

Speech Errors, Phonotactic Constraints, and Implicit Learning: A Study of the Role of Experience in Language Production

Gary S. Dell and Kristopher D. Reed
University of Illinois at Urbana–Champaign

David R. Adams
Washington University

Antje S. Meyer
Max Planck Institute for Psycholinguistics

Speech errors follow the phonotactics of the language being spoken. For example, in English, if [n] is mispronounced as [ŋ], the [ŋ] will always appear in a syllable coda. The authors created an analogue to this phenomenon by having participants recite lists of consonant–vowel–consonant syllables in 4 sessions on different days. In the first 2 experiments, some consonants were always onsets, some were always codas, and some could be both. In a third experiment, the set of possible onsets and codas depended on vowel identity. In all 3 studies, the production errors that occurred respected the “phonotactics” of the experiment. The results illustrate the implicit learning of the sequential constraints present in the stimuli and show that the language production system adapts to recent experience.

We know that “king” is a word of English and that, as far as we can tell, “hing” is not. However, most people would not be greatly surprised to hear that “hing” is a word that they just do not know. This is because [hŋ] is well formed. Each of its phonemes occur in English, and their ordering is consistent with English *phonotactics*, the constraints that define the language’s sound sequences.

Just as it is apparent that “hing” is a possible word, it is even more apparent that [ŋh] is not. It violates phonotactic constraints that [h] must occur in *onset* position (i.e., it must begin a syllable) and that [ŋ] must occur postvocally in a syllable, in the *coda* position. Although syllables in other languages may begin with [ŋ] (e.g., Burmese) or end with [h] (e.g., Hebrew), no English syllables do.

Knowledge of a language’s phonotactics does more than underlie intuitions about possible lexical items. It plays a major role in language acquisition and processing. Adults use their knowledge of possible sound sequences to aid in the identification of speech

sounds (e.g., Massaro & Cohen, 1983; Pitt, 1998) and word boundaries (e.g., McQueen, 1998; Norris, McQueen, Cutler, & Butterfield, 1997). Furthermore, nonwords with common sound sequences can be repeated more quickly than those with low probability phonotactics (Vitevitch, Luce, Charles-Luce, & Kemmerer, 1997). Sensitivity to these sequential constraints begins very early in life. Nine-month-old infants will listen longer to word lists that adhere to the phonotactics of their native language than to lists that do not (Jusczyk, Friederici, Wessels, Svenkerud, & Jusczyk, 1993), and they are sensitive to how phonotactic sequences align with word boundaries (Mattys, Jusczyk, Luce, & Morgan, 1999). Moreover, infants can learn the sequential structure of a continuous stream of speech sounds from as little as 2 min of exposure (Aslin, Saffran, & Newport, 1998; Saffran, Aslin, & Newport, 1996).

One important manifestation of phonotactics can be seen in speech production errors. Phonological speech errors often create nonwords such as “hymn to hing” for “hymn to sing” (from Dell, 1990), but the chance of such an error violating phonotactic constraints has been thought to be small (e.g., Boomer & Laver, 1968; Fromkin, 1971; Wells, 1951). Violations do occur, though. Stemberger (1983) collected several examples, such as [ætk] for “act” or [dlorm] for “dorm,” and these constituted a bit less than 1% of his phonological error corpus (see also Mowrey & MacKay, 1990, and Frisch & Wright, 1998, who provided electromyographic and acoustic evidence, respectively, that some errors elicited from tongue-twister recitation can be phonetically as well as phonotactically ill formed). Despite the existence of violations, it is generally agreed that there is a reasonably strong tendency for errors to obey phonotactic constraints. This tendency, the *phonotactic regularity effect*, suggests that the constraints are actively used during speaking (Fromkin, 1971; Motley & Baars, 1975).

This article presents data that illuminate the mechanisms behind the phonotactic regularity effect. In three experiments, we asked

Gary S. Dell and Kristopher D. Reed, Beckman Institute, University of Illinois at Urbana–Champaign; David R. Adams, Department of Psychology, Washington University; Antje S. Meyer, Max Planck Institute for Psycholinguistics, Nijmegen, the Netherlands.

Antje S. Meyer is now at School of Psychology, University of Birmingham, Birmingham, United Kingdom.

This research was supported by National Science Foundation Grants SBR 94-11627 and 98-73450 and National Institutes of Health Grants HD-21011 and DC-00191. We thank Kathryn Bock, Jennifer Cole, Rina Donchin, Cynthia Fisher, Stefan Frisch, Susan Garnsey, Prahlad Gupta, Harlan Harris, and Natasha Warner for their comments on this work and Linda May for work on the manuscript.

Correspondence concerning this article should be addressed to Gary S. Dell, Beckman Institute, University of Illinois at Urbana–Champaign, 405 North Mathews, Urbana, Illinois 61801. Electronic mail may be sent to gdell@s.psych.uiuc.edu.

speakers to recite sequences of four consonant–vowel–consonant (CVC) syllables, such as “hes feng neg kem.” Each experiment consisted of four sessions per speaker, with the result that each speaker said more than 6,000 syllables altogether. The sequences always included one [h] and one [ŋ], segments that are required to be onsets and codas, respectively, in English. Consequently, we expected that any speech errors in which another sound was replaced with [h] or [ŋ] would nonetheless be phonotactically regular. That is, error outcomes that are [h] would be onsets, and [ŋ] outcomes would be codas. The critical aspect of the experiments, however, concerned errors involving other consonants. For example, in Experiment 1, half of the speakers experienced [f] always as an onset and [s] always as a coda, and half experienced the reverse assignment. More generally, each experiment created experiment-wide phonotactic constraints in regard to the distribution of consonants within the context of the experiment. To anticipate our results, we found that errors adhered to the experiment-wide constraints to a surprisingly strong extent. We show later that this fact, along with other aspects of the data, is informative about the mechanisms responsible for the phonotactic regularity effect and the role of experience in language production.

Two hypotheses motivated our experiments: the *breadth of constraint hypothesis* and the *implicit learning hypothesis*. The former is a claim about the relation between language-wide phonotactic constraints and other less general sound patterns, and the latter concerns the mechanisms for acquisition and representation of these patterns. The next two subsections introduce these hypotheses to provide background for the experiments.

Breadth of Constraint Hypothesis

The breadth of constraint hypothesis is the claim that patterns in language occur at many levels of generality and that the processing system is sensitive to all of these levels. For example, one can speak of constraints that are true for the entire language. The fact that [ŋ] is a coda whenever it occurs in English is an example. There are also patterns that apply to a subset of words. A classic example is the differing stress patterns of English nouns, which tend to have trochaic stress as in *REcord*, and verbs, which tend to be iambic as in *reCORD*. Ultimately, one can speak of constraints that are so specific that they apply to a single word. Thus, one can say that [k] is an onset, in the context of the word *king*.

The claim that language processing is sensitive to patterns at many levels of generality is hardly controversial. Consequently, our goal is not to prove or disprove the breadth of constraint hypothesis. Instead, we seek to reinterpret some well-known speech-error phenomena in light of this hypothesis and relate them to new findings. In so doing, we can identify relations among these phenomena that, to our knowledge, have not been brought out before.

One can characterize a particular speech-error phenomenon in terms of very local constraints. Most phonological speech errors involve the movement of sounds from one syllable to another, such as “leading list” for “reading list” (Fromkin, 1971). These movement errors exhibit what is called the *syllable-position effect* (Boomer & Laver, 1968; Fromkin, 1971; MacKay, 1970; Nootboom, 1969). When a sound moves, it retains its position in the syllable structure. For consonants, this effect reduces to the claim that onset consonants move to other onset positions, and codas

move to other coda positions. In the “leading list” example, the erroneous [l] in “leading” is an onset, as is its source in “list.” The syllable-position effect, like the phonotactic regularity effect, is subject to violations (MacKay, 1970; Stemmer, 1982). Moreover, there is debate about how much of the effect reflects syllable position as opposed to word position (Meyer, 1992; Shattuck-Hufnagel, 1983). Nonetheless, there appears to be some kind of tendency for consonants to retain their positions when they slip. Dell (1986) offered an account of the syllable-position effect by associating words with position-specific phoneme nodes. For example, the [l] in *list* corresponds to an onset-*l* node. When *list* is active, the onset-*l* is activated and thus can appear erroneously in other onset positions in the vicinity of *list*, such as “leading list.”

We suggest that the syllable-position effect and the phonotactic regularity effect are related phenomena. They simply reflect constraints at different levels of generality. Consider the fact that [h] is always an onset when one is speaking English. This is an example of a language-wide constraint, and errors tend to adhere to it. Next consider the fact that [k] is always an onset when one is saying “king.” This is a local constraint, one that is true only for “king” and other words with [k] onsets. An error in the vicinity of “king” involving the migration of [k] will tend to preserve [k]’s onset status (e.g., “red king” spoken as “ked king”). In this way, the syllable position effect can be thought of as adherence to a local constraint in the vicinity of a controlling word. Adherence to language-wide phonotactic constraints is just a more general application of the principle: Misplaced speech sounds tend to reflect the positions in which they correctly occur.

The syllable-position and phonotactic regularity effects can thus be thought of as two ends of a continuum of breadth of constraint, with the former at the narrow end and the latter at the wide end. The distinction between a language’s rules and its exceptions is analogous. A rule applies widely and exceptions locally. The past tense for English is *-ed*, but the past tense for *come* is *came*. Moreover, the regularities in the English past tense fall on a continuum between the wide rule and the local exceptions, including subregularities such as *ring-rang*, *sing-sang*, and *drink-drank*.

In this vein, our experiments constitute an attempt to find a speech error effect that occupies a middle ground between the wide phonotactic regularity effect and the narrow syllable-position effect. The experiment-wide phonotactic constraints (e.g., [f] being an onset throughout this experiment) represent a domain that is broader than a single syllable and its immediate vicinity but considerably more restricted than a lifetime of speaking English. We attempted to separate the contributions of language-wide, experiment-wide, and local constraints on error and, hence, evaluate the breadth of constraint hypothesis.

Implicit Learning Hypothesis

How does the language processing system acquire sensitivity to constraints on the positions of sounds in words and syllables? The simplest answer is that the processing system learns. It experiences sound sequences and stores them in memory. Then it uses those memories in subsequent processing of sound sequences. This is not to deny that an important component of phonotactic knowledge may be independent of experience, that is, inherent in the linguistic, cognitive, articulatory, and auditory systems. But the variabil-

ity of phonotactics across languages and the often arbitrary nature of the phonological forms of lexical items force the conclusion that much of the knowledge of sound patterns comes from experience.

Our experiments examined the learning mechanism that is responsible for the production system's sensitivity to the positions of speech sounds. Specifically, we hypothesized that this mechanism has three properties: It is sensitive to recent experience, it is implicit, and it is capable of generalization. We refer to these properties jointly as the implicit learning hypothesis, in deference to the large literature on the learning of sequential patterns in nonlinguistic domains that associates the properties with an "implicit" or "procedural" learning system (e.g., Cleeremans, 1993; Gupta & Cohen, 1998; Knowlton & Squire, 1994; Reber, 1993; Redington & Chater, 1996; for a recent review, see Neal & Hesketh, 1997). How sensitive the learning mechanism is to recent experience was determined in our experiments by examining whether speakers acquire any special sensitivity to experiment-wide constraints in their errors and, if so, how quickly it is acquired. Implicitness implies that the speakers have no intention to learn the distribution of the speech sounds in the syllables that they are uttering. They are just producing syllables as accurately as they can, and any learning of the sound distribution occurs simply as a result of speaking. Furthermore, their error patterns should be unaffected by explicit awareness of these distributions, something that we manipulated in our experiments.

The issue of generalization concerns the specificity of what is learned and how that knowledge is applied. At one end of the specificity dimension are instance theories (e.g., Goldinger, 1998; Hintzman, 1986; Logan, 1988). Each experience leaves a separate memory trace. For the acquisition of sound distributions within a syllable, an instance would correspond to the utterance of a particular syllable. If a recency mechanism is assumed (e.g., less recent instances are lost or less accessible), an instance theory could explain the sensitivity of errors to experiment-wide constraints. The more accessible instances that there are of a syllable, the more likely that syllable will be produced instead of an intended syllable. Because the recent syllables themselves follow the constraints, the errors will as well. Somewhat less specific than a pure instance theory would be a syllable lexicon (e.g., Levelt & Wheeldon, 1994), a store of syllable types that have been produced before. Augmented with a recency mechanism, a syllable lexicon can explain any tendency for errors to follow experiment-wide constraints for the same reason that instance theories can. Any bias to produce recently spoken syllables necessarily is associated with a sensitivity to sound distributions present in those syllables.

Opposed to instance theories are accounts in which rules are inferred from experience and stored in memory (see Marcus, Vijayan, Rao, & Vishton, 1999). Such rules would, for our experiments, refer to units smaller than the syllable. For example, speakers may induce the experiment-wide rules *[s] is an onset* and *[f] is a coda*, and their acquisition would offer an account of any tendency for errors to respect them. These kinds of rules may refer to more than just the set of possible onsets and codas. Because languages of the world tend to have many restrictions on which codas can accompany which vowels, phonotactic constraints can be usefully described by listing the set of possible onsets and the set of possible *rimes*, or nucleus-coda combinations. Thus, instead of acquiring the rule *[f] is a coda*, our speakers might acquire rules that *[ɛf]*, *[ɪf]*, and so on are legal rimes. Evidence that

speakers are sensitive to the onset-rime distinction comes from speech errors (MacKay, 1972; Shattuck-Hufnagel, 1983; Stemberger, 1983), short-term memory errors (Treiman & Danis, 1988), and a variety of tasks requiring manipulations of pronunciations (Fowler, Treiman, & Gross, 1993; Treiman, 1983, 1986). Accounts of learning based on rule acquisition can also be combined with instance theory, allowing for the storage of both abstractions and specific experiences or items. Such dual-mechanism theories have been proposed to explain categorization data (e.g., Erickson & Kruschke, 1998; Nosofsky, Palmeri, & McKinley, 1994) and aspects of morphology (Pinker, 1991).

A third class of theories, connectionist learning models, can be distinguished from rule and instance theories and even from dual-mechanism theories. Pronunciations of words are learned by the incremental tuning of weights between representational units in a network, in response to experience (e.g., Anderson, Milostan, & Cottrell, 1998; Cleeremans, 1993; Dell, Juliano, & Govindjee, 1993; Elman et al., 1996; Gupta & Dell, 1999; Jordan, 1986; Plaut & Kello, 1999; Plaut, McClelland, Seidenberg, & Patterson, 1996; Seidenberg & McClelland, 1989). Connectionist learning models are associated with flexibility in the specificity of what is learned. Some of the weight changes in the network can be characterized as the induction of "rules" at various levels of generality, whereas other changes may be akin to the storage of a specific instance (McClelland & Rumelhart, 1985). The exact nature of what is learned, however, depends on the training experience, any a priori representations, and the learning algorithm. With respect to our experiments, one would expect connectionist learning models to be able to explain the acquisition of experiment-wide constraints on sound distribution through weight change mechanisms. And speech errors in these kinds of models seem to naturally follow whatever phonotactic constraints are present in the training set provided that output representations are based on well-motivated phonological features (Anderson et al., 1998; Dell et al., 1993).

Our experiments tested the breadth of constraint and implicit learning hypotheses by altering the distributions of speech sounds in a laboratory speech production task and evaluating the extent to which speech errors reflected those distributions. We compared error sensitivity to local constraints (e.g., *[k]* is an onset in this sequence of four syllables), experiment-wide constraints (e.g., *[f]* is always an onset in this experiment), and language-wide constraints (e.g., *[h]* is always an onset in English). Furthermore, we attempted to determine whether any learning of experiment-wide sound distributions is implicit using what we call the *ask-tell technique*. The ask component of this technique is the standard way of assessing participants' explicit knowledge; we asked them at the end of the experiment to report anything they noticed about the syllables that they pronounced. The tell component involves directly telling half of the participants about the sound distributions. If informed and uninformed speakers have similar error patterns, then we have evidence that explicit knowledge is not responsible for those patterns.

In the first experiment, speakers repeated four-syllable sequences such as "feng keg hem nes" in time with a metronome. They repeated 96 such sequences four times each in an experimental session, and they did so for four sessions on separate days. The eight consonants in this example and the vowel *[ɛ]* were the only sounds used in the experiment. Each consonant occurred exactly once in each sequence. *[h]* was always an onset and *[ŋ]*

was always a coda, respecting English phonotactics. The misplacement of these sounds therefore tested for language-wide constraints. The sounds [n], [m], [g], and [k] could occur either as onsets or codas in the experiment's sequences. But because each occurred only once per sequence, each was an onset or coda for that sequence. Misplacement of these sounds allowed for a test of the syllable-position effect, or what we have called local positional constraints. Finally, the sounds [f] and [s] were each confined to a syllable position throughout the four experimental sessions. Half of the speakers experienced only [f] as onset and [s] as coda, which we call the *fes* condition, and half experienced the reverse assignment, the *sef* condition. Errors involving the placement of [s] and [f] therefore tested for knowledge of experiment-wide constraints. The key comparisons concerned the strength of error adherence to the three kinds of constraints. If there is no learning of the experiment-wide constraints, one expects [f] and [s] errors to be "legal," that is, to retain their proper onset or coda status, at about the same rate that errors involving [n], [m], [k], and [g] retain their onset or coda status. In other words, if there is no sensitivity to the experiment-wide distributions, the only factor contributing to the legality of [f] and [s] errors should be the syllable-position effect. Experiment-wide positional constraints then reduce to local positional constraints. If experiment-wide constraints do exert an influence beyond that of local constraints, it will be of further interest to compare their strength with that of language-wide constraints involving [h] and [ŋ].

Experiment 1

Method

Participants. Eight students at University of Illinois at Urbana-Champaign were paid \$20 for participating in the four-session experiment. Each was a native speaker of English. Two participants were randomly assigned to each condition: informed *sef*, informed *fes*, uninformed *sef*, and uninformed *fes*. *Sef* versus *fes* refers to the experiment-wide distribution of [f] and [s]; the two conditions counterbalanced this assignment. Informed versus uninformed refers to whether participants were told about the distribution of these sounds in their materials.

Materials. For each participant, four sets of 96 sequences were generated, subject to the following constraints. Each sequence contained four CVC syllables, with the vowel [ε] for each syllable, and the consonants [h], [ŋ], [m], [n], [k], [g], [f], and [s] appeared exactly once per sequence. [h] was always an onset and [ŋ] always a coda. For speakers in the *fes* condition, [f] and [s] were always onsets and codas, respectively, and speakers in the *sef* condition experienced the reverse. The other consonants appeared equally as onsets and codas.

These constraints defined a vocabulary of 32 possible syllables in each condition. In each set of 96 sequences (with 384 [4 × 96] syllable tokens), the syllable types appeared with the expected frequencies given in the Appendix, which presents all of the types in the *fes* condition. The syllables differed in expected frequency, because each consonant had to occur exactly once in each sequence, and four of the consonants were restricted to be just onsets or just codas throughout the study ([h], [ŋ], [f], and [s]). Each consonant appeared exactly once per sequence to equate for error opportunities among the consonants and allow evaluation of any tendency for unrestricted sounds ([k], [g], [m], and [n]) to obey the syllable-position constraint on movement errors.

The 96 sequences associated with each set were printed, in 16-point bold Helvetica lowercase type, one sequence to a line and 11 lines to a page. We found that the sequences were easier to read when printed than when presented on a computer screen. Syllables were spelled as indicated in the

Appendix. The spelling was straightforward except that [g] onsets in [gem], [geŋ], and [gen] were spelled "gh" to remove the hard-soft "g" ambiguity. The order of the syllables within a sequence, and that of the sequences within a set, was random, with different randomizations used for each set and each participant.

Procedure. Each participant produced four sets of 96 sequences, each set on a separate day. There was, at most, 1 day intervening between sessions. A sequence set was produced as follows. The printed sequences were displayed to the participant via a cutout window so that only a single sequence was visible at a time. After each sequence was made visible, a Franz electric metronome (Model LM FB5, Franz Manufacturing, New Haven, CT) was started at a rate of 1 beat per second, and the participant recited the sequence once, coinciding each syllable with a beat. This slow recitation was designed to ensure proper pronunciation of each syllable and to familiarize the participant with the sequence; the very few errors that occurred were not counted. Then the metronome was set at 2.53 beats per second, and the participant recited the sequence three times at this rate without pause. The printed sequence remained visible throughout recitation. Each day's testing was tape recorded (Marantz cassette recorder PMD201, Marantz Company, Chatsworth, CA). Participants in the informed condition were told at the beginning of each experimental session that "when you see an 'f' (or 's'), it will always be at the beginning of a syllable, and when you see an 's' (or 'f'), it will always be at the end of a syllable." Those in the uninformed condition were told nothing about the distribution of sounds or letters. At the conclusion of the 4th day's testing, all participants were given a sheet of paper with the following instruction: "In the space below, please note any observations concerning the syllables that you were asked to pronounce in this experiment. Example: All of the syllables contained the vowel 'e.'"

Results and Discussion

Tape recordings were scrutinized for errors by Kristopher D. Reed. The reliability of error coding was good. A second coder, who was unaware of the experimental conditions, independently listened to the tape recordings from the 1st day of testing of 3 participants and identified errors. Both coders had access to the correct sequences. Of the 3,456 syllable productions examined, both coders agreed that there was no error on 3,216 syllables. Of the remaining 240 syllables, the coders agreed that there was an error and they agreed on the nature of the error for 191 syllables. For 16 syllables, the primary coder identified an error, and the secondary coder did not. The reverse occurred for 23 syllables. Finally, there were 10 cases in which both coders identified an error but disagreed as to what it was. Thus, the overall coding agreement rate was 98.6% that is, (3,216 + 191)/3,456. If one conditionalizes on syllables in which the primary coder identified an error, the rate was 88.0%. This level of agreement is acceptable, provided that differences of interest are not subtle ones. None of the original coding decisions were changed as a result of this reliability check.

Our focus was on errors in which one of the eight consonants is replaced by another consonant from the sequence. All such errors were tallied. This included cutoff errors such as "kes" → "m . . . kes." All of the tallied errors could be termed consonant movement errors or consonant misorderings because all eight sounds occurred in each sequence. Other error types, for example, vowel errors or substitutions of some other consonant, were rare and were not counted.

We identified 3,065 consonant misorderings. Because each participant produced 4,608 syllables at the faster speech rate (4 syllables × 3 repetitions × 96 sequences × 4 days), the overall

rate of these errors per syllable was 8.3%, that is, 3,065/(4,608 × 8 participants). Errors were more likely on the 1st day (1058) than on subsequent days (Day 2, 799; Day 3, 544; and Day 4, 664).

All misordering errors are presented in the four matrices making up Table 1. Each matrix entry gives the number of times that a particular syllable was produced as an error in a particular condition. For example, in the matrix labeled *fes* legal, the entry 37

Table 1
Error Outcomes in Experiment 1 Classified by Condition and Legality of the Error

fes condition—legal outcomes								
Onset								
Coda	g	k	m	n	f	s	h	ŋ
g	—	37	28	61	32	—	50	—
k	6	—	19	34	37	—	29	—
m	20	37	—	26	41	—	44	—
n	35	30	60	—	69	—	49	—
f	—	—	—	—	—	—	—	—
s	35	28	36	33	65	—	59	—
ŋ	69	36	15	21	66	—	69	—
h	—	—	—	—	—	—	—	—

fes condition—illegal outcomes								
Onset								
Coda	g	k	m	n	f	s*	h	ŋ**
g	17	10	9	23	7	0	21	0
k	4	5	0	10	6	1	5	0
m	4	3	1	5	24	0	9	0
n	9	10	15	24	35	0	27	0
f*	0	0	0	0	0	0	2	0
s	23	5	9	9	—	0	—	0
ŋ	9	4	1	1	—	0	—	0
h**	0	0	0	0	0	0	0	0

sef condition—legal outcomes								
Onset								
Coda	g	k	m	n	f	s	h	ŋ
g	—	24	6	21	—	16	21	—
k	8	—	12	20	—	38	18	—
m	23	16	—	18	—	27	21	—
n	24	27	35	—	—	33	17	—
f	43	21	12	20	—	76	54	—
s	—	—	—	—	—	—	—	—
ŋ	66	44	17	19	—	94	83	—
h	—	—	—	—	—	—	—	—

sef condition—illegal outcomes								
Onset								
Coda	g	k	m	n	f*	s	h	ŋ**
g	15	7	8	5	0	5	8	0
k	2	11	6	9	0	10	5	0
m	6	1	1	4	0	2	3	0
n	11	13	14	9	0	6	16	0
f	13	4	0	3	1	—	—	0
s*	2	0	1	1	0	2	1	0
ŋ	21	13	0	2	0	—	—	0
h**	0	0	0	0	0	0	0	0

* violates experiment-wide constraints.

** violates language-wide constraints.

associated with onset k and coda g means that 37 times a participant in the *fes* condition said "keg" instead of the correct syllable. Thus, the matrices are not confusion matrices (which consonant replaces which); they represent the form of the error outcome syllables. The legal-illegal dimension refers to whether or not the error involved movement that violated a positional constraint. Errors in legal matrices were events in which the spoken erroneous onset (or coda) occupied the same syllable position in the error as it did in the sequence. So, for example, if someone said "keg ken heng fes" instead of "meg ken heng fes," this would count as a legal instance of a "keg" error. The [k] was an onset in both the error and the intended sequence. Note that the s-onset column and the f-coda row are necessarily zero in the *fes* legal matrix, and the converse is true for the *sef* legal matrix. Error outcomes such as "geg," "kek," "nen," and "mem" also cannot occur in the legal matrices. Any of these outcomes would be illegal, because at least one of the consonants would have to violate a constraint.¹

Errors in the illegal matrices reflect events in which a consonant is produced in a different syllable position. For the sounds [k], [g], [m], and [n], an illegal error involves violation of a constraint that is true just for the current sequence, what we call local positional constraints. For [f] and [s], an illegal error violates a constraint that is true for the entire experiment, and, for [h] and [ŋ], the constraint applies to all of English. The labels on the rows and columns of the illegal matrices are augmented by a single asterisk to indicate that the errors in that row or column violate the experiment-wide constraint. Potential language-wide violations are indicated by two asterisks. If there is no asterisk, the row or column in the illegal matrix indicates only a violation of local positional constraints.

Consider the relative strength of these constraints, starting with the language-wide constraints. All 640 of the errors involving misplacement of [h] and [ŋ] obeyed the constraint. Thus, the experiment exhibited the phonotactic regularity effect. Although we are quite certain that this constraint was very strongly upheld in the data, we cannot be certain that there were no phonotactically illegal errors, because error coding can easily be biased against hearing these sounds in illegal positions (see Cutler, 1981; Frisch & Wright, 1998).

Next, consider the local positional constraints, those involving misplacement of [m], [n], [k], and [g]. Of the 1,941 errors in this group, 68.2% obeyed the constraint. That is, more than two thirds of the time, the erroneous sound occupied the same syllable position as it did in the target sequence. Each of the 8 participants produced more errors with these consonants that retained position than errors that changed position ($p < .01$; this p value and all subsequent ones were determined from sign tests). Across the participants, the range was 53.8% to 86.1% legal errors. This result shows that movements are truly constrained by position. The effect, however, is not a dominating one. In fact, our finding is somewhat weaker than what one would expect from reports of the syllable-position effect in the speech error literature. For example, in Stemberger's (1989) analysis of English between-words conso-

¹ Sometimes both the onset and coda of a string are in error. Although such a case counts as two errors, it is a single outcome in the matrices. The error outcome would be counted in the appropriate legal matrix only if both the onset and the coda were legal; otherwise, the outcome appears in an illegal matrix.

nant errors, 90.3% of the errors involved onsets interacting with onsets and codas interacting with codas, 2.5% were clear exceptions to the constraint, and the remainder involved one or more consonants that were ambisyllabic. Insofar as our results are not as strong as Stemberger's, they call into question accounts of the syllable-position effect that do not admit exceptions, such as Dell's (1986) use of different nodes for the onset and coda versions of a consonant. However, the reading task used here may have generated some uncertainty regarding intended syllables, and this could account for a weakening of the constraint. Another point worth noting is that the present experiment confounded word and syllable position, at least insofar as each syllable was construed as a word. Thus, it is uncertain whether the obtained effect was one of word position, syllable position, or both (see Meyer, 1992; Shattuck-Hufnagel, 1986).

The principal findings of this experiment concern the experiment-wide constraints. There were 484 errors in which [f] or [s] was misplaced. Of these errors, only 11 violated the constraint. As can be seen from the illegal outcome matrices in Table 1, there were 7 cases of an onset [s] moving to a coda slot (the [s] row in the *sef* condition), 1 case of coda [s] moving to an onset slot, 2 cases of an onset [f] moving to a coda slot, and 1 case of coda [f] moving to an onset slot. Thus, the experiment-wide constraints were upheld 97.7% of the time (range: 89.3% to 100%). For every participant, the experiment-wide constraint was stronger than the local positional constraint ($p < .01$), demonstrating that error adherence to the experiment-wide distribution was not simply due to the tendency for moving sounds to retain their syllable position.

As far as we can determine, the experiment-wide effect on errors was not associated with verbalizable knowledge of the sound distributions. Participants who were not informed about the distribution of [f] and [s] exhibited the constraint at the same rate (98.2%) as those who were informed (97.0%). Furthermore, none of the uninformed participants wrote down any observations at the end of the experiment about the distribution of [f] and [s]. Although post-experimental assessment of knowledge is an extremely weak test of implicitness (Dulany, Carlson, & Dewey, 1984), the lack of a difference between the errors of informed and uninformed speakers, together with the fact that the key data are errors—responses that are unintended—provides good evidence that the experiment-wide effect was unrelated to explicit knowledge of the sound distributions.

To a first approximation, the experiment-wide effect on errors resembles that of the language-wide effect. It is just not as strong. Thus, one can interpret our effect as a kind of experimentally induced phonotactic regularity effect, one that reflects a learning process that occurred during the experiment. The learning process appeared to be quite powerful—leading to 97.7% adherence to the constraint—and it also seemed to develop very quickly. On the 1st day, only 4 errors (of 184 f-s errors) violated the constraint. The 2nd, 3rd, and 4th days were associated with 3, 2, and 2 violations, respectively. The violations were scattered haphazardly within the 96 trials of each day: Trials 26, 31, 46, and 78 on Day 1; Trials 13, 23, and 31 on Day 2; Trials 83 and 93 on Day 3; and Trials 57 and 76 on Day 4.

Before consideration of possible mechanisms for this learning, it was important to replicate the experiment-wide effect. Something odd about the sounds [f] and [s] may make them less likely to violate positional constraints. To examine this possibility, we

carried out Experiment 2, which was exactly the same as the first experiment except that [k] and [g] functioned as [f] and [s] and vice versa. Thus, half of the speakers experienced [k] only as an onset and [g] only as a coda, the *keg* condition, and half experienced the reverse, the *gek* condition. Throughout the experiments [f] and [s], along with [m] and [n], occurred as both onsets and codas; [h] and [ŋ] were distributed as before.

Experiment 2

Method

The method was the same as that in Experiment 1, except that [k] and [g] were restricted to a particular position for a given speaker, whereas [f] and [s] were unrestricted. To ensure that there was no ambiguity about the pronunciation of the [g], all [g] onsets were spelled "gh." Eight new speakers from the same population as before participated and were assigned 2 each to the informed *keg*, informed *gek*, uninformed *keg*, and uninformed *gek* conditions.

Results and Discussion

Errors were tallied as in Experiment 1, leading to the identification of 3,584 consonant movement errors, a 9.7% error rate per syllable. The 1st day was associated with 1,111 errors, and subsequent days were associated with fewer errors: 960, 789, and 724 for Days 2, 3, and 4, respectively. Table 2 shows all of the legal and illegal error outcomes.

Experiment 2 replicated the first experiment in all important respects. For the language-wide constraints, all of the 1,016 [h] and [ŋ] movement errors were phonotactically legal, a finding identical to that of the earlier study. In contrast, local positional constraints, which pertained to the 1,850 errors in which [f], [s], [m], or [n] moved, were upheld only 77.5% of the time. This compares reasonably well with the 68.2% found in Experiment 1. As in the first experiment, all 8 speakers in Experiment 2 had more legal than illegal outcomes with [f], [s], [m], and [n] errors ($p < .01$), thus replicating the earlier experiment's finding of a significant local positional effect (range: 62.2% to 88.2%).

The principal finding from Experiment 2 was a strong experiment-wide constraint on errors. Of the 718 errors in which [k] or [g] was misplaced, 94.7% kept to their position (range: 90.1% to 100%). All 8 participants exhibited a higher legality percentage for the experiment-wide constraint than for the local positional constraint ($p < .01$), effectively showing a sensitivity to the experiment-wide distribution of [k] and [g] that could not be due to the syllable-position effect. Interestingly, there were somewhat more violations of the experiment-wide constraint involving [k] (30) than [g] (8). This could reflect any of a number of factors, ranging from phonetic differences to the fact that onset and coda [g]s (and not onset and coda [k]s) were orthographically distinct in the printed words. There had also been a hint of an asymmetry in Experiment 1, with more violations of the experiment-wide constraint involving [s] (8) than [f] (3). These differences in numbers of violations for particular consonants may reflect influences of language-wide knowledge. The sound most associated with experiment-wide violations across the two experiments, [k], is quite common as both an onset and a coda. For example, in Kessler and Treiman's (1997) analysis of English CVC words, [k] occurred as an onset in 142 words, as compared with 88, 126, and 92

Table 2
*Error Outcomes in Experiment 2 Classified by
 Condition and Legality*

gek condition—legal outcomes								
Onset								
Coda	f	m	n	s	g	k	h	η
f	—	8	28	22	31	—	58	—
m	8	—	38	17	23	—	69	—
n	30	42	—	25	37	—	45	—
s	24	12	11	—	42	—	25	—
g	—	—	—	—	—	—	—	—
k	24	36	43	36	25	—	97	—
η	55	30	21	43	129	—	90	—
h	—	—	—	—	—	—	—	—

gek condition—illegal outcomes								
Onset								
Coda	f	m	n	s	g	k*	h	η**
f	9	0	0	0	4	1	9	0
m	1	4	4	3	3	3	5	0
n	14	13	30	6	25	3	4	0
s	5	7	2	4	2	4	4	0
g*	0	0	3	2	0	0	0	0
k	6	0	2	0	—	5	—	0
η	4	0	1	1	—	3	—	0
h**	0	0	0	0	0	0	0	0

keg condition—legal outcomes								
Onset								
Coda	f	m	n	s	g	k	h	η
f	—	8	44	27	—	27	32	—
m	20	—	24	59	—	47	81	—
n	86	64	—	32	—	62	42	—
s	19	21	18	—	—	23	36	—
g	42	41	105	40	—	88	77	—
k	—	—	—	—	—	—	—	—
η	58	37	43	87	—	165	134	—
h	—	—	—	—	—	—	—	—

keg condition—illegal outcomes								
Onset								
Coda	f	m	n	s	g*	k	h	η**
f	7	2	3	5	0	2	7	0
m	7	0	11	1	0	8	9	0
n	2	35	13	11	0	13	14	0
s	5	4	10	2	1	7	5	0
g	1	2	7	2	2	—	—	0
k*	1	2	1	4	0	1	2	0
η	3	2	8	3	0	—	—	0
h**	0	0	0	0	0	0	0	0

* violates experiment-wide constraints.

** violates language-wide constraints.

words for [g], [s], and [f], respectively. For codas, [k] was even more dominant: 182 words, as compared with 67 for [g], 116 for [s], and 68 for [f]. We wish to stress, however, that the 30 violations with [k] in the second experiment still represent a small total in comparison with the number of violations of local positional constraints associated with [m] (60), [n] (130), [f] (65), and [s] (50). (Compare the k-coda row in the keg illegal matrix with the

[f], [m], [n], and [s] rows in the same matrix, and do the same for the k-onset column in the gek illegal matrix.)

As in the first experiment, there were more violations of the experiment-wide constraints on the 1st day (17) than on the 2nd (4), 3rd (10), or 4th day (7). In percentage terms, 7.0% of the relevant movement errors on the 1st day violated the experiment-wide constraint, as compared with 4.4% on subsequent days. The distribution of violations across trials within days was unsystematic, as in the previous experiment. Mean trial positions (out of 96) for the violations were 37, 36, 54, and 45 for Days 1, 2, 3, and 4, respectively. Finally, we note that the experiment-wide effect was present for participants who were informed of the sound distributions (94.0%) as well as those who were not informed (95.7%). None of the uninformed group members wrote down anything about the restricted distributions of [k] and [g] in response to the post-experimental query about the stimuli.

Although not as strong as Experiment 1's finding of 97.7% error adherence to experiment-wide constraints, the second experiment's 94.7% adherence adequately replicates the experiment-wide effect. The data further show that this effect was greater in magnitude than the 77.5% local positional (or syllable-position) effect and suggest that it was unrelated to explicit knowledge of the sound distributions. Finally, the similarity of the two experiments' results indicates some generality to the findings.

At this point, we have support for the claim that the participants have learned, that is, that they have been affected by their experience with the stimuli. This learning is revealed in the sensitivity of their errors to the experienced sound distributions. Their errors followed the phonotactics of the experiment, much as normal phonological speech errors follow the phonotactics of the language. However, the experiment-wide effect was not as strong as the language-wide effect and thus seemed to be genuinely intermediate between the local positional effect and the language-wide effect. In this way, the data support the claim that the processing system is sensitive to patterns at different levels of generality, what we have called the breadth of constraint hypothesis.

Before we turn to our final experiment, we need to revisit the theoretical issues raised in the introduction concerning the implicit learning hypothesis. What mechanisms could be responsible for the experiment-wide effect? As stated earlier, one possibility is instance theory. Each experience leaves a memory trace, and later performance depends on the retrieval of one of these traces (e.g., Logan, 1988). One might explain the experiment-wide effect by proposing that each utterance of a syllable constitutes an instance, and more recent instances are more accessible than less recent ones. All that is required, then, is an assumption that the more accessible an instance, the more likely it will occur instead of the correct syllable. The more accessible instances will naturally follow the experiment-wide constraints.

This account, which we call syllable-instance theory, makes a prediction that can be tested with the data of the first two experiments. The prediction is that syllables that did not occur in the experiment should be much less likely as error outcomes than those that did. For example, in Experiment 1, the syllables "kek," "geg," "mem," and "nen" never occurred in the stimuli. Yet, they are consistent with the set of legal onsets and codas in the study, because each of these consonants could appear as both onsets and codas. By consulting the illegal matrices of Table 1 (these errors necessarily are illegal in that a consonant appears in a syllable

position different from its origin in the sequence), one can see that these error outcomes did occur reasonably often. In fact, the average number of such outcomes per syllable type for “kek,” “geg,” “mem,” and “nen” was 20.7. This can be compared with the average number of illegal syllable outcomes involving these consonants for actual occurring syllables such as “kek,” “men,” and so on: 15.3 errors per type, less than that for nonoccurring syllables. The same analysis of Experiment 2’s data from Table 2 shows 17.2 illegal errors per syllable type for nonoccurring syllables (“fef,” “ses,” “mem,” and “nen”) versus 12.8 illegal errors per type for occurring syllables (e.g., “mef” and “nem”). In short, there is no evidence that error outcomes were restricted to occurring syllables. The reason that there were more errors creating nonoccurring syllables could reflect the fact that the nonoccurring outcomes had identical onsets and codas. However, the number of errors per type was highly variable across types, and so one should not put too much stock in the difference.² In any event, it is clear that errors often create nonoccurring syllables. Hence, the syllable-instance theory is not a likely explanation for the experiment-wide effect. Also ruled out are theories that associate the effect with a lexicon of stored syllable types as opposed to instances or tokens.

It is important to note that our ruling out of the syllable-instance theory is predicated on an assumption of indivisible syllable instances. If one allows instances to be blended or averaged, one can account for the generation of nonoccurring syllables. For example, the retrieval of a “gek” and a “keg” could, if the instances are properly structured, blend to a “geg” or a “kek” output. Such blended output-instance theories (e.g., Hintzman, 1986) can be distinguished from pure instance theories (e.g., Logan, 1988) and are difficult to distinguish from theories that allow for the storage of abstractions, such as rule-learning or connectionist theories.

The presence of a strong experiment-wide constraint on the location of erroneous consonants, coupled with the production of nonoccurring syllable outcomes, suggests that what is being learned involves smaller units than the syllable. The most intuitive hypothesis is that speakers were learning some kind of implicit generalization or rule about the set of possible onsets and possible codas in the study. So, in Experiment 1, the fes-condition participants learned that [f] was a possible onset. That is, they learned the set of possible onsets and the set of possible codas. This information could also be stated in a negative fashion ([s] *cannot be an onset*) or conditionally (*if the segment is [f], then it must occur as an onset*). These kinds of generalizations can be referred to as positional, in that they refer to the position of a consonant in the syllable but not to other characteristics of the syllable.

If the learned generalizations about consonants are solely positional, they would be indifferent to the identity of the vowel. We tested this in a third experiment. This experiment was similar to the first one in that the consonants [f] and [s] had restricted distributions and [m], [n], [k], and [g] did not. However, the experiment-wide constraint was a second-order one that depended on the identity of the vowel. Specifically, speakers were given two kinds of sequences: sequences with the vowel [æ], such as “fas hang kag nam,” and sequences with the vowel [ɪ], such as “sif hing kig nim.” Every participant received both [æ] and [ɪ] sequences. But all of the sequences for the participants in the fas-sif condition obeyed the following second-order constraints: *If the vowel is [æ], then [f] must be an onset and [s] must be a coda, and if the vowel is [ɪ], then [s] must be an onset and [f] must be a coda.* The other

condition, saf-fis, reversed the consequents of the two conditional statements. So, each speaker in both conditions experienced [f]s and [s]s as both onsets and codas. If all that is learned is positional—that is, which consonants can be onsets and codas—then movements of [f] and [s] should adhere to their syllable positions at the same rate as that for the other consonants. On the basis of the two previous experiments, we would expect this rate to be between 65% and 80% of the time. However, if speakers can learn the second-order constraints inherent in the stimuli, then we would expect [f] and [s] movements to maintain their syllable position at a higher rate. If the sensitivity to these constraints is as powerful as that of the purely positional experiment-wide constraints tested in the first two experiments, we would expect [f] and [s] movement errors to have 94%–98% adherence.

Experiment 3

Method

The only difference between this experiment and Experiment 1 concerned the makeup of the stimuli (and the information about the stimuli given to those participants in the “informed” conditions). There were 8 participants, as before, with 2 each assigned to the informed saf-fis, informed fas-sif, uninformed saf-fis, and uninformed fas-sif conditions. Each did 96 sequences a day for 4 days using the same procedure as the previous experiments. The sequences for each day for each participant in the fas-sif conditions were composed as follows. A 96-sequence set for a single day of the fes condition of Experiment 1 was generated. However, every other sequence in the set was transformed by replacing the occurrence of onset [f] with [s], and the occurrence of coda [s] with [f]. Then the vowels in the untransformed sequences were changed to [æ], spelled “a,” and the vowels in the transformed sequences were changed to [ɪ], spelled “i.” (The spelling of the “g”s from Experiment 1 was maintained.) The resulting list of 96 sequences alternated [æ] and [ɪ] sequences as in the following: fas nag hang mak ([æ]: f onset and s coda), nif sig kim hing ([ɪ]: s onset and f coda), mak hag nas fang ([æ]: f onset and s coda), and ghing sin mif hik ([ɪ]: s onset and f coda). Sequence sets for the saf-fis condition were generated in the same way but with the opposite assignment of vowels to sequences. These procedures resulted in stimulus lists in which [f] and [s] appeared as both onsets and codas, as was also true for [m], [n], [k], and [g]. But whether [f] or [s] was an onset or a coda for a particular participant was completely predictable from the vowel. In contrast, the syllable position of [m], [n], [k], or [g] was unrelated to the vowel. Thus, the lists embodied a second-order experiment-wide constraint. If participants are sensitive to this constraint, one would expect their [f] and [s] misorderings to maintain their syllable position to a greater extent than the other consonants that occurred as both onsets and codas.

Results and Discussion

We tallied 1,769 consonant movement errors, a rate of 4.8% errors per syllable. This rate was somewhat lower than the rates of the two previous experiments (8.3% and 9.7%). But with the large individual differences in error rates in these studies—the most inaccurate speaker had more than 10 times as many errors as the most accurate one—the mean rate would be expected to vary

² Any of several mechanisms could achieve this effect. One is that the violations of local positional constraints that result in something such as “kek” often involve within-syllable movement, such as “gek” → “kek.” Or there could be a bias for onset-coda harmony in the output, an effect that is often present in early child phonology.

across studies with only 8 participants. Errors declined over days: 635, 489, 355, and 290 errors for Days 1, 2, 3, and 4, respectively. All of the error outcomes are presented in Table 3.

The effects of language-wide and local constraints were the same as in the two previous experiments. All of the 484 misorderings of [h] and [ŋ] maintained their positions, replicating the finding of a language-wide phonotactic regularity effect. For the 1,026 errors involving the movement of [m], [n], [k], or

[g], 76.8% obeyed the local positional constraint. All 8 participants had more legal than illegal errors with these four consonants ($p < .01$; range: 54.6% to 86.7%). Thus, this experiment produced a syllable position effect of a similar size and reliability as in the previous two experiments.

The results for the experiment-wide constraints showed that these second-order constraints were indeed learned. The 259 misorderings involving [f] and [s] maintained their positions at a high

Table 3
Error Outcomes in Experiment 3 Classified by Legality and Condition

		fas-sif condition—legal outcomes with [æ] Onset							saf-fis condition—legal outcomes with [æ] Onset									
Coda		g	k	m	n	f	s	h	ŋ	Coda	g	k	m	n	f	s	h	ŋ
g	—	16	6	6	15	—	9	—	—	g	—	10	0	5	—	3	10	—
k	2	—	3	5	9	—	4	—	—	k	2	—	7	8	—	10	5	—
m	16	20	—	6	8	—	15	—	—	m	8	13	—	2	—	13	14	—
n	20	12	10	—	23	—	17	—	—	n	6	5	16	—	—	10	2	—
f	—	—	—	—	—	—	—	—	—	f	11	5	1	2	—	7	25	—
s	9	8	0	1	17	—	23	—	—	s	—	—	—	—	—	—	—	—
ŋ	27	32	7	10	25	—	41	—	—	ŋ	17	12	4	8	—	21	30	—
h	—	—	—	—	—	—	—	—	—	h	—	—	—	—	—	—	—	—

		fas-sif condition—legal outcomes with [ɪ] Onset							saf-fis condition—legal outcomes with [ɪ] Onset									
Coda		g	k	m	n	f	s	h	ŋ	Coda	g	k	m	n	f	s	h	ŋ
g	—	8	9	22	—	18	6	—	—	g	—	5	6	5	17	—	5	—
k	5	—	15	11	—	18	4	—	—	k	4	0	13	6	2	—	5	—
m	6	13	—	3	—	12	24	—	—	m	3	4	—	11	4	—	11	—
n	9	7	7	—	—	4	6	—	—	n	10	1	15	—	4	—	5	—
f	21	4	4	23	—	31	17	—	—	f	—	—	—	—	—	—	—	—
s	—	—	—	—	—	—	—	—	—	s	4	7	4	6	13	—	14	—
ŋ	7	22	6	14	—	53	25	—	—	ŋ	11	25	12	7	27	—	25	—
h	—	—	—	—	—	—	—	—	—	h	—	—	—	—	—	—	—	—

		fas-sif condition—illegal outcomes with [æ] Onset							saf-fis condition—illegal outcomes with [æ] Onset									
Coda		g	k	m	n	f	s*	h	ŋ**	Coda	g	k	m	n	f*	s	h	ŋ**
g	1	1	0	2	3	1	6	0	0	g	3	7	0	6	0	8	0	0
k	0	1	0	0	1	1	0	0	0	k	1	1	2	1	1	6	2	0
m	0	0	2	2	0	2	3	0	0	m	1	0	0	2	0	2	3	0
n	0	1	0	1	2	1	2	0	0	n	0	2	0	4	0	1	1	0
f*	0	0	0	0	0	0	2	0	0	f	0	0	0	0	0	—	—	0
s	1	1	3	0	—	1	—	0	0	s*	1	0	1	0	1	0	0	0
ŋ	5	1	0	0	—	1	—	0	0	ŋ	1	1	1	0	0	—	—	0
h**	0	0	0	0	0	0	0	0	0	h**	0	0	0	0	0	0	0	0

		fas-sif condition—illegal outcomes with [ɪ] Onset							saf-fis condition—illegal outcomes with [ɪ] Onset									
Coda		g	k	m	n	f*	s	h	ŋ**	Coda	g	k	m	n	f	s*	h	ŋ**
g	6	1	6	0	0	9	2	0	0	g	1	1	2	1	3	0	1	0
k	1	6	0	3	0	5	3	0	0	k	0	5	0	4	3	0	3	0
m	1	0	2	1	0	3	8	0	0	m	0	0	0	0	3	0	3	0
n	2	5	5	3	0	4	3	0	0	n	0	1	1	1	5	0	6	0
f	1	1	0	2	4	—	—	0	0	f*	4	0	0	0	0	0	1	0
s*	0	0	0	0	0	0	1	0	0	s	0	2	0	0	—	2	—	0
ŋ	3	2	0	0	0	—	—	0	0	ŋ	0	3	0	0	—	0	—	0
h**	0	0	0	0	0	0	0	0	0	h**	0	0	0	0	0	0	0	0

* violates experiment-wide constraints.

** violates language-wide constraints.

rate, 90.3% of the time (range: 82.9% to 100%). This occurred even though [f] and [s] appeared as both onsets and codas for each participant. All 8 participants individually followed the experiment-wide constraint to a greater extent than the local positional constraints ($p < .01$). Thus, error adherence to the experiment-wide distributions seemed to reflect more than just the syllable-position effect.

Errors that violated the experiment-wide distribution were more likely on the 1st day (13 errors; mean trial position = 49) than on the 2nd (4 errors; mean position = 80), 3rd (3 errors; mean position = 73), or 4th day (5 errors; mean position = 72). In percentage terms, the rate of violations was greater on the 1st day (14.3%) than on subsequent days (7.1%). This same pattern of a greater percentage of experiment-wide violations on the 1st day was also seen in Experiment 2, although in neither experiment were there enough errors for a robust statistical test.

The information that was given to informed participants had no effect on the error pattern. Informed speakers' [f] and [s] errors held their positions 90.0% of the time, as compared with 90.7% for uninformed speakers. Uninformed speakers did not report anything about the distribution of [f] and [s] in response to the postexperiment query. Thus, as in the previous two experiments, there is little evidence that explicit knowledge was responsible for error adherence to the experiment-wide constraints.

In summary, the speakers in Experiment 3 appeared to learn implicitly about the constraint in the materials involving the association between the position of [f] and [s] and the particular vowel in the sequences. This learning was expressed as a strong tendency for their errors to adhere to the constraint, a tendency that was stronger than that for the unrestricted consonants ([k], [g], [m], and [n]) to maintain their positions in movement errors. However, Experiment 3's experiment-wide effect (90.3% adherence) was somewhat smaller than that found in the first two experiments (97.7% and 94.7%), suggesting that the second-order constraint may not be registered as strongly as purely positional constraints.

Consider now some possible mechanisms for the experiment-wide effect found in Experiment 3. As we did for the earlier two experiments, we can rule out an explanation based on a tendency for errors to create whole syllables that are present in the stimuli. As the illegal matrices in Table 3 show, errors violating local positional constraints that create nonoccurring syllables (e.g., "nin") were at least as common as those creating occurring syllables (e.g., "kam"). However, the finding that [f] and [s] movement errors were sensitive to the identity of the vowel reveals that speakers were acquiring more than just purely positional rules of the form [f] is an onset for this experiment. Thus, it appears that what was being learned is some kind of consonant position-vowel link.

One hypothesis regarding these consonant-vowel links is that speakers are learning the sets of onsets and rimes that occur in the stimuli (see Treiman, 1983, 1986). For example, in the fas-sif condition, they learn that the possible rimes include [æs] and [ɪf], but not [æf] and [ɪs]. Along with this, they learn that the possible onsets include both [f] and [s]. It is further assumed that errors have a strong bias to create only occurring onsets and rimes. This onset-rime hypothesis makes a clear prediction: Violations of Experiment 3's experiment-wide constraints should be largely confined to movements to the onset position. This is because the onset-rime hypothesis links vowels and codas but not onsets and

vowels. Consequently, one should find illegal error outcomes such as "mak" spoken as "sak" in the fas-sif condition, in which the onset [s] illegally surfaces in an [æ] syllable. This violation would not be prevented, however, because [s] was an acceptable onset in the experiment. In contrast, one would not find "mak" spoken as "maf," in which the coda [f] illegally appears in an [æ] syllable; [æf] is a nonoccurring rime in this condition and thus would be avoided.

We examined all 25 outcomes that violated the experiment-wide constraint in this experiment (see the rows and columns in the illegal matrices in Table 3 labeled with a single asterisk). These errors involved 15 cases of onset movement and 11 of coda movement. (The reason that this sums to 26 is that one outcome involved the movement of both an onset and a coda.) So, although there were 4 more onset than coda violations (directionally consistent with the onset-rime hypothesis), it appears that error adherence to the experiment-wide constraint was not at all confined to coda movements. Whether there is any reliable asymmetry between onset and coda violations with our methods can be determined only in a much larger experiment. Nonetheless, it seems likely that the consonant-vowel links that speakers acquire in our procedure involve onsets as well as codas. This result goes along with other evidence that listeners' sensitivity to the frequency of diphones in CVC syllables includes the onset-vowel and vowel-coda combinations (Treiman, Kessler, Knewasser, Tincoff, & Bowman, 1999).

General Discussion

Speech errors reflect a speaker's knowledge of where sounds occur in syllables and words. Errors obey language-wide phonotactic constraints; one rarely finds an error whose outcome contains an impossible sound sequence. At the same time, errors exhibit the syllable-position effect; when a sound moves from one syllable to another, it tends to stick to its position. We demonstrated these two phenomena experimentally and found, as well, a sensitivity to the distribution of sounds within an experiment.

In Experiment 1, we found that an experiment-wide constraint regarding the syllable positions of [f] and [s] was obeyed by [f] and [s] movement errors to a very strong extent, 97.7% of the time. Misplacements of consonants not subject to experiment-wide restrictions maintained their positions at a lower 68.2% rate. Experiment 2 replicated this pattern, except with [k] and [g] as the sounds subject to experiment-wide constraint; [k] and [g] movements kept to their position 94.7% of the time, in contrast to the other unrestricted consonants, whose errors kept their syllable position at a 77.5% rate. The final experiment provided evidence that second-order experiment-wide constraints, such as [f] being limited to the onset position when the vowel is [æ], are revealed in errors. Of the relevant movement errors, 90.3% obeyed this constraint, a rate significantly higher than the 76.8% with which unrestricted consonants maintained their syllable positions when they were misplaced.

These results, we claim, say something about both language and learning. With respect to language, they support what we called the breadth of constraint hypothesis, a claim that patterns in language occur at many levels of generality, and the processing system is sensitive to all of these levels. Thus, we see error adherence to language-wide constraints (the phonotactic regularity effect) and

to local positional constraints (the syllable-position effect) as two ends of a continuum. The former applies to the whole language, and the latter applies to a single word or syllable. Our finding of sensitivity to experiment-wide constraints is intermediate between these extremes.

With respect to learning, our data illustrate the implicit learning of sequential patterns, akin to the many experiments in which participants push buttons in response to rule-governed sequences of stimuli (e.g., Cleeremans, 1993; Poldrack, Selco, Field, & Cohen, 1999). The learning that occurred in our studies was implicit, as this term is normally used. The error patterns were unaffected by informing the participants of the sound distributions in their syllables, and those who were uninformed did not report any knowledge of the distribution when asked. Furthermore, the task required of the participants, repeating syllables, can be performed quickly and does not necessarily require any knowledge of the experimentally manipulated patterns. Finally, the data of interest, errors, were not intended responses by our speakers. Hence, it is difficult to see the error pattern resulting in a simple way from speakers' conscious intentions.

An illuminating characteristic of our data is that sensitivity to experiment-wide constraints occurred very quickly. Although the effect seemed to strengthen across days, it was clearly present during the 1st day. In Experiment 1, for example, only 4 of 184 misplacements of [f] and [s] were violations, and these violations were haphazardly distributed across the trials. These findings indicate very rapid learning. With this in mind, we offer the following general mechanism: Each utterance of a syllable tunes the language production system to favor the production of that and similar syllables. The effect of this tuning endures longer than a single trial, and it accumulates with the tuning associated with other utterances. The overall effect is to adapt the production system to recent experience. It is noteworthy that this kind of mechanism has also been offered for the grammatical encoding component in language production. Bock and Griffin (2000) have showed that structural priming, the tendency for speakers to use the same syntactic structure as a recently experienced prime sentence (e.g., Bock, 1986), is a form of implicit learning. The effects of a single prime are relatively long lasting, and hence the tunings associated with individual utterances can accumulate to create general biases for particular structures.

Specific mechanisms for this kind of adaptive tuning have been proposed in recurrent network theories of sequence learning, which associate implicit learning with connection weight changes between representations of potential responses and distributed representations of prior contexts (e.g., Chang, Dell, Bock, & Griffin, 2000; Cleeremans, 1993; Elman, 1990; Gupta & Cohen, 1998; see Christiansen & Chater, 1999, for a review). However, to account for our data, very small weight changes will probably not do the trick. All that is required is the absence of any [f] codas over a few sequences to make the system considerably less likely to generate an [f] coda as an error than when there are [f] codas in these sequences. One may need a large short-term component to the weight change process to account for the alacrity and size of our effects. At least one account of implicit sequence learning, that of Cleeremans (1993), includes both a short-term and a long-term component to adaptive weight change. Each response creates a sizable change in weights, but only a fraction of that change is permanent. Similarly, the node structure theory of MacKay (1987)

allows for both short-term (recency) and permanent changes in the transmission of priming among representational units responsible for producing linguistic sequences. These would be analogous to the weight changes in distributed connectionist models.

Another issue related to learning concerns the representations over which learning occurs. We have suggested that the learning is not based on indivisible syllables. That is, the speakers did not simply inventory which syllables occurred in the experiment and which did not with the sole effect that errors were biased toward creating occurring syllables. Furthermore, we have shown that speakers can learn more than just the set of occurring onsets and codas. Experiment 3 suggests that the experiment-wide error effect can be sensitive to consonant position–vowel combinations, as well as straight consonant positions. However, this was just a single study, and so we do not wish to make any strong claims about what units and structures the learning occurs over. At this point in time, the available data are consistent with any number of representational proposals, provided that they allow for the registration of word and/or syllable boundary information and the co-occurrence of, at least, adjacent speech sounds. For example, if a recurrent network's output consisted of sequences of phonemes or features and included boundaries, it would naturally tend to register these kinds of co-occurrences. And recently experienced co-occurrences would have privileged effects on output. Thus, recurrent network models might be able to explain the learning effects on errors. However, we are a long way from being able to offer this kind of model of our data. Leaving learning aside, any model will, first, have to create the right kinds of phonological errors, particularly errors in which phonemes exchange or otherwise move between syllables. None of the existing recurrent network models of the production of word forms (e.g., Dell et al., 1993) can give a good account of phonemic movement errors, particularly exchanges (see Dell, Chang, & Griffin, 1999, for a discussion).

To conclude, we return to the phonotactic regularity of speech errors. Why does this effect occur? Our experiments invite the conclusion that at least part of the effect is the result of recent experience. You may say "blug" instead of "bug" because you have said words beginning with [b] in your recent past. But you would never say "lbug." Perhaps that is because your phonology does not allow syllable initial [lb]. But perhaps it is instead because you do not have any recent experience saying this cluster. Some modern phonological theories tend to see the phonology of a language as being projected from its lexicon (see Broe & Pierrehumbert, 1999). That is, there is no independent abstract representation of phonological patterns aside from what is stored in the lexicon. Or, if there is some abstraction, it is computed from the contents of the lexicon. We would only add to this notion that the phonology is projected preferentially from those parts of the lexicon that are most accessible, such as recently experienced sound forms.

References

- Anderson, K., Milostan, J., & Cottrell, G. W. (1998). Assessing the contribution of representation to results. In *Proceedings of the Twentieth Annual Meeting of the Cognitive Science Society* (pp. 48–53). Mahwah, NJ: Erlbaum.
- Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of

- conditional probability statistics by 8-month-old infants. *Psychological Science*, 9, 321–324.
- Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology*, 18, 355–387.
- Bock, J. K., & Griffin, Z. M. (2000). The persistence of structural priming. *Journal of Experimental Psychology: General*, 129, 177–192.
- Boomer, D. S., & Laver, J. D. M. (1968). Slips of the tongue. *British Journal of Disorders of Communication*, 3, 1–12.
- Broe, M., & Pierrehumbert, J. (Eds.). (1999). *Papers in laboratory phonology V: Acquisition and the lexicon*. Cambridge, England: Cambridge University Press.
- Chang, F., Dell, G. S., Bock, J. K., & Griffin, Z. M. (2000). Structural priming as implicit learning: A comparison of models of sentence production. *Journal of Psycholinguistic Research*, 29, 217–229.
- Christiansen, M. H., & Chater, N. (1999). Connectionist natural language processing: The state of the art. *Cognitive Science*, 23, 417–437.
- Cleeremans, A. (1993). *Mechanisms of implicit learning: Connectionist models of sequence processing*. Cambridge, MA: MIT Press.
- Cutler, A. (1981). The reliability of speech error data. *Linguistics*, 19, 561–582.
- Dell, G. S. (1986). A spreading activation theory of retrieval in language production. *Psychological Review*, 93, 283–321.
- Dell, G. S. (1990). Effects of frequency and vocabulary type on phonological speech errors. *Language and Cognitive Processes*, 5, 313–349.
- Dell, G. S., Chang, F., & Griffin, Z. M. (1999). Connectionist models of language production: Lexical access and grammatical encoding. *Cognitive Science*, 23, 517–542.
- Dell, G. S., Juliano, C., & Govindjee, A. (1993). Structure and content in language production: A theory of frame constraints in phonological speech errors. *Cognitive Science*, 17, 149–195.
- Dulany, D. E., Carlson, R. A., & Dewey, G. J. (1984). A case of syntactical learning and judgment: How conscious and abstract? *Journal of Experimental Psychology: General*, 113, 541–555.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14, 179–211.
- Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.
- Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, 127, 107–140.
- Fowler, C. A., Treiman, R., & Gross, J. (1993). The structure of English syllables and polysyllables. *Journal of Memory and Language*, 32, 115–140.
- Frisch, S., & Wright, R. (1998). *The phonetics of phonological speech errors: An acoustic analysis of /s/ and /z/* (Progress Rep. No. 22). Bloomington: Indiana University.
- Fromkin, V. A. (1971). The non-anomalous nature of anomalous utterances. *Language*, 47, 27–52.
- Goldinger, S. D. (1998). Echoes of echoes? An episodic theory of lexical access. *Psychological Review*, 105, 251–279.
- Gupta, P., & Cohen, N. J. (1998). *Skill learning, repetition priming, and procedural memory: Theoretical and computational analysis*. Manuscript submitted for publication.
- Gupta, P., & Dell, G. S. (1999). The emergence of language from serial order and procedural memory. In B. MacWhinney (Ed.), *Emergentist approaches to language* (28th Carnegie Mellon Symposium on Cognition) (pp. 447–481). Mahwah, NJ: Erlbaum.
- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, 93, 411–428.
- Jordan, M. I. (1986). Attractor dynamics and parallelism in a connectionist sequential machine. In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society* (pp. 531–546). Hillsdale, NJ: Erlbaum.
- Jusczyk, P. W., Friederici, A. D., Wessels, J. M. I., Svenkerud, V. Y., & Jusczyk, A. M. (1993). Infants' sensitivity to sound patterns of native language words. *Journal of Memory and Language*, 32, 402–420.
- Kessler, B., & Treiman, R. (1997). Syllable structure and the distribution of phonemes in English syllables. *Journal of Memory and Language*, 37, 295–311.
- Knowlton, B. J., & Squire, L. R. (1994). The information acquired during artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 79–91.
- Levelt, W. J. M., & Wheeldon, L. (1994). Do speakers have access to a mental syllabary? *Cognition*, 50, 239–269.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492–527.
- MacKay, D. G. (1970). Spoonerisms: The structure of errors in the serial order of speech. *Neuropsychologia*, 8, 323–350.
- MacKay, D. G. (1972). The structure of words and syllables: Evidence from errors in speech. *Cognitive Psychology*, 3, 210–227.
- MacKay, D. G. (1987). *The organization of perception and action: A theory for language and other cognitive skills*. New York: Springer-Verlag.
- Marcus, G. F., Vijayan, S., Rao, S. B., & Vishton, P. M. (1999). Rule learning in 7-month-old infants. *Science*, 283, 77–80.
- Massaro, D. W., & Cohen, M. M. (1983). Phonological constraints in speech perception. *Perception & Psychophysics*, 34, 338–348.
- Mattys, S. L., Jusczyk, P. W., Luce, P. A., & Morgan, J. L. (1999). Phonotactic and prosodic effects on word segmentation in infants. *Cognitive Psychology*, 38, 465–494.
- McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114, 159–188.
- McQueen, J. M. (1998). Segmentation of continuous speech using phonotactics. *Journal of Memory and Language*, 39, 21–46.
- Meyer, A. S. (1992). Investigation of phonological encoding through speech-error analyses: Achievements, limitations, and alternatives. *Cognition*, 42, 181–211.
- Motley, M. T., & Baars, B. J. (1975). Encoding sensitivities to phonological markedness and transition probability: Evidence from spoonerisms. *Human Communication Research*, 2, 351–361.
- Mowrey, R. A., & MacKay, I. R. A. (1990). Phonological primitives: Electromyographic speech error evidence. *Journal of the Acoustical Society of America*, 88, 1299–1312.
- Neal, A., & Hesketh, B. (1997). Episodic knowledge and implicit learning. *Psychonomic Bulletin and Review*, 4, 24–37.
- Nooteboom, S. G. (1969). The tongue slips into patterns. In A. G. Sciarone, A. J. van Essen, & A. A. Van Raad (Eds.), *Leyden studies in linguistics and phonetics* (pp. 114–132). The Hague: Mouton.
- Norris, D., McQueen, J. M., Cutler, A., & Butterfield, S. (1997). The possible-word constraint in the segmentation of continuous speech. *Cognitive Psychology*, 34, 191–243.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53–79.
- Pinker, S. (1991). Rules of language. *Science*, 253, 530–535.
- Pitt, M. A. (1998). Phonological processes and the perception of phonotactically illegal clusters. *Perception & Psychophysics*, 60, 941–951.
- Plaut, D. C., & Kello, C. T. (1999). The emergence of phonology from the interplay of speech comprehension and production: A distributed connectionist approach. In B. MacWhinney (Ed.), *The emergence of language* (pp. 381–415). Mahwah, NJ: Erlbaum.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. E. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, 56–115.
- Poldrack, R. A., Selco, S. L., Field, J. E., & Cohen, N. J. (1999). The relationship between skill learning and repetition priming: Experimental

- and computational analyses. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 208–235.
- Reber, A. S. (1993). *Implicit learning and tacit knowledge: An essay on the cognitive unconscious*. New York: Oxford University Press.
- Redington, M., & Chater, N. (1996). Transfer in artificial grammar learning: A reevaluation. *Journal of Experimental Psychology: General*, 125, 123–138.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926–1928.
- Seidenberg, M. S., & McClelland, J. L. (1989). A distributed developmental model of visual word recognition and naming. *Psychological Review*, 96, 523–568.
- Shattuck-Hufnagel, S. (1983). Sublexical units and suprasegmental structure in speech production planning. In P. F. MacNeilage (Ed.), *The production of speech* (pp. 109–136). New York: Springer-Verlag.
- Shattuck-Hufnagel, S. (1986). The representation of phonological information during speech production planning: Evidence from vowel errors in spontaneous speech. *Phonology Yearbook*, 3, 117–149.
- Stemberger, J. P. (1982). *The lexicon in a model of language production*. Unpublished doctoral dissertation, University of California, San Diego, La Jolla.
- Stemberger, J. P. (1983). *Speech errors and theoretical phonology: A review*. Bloomington: Indiana University Linguistics Club.
- Stemberger, J. P. (1989). Speech errors in early child language production. *Journal of Memory and Language*, 28, 164–188.
- Treiman, R. (1983). The structure of spoken syllables: Evidence from novel word games. *Cognition*, 15, 49–74.
- Treiman, R. (1986). The division between onsets and rimes in English syllables. *Journal of Memory and Language*, 25, 476–491.
- Treiman, R., & Danis, C. (1988). Short-term memory errors for spoken syllables are affected by the linguistic structure of syllables. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 145–152.
- Treiman, R., Kessler, B., Knewasser, S., Tincoff, R., & Bowman, M. (1999). English speakers' sensitivity to phonotactic patterns. In M. Broe & J. Pierrehumbert (Eds.), *Papers in laboratory to phonology V: Acquisition and the lexicon* (pp. 269–282). Cambridge, England: Cambridge University Press.
- Vitevitch, M. S., Luce, P. A., Charles-Luce, J., & Kemmerer, D. (1997). Phonotactics and syllable stress: Implications for the processing of spoken nonsense words. *Language and Speech*, 40, 47–62.
- Wells, R. (1951). Predicting slips of the tongue. *Yale Scientific Magazine*, 3, 9–30.

Appendix

Syllable Types in the fes Condition in Each Set of 96 Sequences

Two restricted sounds (expected frequency: 24)	One restricted sound (expected frequency: 12)	No restricted sound (expected frequency: 8)
heng	keng	kem
hes	gheng	ken
fes	meng	keg
feng	neng	mek
	hek	men
	heg	meg
	hem	ghem
	hen	gek
	fek	ghen
	feg	nem
	fen	nek
	fem	neg
	kes	
	ges	
	nes	
	mes	

Received August 25, 1999

Revision received March 14, 2000

Accepted March 21, 2000 ■