among a given child's words, that there are generally systematic relations between the child's word and the adult model word, and that it is possible, by comparing children who have very different ways of dealing with adult words, to come up with a general theory of why and how these "generally systematic" relations exist. These three realizations will be developed in this section and in the two which follow.

Note: Beginning in this section, I will occasionally draw small flow-chart diagrams in order to keep track of the various capacities for processing and storage that we postulate in order to account for the child's language behaviour. It is important to keep in mind that the entities and processes represented by these boxes and arrows are only hypothetical constructs, and that even the best guesses among them must be grossly oversimplified compared to whatever it is that we have in our heads.

A. The Input Lexicon: Representation of the Adult Word

"Lexicon" is a word whose precise meaning varies from user to user, but it at least denotes a collection of stored, accessible, memorized bits of information about the sounds and meanings of words and/or their component meaningful parts. We must grant that something which should be called a lexicon exists in the human individual; that is, there must be some form of long-term storage containing at least a sketchy encoding of the sound-pattern and meaning which is accessible when we recognize and understand a word.

In order to say a word spontaneously and meaningfully, one must also have access to stored information about how it sounds and what it means; a standing controversy is whether this knowledge is best represented by postulating a separate "output lexicon" or whether both recognition and production information are best conceived of as being in a single lexicon. [See Chapters 6 and 7.]

To advocate a single lexicon in a psycholinguistic model of child phonology is to hypothesize that the rules which create the child's output form from her

input form operate in real time; to advocate a two-lexicon model is to claim that a form "closer" to the output form is also stored and that this second form is used as a basis for production. Much of the data that we will consider can be handled more gracefully in a two-lexicon model than in a one-lexicon model; I think the two-lexicon model is likely to be a better approximation to what we really utilize in speaking, and so I will use it in this chapter. It is by no means universally accepted as the superior model (cf. N. V. Smith, 1978), however, and formally all the data that it handles can be managed in a one-lexicon system, by the use of markings on each lexical entry specifying which rules apply to it in the event of competing rules applying to the same domain.

We shall say, then, that two forms may be stored for each word: a recognition form and a production form. The collection of words (form-meaning pairs) that a speaker can recognize and understand is called the "input lexicon"; it could equally well be called the "recognition lexicon" or the "passive lexicon". The collection of words that a speaker can use (that is, the information necessary to use them meaningfully and to pronounce them) is referred to as the "output lexicon", but could also be thought of as the "active lexicon". (This active/passive dichotomy is usually thought of as a matter of knowledge of word meaning rather than pronunciation, but the extension of it to include knowledge of pronunciation seems to capture the right distinction.) So far, then, we have the rudimentary diagram shown in Fig. 1.




FIG. 1

Let us explore the properties that can be ascribed to the input lexicon. We know that speech perception is an active process: the hearer filters and structures the incoming sound. Several researchers, including Waterson (1970, 1972), Ingram (1974), Hawkins (1973), Macken (1979) and Wilbur (1980), have called attention to the possibility that a child may not succeed at first in getting a complete picture of a word he has begun to learn. Therefore, we may be more accurate in particular cases if we represent the child's knowledge of some part of the word's sound pattern by "noise" (Ingram, 1974) or by underspecified phonemes (archiphonemes, macrophonemes). These are useful notational devices whenever we have reason to believe that, for example, a child has not figured out what sounds are present in the unstressed syllable of a word or has been unable to tell which of several fricatives a word ends with. To be more explicit, these devices are useful notations whenever the child apparently cannot distinguish perceptually among particular sets of similar words.

Note that we cannot rely on the child's pronunciation to let us know what perceptual distinctions she is making, for children can in fact frequently tell the difference between two words while they are still unable to pronounce either one of them. (Wilbur points out in personal communication that in adults, cross-dialect phenomena continue to give examples of perception outstripping production: American Midwesterners who do not distinguish among /ɛ,æ,e/ before /r/ in their speech nevertheless can reliably distinguish "merry", "marry", and "Mary" in the speech of those who do make the distinction.)

To give two simple examples of the use of these notations for incomplete phonetic input information: suppose that a certain child appears unable to distinguish between two words which differ only in the shape of a pretonic syllable, such as "along" and "belong", but that she can distinguish them from "long". Then "noise" marker would be appropriate to represent the first syllable of iambic words in the input lexicon.

Now suppose that we have a child who cannot distinguish /bæs, bæθ, bæf/ from one another at an above-chance level in an appropriate test situation, but who can tell them from /bæt, bæv/. Here, the child has some knowledge of the final sound of, say, "bath", so we would not use a noise-marker. Instead, we would say that "bath" is entered as /bæ (unvoiced fricative)/ in the child's input lexicon.

So, what we have been saying is that the child's ability to use acoustic features to discriminate meaningful words is typically well ahead of her ability to control those features for making contrasts in production, but may well be inferior to the linguistic discrimination ability of the adult. Some discrimination which the child appears to make may in fact be carried out partly on the basis of extra-linguistic information and linguistic context. For, like all of us, a child's ability to "hear" is conditioned by her expectations of what she is about to hear. This factor is important to emphasize for two reasons. One will be discussed in Section VI, where we will explore some implications of Barton's (1976) work which shows that unfamiliar words in minimal-pair tests of discrimination ability tend to be misheard as familiar ones. This biases the tests and increases the difficulty of ascertaining what the child's input lexical representation of a word "really is".

The other reason for bringing up the notion of the child's expectations is the following phenomenon: Macken (1979) and Platt and MacWhinney (1983) have argued that we sometimes have good evidence for the following sequence of events. First, a child learns to recognize the sounds of a word adequately but cannot produce it very well: we say that the input representation is good, but not the output representation. Usually, the child will then slowly bring the production into line with the target, but in certain cases, expected improvements fail to occur in particular words or sets of words. The child maintains his old pronunciation in such a way that it seems that he is no longer even trying to match the adult model. Instead, it seems that he has replaced his original input representation with a new one which is based on his own output. For example, Macken (1979) gives this analysis for

certain events reported by Smith (1973). Amahl, his subject, produced the word "take" as [Geik] at an early stage, using a general velar assimilation rule (a type of rule which we will shortly be discussing in some detail). The rule stopped operating for all other words by Smith's "stage 14", but Amahl retained a velar-harmonized form for "take" until "stage 22", and even created a participle [kukən] for "taken" at "stage 18".

Now, if a child maintains his own form when he is capable of improving it, it must mean that he has temporarily stopped monitoring, stopped really listening to himself and/or to the adult model. He expects that he is correct, and does not bother to check up. Indeed, many of us have adult acquaintances who have an idiosyncratic pronunciation of some word, and who seem quite unaware that they are not speaking as other people do. Many irregularities in children's phonological behaviour thus seem to be explainable in terms of the biasing of perception by expectation.

B. Segments, Phones, and Phonemic Contrasts

Now we will consider the early stages of the production of proto-words and words. Early child speech is often called pre-phonemic (Nakazima, 1972; Menyuk, 1977). There are very good reasons for this. One is that phonemic contrast and phonetic control do not develop in synchrony. One example of this sort of uneven development can occur when a child honours a contrast without being able to handle the relevant phonetics at all. So we may find a child who renders the voicing contrast in word-final position by deleting voiced final stops and producing the unvoiced stops as a glottal stop. In such a case, for example, the pair "bead, beat" would be rendered as the pair [bi, biʔ]. This hypothetical child has preserved a phonemic contrast without being able to produce either adult phone involved.

The converse case can occur as well: phonetic control can develop ahead of phonemic contrast. It is very common for all initial stops, regardless of target, to be produced by a child learning English as "voiced" (more precisely, to have voicing onset time between 0 and 20 msec; see Macken and Barton, 1980). In such a case, the phonetics of voiced (short-lag VOT) initial stops could be under control, but not the phonetics of unvoiced (long-lag) initial stops. One could correctly say that the child at this stage had acquired the *phones* [b,d,g], but it would be quite wrong to say that she had acquired the *phonemes* /b,d,g/ since she does not have the contrast between them and /p,t,k/. (For further discussion with examples, see Moskowitz, 1975).

The second reason why the concept of phoneme is difficult to apply in the early stages of language development is that for many children, minimal pairs (pairs of words differing only by the contrast in question) are so rare as to make statements about the presence or absence of contrast impossible (see Itkonen, 1977).

And the third good reason for calling early speech pre-phonemic is even

more linguistically unsettling. At least we can speak of phones in the first case above, and nothing prevents us from doing so in the second case. That is, we appear to have phonetic targets which are comparable to one another, independent of the lexical items – the particular words – in which they are located. In adult language, we expect that any difference between, say, an /a/ in one word and an /a/ in another will be completely due to the sounds surrounding them, the stress pattern, and possibly to some kinds of morpho-syntactic factors (e.g. being used as a clitic) or more social factors (formality, rate). We are not prepared to see arbitrary variation in phonetic targeting between one lexical item and the next. Yet it does happen; it even occurs in adult language in special marginal cases.

Let us first consider a special case in adult English where a segment fails to satisfy the criteria for being a phone. The “o” of “no” is subject to a huge amount of variation in realization because of the expressive roles it plays; it can occupy almost all positions in the English vowel space “below” a diagonal from [æ] to [o], including for example [ɔ, a, æ], and [ə] as well as the citation form [o]. We must therefore record as a lexical fact about the word “no” the colours its vowel would take – in other words, we cannot describe the vowel in “no” as the phone [o], and if we insist on saying (for good reasons outside the scope of this paper) that it is still the phoneme /o/, we must have a special marking in the lexicon preventing this /o/ from having its usual phonetic spelling-out as [o] in certain usages.

The child phonology case to be cited here, from Jacob, parallels the adult one: the problem is caused by inconsistencies in the amount of variation found for what should be two instances of the same phone. Jacob produced many tokens of the targets “down” and “round”; both favourite action words. The vowels of the two words differed in output: the renditions for “down” were much more variable than those for “round”. But there was no reasonable way to ascribe this difference to phonetic conditioning or to any of the other factors just cited as causing variation. Thus, these two segments could not be considered tokens of the same phone.

Similar problems in the definition of consonant phones were noted by Ferguson and Farwell (1975), and contribute to Ferguson’s repeated suggestions that the earliest productive stage of language acquisition should be considered a lexical acquisition period rather than a period of acquisition of primitive phonemes. In this chapter we will be working towards a compromise model that allows for both the idiosyncratic properties of segments in particular words and the general properties of those segments which do seem to be comparable from one word to the next.

C. Strategies for Dealing with Words and Sounds

There seem to be a number of strategies that children may draw on as they try to render adult words within their limited articulatory abilities. Two types of strategies have been clearly identified in the literature to date. The

first type induces little distortion in the model word, while strategies of the second type tend to modify it considerably. Most children probably draw on all of these strategies to varying degrees. However, some of them rely quite heavily on those which do little violence to the model word, while other children show no compunction about making gross changes in a fair number of the words that they attempt. (It has often been speculated that this is a matter of cautious v. bold temperament on the child’s part, but to date there has been no systematic attempt to compare phonological behavior with any aspect of personality, or even with the strategies chosen for acquiring any other aspect of language.)

1. *Non-distorting strategies: avoidance and exploitation*

The non-distorting strategies, which may also be termed “selection strategies”, are (a) avoidance and (b) exploitation of favourite sounds.

(a) *Avoidance.* By avoidance we mean that the child does not even attempt to say words containing certain adult sounds. The confirmation that this phenomenon can exist in normal children as young as 15 months old, and not merely in the older child who has required articulation therapy, is a matter of major importance on both linguistic and psychological grounds. Linguistically, it is important because it lies entirely outside the range of behaviour considered by Jakobson and requires the construction of additional acquisition theory (see Ferguson and Macken, to appear). Furthermore, it provides one of the clearest demonstrations of the fact that perceptual discrimination can precede production by many months; if there are two similar sounds and one is avoided while the other is attempted, the child must be able to discriminate between the two sounds while being able to make only one of them. Psychologically, avoidance is a stunning phenomenon because it implies considerable metalinguistic awareness on the part of a child who has only recently begun to speak. After all, avoidance must be the result of a kind of decision.

Consider a child who imitates and uses a set of words beginning with, say, /d/, but who will not attempt any with /b/ even though he has demonstrated comprehension of b-initial words like “ball”, “block”, “box”, and so on. At the very least, such a child must have the feeling that there is something special the matter with b-initial words, some reason why he does not want to say them. Ferguson and Farwell (1975) suggested that this might be happening in some of their subjects; Menn (1976a) was able to demonstrate the b/d case just cited for Jacob, including showing that the child knew the meanings of a good number of b-initial words; and Schwartz and Leonard (1982) showed that avoidance could be demonstrated experimentally in children near the onset of speech (having less than 50 words), although not in somewhat older ones (Leonard, Schwartz, Folger, and Wilcox 1978).

(b) *Exploitation of favourite sounds.* Some children early in their speaking-lives seem to seek out adult words that contain particular sounds and add

these words preferentially to their output, although they learn other words as well. Farwell (1976) is the first study to document this strategy; her case, from the collection of the Stanford Child Phonology Project, was a little girl who apparently especially liked fricatives and affricates, for her output was loaded with words like "juice", "choo-choo", "shoes".

It is clear that both avoidance and exploitation are strategies that we should expect to find if a child is, in fact, treating the mastery of pronunciation as a problem to be solved, and is capable of avoiding perceived areas of difficulty and of capitalizing on perceived areas of success.

2. Modification strategies: rule use

Now let us consider modification strategies, those which result in changes to the shape of the word. One case has become familiar: the case of rule-use. Here, the child has a systematic method of dealing with adult words, one that can be described by a set of rules for substitution, omission, and occasional metathesis of the sounds of the adult word. First we will consider some typical examples of this well-studied type of modification strategy, and then, in Section IV.4.C, we will study some more unruly modifications.

Child-phonology rules represent the child's modifications of the adult model word in a segment-by-segment fashion. They are usually written as direct maps from the adult sound to the child's sound. When the rules are written this way, of course, a step is left out: the psychologically intermediate but inaccessible step of the child's internal recognition encoding of the adult model word which we just discussed in Section III.A. For the present, we will write rules without that intermediate step; when we discuss the construction of a psychological model for child phonology in Section V., we will put it back in again, and also hypothesize some other intermediate processing levels.

To begin with, let us consider a hypothetical child near the beginning of speech who has the following list of words:

hat [æ']; boy [bɔj]; cat [kæ]
 nice [naj]; house [æw]; dog [da]
 please [pi]; blue [bu]; clock [ka]
 drum [dʌ]; up [ʌ]; down [dæw]

This "child" would appear to substitute glottal stop for final /t/ and to delete other final consonants. Initial /h/ is also deleted. Liquids are dropped from consonant clusters. These statements may be translated into formal terms like this:

$t \rightarrow \text{ʔ} / _$ (/t/ becomes glottal stop and then
 $C \rightarrow \emptyset / _$ all other consonants are deleted word-finally)
 $[+ \text{cons}] \rightarrow \emptyset / _$ (liquids are deleted from initial clusters)
 $[+ \text{voc}] \rightarrow \emptyset / \# C _ V$
 $/h/ \rightarrow \emptyset / _ \# _$ (h is deleted word-finally)

The reader may have noted that these four rules are not the only ones that can be devised to describe the observed behaviour. It is important to understand that in most cases we do not get enough different words from a young child to determine her set of rules fully. Rules are always to be regarded as the analyst's tentative hypotheses about the child's mental operations. And it is also important to remember that a rule is no more than a description of a hypothesized regularity of behaviour. It is not an explanation of anything to say that a child "has" a deletion rule or a substitution rule, just as it is no explanation to say that an apple falls because of gravity.

Now let us examine in more detail two of the best-known rule types of child phonology: assimilation and voicing/devoicing. Towards the end of this section we shall also see that there are other strategies that children use which produce the same effects that these rules do.

(a) *Assimilation and consonant harmony: assimilation rules.* We often notice that young children have rules which change the consonants in a word to make them more similar to one another. As in general phonology, these are called consonant assimilation rules. For example, a child who can say "daddy" with a good initial [d] and "egg" with a good final [g] may yet say [gɔg] for "dog". Such a child usually also says [gʌk] for "duck" and "truck", etc. These rules may be so strict for a time that all the consonants in any given output word must be homorganic – that is, made with the same position of the articulators. "Boat", for example, would have to be produced as either [bowp] or [dowt].

Assimilation involving the feature [nasal] is common, too, in child phonology: "dance" may become [næns], with the /d/ assimilating in nasality to match the following nasal; or "meat" may become [dit], with the /m/ losing both its nasality and its labial position as it assimilates to the final /t/. (Both of these forms are from Daniel, Menn 1971).

Sometimes a child may produce some non-harmonic sequences and yet apparently require harmony in other words: he may say "gate" correctly, but produce [gtg] for "big" and [gejk] for "take". In this case, the assimilation of labials or dentals to velars occurred only if the velar was word-final; if it was word-initial, both stops were produced correctly. Relative position of the consonants in a word is often a factor when some sort of asymmetry of consonant harmony is found (Ingram 1974). Vihman's (1978) survey suggests that sounds at the beginning of a word are somewhat more likely to be the ones which are changed when there is an assimilation rule, but this is merely a tendency.

Assimilation rules can be found in great numbers in adult language as well, but there is an important difference. In adult language, the usual type of consonant assimilation is contact assimilation: a segment changes and becomes more like one that is next to it. Although many adult languages have vowel harmony, which occurs even when consonants lie between the vowels, very few adult languages have consonant assimilation at a distance; Vihman finds it in only three of the 88 languages in the Stanford Phonology

Archive (not including some cases in which the intervening vowel is "coloured" by nasalization in nasal harmony, or by pharyngealization in pharyngeal harmony; these cases are called "prosodies" by Vihman). Something special is taking place in child phonology.

When we find deletion rules, as in our initial examples for this section, or contact assimilations like the change of "ask" to [æst], we usually feel that mechanical "ease of articulation" should account for them. But when we contemplate distance assimilation, we find our intuitive notion of simplicity challenged. Why should a child who can say "dad" and "egg" find [gɔg] easier than [dɔg]? Is this to be explained in "natural" terms? In a sense, yes – but in terms of a different kind than we have previously considered, terms which are very important to the construction of a theory of child phonology.

In trying to understand distance assimilation, we can get some help from considering general motor behaviour. Under what circumstances is an ABA pattern of behaviour "easier" than an ABC pattern (assuming that A and C are equally easy to carry out in themselves and as sequels to B)? The only way that doing A again can be "easier" is if the sequence is to some extent pre-assembled or preprogrammed, for in a memoryless series of events it would not matter whether an element is one that has recently been used. In other words, doing A the second time is easier than doing "C" only if we know that we are going to do "A" again and can make use of that information. So the argument goes as follows: young children often use distance assimilation. We take as a working assumption that this must make words easier for them. It cannot make words easier for them unless there is a stage of production at which a word is programmed or assembled before it is spoken. Therefore, I think that a model of how words are produced by young children must have such a stage in it. Later on, we will come back to this point and try to deduce more about the properties of this stage from the data that we have available.

(b) *Other strategies: consonant harmony as a goal.* Assimilation rules are not the only way that children deal with disharmonic sequences. Some children omit one of the offending consonants: Daniel, who used assimilation on "dog", "boat", and a good many other words, said [gej] for "gate", rather than [gejk] or [dejt]. Other children use a glottal stop in place of one of the adult sounds. Such patterns of rule use linked by similar input and similar output strongly suggest that we should take a functional approach to child phonology rules; that is, they make more sense if we think of them as means to some end. And in fact, we have been doing just that: we have been assuming that these rules are somehow designed to eliminate disharmonic sequences.

3. *Output constraints and conspiracies: first mention*

At this point it will help to develop some terms for dealing with sets of rules which appear to serve some common function. Suppose none of the forms produced by a child contain consonant clusters, for example, or that none

have final stops, or that none have disharmonic sequences. A statement that a particular sound-pattern does not appear in a corpus and is not expected to appear if we get a larger sample is a statement of an output constraint. Adult languages have output constraints as well; consonant clusters are absent from many languages, and every language has restrictions on how many and what kind of consonants form a pronounceable cluster (Bell, 1971). Vowel harmony, present in quite a number of languages, is also describable as an output constraint.

Following Kisseberth (1970), when we have a set of rules that all contribute to eliminating sound patterns which would violate a particular output constraint, we say that those rules form a conspiracy. In the example from Daniel, assimilation rules and a (limited) deletion rule were part of the conspiracy to eliminate disharmonic sequences.

Conspiracies of rules are not the only devices that children use to maintain output constraints, however. Selection strategies may also contribute – children may avoid adult words which violate a constraint. Sometimes, this may be a very minor strategy for a particular child (Daniel probably avoided the word "cup"), but sometimes it is a major contributor to the maintenance of an output constraint.

Let us now look at some cases involving another very common output constraint in young children. This one actually involves a pair of phenomena collectively referred to by Ingram (1976) as "voicing": the constraint that initial stops be voiced and final stops be unvoiced. A child may have only one of these or neither, but the pair is very common for English-learning children.

At the acoustic-phonetic level, the statement is slightly different: initial stops tend to be voiceless-unaspirated (short-lag VOT) and final stops to be partially devoiced (see again Macken and Barton, 1980; N. V. Smith, 1973; B. Smith 1979). This difference in statement is not important within English phonology, but it becomes very important cross-linguistically, since voiceless unaspirated stops count as "voiced" in English phonology, but as "unvoiced" in Spanish, French, and many other languages. An explanation for this pair of phenomena should be in terms of the regulation of glottal air-flow – for discussion see Flege and Massey (1980) and Westbury and Keating (1980). If there is any rule which deserves to be called a "natural process", surely it is the rule of final devoicing: it is not only found in child language, but is one of the most frequent rules in adult language, appearing in many forms from a low-level tendency (as in American English) to the familiar German and Russian final devoicing rule and Turkish syllable-final devoicing.

So, many children use the natural-process rule of devoicing final stops, and many also use the natural-process rule of voicing initial stops; Joan Velten is undoubtedly the best-known example. She said [bat] for "pocket", [ba] for "pie", [bat] for "bad", [ap] for "up" and [zas] for "sauce", to choose from a long list (pp. 86–87 in Velten 1941). There are no examples involving velar stops in output, for at this age (23 months) Joan changed all adult velars to coronals (except for [bup], "book"). Other children who have the

same voicing constraint may use a selection strategy: words beginning with /p,t,k/ or words ending with /b,d,g/ may be avoided, and words which begin and end with the preferred sounds may be selected.

Now let us look at a more complicated case, one in which all the three principal stop positions of English were being produced by the child. Here the voicing constraint is in full force in final position: final [p,t,k] have been mastered, while the final voiced stops /b,d/ are avoided, and final /g/ is modified by being devoiced or deleted.

The constraint has been overcome in initial position: the contrast between initial /d/ and /t/ has been mastered and initial [k] has been acquired. Initial /p/ is avoided, but so is initial /g/. (Ferguson (1975) has commented on similar asymmetries of consonant distributions in child phonology and across adult languages.) These statements are summarized in tabular form (Table I). Another important point is exemplified by these data; notice that the voicing contrast has been mastered for initial dentals, but not for initial velars or labials, and that in this case we cannot even say that one value of the feature is present for all three initial stops. A feature that has been mastered (in either the control sense or the contrast sense) in one phoneme may or may not spread to other phonemes in the same word-position. We presently do not know whether it is possible to explain the difference between the cases in which a feature generalizes and the cases in which it remains "bound" to a particular phone.

TABLE I
Jacob's consonants

Initial		Final	
p absent	b mastered	p mastered	b absent
t mastered	d mastered	t mastered	d absent
k mastered	g absent	k mastered	g devoiced deleted

Other rule strategies besides the use of voicing or devoicing rules can be found in children obeying the voicing constraint. We have just mentioned Jacob's occasional deletion of final /g/, but there are much more interesting cases to be found. These are the children who add extra segments in order to render a voicing contrast. It has been claimed that some children add a vowel to the end of a word with a final voiced stop; this brings the sound into the interior of the word where it could be managed. "Bag" might be produced as [bægə] or [bægæ].

Also, two cases are now reported in which children added nasals rather than vowels in their apparent efforts to preserve the voicing contrast in final position. Fey and Gandour (1979) presented a study of a child who found that he could preserve the voicing of adult final stops by adding a final homorganic nasal: "bag" became [bægŋ]. (Phonetically this is rather less

exotic than it looks written out; the effect is just produced by releasing the velar closure before releasing the stop articulation. However, this cannot well be considered a natural process; there is no evidence that there is a general tendency for speakers attempting to maintain voicing through a final closure to fail with this result.) Bowerman (pers. commun.) reports a different use of added nasal segments: one of her daughters added a homorganic nasal before final voiced stops, so that for example "Bob" became [bamp]. The stops themselves were still devoiced, but contrast was maintained (and the insertion of the nasal should have helped to maintain the vowel-lengthening which precedes final voiced stops in English and which in fact serves to carry the final voicing contrast in some dialects).

Now that we have seen how the notion of output constraint can serve to bring together several rules and/or strategies under the observation that they all "serve to maintain the same output constraint", it is time to take a critical look at the notion itself. So far, all we have is description, not explanation. To say that a rule "serves an output constraint" or "is part of a conspiracy" is only organization of data. But once we organize the data in this way, a plausible explanation jumps out at us: the child is modifying unfamiliar sound patterns to make them like the ones he has already mastered. *And that means that the child has to learn sound patterns, not just sounds.* Again, output constraints are only descriptive devices; what they describe are those sound patterns which a child has mastered v. those that he has not. That is why words which do not fit the constraints are almost all avoided or modified. This is the central thesis of this chapter; we shall explore its empirical support and its implications in many of the remaining sections.

4. Another modification strategy: template matching

Now let us consider another type of modification strategy, one evidenced primarily in work done by Vihman (1976, 1981), Macken (1979), and Priestly (1977). These cases involve fairly violent rearrangements of sounds of adult words to match "templates" of preferred sound patterns. The simpler cases can just as well be considered cases of rule use, and usually are described in terms of metathesis (place-exchanging) rules. The more complex cases, however, cannot be described by rules without a lot of artificial special-case magic, for what makes them so complex is the fact that the child's attack on the adult word is not fully systematic.

A good simple case to begin with is Vihman (1976). A child learning Estonian as her first language seemed to have learned to say words containing two different vowel sounds only if the first vowel was lower than the second. The Estonian words for mother, /ema/, and for father, /isa/, do not happen to follow this pattern. For a little while, the child said just [sa] for "father"; then for four months she failed to attempt either word, although "both father and mother made earnest attempts to elicit the words /ema/ and /isa/". At fifteen and one-half months, the child began to rearrange those words to conform to her output constraint: "/ema/ emerged as [ami] or [ani]

... at which time /isa/ also reappeared, now pronounced [asi], and the word /liha/, 'meat', was reproduced, following the same rule, as [ati].

An example of a case where the child was less systematic about the map from the adult word to output is given in Priestly (1977) (also discussed in Ingram, 1979). Priestly's son Christopher treated virtually all stop-final adult two-syllable words and a fair number of vowel/sonorant-final two-syllable words according to the following patterns: Consonant selection:

$$C_1 - C_2 \rightarrow C_1 - j - C_2$$

examples: pillow [pijal]; Brenda [bajan]; tiger [tajak]

or

$$C_1 - C_x - C_2 \rightarrow C_1 - j - C_2$$

examples: rabbit [rajat]; melon [majan]

with a few cases of idiosyncratic rearrangements, such as "streamer" being produced as [mijat]. There was also a choice of vowel treatments; sometimes Christopher was able to match two vowels of the target, but at other times he replaced one or both by [a]. In addition to the cases already listed, consider the apparent metathesis of vowel features involved in his rendition of "woman" as [wajum].

Other two-syllable words which ended in a vowel or sonorants were treated without these special medial-[j] rearrangements: examples are "bacon", produced almost correctly as [bejkan], "kitchen", where the medial affricate apparently caused the only problem, rendered [kɪkɪn, kɪtɪn], and "scissors", [stɪzɪz].

While it is possible to discern some tendencies in Christopher's assignments of particular adult forms to particular outputs, Priestly makes it clear that there is considerable arbitrary variation from word to word. This fact of lexical variation is further emphasized by Christopher's variation across tokens of the same word: "monster" was recorded as [majos] in weeks 4 and 6 of the study, but as [mejan] in week 5; "dragon" was given as both [dajan] (week 3) and as [dajak] (week 4).

In Priestly's case, then, the child had a favourite output shape to fill, but only a few constraints on which consonants and vowels he picked to fill it with. Macken's 1979 subject Si, acquiring Spanish, shows us a much more constrained output template – that is, one which allowed a very limited set of consonants – and a much greater abandon in her treatment of the model word. (The latter fact probably also reflects the much greater proportion of polysyllabic words among her targets.)

Si could produce disharmonic sequences in a word only if one target consonant was labial and another was dental. Adult words which met this criterion were produced so that the labial preceded the dental; much deletion and occasional metathesis occurred.

examples: manzana [mana] pelota [patda]
zapato [patda] elefante [batte]
Fernando [wanno] sopa [pwæta]

In Si's case, the details of what is deleted and what is selected defy organized statement in terms of rewrite rules. As Macken says, this is goal-directed behaviour: the child is looking for consonants that she can fit into her output template and ignoring the rest.

IV. Rule Creation

A. Extending the Notion of "Ease of Articulation" One Key to a New Theory

When a child's production of a word fails to match the adult model, we cannot help assuming that there must be some sense in which what he does produce is easier than what he has failed to produce. But what sense is this? How can [bada] be easier for Macken's Si than [daba]? Why will some children use [l] for /y/ and others use [j] for /l/? Why do some children exploit fricatives while others delete them, avoid them, or replace them with stops? Clearly, if we stick to our common-sense starting assumption, then it must be the case that what is easier for one child can be harder for another. Perhaps a little of the variation is due to anatomical differences, but we simply do not have the means to investigate that hypothesis. A much more fruitful approach is to assume that a great deal of "ease" and "difficulty" is not a matter of physiology at all – or, to put it another way, that physiological causes are only one factor in determining "ease of articulation" for the individual child. The other factor, and I propose that it is the major factor, is the state of a child's knowledge at a given time.

Let me give an example. A child may, as we have said, discover "how to say [l]" before "how to say [j]", or the reverse may be true. Suppose a particular child has discovered [l] first, by chance. We notate this discovery as the invention of a rule taking /l/ into [l]. Now this child may slip into her [l] while trying to say [j], either accidentally or on purpose. If she finds the approximation good enough, she will continue to use it: she will have thus discovered or invented a modification rule. Again, in this case, [l] is "easier" than [j] only because this child happens to have found out how to make an [l] first.

I suggest, in short, that a two-stage discovery process is probably involved in a child's establishment of a new articulatory gesture as her way-of-saying a particular target sound. The first stage is a matter of trial-and-error attempts to match the sound sequence; the second stage is one of deliberate or accidental overgeneralization of the success of that articulatory gesture – that is, the use of it to render similar adult targets.

Let us consider the hypothesized scenario here in more detail, for it is the heart of this chapter's proposal for dealing with one of the fundamental problems of child phonology, namely, how can there be so much individual variation and yet such strong general tendencies? We suppose, then, that

variability across children originates with each child making trial-and-error starts at matching adult sound patterns. For each given sound or pattern, some children will succeed and some will fail. "External" factors, such as the frequency and salience of the sound in the speech of others, may contribute to the likelihood of success; so will "internal" factors: the probability of accidentally hitting on an acceptable way to produce it and the salience of the sound in one's own speech.

We frankly do not know why some sounds are more probable than others; Stevens' (1972) notion that favoured phones are those which are acoustically stable (i.e. permit a certain sloppiness in articulation without showing appreciable acoustic change) is certainly an attractive idea, but we cannot yet simulate the child's vocal tract accurately enough to test this idea with acoustic modelling. (However, progress has recently been made in this area – see Goldstein, 1980). The accidental aspect of learning to produce target sounds is a principal source of individual variation, but it is also a principal source of the probabilistic universals of order of acquisition; roughly and with all due caveats, stops usually are acquired before fricatives, labials usually before velars, nasals usually early, liquids usually late. (See Sander, 1972, both for data on English and for methodological considerations.) If the reader will permit me some licence in the statement of probabilities, we might say that a [b] is a low pair, [k] is jacks or better, [l] is a flush, [θ] is a straight flush, and the fricative [ʃ] which Jakobson dwelt on as the latest-acquired Czech phoneme is a royal flush in spades: some kid somewhere in Czechoslovakia is going to get it phonetically right in her first ten words, but don't bet on her being in your data sample.

We should stress one more thing about this proposed initial trial-and-error stage of discovery: a child may accept her rendition of a sound even when it is quite inaccurate. Some rules that give inaccurate renditions of adult targets therefore arise at this first stage. But many more may arise in the second stage, as the child makes use of her initial accomplishment.

B. "Natural Processes"

It is quite reasonable to say that both $/l/ \rightarrow [j]$ and $/j/ \rightarrow [l]$ are "natural phonetic processes", in that articulatory factors make it quite likely that a clumsy attempt at either of them will produce the other, rather than, say, a [t] or a [b]. Put another way, a child with a certain amount of experience at making speech sounds with his mouth is likely to get some of the properties of, say, [l], correct (in a word that does not present a host of other problems): perhaps the voicing, the continuancy, the central tongue placement, or the lack of rounding. [l] and [j] share all of these properties, so a child who is doing well at approximating one of these two phones is quite likely to end up with the other as his approximation to it.

Informal observation suggests that [l] and [j] are roughly equally likely to be found substituting for one another ([w] or a similar sound is also found frequently for "dark L" [ɫ], of course). In other cases, there is a heavy bias in

favour of one of a pair of phones. For example, in word-initial position, stops are much more likely to be discovered before fricatives and then to be used to substitute for them. Similarly, voiced stops are likely to be used for unvoiced stops in initial position, as we have already seen. We certainly have enough reason to say that "stopping" (use of stop for fricative) and "voicing" are natural in initial position; that is, we have reason to believe that there is a high, physiologically-governed probability that the child making a first attempt at an initial fricative will produce an initial stop, and that the child first attempting an initial unvoiced stop will produce a voiced stop instead. This, I think, is the only coherent interpretation of the notion "natural process", although other views certainly appear to be held (see Stampe, 1969; Ingram, 1976, but also Ingram, 1979).

In summary, I propose that "natural processes" are really descriptions of those pitfalls of learning to articulate which are commoner and more heavily determined by physiology. To build a rigorous theory of the acquisition of phonology, one must also be able to explain why children fall into those particular pits.

And that step would still be only a beginning, for physiology only dictates what articulatory goals are likely to be surrounded with what traps. To explain how children succeed in avoiding or climbing out of them, we need a problem-solving theory, a cognitive theory. The essence of such a theory for the acquisition of phonology, again, is the trial-and-error discovery followed by application of the discovered skill to new cases – a model which will be very unsurprising to any developmental psychologist.

C. Non-natural Rules

There remain some kinds of rules that are at a considerable remove from the solution of particular articulatory problems.

A very important kind of non-natural rule arises as the child begins to attend to the fact that what appears to be the same morpheme is not always produced in the same way by adults. Sometimes that child is correct in interpreting her observations this way – that is, sometimes she has indeed run into a case of allomorphy or of stylistic variation. However, sometimes she is incorrect; what appears to be variation in the shape of a single morpheme is in fact a case in which the adult is sometimes using one morpheme and sometimes using two which the child has failed to segment. For example, if a child notices the "Z"-morpheme of the English possessive and plural appearing on certain nouns but does not yet understand that the final sibilant has one or both of those meanings, he may develop his own phonological "hypothesis" about where those final sibilants are supposed to appear. Daniel (Menn, 1971) created a rule adding [s] to the end of all English words ending in /r/, apparently because there was an accidental abundance of plurals and possessives on names and objects in /r/ in his environment. He may have figured that the sibilant-final forms which he heard were the full and correct forms of the words which he also heard with

final /r/ – that is, he took “pears” as the full form of “pear”, “Peter’s” as the full form of “Peter”, etc.

It is also the case that rules which once had an articulatory base, after they have been invented, seem to acquire considerable autonomy and may generalize without any further articulatory motivation. A child may apply a rule for one segment or (sequence of segments) to a similar one *even though he could have produced the latter correctly*. This seems to be the case for several rules used by Amahl (see Smith, 1978). Rules are much more than articulatory habits, then; they are transduction habits, habits of rendering perceived targets in particular ways. Illustrations and further discussion will be presented in Section IV.E., “Overgeneralization”.

It is too early to make strong generalizations about the ages at which transduction rules of different kinds can be found, but roughly, it seems that the very youngest children’s rules are mostly those which lend themselves to explanations in terms of seeking solutions to articulatory problems; as these problems are overcome, we begin to see more instances of rules that arise from overgeneralizations of other rules, and more rules which reflect the child’s guesses about the reasons for variation in words of the adult language.

D. Rule Origin and Growth

We have already found ourselves considering the topic of rule origin; let us now do so in more generality and depth. We have characterized transduction rules as systematic correspondences between adult and child sound-patterns, ranging from correct renditions (/d/ → [d]), omissions, and natural substitutions (/ð/ → [d]) to the idiosyncratic rule inserting [s] after word-final /r/ that we have just discussed. There is also a range in how systematic a rule is. Some are exceptionless; most have a few lexical exceptions which typically consist of forms that were learned before the child invented the rule in question, or of forms which are the forerunners of a new rule. And some rules have so many exceptions that they reach the point where we are better off abandoning the attempt to write them; the Priestly case was one example of such a state of affairs.

The evidence for the nature of rule-change is somewhat sketchy, because rule-changes can take place in a short time, sometimes within a few hours. Fine-grained longitudinal study is needed to give a picture of Before, During, and After in such cases. This is emphatically not to say that all rule-change is rapid. Replacement of one well-established rule by another may take place over a period of weeks (and fossil forms created by the old rule may survive indefinitely).

1. Rule origin

We have already discussed trial-and-error experimentation as a source for correct transduction rules (/d/ → [d]) and for natural transduction rules. But

it should be noted that a child’s trial-and-error sessions do not always lead to the formation of a rule. Even if the child manages a perfect rendition of some sound pattern, she may be unable to capture the trick of doing it at will. For example, Daniel (Menn, 1971) made dozens of attempts at the word “peach” during the period when his consonants were subject to assimilation.

If he had been able to make the beginning of the word affricate to match the end, he presumably would have had no problem. But he had not learned to produce any initial affricates, and his versions of the word included [dits, citʃ, nits, its, pipʃ] and [pitʃ] itself at various times. He settled on none of them.

Yet sometimes a rule actually emerged within hours: Daniel tried [pəqs] and [dæts] for “box” at 10:16, and later the same day his assimilation rule made its first true appearance, with “dog” as [gVg], a form it kept stably for months (as far as the consonants were concerned).

The other case of rule origin in the literature has been called consolidation (Menn, 1976a). This term is used to describe the situation in which two similar adult target sound patterns are involved in very similar trial-and-error sequences, and end up being handled in the same way. Correct versions of both of the patterns may be produced in the course of the trials. Jacob varied between [ei] and [i] for the vowel of both “tea” and “table” for some weeks before settling on [i] for both. The mutual influence of similar sound patterns is clearly demonstrated in such cases. Template matching can also originate in this fashion – see Vihman (1981).

2. Rule generalization

Rule origin can occur through rule generalization, for of course dividing a rule from its predecessor is often difficult or arbitrary – there is often no sense to the question “is this a new rule or a generalized version of an old one?” Rule generalization basically means the extension of a rule to new cases, and this covers two different kinds of events. To discuss them, we need the concept of the domain of a rule. The domain of a rule is simply the set of cases to which it is actually applied. For example, the domain of a rule that applies to all English voiced obstruent is just the set of all instances of /bdgvðzʒʒ/.

Formally, if we have an exceptionless rule, its domain is specified in its structural description. In the example given, the structural description could be written [+ obstruent, + voice].

If a rule has lexical exceptions, sounds in the excepted words are not in its domain even if they meet its structural description. Thus, if the word “bad” were simply listed as a lexical exception to a rule otherwise applying to all voiced obstruents, the /b/ and /d/ in it would be outside the domain of the rule. If, at a later time, “bad” ceased to be an exception, it would by definition have been brought into the domain of the rule and thus, the rule would have become more general without any change in its structural description at all. We might term this type of rule generalization “lexical smoothing”. Lexical smoothing is important in child phonology because

lexical exceptions to rules are so frequent. Yet it is not really a change in the rule; it is only a change in the set of exceptions to it.

The other type of rule generalization is formally expressible as a relaxation of the structural description, allowing additional phonologically-defined sets of words to be operated on by the rule. For example, a rule which at some point applies only to final /b/ might at a later time apply to all final labials, or to all final obstruents, or to all instances of /b/. Any of those changes would bring new sets of sounds into the domain of the rule, thus generalizing it. A relatively technical note: in child phonology, we often have trouble determining the domain of a rule for various reasons. Here is one interesting problem: Consider the data from Joan Velten given above (Section III.C.3). She had no velars in her output; she had initial voicing and final devoicing of other stop consonants. Should velars be considered to be in the domain of the voicing and devoicing rules? It is easy to write the rules either way (with voicing and devoicing rules applying directly to all stops before the conversion of velars to dentals, or with "fronting" preceding voicing and devoicing). Only in the latter order can the voicing rules be written excluding velars and still give us the observed distribution of forms. Now the fact is that when velars show up, they may not be subject to either of the rules obeyed by the other stops, so it is preferable to write the rules the second way, and thus to make no vacuous claims about the velars. If the velars do show up obeying the voicing and/or devoicing rules, that would then count as a generalization of the two rules.

E. Overgeneralization

Just as in the acquisition of morphology or syntax, rule generalization can create incorrect forms, and thus, from the adult point of view, be overgeneralization. The term is used loosely; typically it is used when a rule produces some "good" results and some "bad" ones. If a rule always produces modified forms ("bad results"), we do not bother to call extensions of it "overgeneralizations" except when they make a child's approximations worse than they were before the rule affected them. Let us consider some examples.

Daniel (Menn, 1971) had the two words "down" and "stone" rendered as [dæʊn] and [don] from the time of his first attempts at them. Then he developed a rule of nasal harmony – he made all of the stops in a word nasal if the final stop was nasal. "Down" and "stone" remained lexical exceptions to this rule; that is, after he had been saying [næns] for "dance" and [neɪn] for "train" for two weeks, he still maintained the two older words in their unassimilated form. Eventually, however, there was a period of time in which he varied between [næʊn] and [dæʊn] for "down", and between [non] and [don] for "stone". Finally, the assimilated forms for these two words took over completely and they were no longer lexical exceptions to the rule.

From the adult point of view, these two words were poorer approximations to the adult model after the rule had applied to them than before (indeed, "down" had been perfect). Therefore, the generalization involved in extending the domain of the assimilation rule to include "down" and "stone" (a case of lexical smoothing, to use the term introduced above) is an overgeneralization of the assimilation rule.

A change in the structural description of a rule can also produce overgeneralization ("recidivism" in N. V. Smith's terminology). Here is his example from Amahl (1973, pp. 152–53): "At stage 1, /s/ and /l/ were normally neutralised as [d], together with all the other coronal consonants . . ." (I omit his description of exceptions to this rule, which generally made coronals into [d].) Then "/l/ began to appear in A's speech before any coronal consonant" – for example, "lady" was rendered either [de:di] or [le:di] So /l/ was optionally excepted from the general treatment of coronals in certain environments. "Then at stage 5 /s/ (and shortly thereafter /l/ became [l] before any coronal consonant . . . : sausage [lɔdid]; shade [le:t] . . ."

Here, the new rule for realizing /l/ as [l] in some environments had added /s/ and /f/ to its domain. So it had generalized by a change in the structural description: the input to the rule had originally been /l/, but later included /l,s,f/. What makes this an overgeneralization? Smith says: "Now originally two words such as "side" and "light" were both [dait], but after the appearance of /l/ before any coronal consonant they became distinct as [dait] and [lait] respectively. However, once /s/ was "liquidised" the two words fell together again – perfectly regularly, as [lait]."

What is lost when this /l/-realization rule is generalized, then, is the contrast between /l/ and the sibilants /s,f/. (Of course there is a compensating gain in this case, because there is contrast of /s,f/ with /d/ only after the /l/→[l] rule generalizes.)

Reviewing this section: we have seen that rule creation can take place through probable or natural failures, such as the production of a stop for an initial fricative, or through the consolidation of similar forms. It is not to be forgotten that the discovery of a correct articulation for an adult sound is also a rule in the sense of a connection between what is heard and what is produced. The child's existing repertoire has a great deal to do with what form new rules may take.

Non-natural rules can arise when a child misapprehends an allomorphic variation and treats it as a purely phonetic rule without semantic significance, or when a child performs major alterations to get a target to fit a canonical form.

Rules can grow and generalize in two ways: by overcoming lexical exceptions (lexical smoothing) or by generalizing the class of sound patterns to which they apply; overgeneralizations can occur as a result of either of these kinds of rule growth.

V. Towards a Psychological Model of Phonological Development

A. The Theoretical Importance of Lexical Exceptions and Overgeneralizations

Lexical exceptions and overgeneralizations are important data for developing a psycholinguistic theory of language acquisition. To begin with, overgeneralizations are inexplicable if one holds the view that the child makes word-by-word progress towards correct productions; that is obvious.

Lexical exceptions are also inexplicable on the neo-Jakobsonian view that the acquisition of phonology is purely a matter of acquiring distinctive features (Menn, 1981). After all, Daniel was able to make the distinction between nasal and non-nasal dentals in production before, during, and after the time his nasal assimilation rule applied: he had no problem producing "daddy" with initial [d] and "no" with initial [n] during the time that he said [næns] for "dance" and so on. So the fact that these words were originally exceptions to the nasal assimilation rule cannot be described in terms of distinctive features. Overgeneralizations cannot be accounted for in terms of acquisition of distinctive features either. "Lexical smoothing" – e.g. the overgeneralization of the nasal assimilation rule to "down" and "stone" is certainly not a matter of learning to make a new distinction, and neither is the loosening of structural descriptions. If we re-examine Smith's "recidivism" case, we see that it only involves a shift in mapping input distinctions onto output ones, not the introduction of new output features. (Amahl mapped /l/ onto [l] in certain environments and all other coronals onto [d]; the overgeneralization which then took place resulted in his also mapping two other coronal continuant consonants, [s] and [ʃ], onto [l] in those environments.)

Similarly, one cannot explain lexical exceptions or lexical smoothing (although one can handle recidivism) within a theory which says that the acquisition of phonology is purely a matter of overcoming natural processes. Consider: if nasal harmony is not a natural process, then the natural process approach is not able to deal with one of the commoner rules of child phonology. On the other hand, if it is a natural process, one has to explain why it did not apply to "down" and "stone" (i.e. why it was "suppressed", in Stampe's terms, for these two words) initially, and then began to show up on other words and eventually on these two themselves.

Finally, one cannot explain lexical exceptions or overgeneralizations within a theory which might claim that the acquisition of phonology is purely a matter of overcoming output constraints, as I might have tempted you to think in Section III.C.3, "Output Constraints and Conspiracies". Such a theory would be subject to exactly the same inadequacies as Jakobson's in these cases – for example, it could not deal with the existence of lexical exceptions to rules.

Summarizing, if we want a functional, explanatory theory of the acquisition of phonology – a theory that does more than say "children have rules, but the rules sometimes have exceptions", we need a theory which is more complicated.

B. Phonological Idioms

One thing that we have just seen is that articulatory success on particular sound patterns sometimes cannot be extended to new instances of very similar patterns. The ability to say "down" and "stone" without nasal harmony apparently was not generalizable to "dance" (let alone to "prune" or to "jump").

The most spectacular cases of non-generalizable articulatory accomplishments were analysed by Moskowitz (1970b); she aptly named them "progressive) phonological idioms". By this she meant words which are pronounced quite well, sometimes perfectly, and, crucially, much better than words of similar adult sound pattern. These are, in short, words which are exceptions to the child's modification rules and/or output constraints. The classic example is Hildegard Leopold's "pretty". She produced this word quite accurately as one of her first words at about 9 months of age. However, then and for many months thereafter, she produced no other consonant clusters and only one other word violating consonant harmony, "tick-tock". Finally, at a point after she had learned to break the consonant harmony constraint in general, "pretty" was changed to roughly [bidi], thus becoming part of the system in effect at that time. (See also Moskowitz, 1970a.)

A good many of the children studied have a few progressive phonological idioms among their early words. These phenomena as well as the less spectacular lexical exceptions discussed in the preceding section are clearly material which must be explained. Note that such lumpy pattern-and-exception landscapes are characteristic of the most closely related psychological areas that we know of: adult language is full of idioms, and cognitive development is full of instances in which the mastery of special cases long precedes the mastery of general skills. It seems that child phonology is more complicated than was once thought, but it still appears to be no more complex than adult syntax or cognitive development. (This rather silly-sounding remark is provoked by those who complain that if one introduces all these complexities, there is no elegant theory left any more. I believe it is one of the corollaries to Murphy's law, however, that nothing is as simple as it originally appears to be.)

C. Canonical forms

Ingram (1974) and Waterson (1971, 1972) have both shown that a young child's output forms can be sorted into sets of canonical forms (Ingram) or

prosodies (Waterson). Prosody is here used in the Firthian sense of a sequence of several archiphonemes (partially specified phonemes), and is exactly equivalent to the notion of canonical form. The members of such a set of forms have some strong syllable-structure restrictions in common: a set will be, say, just CV words, or just CVC and VC words, etc. What makes them interesting, indeed surprising, is that these sets are also restricted as to what phones can appear in them.

For example, taking Waterson's data, one set consists of forms for "fly", "barrow", and "flower"; these are all realized by forms consisting of an open syllable with voiced, continuant, labial onset: [wæ], [bβæ]. Another set consisted of "fish", "fetch", "vest", "brush", and "dish"; these were rendered as [(C)Vf], with the vowel always mid-high as it is in the targets. A third set was made up of CVCV forms in which the C's were stops and the second syllable was an exact reduplication of the first; the targets mapped into this canonical form included "Bobby", "biscuit", "kitty". Another set, which allowed the vowels to differ, was of the form [ɲVɲV], used for "Randall", "finger", "window", and "another".

Such sets may be maintained by any of the strategies that we have discussed: by selection of adult words that "fit" a form, by use of a rules, or by template-matching. We can thus see in phonological development a gradual weakening of restrictions on the co-occurrence of phones and the realization of more combinations of syllable structure with phonetic content, until we can no longer sort the child's output into these neat sets. In this progression, phonological idioms represent the most primitive level in the sense that they are the forms with the tightest relationship between phonetic content and syllable shape. A little set of lexical exceptions to a rule like "down" and "stone" represent a slight weakening in that relationship—they were produced, remember, as [dæon] and [don], two forms differing only in the vowel.

D. Motor Programming: A Psycholinguistic Account of Output Constraints and Canonical Forms

In the preceding section, we implied an interpretation of early output constraints and their gradual relaxation: it is as though the beginning speaker cannot vary some feature values in the course of a single word even though he can make the different sounds in separate words. To take a familiar example, a child with a consonant harmony constraint may be able to make consonants at two or more positions of articulation, e.g. be able to say "toy" and "boy", yet be able to say only [bʌb] for "tub". As Waterson says (1972, p. 13), there is "difficulty in the planning and production of rapid changes of articulation in a short space of time". There is a sense in which the whole word, for a child such as this, can be thought of as bearing a single specification for place of articulation. (This idea has antecedents in several theories of vowel harmony in adult language, e.g. Weilmers and Harris

1942, Waterson 1956/1970.) For a child like Waterson's P, an output word must conform to one of the given canonical forms, and within that restriction, only few degrees of freedom are left for the individual word.

We can tie all of these phenomena together and understand how they fit into an acquisition process if we make an analogy with computer programming. Suppose that learning to pronounce a sequence of sounds is like creating a program that the articulators and the larynx execute. A phonological idiom would then be like an invariant program, one which has no variable parameters that the user is free to set. A canonical form would be like a program in which some parameters are fixed, but others are settable. Let us consider some examples using this metaphor. Assume a child has CV(CV) as the canonical form subsuming, say, "bye-bye" and "baby" as [baba], "ball" as [bɔ], "doggie" as [dʌdʌ], and "there" as [de]. In this hypothetical case, the "program" can either stop after one CV cycle or produce a second CV. The only stops are [b] and [d], which means that there are two choices for consonant position: labial or dental; this choice is made once for the whole word. It also means that there is no choice for voicing or nasality within this program (which means that the canonical form should in fact have been written out as C[+voice, -nasal] V (CV)). Note that there is considerable freedom of choice for the vowels, but that the vowel is also specified once for the whole word.

This child might also have another canonical form, say (C)Vf, like Waterson's child. This form is like a program that allows some leeway for specifications of the initial consonant and the vowel, but always finishes the word with an [f]. Such forms have always been puzzling before—it is easy to imagine why assimilated forms are simpler than non-assimilated forms, but what good are canonical forms like CVf?

If the "programming" metaphor is roughly accurate, we now have an answer to that question. Even though a form like CVf requires a change in the articulatory position for the production of "bush" or "fish", it has very few variable parameters. Therefore, once it has been learned, it can be highly automatic to "run". The program is called up, the initial consonant is chosen, the vowel is chosen, and it runs with no more attention than would have been necessary to produce an open syllable. Waterson (1972, p. 17) noted: "each word appeared to be learned as an individual item . . . at first there were only one or two examples of a particular pattern and then there would be quite a sudden increase."

So now, we can describe phonological development by saying that the child gradually learns to improve in three areas of production control: (1) she learns to increase the number of parameters that can be freely assigned values in a given word; the consequence of this is that more of the segments in a word can vary (2) she learns to increase the number of values that each parameter can take on; this increase means that there is a wider range of possible phones that can be put into each segmental position in a word (3) she learns to link up short programs to make longer ones which can generate polysyllabic words.

In summary, the patterns of language behaviour that we have surveyed suggest that the child must initially discover (by trial and error) how to make sequences of sounds, not merely how to make segments in isolation. Some of these sequences she learns to vary systematically in one or two respects; these we see as groups of similar words, that is, as sets of words belonging to canonical forms. Other sequences she does not learn the trick of varying for a long time, possibly because they were among the most complicated to begin with; these remain phonological idioms. Some canonical forms run into developmental dead ends: Daniel learned only to vary the vowel in his [dVn] canonical form, producing only "down" and "stone" ([non]) with it. But apparently he could not go on from there to learn to vary the place of articulation of the consonants; he had to abandon his temporary conquest of nasal disharmony and make a fresh start.

E. The Articulatory Program and the General Model

1. *The Output Lexicon*

We have described many typical rules of child phonology, we have considered what might be difficult about certain sound sequences that children seem to avoid producing, and we have seen that many rules may be explained as devices which children invent to get around those difficulties. We have found rules that get rid of consonant clusters, of consonant disharmony, and of particular sounds in particular environments. We have also seen that there are some rules and looser strategies that cannot be explained in terms of articulatory simplification, at least not in the usual sense. Instead, we have had to invoke the idea that getting a word out involves the assembly of some sort of articulatory program.

Let us now go back to another aspect of psycholinguistic modelling. There is another important property of children's output that we have mentioned but not really discussed: the fact that some rule changes are carried out gradually. Sometimes this can be explained, following Macken (as we did earlier), by postulating that the child has misheard some word to begin with or has replaced an originally correct encoding of the word by an erroneous version based on his own output. In either case, the result can be that when a new rule comes in which should apply to the word, it will fail to do so because the word has the wrong stored form. Recall that in Macken's example, taken from N. V. Smith, the child had apparently stored "take" as [geik], because when all other velar-final words had broken free of the consonant-harmony rule, that word remained harmonized as an exception.

But often enough, there is quite a delay in applying a new rule to a word that is already established in the output vocabulary, and this can happen even when it is quite unlikely that there has been any miscoding of the word. For example, we mentioned that it took about two weeks for the nasal assimilation rule to begin to affect Daniel's "down" and "stone", and several more before the new forms replaced them entirely. What accounts

for the persistence of these forms? The most straightforward account, I think, is given by the two-lexicon model. What we can say with this model is that ways-to-say words are stored, too, in an output lexicon; application of a new rule to a word that is already in a child's active vocabulary involves the ouster of the old form which was stored in the output lexicon and its replacement by the new form. In this model, rules are the links from the input lexicon to the output lexicon. To show this our original figure is relabelled in Fig. 2.




FIG. 2

Lags in the adoption of a rule, in this model, simply are cases in which a child has formed the habit of saying a certain word a certain way and maintains that habit instead of "updating" it.

Now we need to fit the notion of articulatory programming, which we developed in the previous section, into the two-lexicon model. This proves to be very easy to do. What we did in that section was to factor the stored information about how to pronounce a word into two parts: (1) information as to which canonical form it belongs to, and (2) information on how the variable parameters in that canonical form should be chosen in order to produce the word. For example, suppose that the child has an accurate rendition of "dish" as part of a C[+voice]V[-tense][f] canonical form. We view its entry in the output lexicon as consisting of the information that (1) it belongs to the canonical form just mentioned and (2) the variable consonant parameters should be set at [+dental, -continuant], giving [d] since the voicing parameter has been fixed at [+voice]; meanwhile the variable vowel parameters should be set at [+front, +high], giving [ɪ] since there is already a fixed vowel parameter of [-tense].

The actual production of a word that belongs to a canonical form thus takes place in two stages. The first is recall of the canonical form and the stored variable-parameter values from the output lexicon, and the second stage is plugging the values into the articulatory program specified by the canonical form. Figure 3 shows this elaboration of the two-lexicon model.

Phonological idioms remain as output lexical entries that cannot be factored – that is, as entries in which there are no variable parameters to be set. This means that in our model, the output lexicon contains only the specification of the program; when it is called up, there is no plugging in of settings to be done – the articulatory program (alias the canonical form) has been stored fully specified.




FIG. 3

2. Rules in the Two-lexicon Model

We have occasionally used the cover term "transduction" to mean all the steps from hearing to speaking a word. As we have analysed this process in terms of perception, storage, and production, we have steadily been breaking it down into finer steps. We have said that one of those steps is the connection between the input lexicon and the output lexicon, and that step is mediated by rules. But we have really only talked about rules in the usual informal mode of relating the adult model word to the child's output word. We need to go back and see what we can deduce about the nature of the rules that would fit into our model.

These rules must account for the difference between what the child knows about the sound of a word as stored in the input lexicon and what is stored as canonical form membership plus variable parameter settings in the output lexicon. In the immature speaker, there is generally a loss of information at this step – that is, kids do not make in production all the distinctions that they can make in perception. The major function of the rules, then, is the selection of which pieces of information about the adult word will be preserved in the output lexicon and which will be abandoned; for this reason, we will refer to the rules in our model which link the input lexicon with the output lexicon as "selection rules".

Let us first consider how selection rules should look for a child who has developed beyond the stage of having obvious canonical forms. For such a child we gain very little by introducing the theoretical complexity of the factored output lexicon, and we make our work easier if we go back to the older model in which the output lexical entry for a word contains all the information needed to say it (see Fig. 2 again).

The notion of selection rule is especially convenient in discussing different children's treatment of consonant clusters, so we will use that topic as an example. The commonest pattern of initial cluster reduction for children acquiring English seems to be the one used in baby talk: stops and nasals are retained, liquids and fricatives lost; /sl/ and /sw/ clusters seem to be indeterminate. (Incidentally, parents tend to perceive their children as adhering to this stereotypic pattern even when the child actually uses a different one; see Menn, 1977; Menn and Berko Gleason, forthcoming.)

Some children find ways of breaking clusters apart, inserting [ə] or moving one of the segments to another part of the word (e.g. saying [nos] for "snow" (Hamp, 1974; also Waterson, 1971). In this discussion, however, we need to focus on those children who do reduce an adult initial consonant cluster to one segment, but who do not do it just by omitting one of the segments. For /sp, st, sk/ we can find some children who use the roughly corresponding fricatives [ʃ or f, s, x or s] to represent the cluster (also [fw] for /skw/); for /sm, sn, sl/ some children use the devoiced counterparts [m̥, n̥, l̥]. It is easy to describe what is happening here: the child is mapping the cluster into one segment by selecting some of the features belonging to the first adult segment and some to the second one. This is usually done with considerable regularity; that is, a given child will preserve either the fricative character or the stop character of all s+stop clusters. (The treatment of s+nasal clusters may differ from the treatment of s+stop clusters, however.)

Selection rules which produce effects such as these can be considered as selecting features from a particular portion of a word in the input lexicon and then putting them in a designated slot in the output lexical entry. Here, certain position and manner features from initial consonant clusters are taken and "put into a slot" so that they will designate the initial consonant of the output word.

Now let us consider briefly the character of the selection rules that would have to be written to characterize the behaviour of a child who is still operating with strict canonical forms. These rules must map the input lexical entry onto the two-part output lexical entry which we constructed in Section V.D. Therefore, they must be able to take each word in the input lexicon and specify both the canonical form (which articulatory routine will be used) and how any variable parameters are to be set.

A great deal of the variation from one child to the next is reflected right here. Take the word "snow"; for some young children, this will be treated as a CV word and most likely be produced as [no], [do] or [no]. Other children may put any target word containing a sibilant into a (C)V's class, and so produce "snow" as [nos] or [dos]. A child who tries to break up the cluster with an inserted vowel, giving [səno'] would probably have a CVCV canonical form to map it onto (but this raises problems of stress, which is clearly fixed for some polysyllabic canonical forms). A syllabic [s] for the first syllable is another possibility.

It is by no means clear how a child goes about picking what canonical form to assign an adult word to. She may be quite systematic about it – say assigning all two-syllable words with initial consonants to CVCV and all fricative-final monosyllables to CVs. But her assignments may seem rather more haphazard, especially for words which could fit equally well into either of two forms and for words that do not fit well into any form.

When we consider children whose transduction patterns are less regular and more like template matching, it is no longer possible to write selection rules; we must be content with guidelines. Note, however, that it is possible for there to be a fairly reliable rule for the choice of canonical form coupled

with some roughness in the way that variable parameter values appear to be selected (Macken's Si; Priestly's child – both discussed in Section III.C.4). I do not know if any case has been analysed as having irregularity in the choice of canonical forms coupled with regular rules for setting the variable parameters once the form has been chosen.

3. *Caution: the limitations of the programming metaphor*

We set up all this apparatus because it does a nice job of rationalizing the transduction patterns that we seem to find, although there are some data that do not fit as easily as one would like. This model is valid only to the extent that producing a word is like running off some fairly simple sort of speech synthesis program. I enjoin the reader to consider how the theory presented in this chapter might be modified so that it simulates the behaviour of real children better than it presently can.

VI. Saying What One Hears: Task Variables

A. Imitation, Self-monitoring, and Spontaneous Speech

We have been using the term "transduction" occasionally as a cover term for the whole process of hearing and then saying a word (regardless of the time delay between those events). One of the major phenomena of child phonology is the great variability that can be found in the accuracy of a single child's transductions, and the apparent relation of that accuracy to the conditions under which the word was produced. There are three reasons why we must be able to deal with this variability: first, obviously, since it exists we must be able to incorporate it in our theories; second, we must take it into account in data collection so as to get a proper sample of a child's performance, and third, in the assessment of phonological development for clinical purposes, we face the same sampling problem as in research data collection but with much greater urgency because of the need for efficient use of time and because of the consequences for the child.

In this section, we shall review some of the factors that are believed to be involved in the observed variations. It is well-known that imitated productions of words may be much better than spontaneous ones; it is also known that they may be just the same (Korte and Bond, 1979) or simply different (Moskowitz, 1975); and under some conditions, imitations can be worse than spontaneous productions. This means that the factor of being imitated cannot be the only one which produces variations in accuracy of transduction; other factors must be interacting with it if the relation of spontaneous to imitated production is unstable. We shall see that one of these factors is whether the target is already in the child's output vocabulary. However, when this is taken into account there is still a large residue of variation which does not seem to be a matter of the choice of test words and tasks at all. Perhaps it is truly random, but there is some evidence to suggest that another possible factor is the child's own moment-to-moment appraisal of what task she is really being asked to do.

Let us first consider why imitation, so often used as a research or assessment tool, is expected to improve a child's performance, and then why it may fail to do so. Recall that "spontaneous" is actually used to describe utterances elicited from the child by any (humane) means, as long as no-one says the target word itself within several minutes prior to the child's attempt at saying it. The intended essence of this distinction between "spontaneous" and imitated speech is that spontaneous utterances require retrieval of some encoding of the sound pattern of a word from long-term memory, while imitation is supposed to rely on short-term auditory memory. Thus imitation should be able to reflect the child's perceptual and articulatory capacity unencumbered by incorrect stored information.

But careful consideration of this supposition and of the data that we actually have about perception and about imitated production shows that it is, in general, false, especially for very young children. One always relies to some extent on old knowledge in both perception and production, and no imitation task can be assumed to escape this reliance. Perhaps it is minimized when the subject succeeds in categorizing and imitating novel sounds, as in Kent's task of imitating foreign vowel sounds (1978). But in general, imitation does not mean listening and reproducing without the interference of old habits; and if imitation relies entirely on old knowledge, as it may when the child is asked to repeat a familiar word, then imitated and spontaneous tokens of a word should be identical.

Now let us see how imitated tokens might be worse than spontaneous tokens. Barton's extensive work (1976, 1980), shows that children aged two years or under have a strong tendency to mis-hear unfamiliar words, re-interpreting them as familiar words which are phonetically similar. Barton attempted to do minimal-pair word-discrimination tasks with very young subjects, and often they had to be taught one of the words – for example, a twenty-month-old might know "coat" but not "goat". Such a child could learn to choose pictures correctly when one picture was a goat and the other was a bull, but given the minimal pair "coat – goat" to discriminate, the child tended to pick "coat" regardless of whether he heard "coat" or "goat". (The bias depended only on this familiarity factor, not on any phonemic factors.)

The implication of Barton's comprehension study for our consideration of imitation should be clear: he was running into a perceptual bias, and the same bias should be present in imitation tasks, unless they use words which are firmly in the child's passive vocabulary. Even tasks using all nonsense syllables may be affected, since any of them may be misperceived as a familiar word. (It is worth remembering that this bias for hearing the novel as the familiar remains throughout life, as anyone who has an uncommon name that resembles a common one can testify.)

Another variable which is involved in transduction is self-monitoring. Conscious monitoring is likely to improve the quality of one's output, and in adults such self-monitoring seems to be maximal when other cognitive loads are reduced; we assume that the same is true for children. Waterson (1978) has certainly shown that for one child's spontaneous speech, phonetic

quality declined as the length of the phrase produced increased; this is suggestive, but the "cognitive load" variation producing the phonological variation was a very highly linguistic one, namely the length of the utterance the word was embedded in, so one must be cautious about generalizing from it.

At any rate, children do indeed go off and quietly practice new sounds (Weir, 1962, is the classic reference). Some children are also observed to whisper new words and sounds (Leopold, 1939). As an aside, we might reflect that the observations of children practising, whispering, and showing off new sounds make the problem-solving theory of child phonology more credible (though they in no way make the old embryonic development theory less credible). It would be difficult to make sense of the claim that the acquisition of the ability to pronounce is a matter of problem-solving, if children never acted as though they were trying to solve a problem.

However, a large part of self-monitoring must also take place below the level of consciousness, for the amount of feedback that must be involved in the achievement of the fine control of the native speaker's accent is immense.

Returning to the main topic, now we must ask whether self-monitoring is improved during imitation. It could be; during imitation, a child has an opportunity to compare her short-term phonetic memory of a word with her production of it and/or with her long-term stored memory of how it should sound. Sometimes such comparisons are made, sometimes they obviously are not. There are plenty of recorded instances in which a child imitates an adult, produces a deviant pronunciation ([ote] for "okay", [fis] for "fish"), shows no signs of dismay, and denies hotly that what she has said differs from what the adult has said. Why the failure to spot such major discrepancies? In a good many of these cases, we are sure that it is cannot be ascribed to perceptual difficulty, for if the adult goes on to imitate the child's faulty pronunciation, the child also hotly rejects the adult's parody. The child can tell the difference. True, sometimes the adult produces a very crude parody; in such cases one can argue that the child would have accepted an accurate one. But in many cases, this dodge is not available; the child would have been capable of distinguishing her version from the correct version, if only she were paying attention. An unpublished observation of Daniel will illustrate this point. During the period in which all initial dental and labial stops were assimilated to following velar stops, Daniel was requested to get his toy duck; the toy was out of sight, and the adult simultaneously pointed in the direction in which it lay. Daniel echoed the word as [dʌk], went and looked in the indicated place, said [gʌk], and toddled away with no further interest. I suggest the following interpretation for this sequence of events: in the absence of non-verbal supporting context, Daniel failed to comprehend the word "duck". He repeated it correctly under this condition, clearly demonstrating that he had no perceptual problem with the adult word whatever. Then he went and found it, and repeated his familiar form for it.

It must be the child's self-monitoring that is at fault when he is at this level

of discriminatory ability and still fails to recognize that there is a difference between his output and the adult's. I suggest that there are two factors contributing to this lack of attention. One has been mentioned above; if the child has somehow arrived at the opinion of his word is adequate, he is likely to believe that he always does it right. (See Zwicky, 1982 on classical malapropisms.)

The other factor is a problem which also besets many Piagetian-style interviews of children: the problem of making sure that the adult and the child are actually directing their attention to the same phenomenon. Suppose that we consider an adult correcting a child's production of "fish" (to take a frequently-used example). The child says [fis], the adult says "No, say [fɪʃ:]", and the child indignantly responds "But I did say [fɪs:].". The adult wants the child to attend to the pronunciation, but how is this desire communicated to the child? Language is usually used, not contemplated; children expect to listen for meaning, not for sound. The child is more often disposed to understand the request to "say fish" as "say the sound pattern that designates the object with fins and scales that swims" than as "pronounce the word 'fish' accurately".

At the beginning of this section, we said that these transduction variables are of theoretical, methodological, and practical importance. By now these claims are obvious. First, as for theory, what we have seen is that the variables of attention and task orientation need to be incorporated into any model of child phonology. We have developed the outlines of a model of child phonology without taking account of these variables. We will not discuss in any detail how it can be modified to allow for them, but for example some more boxes and arrows need to be introduced into Fig. 3 to represent the following statement: spontaneous productions come from storage in the output lexicon, but imitated ones, to the extent that they are better than spontaneous ones, bypass the output lexicon and draw on less-automatic production mechanisms (Menn, 1979). Second, if we want to assess a child's ability to pronounce, whether our goals are research or remediation, we want to know whether the child is using her "best" pronunciation or an old familiar one. Recognizing that we have no control over this variable, we need to think of tasks that would make attention to pronunciation instead of meaning more or less likely.

Some speech pathologists and researchers attempt to test a child's best articulatory capacities by asking her to imitate nonsense words. This is intended to reduce the child's reliance on her habitual ways-of-saying known words; it is a very reasonable procedure, but we have seen that a child may assimilate a nonsense word to a known word in perception.

Her target word would then be different from the one that the examiner said, and thus there would be two sources of error: the misperception of the target and the effects, if any, of drawing on established articulatory habits. A very useful discussion of these and other task variables in assessment is Menyuk (1980).

B. The Word as Means to an End

We have introduced the problem of task variables in the context of observing and testing the child who has begun to talk. We should also consider the role of task variables during the transition from babbling to speech.

One of the perennial puzzles of child phonology is the phenomenon that got exaggerated into the legend of the universal "silent period": frequently a child will be able to produce a sound in babbling even though she cannot put it into any words. How can this come about? Like any other voluntary motor performance, the production of a sound or sound-sequence is easier in some contexts of action than others. Consider producing a given sequence of sounds under each of the following conditions:

- (i) having just made the sound(s) by accident (the context for "circular" babble);
- (ii) having just heard someone else make the sound(s) (the context for imitation);
- (iii) having decided to make the sound for its own sake or to execute the motor sequence that will produce the sound(s) (the context for sound-play);
- (iv) having decided to obtain a goal which requires the use of the sound(s) as a subgoal (the context for meaningful utterances).

Observation tells us that (i) is the earliest, and therefore the easiest of these four conditions, while (iv) is the hardest. It is not clear whether (ii) is easier than (iii), however. But the important point is that (iv) requires the ability to carry out (ii) or (iii) plus attention to the goal of the act of speaking. We might hypothesize that the means-ends gap found here is the reason why sounds can appear in babble before being used in speech, drawing on general principles of cognitive development.

But there may be some other factors involved in this delay. For example, a child might fail to realize that a sound made in play is just the one needed for certain words. This might happen because the recall memory for the sounds in those words is not strong enough to bring them to mind without supportive contexts, even though the child can recognize them when others say them. Second-language learners will certainly recognize this kind of recognition/recall disparity.

In conclusion, we cannot say with certainty why a child is unable to use in words a sound that he can produce in play, but there are many possible cognitive reasons why this might happen, so there is no point to invoking some mystery of the "language faculty" until it is shown that none of these possible reasons is plausible.

Note: It is important to make one's analogies carefully when comparing language with other cognitive abilities. There is possible confusion about my use of the terms "means" and "end". Children can indeed learn to produce words for various social and personal ends well before they show the

innovative means-ends behaviour that is called for on Piagetian developmental scales. But the kind of means-ends behaviour required for the onset of meaningful speech is of the most primitive variety; early words are acquired by plenty of practice and are deployed in familiar situations for familiar purposes.)

VII. The Acquisition of Allophones and Allomorphs

So far, we have concentrated on the development of the child's ability to go from a shallow phonemic input representation of the adult's word to some tolerable output approximation of it. But this, of course, is only the surface of the acquisition of phonology. How do children begin to dig below the phonetic surface?

This is a major topic, and in this section we will only discuss some theoretical issues and cite some of the recent studies in this area. To begin with, there are terminological problems that I would like to avoid, so I will specify the terms I will use in this section. A morphophonological or morphophonemic rule is one which requires morphological information for its operation, e.g. a rule which applies to verb stems, to plural morphemes, to members of a declensional class. An allophonic rule is one which requires only phonological information: the identity of neighbouring sounds, boundaries, assigned stresses, etc. (Boundary and trace markers are essentially devices for recasting morphophonological rules as allophonic rules.)

A productive rule is one which would apply to new words coming into a language and which can therefore be tested on nonsense words of properly chosen shapes. The effects of non-productive allophonic rules may persist for a long time in redundancy rules, which specify possible output shapes of morphemes without giving directions as to how aberrant morphemes are to be rearranged.

A rule of any type, morphophonological or allophonic, productive or not, may produce allomorphy: the appearance of a given morpheme in two or more shapes that would be written distinctly in phonetic transcription. (Examples will be supplied in the text as needed, rather than being given here.)

The distinction between superficial and cognitive aspects of acquisition has been kept in clear focus – in fact, has been the focus of debate – in studies of the acquisition of morphophonology. Berko Gleason's "wug test" (Here is a wug. Now there are two of them; there are two . . .) (Berko, 1958/1971) contrasted the child's ability to produce forms which might have been memorized (one glass, two glasses) with non-word forms which could not have been heard before (tass, tasses; gutch, gutches). Here, the pattern of /-s, -z, -əz/ allomorphy is productive in the adult language, and the test distinguishes between the child who can produce the correct allomorphs only on familiar words and the more advanced child who can supply them for novel words and therefore must know the underlying pattern.

When a pattern does not reach productivity in the adult language, as is the case with many of the alternations in the late-acquired "learned" vocabulary in English, it is more difficult to assess the degree to which a speaker has acquired a pattern rather than a list of surface forms. As McCawley (ms) has pointed out, when a pattern is non-productive, it is probably not necessary to go beyond memorization of a short list of words to be a competent user of the language.

However, some techniques show that a degree of awareness of such patterns does develop in many speakers. It should be noted that the cognitive demands of the acquisition of the common non-productive rules of English (tri-syllabic laxing, various stress-shift rules, velar softening) are no greater than the demands of the acquisition of the complex productive morphophonemics of German or Russian. (Review of the acquisition of complex morphophonologies is beyond the scope of this chapter; the reader should consult MacWhinney 1978.)

Several techniques have been developed for studying knowledge of non-productive morphophonemic rules. There are the memory-reversion technique of Myerson (1975), the meaning-guessing technique of Wilbur and Menn, and the concept-forming technique most recently used by Jaeger. The Wilbur and Menn (1974, 1975) technique is the simplest: here subjects were given pseudo-words created from Latin or English morphemes according to regular non-productive patterns, and asked to pick among three possible meanings – for example, for "chibble" the choices were (a) light rain, (b) a kind of smooth cloth, (c) coarse sawdust. Responses of experimental subjects showed that attenuated sound-meaning correspondences were indeed available to the subjects for most of the obsolete allomorphic patterns: for "chibble", 65% of the subjects chose "coarse sawdust", 22% chose "light rain", and only 12% chose "smooth cloth"; for the test word "abducive", 72% chose "distracting", 10% "conserving", and 18% "informing".

But as Linell (1979) correctly warns, one cannot infer awareness of particular rules (e.g. the rules postulated by SPE) just by showing awareness of the allomorphy that those rules describe. Much more work is required in this area, and Jaeger's, which is too complex to discuss here, is a good start.

So far we have been discussing allomorphic relations that clearly go across phoneme boundaries: equivalences of /s,z, əz/; or of /p/ and /b/ (chip, chibble). How do we study the acquisition of strictly allophonic rules, that is, rules which have purely phonological conditioning? Some of these also go across phoneme boundaries (i.e. produce neutralization) and some do not. For example, final devoicing of consonants in a language with a voiceless distinction produces neutralization (e.g. "Hund", /hʊnt/; "Hunde" /hʊnd/, where the underlying /d/ in the singular cannot be distinguished from an underlying /t/ unless one looks at the plural or another inflected form). On the other hand, the lengthening of vowels before voiced segments in English does not cause neutralization – there is no problem reconstructing an underlying segment different from the surface form.

Allophonic rules are easy enough to study if their context can be manipulated – if the segment can be made to appear both in the conditioning context and out of it. Vogel (1975) studied nasal assimilation to following stops in Spanish and Drachman and Malikouti-Drachman (1973) studied the same phenomenon in Greek. The overall impression from such studies is that stages of acquisition can be understood only from the perspective that the child is trying to work his way "back" from the surface to account for the patterns he observes; intermediate stages of rule acquisition need not look like simplified versions of the rules written for adult phonology.

These two nasal assimilation rules, incidentally, both function in two kinds of contexts: across morpheme boundaries, where they are easy to study by manipulation of context, and within morphemes, where their productivity is much harder to demonstrate. The techniques for appropriate tasks now exist, however, such as repetition of synthetic stimuli which violate the rules, and we can expect considerable progress on this front. In the meantime, some studies are available on the achievement of adult-like control of the surface manifestations of these rules: for example, Hawkins (1973, 1979a) on the acquisition of proper stop duration within consonant clusters in English, and Naeser (1970) on the duration of children's vowels before voiced and unvoiced stops in English.

VIII. Summary and Conclusion

It would be pleasant to say: these are the facts about the acquisition of phonology. However, we must hedge, this being a human science, and say instead: these are the major conclusions about the early stages of the acquisition of phonology that appear to be justified at the present time.

(1) Some children take a very "holistic" approach to the acquisition of phonology; their speech is so hard to transcribe and describe that we can say little about them in existing theoretical frameworks (Peters, 1977). Even the more analytical children sometimes resort to holistic approaches to varying degrees.

(2) The child's early acquisition of phonology has two aspects: the acquisition of phonetic control and the acquisition of phonemic contrast. Later, the same dichotomy extends to the acquisition of the surface forms of words v. the acquisition of the patterns that they are instances of.

(3) For most of the children whose approach we currently can handle, we find a rough division into an early period of very slow growth of the output lexicon, and then a period of more rapid growth. However, some children never show such a marked point of acceleration.

(4) Most of the words of the early period will have alternating vowels and consonants. Some words will probably be very well controlled and be more complex in structure (progressive phonological idioms); others may be extremely vague and variable in their output token forms.

(5) During the early period and for a while thereafter, most words will fall into groups. The words in such a group will be similar in both syllable structure and phonetic content; they will be describable as instances of a canonical form.

(6) The acquisition of phonemic oppositions can be studied only in terms of syllable structures: the typical picture is for a child to have a particular feature contrast in one position (initial, intervocalic, preconsonantal, final . . .) well before it appears in others. Within a given position, phonemic contrast may be evidenced indirectly before the child achieves good control of the pair of adult phonetic features involved in the contrast, but on the other hand there may be good phonetic control of one value of a feature without the presence of a contrasting phone.

(7) The mismatches between adult model and child word are the results of the child's trial and error attempts; they are shaped by the child's articulatory and auditory endowments (and thus to that extent are "natural") and by the child's previous successes at sound production. All rules of child phonology are learned in the sense that the child must discover for herself each correspondence between the sounds that she hears and what she does with her vocal tract in an attempt to produce those sounds.

(8) Knowledge gained by articulatory success on a particular sound or sound-pattern does not always generalize to cases which we phonologists feel to be similar: a feature or a phone or a string mastered may remain an isolated success for a long time.

(9) Regular mapping patterns (rules) grow, generalize, and often over-generalize, even to the point of diminishing the child's accuracy of production of some words.

(10) Whole-word mapping strategies are used to varying degrees and are a major type of irregular mapping. Even in later stages of acquisition, such strategies can be found on the more difficult polysyllables.

(11) Instead of modifying adult words which are not within her capacity to produce accurately, a child may use selection strategies, avoiding problematic sounds and sound sequences and/or exploiting favourites.

(12) As implied by the phrasing of all these statements, individual variation among children is considerable. A deterministic theory would therefore have to be so weak as to be meaningless. Yet typical patterns emerge. The prevailing theories allow for individual variation by considering the child to be experimenting with solutions to the problem of how to say words. As we look across children, trying to discover what tends to be earlier, and therefore presumably easier, and what tends to be later, and therefore presumably harder, we find three articulatory sources of difficulty for the young speaker: the articulation of certain phones (e.g. [ɪ, θ]), the sequencing of dissimilar consonantal targets, and departures from CVCV . . . alternations. Difficulties also arise from perceptual sources, including a tendency to perceive unfamiliar forms as similar familiar ones, and (probably) an inability to take in all the information about a relatively long word until its most salient sounds have already been well-learned.

Finally, unexpected hindrances and aids may arise from the child's current array of strategies: a sequence which "should" be easy may be difficult for a particular child because it does not fit into the rules or prosodic strategies that she happens to have developed up to that time.

(13) This chapter presents the view that the child's mastery of production mechanisms can be described as learning to (a) control the accuracy of articulatory movements, (b) specify more contrasting articulatory targets in a given sequence position, (c) produce more different sequence types, and (d) concatenate sequence types.

Let us conclude by considering the assertion made in the introduction to this chapter: that evidence from the studies which have become available in the last ten or fifteen years has forced a change in our basic conception of the nature of phonological development. We can no longer sustain the developing-embryo model; we need problem-solving models to make sense of peculiarly skewed output distributions such as we find in children who "avoid" or "exploit" heavily. Just as in the acquisition of morphology and syntax, what has been called the "implicit defining question" of our research has changed. We used to ask: What linguistic theory will explain the order in which the various language behaviours develop? This question assumed that there is such an order, and that it should be explainable by linguistic theory. The new question is roughly: What behavioural predispositions and abilities does the child bring to the task of learning to communicate with language, and how does the individual go about solving the articulatory and phonological problems posed by the language to be learned?

The presuppositions of the second question differ markedly from the first. We now presuppose that there are a variety of "predispositions and abilities" of memory, motor control, perception, etc. – including perhaps some "purely linguistic" predispositions which might have evolved just for handling the special rapid information processing and complex pattern learning involved in the acquisition and use of language. We also presume that the notion of problem solving is the best heuristic for explaining the kind of very rough consensus of developmental order that we find in the data. As for the old assumption that linguistic theory can explain what we find in acquisition, we have seen that the more likely scenario is that linguistic theory and acquisition data will have to come to terms with one another. A theory based only on the performance of the mature skilled user cannot anticipate the temporary learning devices and detours of the unskilled learner.

Note: In this chapter, the feminine form has often been used for the indefinite pronoun. The reader may not realize it, but my female colleagues and I are still receiving professional form correspondence – for example, reprint requests – that address us as "Dear Sir / Sehr geehrter Herr / Cher monsieur". At least until I have evidence that more scientists in this field can conceive of their fellows in two sexes as well as in three languages, I think it well to jog their sense of "markedness" a bit.

Acknowledgements

My grateful thanks to Sarah Hawkins, Paula Menyuk, and Ronnie Wilbur, who spent considerable time and effort working over the first draft of this chapter. At the urging of Charles A. Ferguson, and with the help of Prof. Hawkins, I have attempted to overcome old habits and use IPA consistently throughout it.

References

- Barton, D. P. (1976). The role of perception in the acquisition of speech. Doctoral dissertation, University of London. Circulated by Indiana University Linguistics Club.
- Barton, D. P. (1980). Phonemic perception in children. In "Child Phonology: Perception and Production" (G. Yeni-Komshian, J. F. Kavanagh and C. A. Ferguson, eds) Vol. 2. Academic Press, New York and London.
- Bell, A. (1971). Some patterns of occurrence and formation of syllable structures. In "Working Papers on Linguistic Universals", Linguistics Department, Stanford University, 6, 23-137.
- Berko, J. (1958). The child's learning of English morphology. *Word* 14, 150-157. Reprinted in A. Bar-Adon and W. Leopold, (eds) (1971). "Child Language: A Book of Readings". pp. 153-167.
- Branigan, G. (1979). Sequences of words as structured units. Doctoral dissertation, Boston University School of Education.
- Clumeck, H. (1977). Studies in the acquisition of Mandarin phonology. Doctoral dissertation, University of California at Berkeley.
- Drachman, G. and Malikouti-Drachman, A. (1973). Studies in the acquisition of Greek as a native language. In "Ohio State University Working Papers in Linguistics" 15, 99-114.
- Farwell, C. B. (1976). Some strategies in the early production of fricatives. In "Papers and Reports in Child Language Development" (No. 12). Stanford University Linguistics Department.
- Ferguson, C. A., Weeks, T. and Peizer, D. B. (1973). Model-and-replica grammar of a child's first words. *Lingua* 31 (1), 35-65.
- Ferguson, C. A. and Farwell, C. B. (1975). Words and sounds in early language acquisition. *Language* 51, 491-439.
- Ferguson, C. A. and Macken, M. A. (To appear). Phonological development in children's play and cognition. In "Children's Language", (Keith E. Nelson ed.), Vol. IV. Gardner Press, New York.
- Fey, M. and Gandour, J. (1979). Problem-solving in early phonology acquisition. Paper read at the Annual Meeting of the Linguistic Society of America, Los Angeles, 1979.
- Fey, M. and Gandour J. (1982). Rule discovery in early phonology acquisition. *Journal of Child Language* 9, 71-82.
- Fliege, J. E. and Massey, K. P. (1980). English prevoicing: random or controlled? Paper read at the Summer Meeting of the Linguistic Society of America, 1980.
- Goldstein, U. (1980). An articulatory model for the vocal tracts of growing children. Doctoral dissertation, Electrical Engineering Department, MIT.
- Halliday, M. A. K. (1975). "Learning How to Mean: Explorations in the Development of Language". Edward Arnold, London.
- Hamp, E. H. (1974). Wortphonologie. *Journal of Child Language* 11, 287-288.
- Hawkins, S. (1973). Temporal coordination of consonants in the speech of children: preliminary data. *Journal of Phonetics*, 1, 181-217.
- Hawkins, S. (1979a). Temporal coordination of consonants in speech of children: further data. *Journal of Phonetics*, 7, 235-267.
- Hawkins, S. (1979). The control of timing in children's speech. *Proceedings of the Ninth International Congress of Phonetic Sciences, Copenhagen, 1979*.
- Ingram, D. (1974). Phonological rules in young children. *Journal of Child Language* 1, 49-64.
- Ingram, D. (1976). "Phonological Disabilities in Children". Elsevier, New York.
- Ingram, D. (1979). Phonological patterns in the speech of young children. In "Language Acquisition" (P. Fletcher and M. Garman, eds). Cambridge University Press, Cambridge.
- Itkonen, T. (1977). Notes on the acquisition of phonology. English summary of Huomioita lapsen äänteistön kehityksestä. *Virittäjä*, 279-308.
- Jaeger, J. J. (1980). Categorization in phonology: an experimental approach. Doctoral dissertation, University of California/Berkeley.
- Jakobson, R. (1968). "Child Language, Aphasia, and Phonological Universals" (Trans. A. Keiler). Mouton, The Hague.
- Jones, L. G. (1967). English phonotactic structure and first-language acquisition. *Lingua* 19, 1-59.
- Kent, R. D. (1978). Imitation of synthesized vowels by preschool children. *Journal of the Acoustic Society of America*, 63, 1193-1198.
- Kisseberth, C. W. (1970). On the functional unity of phonological rules. *Linguistic Inquiry* 1, 291-306.
- Korte, S. S., and Bond, Z. S. (1979). Children's spontaneous and imitative speech: An acoustic analysis. Paper read at meeting of American Speech and Hearing Society, November, 1979.
- Leonard, L., Schwartz, R., Folger, M. K. and Wilcox, M. J. (1978). Some aspects of children phonology in imitative and spontaneous speech. *Journal of Child Language* 5 (3), 403-416.
- Leopold, W. F. (1939-1949). "Speech Development of a Bilingual Child". Vols I-IV. Northwestern University Press, Evanston.
- Linell, P. Psychological reality and the concept of phonological rule. *Proceedings of the Ninth International Congress of Phonetic Sciences, Copenhagen, 1979*.
- Macken, M. A. and Barton, D. (1980). The acquisition of the voicing contrast in English: a study of voice onset time in word-initial stop consonants. *Journal of Child Language* 7 41-75.
- Macken, M. A. The child's lexical representation: evidence from the 'puzzle-puddle-pickle' phenomenon. Stanford University Papers and Reports in Child Language Development (No. 16), 1979; also to appear in *Journal of Linguistics*.
- Macken, M. A. (1979). Developmental reorganization of phonology: a hierarchy of basic units of acquisition. *Lingua* 49 11-49.
- MacWhinney, B. (1978). The acquisition of Morphophonology. *Monographs of the Society for Research in Child Development* 43, 1-2.
- Menn, L. (1971). Phonotactic rules in beginning speech. *Lingua* 26, 225-241.
- Menn, L. (1976a). Pattern, control, and contrast in beginning speech: a case study in the acquisition of word form and function. Doctoral Dissertation, University of Illinois. Circulated by Indiana University Linguistics Club.

- Menn, L. (1976b). Semantics of intonation contour in late babble and beginning speech (English). Paper read at the Summer Meeting, Linguistic Society of America, 1976.
- Menn, L. (1977). Parental awareness of child phonology. Paper read at the Annual Meeting of the Linguistic Society of America, 1977.
- Menn, L. Transition and variation in child phonology: modelling a developing system. *Proceedings of the Ninth International Congress of Phonetic Sciences, Copenhagen, 1979*.
- Menn, L. (1981). Review of S. P. Blache, The acquisition of distinctive features. *Language*, 57, 953-958.
- Menn, L. and Haselkorn, S. (1977). Now you see it, now you don't: tracing the development of communicative consciousness. In "Proceedings of the Seventh Annual Meeting, NorthEast Linguistic Society", (Judy Kegl, ed.) 1977.
- Menn, L. and Berko Gleason, J. (to appear). Babytalk as folk phonology: Bias in parents' reports of children's pronunciations. (forthcoming.)
- Menyuk, P. (1977). "Language and Maturation". Cambridge, MIT Press, Cambridge, Massachusetts.
- Menyuk, P. (1980). The role of context in misarticulations. In "Child Phonology", (Yeni-Komshian, G. Kavanagh, J. and Ferguson, C. A., eds) Vol. I. Academic Press, New York and London.
- Moskowitz, A. (1970a). The two-year-old stage in the acquisition of English phonology. *Language* 46, 426-441.
- Moskowitz, A. (1970b). The acquisition of phonology. Working paper (No. 34), Language-behavior Research Laboratory, University of California, Berkeley.
- Moskowitz, A. (1975). The acquisition of phonetics: a study in phonetics and phonology. *Journal of Phonetics* 3, 141-150.
- Myerson, R. (1975). A developmental study of children's knowledge of complex derived words of English. Doctoral dissertation, Harvard Graduate School of Education.
- Naeser, M. A. (1970). The American child's acquisition of differential vowel duration. Technical report No. 144 (in two parts), Wisconsin Research and Development center for Cognitive Learning, University of Wisconsin, Madison.
- Nakazima, S. (1972). A comparative study of the speech development of Japanese and American children, part IV. *Studia Phonologica* VI, 1-37.
- Peters, A. M. (1977). Language learning strategies. *Language* 53, 560-573.
- Platt, C. and MacWhinney, B. (1983). Solving a problem vs. remembering a solution: error assimilation as a strategy in language acquisition. *Journal of Child Language*, 7, 41-75.
- Priestly, T. M. S. (1977). One idiosyncratic strategy in the acquisition of phonology. *Journal of Child Language* 4, 45-66.
- Sander, E. K. (1972). When are speech sounds learned? *Journal of Speech and Hearing Disorders* 37, 55-63.
- Schwartz, R. G. and Leonard, L. B. (1982). Do children pick and choose? An examination of phonological selection and avoidance in early lexical acquisition. *Journal of Child Language*, 9, 319-336.
- Slobin, D. I. (1966). Comments on "Developmental Psycholinguistics". In "The Genesis of Language" (F. Smith and G. A. Miller, eds.), MIT Press, Cambridge, Massachusetts.
- Slobin, D. I. (1973). Cognitive prerequisites for the development of grammar. In "Studies of Child Language Development" (C. A. Ferguson and D. I. Slobin, eds), Holt, Rinehart and Winston, New York.

- Smith, B. L. (1979). A phonetic analysis of consonantal devoicing in children's speech. *Journal of Child Language* 6, 19-28.
- Smith, N. V. (1973). "The Acquisition of Phonology: A Case Study". Cambridge University Press, Cambridge.
- Smith, N. V. (1978). Lexical acquisition and the acquisition of phonology. Summer Forum Lecture, Linguistic Institute of the Linguistic Society of America.
- Snow, C. (1977). The development of conversation between mothers and babies. *Journal of Child Language* 4, 1-22.
- Stampe, D. (1969). The acquisition of phonemic representation. *Proceedings of the Fifth Regional Meeting of the Chicago Linguistic Society*, 433-444.
- Sterne, D., Jaffe, T., Beebe, B. and Bennett, S. L. (1975). Vocalizing in unison and in alternation: two modes of communication in the mother-infant dyad. In "Annals of the New York Academy of Sciences." Vol. 263: Developmental Psycholinguistics and Communication Disorders. (Aaronson, D. and Rieber, R. W. eds).
- Stevens, K. N. (1972). The quantal nature of speech: evidence from articulatory-acoustic data. In "Human Communication, A Unified View". (Denes, P. B. and David, E. E. eds) McGraw-Hill, New York. pp. 51-66.
- Velten, H. V. (1941). The growth of phonemic and lexical pattern in the infant. *Language* 19, 440-444. Reprinted in A. Bar-Adon and W. Leopold (eds.) (1971). "Readings in Child Language". Prentice-Hall, Englewood Cliffs, New Jersey.
- Vihman, M. M. (1976). From prespeech to speech: On early phonology. "Papers and Reports on Child Language Development", (No. 12), Stanford University Linguistics Department.
- Vihman, M. M. (1981). Phonology and the development of the lexicon: evidence from children's errors. *Journal of Child Language* 8, 239-264.
- Vihman, M. M. (to appear). Homonymy and the organization of early vocabulary. Child Phonology Project, Stanford University.
- Vihman, M. M. (1978). Consonant harmony - its scope and function in child language. In "Universals of Human Language" (J. H. Greenberg, ed.), Vol. III. Stanford University Press, Stanford.
- Vogel, I. (1975). Nasals and nasal assimilation patterns in the acquisition of Chicano Spanish. In "Papers and Reports on Child Language Development" (No. 10), Stanford University Linguistics Department.
- von Raffler-Engel, W. (1973). The development from sound to phoneme in child language. In Slobin, (eds.), "Studies of Child Language Development". (Ferguson, C. A. and Slobin, D. I., eds) pp. 9-12. Translated from *Proceedings of the Fifth International Congress of Phonetic Sciences, Munster, 1964*.
- Waterson, N. (1970). Some aspects of the phonology of the nominal forms of the Turkish word. In "Prosodic Analysis" (F. R. Palmer, ed.) pp. 174-187. Oxford University Press, London.
- Waterson, N. (1970). Some speech forms of an English child: a phonological study. *Transactions of the Philological Society*.
- Waterson, N. (1971). Child phonology: a prosodic view. *Journal of Linguistics* 7, 179-221.
- Waterson, N. (1972). Perception and production in the acquisition of language. *Proceedings of the International Symposium on First Language Acquisition, Florence, 1972*.
- Waterson, N. (1978). Growth of complexity in phonological development. In "The Development of Communication". (Waterson N. and Snow, C. E., eds) Wiley, New York.

- Weir, R. (1962). "Language in the Crib" Mouton, The Hague.
- Wellmers, W. E. and Harris, Z. S. (1942). The phonemes of Fanti. *JAOS* 64 (No. 4), 318-333.
- Westbury, J. R. and Keating, P. A. (1980). A model of stop consonant voicing and a theory of markedness. Paper read at the Annual Meeting of the Linguistic Society of America, 1980.
- Wilbur, R. B. and Menn, L. (1975). Psychological reality, linguistic theory, and the internal structure of the lexicon. San Jose State University Occasional Papers in Linguistics.
- Wilbur, R. B. (1980). Theoretical phonology and child phonology: argumentation and implications. In "Phonology in the 1980's". (Goyvaerts, D. ed.) Story-Scientia, Ghent.
- Yeni-Komshian, G., Kavanagh, J. and Ferguson, C. A. (eds) (1980). "Child Phonology: Perception and Production". Academic Press, New York and London.
- Zwicky, A. M. (1982). Classical malapropisms and the creation of a mental lexicon. In "Exceptional Language and Linguistics"; (L. Menn and L. K. Obler, eds), pp. 115-132. Academic Press, New York and London.