Published in final edited form as:

J Exp Psychol Learn Mem Cogn. 2008 September ; 34(5): 1289–1295. doi:10.1037/a0013033.

Limits on Learning Phonotactic Constraints from Recent Production Experience

Jill A. Warker, Gary S. Dell, Christine A. Whalen, and Samantha Gereg

University of Illinois at Urbana-Champaign

Abstract

Adults can learn new artificial phonotactic constraints by producing syllables that exhibit the constraints. The experiments presented here tested the limits of phonotactic learning in production using speech errors as an implicit measure of learning. Experiment 1 tested a constraint in which the placement of a consonant as an onset or coda depended on the identity of a nonadjacent consonant. Participants' speech errors reflected knowledge of the constraint, but not until the second day of testing. Experiment 2 tested a constraint in which consonant placement depended on an extralinguistic factor, the speech rate. Participants were not able to learn this constraint. Together, these experiments suggest that phonotactic-like constraints are acquired when mutually constraining elements reside within the phonological system.

Keywords

phonotactic learning; speech errors

Phonotactic constraints characterize possible phoneme sequences in a language. For example, in English, /ŋ/ may occur postvocally in a syllable, that is, in the *coda* position, but never in the prevocalic, or *onset* position. Yet in other languages such as Vietnamese, /ŋ/ appears often as an onset, demonstrating that these kinds of constraints must be learned. Here, we investigate learning mechanisms by taking an established paradigm for studying the effects of phonotactic learning on speech production and asking whether participants can learn something a little out of the ordinary. In one case we find success; and in another, failure. These contrasting results allow us to set limits on the adaptability of the phonological processing system.

It seems likely that knowledge of native-language phonotactics is acquired early in life, as young as 6 to 9 months of age (Jusczyk, Friederici, Wessels, Svenkerud & Jusczyk, 1993). Moreover, older infants rapidly learn artificial phonotactic constraints. Chambers, Onishi, and Fisher (2003) exposed 16-month old infants to syllables adhering to constraints, such as /f/ is *always an onset*, for 2 minutes. Later, they were able to discriminate between syllables obeying those constraints and syllables violating them. This ability to learn new artificial constraints from perceptual experience continues into adulthood. Adult participants learned these same constraints from the same set of syllables, as measured with a response time task (Onishi, Chambers & Fisher, 2002).

What about production? We believe that native-language phonotactics affect production because everyday speech errors appear to be phonotactically legal. One might slip up and pronounce “nun” as “nung,” but not “ngung” (Fromkin, 1971)¹. Can the sensitivity of slips to

Correspondence may be directed to Jill A. Warker (warker@uiuc.edu), Beckman Institute, 405 N. Mathews Ave., University of Illinois at Urbana-Champaign, IL 61801. 217-244-1294.

phonotactics be altered by recent experience? Dell, Reed, Adams, and Meyer (2000) reasoned that speech errors would be a good tool to investigate the acquisition of new phonotactic constraints because people do not intend to misspeak; their errors are thus an unobtrusive, implicit measure of their knowledge. In their studies, participants produced sequences of syllables in four sessions on separate days and errors were recorded. Each sequence consisted of four consonant-vowel-consonant (CVC) syllables containing the vowel /ε/ and the eight consonants in this example: *hes feng meg ken*. The positions of the consonants varied but some were constrained. /h/ and /ŋ/ were subject to a *language-wide constraint* in accordance with English phonotactics, i.e. /h/ always began a syllable and /ŋ/ always ended a syllable. The key part of the study, however, concerned two consonants subject to an *experiment-wide constraint*: e.g. during the experiment /f/ was always an onset and /s/ was always a coda. The remaining consonants (e.g. /k/, /g/, /m/, /n/) were *unrestricted* and could appear as onsets or codas. Participants repeated each sequence three times quickly to induce speech errors. The resulting errors were coded as legal (a moving consonant maintained its original syllable position when it slipped to a different syllable) or illegal (slipped to a different position). Dell et al. found that slips of /h/ and /ŋ/ were always legal, demonstrating the effect of native phonotactics on errors. More importantly, slips of experimentally constrained consonants were also overwhelmingly legal (97% legal across two experiments), far more so than those of unrestricted consonants (73% legal). The slips reflected the phonotactic distribution within the experiment and hence demonstrated learning of that distribution. Moreover, the learning was rapid. Participants' errors adhered to the experimental constraints as strongly on the first day of testing as on the fourth day (see also Taylor & Houghton, 2005 for evidence for rapid learning in this paradigm).

The experimental constraint in Dell et al. (2000) was a *first-order constraint*, meaning that it did not depend on context, e.g. /f/, when it occurs, is always an onset. Using Dell et al.'s (2000) paradigm, Warker and Dell (2006) embedded a *second-order constraint* in their stimuli, that is, the syllable position of a constrained consonant depended on another aspect of the syllable. For example, Warker and Dell tested *if the vowel is /æ/, /k/ is an onset and /g/ is a coda but if the vowel is /I/, /g/ is an onset and /k/ is a coda*. Participants were able to learn this constraint as evidenced by their speech errors. Errors involving restricted consonants were legal more often (86%) than errors involving unrestricted ones (75%).

Three additional experiments by Warker and Dell (2006) replicated the learning of second-order constraints. In all four experiments, though, the learning contrasted with the learning of first-order constraints found by Dell et al. (2000) in that there was no evidence of learning on the first day. Despite reciting over a thousand syllables in the first session, participants' slips of restricted consonants were no more legal than their slips of unrestricted consonants. Figure 1 presents the data from these four second-order studies, separating the effects of Day 1 from the other three sessions. Warker and Dell suggested that second-order constraints take time to learn because they are *self-interfering*, and implemented this idea in a connectionist model. For example, learning to map /k/ to onset position when the vowel is /æ/ interferes with mapping /k/ to the coda position when the vowel is /I/. This interference can only be overcome through learning new context-sensitive representations in the model's intermediate or "hidden" layer of processing units, which takes many additional trials.

In this paper, we explore the limits of learning new second-order constraints using the speech-error paradigm. Experiment 1 looks at whether the conditioning context must be adjacent to the restricted consonants (as vowels are next to onsets and codas in CVCs) whereas Experiment 2 tests whether conditioning contexts must be phonological in nature.

¹Although slips obey phonotactics to a great extent, it should be noted that experimental studies of tongue-twister production have documented the production of ill-formed consonants (see Goldstein et al., 2007, for review).

Experiment 1

Second-order constraints are common in the world's languages, particularly constraints in which mutually constraining elements are adjacent (e.g. in English, onsets /s/ and /dr/ are legal, but /sr/ and dl/ are not). Although less common, nonadjacent phonotactic constraints do occur in several languages. They most often involve mutually constraining vowels (e.g. vowel harmony as in Finnish), but constraints between nonadjacent consonants have been described (e.g. Koo & Cole, 2006; Hansson, 2001, Rose & Walker, 2004). Can nonadjacent second-order phonotactic constraints be readily learned in the laboratory? One important study showed that adults are able to detect that a nonadjacent sequence of three phonemes recurs in a continuous list of syllables, showing that the perceptual system can register and remember such dependencies (Newport & Aslin, 2004). Here we ask a related question for production.

Although the existence of nonadjacent phonotactic constraints in some languages is, ipso facto, evidence for the learnability of such constraints in both perception and production, it is important to study learning in an experimental setting for two reasons. First, principles of sequence learning under controlled conditions can be examined and so the nature and limits of learning can be determined. Second, any such principles that are discovered may, in fact, be related to true language learning and so one can use the experimental findings to develop testable hypotheses about acquisition.

Two of Warker and Dell (2006)'s speech error studies (the final two in Figure 1) demonstrated learning of constraints that, at first glance, appear to be nonadjacent. For example, in their Experiment 2a, participants recited sequences of /r/-medial CVCVC disyllables such as *kerem nereg hereng feres* and /l/-medial ones such as *gelen melek feleng heles*. All disyllables conformed to the constraint: *if the middle consonant of a CVCVC is /r/, /k/ is an onset and /g/ is a coda but if it is /l/, /g/ is an onset and /k/ is a coda*. Described in this way, the positioning of /g/ and /k/ depends on the nonadjacent /r/ or /l/.

These studies, however, did not test a truly nonadjacent dependency because postvocalic /r/s and /l/s greatly color the preceding /ε/ (e.g. “kereg” as in “Kerry”, but “keleg” as in “Kelly”). To show this, we examined utterances by a naive male speaker of American north midlands dialect of 8 disyllables from Warker and Dell (2006) (4 with ‘ere’ matched to 4 with ‘ele’). Acoustic analysis of the formant midpoints found that the preceding /ε/s consistently and significantly differed as a function of medial consonant at every formant measured (F1: $t(3) = 6.29, p = .008$; F2: $t(3) = -5.73, p = .011$; F3: $t(3) = 4.59, p = .019$). The differences were particularly dramatic in F2 and F3 with medial /l/s leading to /ε/s with an F3-F2 difference of 1062 Hz, while those preceding medial /r/s had a difference of only 551 Hz. These “/ε/”s are quite different phonetically. Thus, Warker and Dell's participants may have learned a contingency between an *adjacent* vowel and the onset consonant, just as in the first two experiments in Figure 1.

The first experiment was designed to see if participants could learn a truly nonadjacent constraint through their production experience by replacing the medial /r/ and /l/ with /v/ and /b/. An analysis of utterances with the same disyllables and speaker as that described above, except with medial combinations of “eve” and “ebe,” showed that the initial /ε/s before /v/ and /b/ were nearly identical in their formant frequencies, (F1: $t(3) = .24, p = .83$; F2: $t(3) = -.57, p = .61$; F3: $t(3) = .23, p = .83$). The mean differences (“ebe” – “eve”) were 2 Hz, –22 Hz, and 6 Hz, for F1-F3, respectively. Thus, the /v/-/b/ contrast does not affect the identities of their preceding vowels the way /r/ and /l/ do.

Because of the second-order nature of the medial-consonant conditioned constraint, we expect any learning that occurs to be subtle and to only appear after the first session. Consequently, this experiment will test twice as many participants as all previous studies using this paradigm

(16 rather than 8) and will do so over four separate days. Furthermore, two, rather than one, sets of restricted onset and coda consonants will be employed, an /f/-/s/ set and a /k/-/g/ set.

Participants

Sixteen native English speaking students from the University of Illinois at Urbana-Champaign received \$40 for participating in 4 testing sessions over the course of 4 days. Half were randomly assigned to the experiment-restricted /f/-/s/ condition and half to the experiment-restricted /k/-/g/ condition.

Materials

Each day, participants received 96 sequences, each containing four CVCVC disyllables. The medial VCV combination of all four disyllables in each sequence was either “eve” or “ebe” and the combinations alternated by sequence. Throughout the study, the same set of eight consonants appeared once per sequence in the onset and coda positions of the disyllables. These consonants were classified into three groups: language-restricted (/h/ & /ŋ/), experiment-restricted (/f/ & /s/ or /k/ & /g/, depending on condition), and unrestricted (/k/, /g/, /m/, /n/ or /f/, /s/, /m/, /n/, depending on condition). As per English phonotactics, /h/ always occurred in an onset position and /ŋ/ always occurred in a coda position. The unrestricted consonants could appear in onset or coda position. The experiment-restricted consonants differed in placement according to condition. For the eight participants in the /f/-/s/ condition, half experienced the constraint: *if the medial consonant is /v/, /f/ is an onset and /s/ is a coda but if the medial consonant is /b/, /s/ is an onset and /f/ is a coda.* (That is, when /v/ is medial, any /f/ in the disyllable must be an onset and any /s/ must be a coda, and analogously for medial /b/). The remaining half in this condition received /f/ and /s/ in the opposite position assignments. For the eight participants in the /k/-/g/ condition, half received the constraint: *if the medial consonant is /v/, /k/ is an onset and /g/ is a coda but if the medial consonant is /b/, /g/ is an onset and /k/ is a coda.* The other half received the opposite assignment of /k/ and /g/.

A computer program randomly generated four sets of 96 sequences following the appropriate constraints for each participant. The lists were printed in 16-point bold Arial font with 1 sequence per line and 16 lines per page, with /v/-medial and /b/-medial sequences alternating as illustrated below for the condition in which /f/ is an onset if /v/ and coda if /b/, and /s/ is an onset if /b/ and coda if /v/:

hevek	gevang	feven	meves
mebef	keben	hebeng	sebeg
nevek	fevang	hevem	neves
etc.			

Procedure

Participants first recited each sequence once slowly at a rate of 1 disyllable/second and then repeated the sequence three times without pause at 2.53 disyllables/second in time to a metronome. Each disyllable was to be given trochaic stress so that the first vowel was /ε/ and the second a schwa. Participants did one set of 96 sequences per day on 4 separate days. Each session was digitally recorded. Before beginning the experiment, participants received four sample sequences to provide practice with pronunciation and the overall procedure.

Results & Discussion

The fast repetitions were transcribed for errors involving movement of the non-medial consonants to non-medial positions. Errors were classified as either legal or illegal depending on what position consonants moved to in an error. An illegal error occurred when a consonant

moved from onset to coda, or coda to onset. A legal error occurred when a consonant kept to its original onset or coda slot in the move. For example, if the sequence *fevek heves mevang neveg* from the /f/-/s/ restricted condition was pronounced *feveng hevek meves nevem*, the errors were scored as follows: For language-restricted consonants, there is one legal error because the /ŋ/ coda moved from *mevang* to *feveng*. For experiment-restricted consonants, there is one legal error because the /s/ coda moved from *heves* to *meves*. For unrestricted consonants, there is one legal error because the /k/ coda moved from *fevek* to *hevek* and one illegal error because the /m/ onset moved from *mevang* to *nevem*. Cutoff errors were included in the error analyses, such as *h...fevek*, but errors involving undistinguishable consonants, consonants outside the original set of eight (such as /d/), and incorrect medial VCV combinations were not included.

The reliability of error transcription was determined by having a second coder who was naive to the conditions transcribe a subset of the recordings. Overall, the reliability was good. Of the 5760 syllables doubly transcribed, both coders agreed there was no error on 5400 syllables and on the presence and nature of 249 errors, an overall agreement rate of 98.1%. Considering only the syllables where the original coder found an error (321 errors), the conditionalized agreement rate was 77.6%. These values are comparable to what has been found in similar studies.

When participants made an error involving /h/ or /ŋ/, it was always legal (mean legality = 100%, SE = 0, based on 1328 errors). However, we were primarily interested in comparing the legality of errors of experiment-restricted and unrestricted consonants. In keeping with previous studies of second-order constraints, we tested three planned contrasts oriented around the distinction between the first day of testing and the remaining three days: (a) the legality of slips of restricted versus unrestricted consonants on Day 1 (no difference is expected) (b) the restricted-unrestricted contrast for Days 2–4 (restricted slips should be more legal if learning occurs), and (c) an interaction contrast comparing the size of the restricted-unrestricted difference on Day 1 to that on Days 2–4 (the difference should be significantly greater on Days 2–4 compared to Day 1). All of the planned contrasts use the normal approximation of the Wilcoxon signed-rank test. The mean legality percentages for Day 1 and Days 2–4 are given in Table 1. On Day 1, there was no significant difference between experiment-restricted and unrestricted errors (Wilcoxon $Z = -.83$, $p = .41$). However, the predicted difference emerged on Days 2–4 (Wilcoxon $Z = 2.38$, $p = .02$) and there was a significant interaction between day of testing and restrictedness (Wilcoxon $Z = 2.33$, $p = .02$). It is worth noting that the effects of learning were quite similar for the f-s and k-g conditions. Over Days 2–4, the legality of restricted slips exceeded that of unrestricted ones by 6.2% and 7.4% in the two conditions, respectively. These differences are comparable to those found in Warker and Dell's (2006) two final experiments, which also used CVCVCV stimuli (Figure 1). Overall, the results demonstrate that participants were able to learn this nonadjacent constraint but, as with other second-order constraints, this learning took until the second day to manifest itself in their speech errors.

Experiment 2

In all previous speech-error experiments testing the learning of second-order constraints, the constraining contexts were phonological. That is, the constraints on possible onsets and codas were governed by other phonological aspects of the strings, such as the identity of nearby vowels and consonants. Is this a necessary condition for learning? Onishi et al. (2002) tested whether an extralinguistic property can serve as the constraining context for learning an artificial second-order constraint. In their study, which involved perception rather than production, participants heard syllables spoken by two different speakers, a male and a female. The syllables exhibited second-order constraints such as, /f/ is an onset if the syllable is spoken by speaker A; /f/ is a coda if spoken by speaker B. Thus, the distribution of consonants depended on the speaker's voice. Speaker identity is an example of an indexical (non-linguistic) speech

property; it greatly affects the acoustics of a syllable, but not its phonological form. Onishi et al. found that adult listeners exhibited no significant sensitivity to this contingency in their response-time measure, and suggested that the contexts that constrain phoneme distributions must be phonological to be learnable. We call this the *modularity hypothesis*, and Onishi et al.'s negative results constitute its principal supporting evidence. It should be noted, however, that their study was just a single session. If second-order learning in perception is, like production, expressed more strongly in subsequent sessions, one can question whether these data are strong support for the hypothesis. It should also be noted that extralinguistic properties, such as speaker identity, do affect many aspects of language processing. For example, they have been shown to affect lexical disambiguation (Creel, Aslin, & Tannenhaus, 2008) and spoken word recognition (Sommers & Barcroft, 2006) in perception as well as spontaneous speaker imitation in production (Goldinger, 1998). Social factors have also been shown to play a role in learning allophonic variation (Hay, Warren, & Drager, 2006). Thus, the modularity hypothesis, if true, must be recognized as a restriction on constraints that can specifically affect the implicit learning of phoneme distributions, not a general constraint preventing the influence of extralinguistic properties on allophonic variation, speech processing, and the retention of these properties in memory.

Experiment 2 uses the speech-error method to test for learning of a second-order constraint in which the constraining context is extralinguistic. We chose speech rate as the constraining context. Participants recited half of their sequences at a fast rate and half at a slower rate and whether particular consonants were onsets or codas was contingent on the rate. Speech rate has a substantial effect on production (as we will demonstrate). Rate, however, is represented separately from the segmental properties of words in production theories; it is a global property of production that controls when output is selected, but is not thought to be part of the relevant abstract phonological representations (e.g. Dell, 1986; MacKay, 1982; see also Sommers & Barcroft, 2006, for discussion about how global extralinguistic properties such as rate may affect perception differently from other more frequently changing properties such as amplitude, and Newman & Sawusch, 1996, for evidence that speech-rate normalization in perception may not interact with phonotactics). If the modularity hypothesis is correct, speech errors should be insensitive to a second-order speech-rate contingency, even over multi-day training.

Participants

Sixteen new participants from the same population as Experiment 1 were compensated similarly for their participation. As in Experiment 1, eight were randomly assigned to the /f-/s/ condition and the remaining eight were in the /k-/g/ condition.

Materials

The materials were similar to Experiment 1 with the exception that each sequence contained four CVC syllables. In all syllables, the vowel was // . The eight consonants were distributed within each sequence as in Experiment 1. In the /f-/s/ condition, half of the participants received sequences obeying the following constraint: *if the speech rate is fast, /f/ is an onset and /s/ is a coda but if the speech rate is slow, /s/ is an onset and /f/ is a coda*. The other half received the opposite assignments of /f/ and /s/. For the eight participants in the /k-/g/ condition, half experienced the constraint: *if the speech rate is fast, /k/ is an onset and /g/ is a coda but if the speech rate is slow, /g/ is an onset and /k/ is a coda*. The remaining participants received /k/ and /g/ in the opposite position assignments.

The lists of 96 sequences for each session were generated and printed as in Experiment 1 with the exception that the word “fast” or “slow” appeared to the left of each sequence, serving as an indicator of the tempo for a particular sequence. The sequences alternated between fast and

slow throughout a given testing session. Thus, the first two sequences for a session in the f-s condition might be:

<i>fast</i>	hes	fek	meng	gen
<i>slow</i>	mek	seng	heg	nef

Procedure

The procedure was the same as Experiment 1, except that the three rapid recitations of each “fast” and “slow” sequence were done at 2.67 and 1.87 syllables/second, respectively.

Results & Discussion

Errors were scored as in Experiment 1. Testing for reliability of the error transcription yielded an overall agreement rate of 98.9% and a conditionalized agreement (based on 122 errors) of 82%. As before, when participants made an error involving /h/ or /ŋ/, it was always legal (mean = 100%, SE = 0, based on 587 errors). Again, the focus was on planned contrasts involving experiment-restricted and unrestricted consonant slips (Table 2). Unlike Experiment 1, there was no evidence of learning. There was no significant difference between the legality percentages of restricted and unrestricted errors for either Day 1 (Wilcoxon $Z = -1.48$, $p = .14$) or Days 2–4 (Wilcoxon $Z = -.21$, $p = .84$), and no interaction between day of testing and the effect of restrictedness (Wilcoxon $Z = 1.60$, $p = .11$). On the critical Days 2–4, the effect is slightly in the wrong direction. Broken down by speech rate, the restricted – unrestricted contrast was 1.8% for errors at the fast rate and –3.1% (wrong direction) for errors at the slow rate.

Two other aspects of the data are noteworthy. First, there were many more errors in the fast condition (1376) than the slow condition (730), attesting to the power and salience of the speech-rate manipulation. It is thus unlikely that the rate difference was too subtle to guide learning if it could. Second, there were fewer errors in Experiment 2 (2106 total errors) than Experiment 1 (5460 total errors). Our explanation for this is that the slow rate is considerably slower than the constant rate used in Experiment 1. Thus, there are fewer errors on slow trials and, on fast trials, perhaps participants benefited from the “break” they received on the intervening slow trial. A consequence of fewer errors is greater variability in legality percentages (compare SE’s in Tables 1 and 2). Because of this greater variability, we cannot rule out that there is a small effect that we did not detect.² Nor can we rule out that additional days of testing would not demonstrate an effect, although the average legality percentages for restricted and unrestricted on Day 4 were 88.9 (SE = 3.15) and 88.7 (SE = 2.89), respectively. We simply found no evidence that restricted slips were more legal than unrestricted ones. As these results are consistent with Onishi et al.’s (2002) findings and the modularity hypothesis, we conclude that participants do not implicitly learn constraints where an extralinguistic feature such as speech rate determines consonant position.³

²The power to detect a restricted-unrestricted difference in the predicted direction on Days 2–4 whose size equals the mean of that obtained in the three previous second-order experiments with CVC stimuli (Dell et al., 2000; Warker & Dell, 2006) is .99. The power remains high (.88) even if the true effect size is equal to the smallest of the three obtained effects. This assumes an estimated standard error of the difference from the current experiments. Please note that these are power calculations based on the parametric analogue to the Wilcoxon test (paired t-test), and so the applicability of this power determination to our situation is approximate.

³To ensure that phonotactic-like learning could occur in the face of variation in speech rate, we tested four new participants on a first-order constraint (e.g. /k/ is onset; /g/ is coda for the entire experiment, with two participants getting the reverse constraints) where the trials alternated between the fast and slow tempos. Robust learning was demonstrated, similar in magnitude to other studies (93.6% legal for restricted consonants, 80.2% legal for unrestricted consonants). The effects were sufficiently strong that they were significant for 3 of the 4 participants, individually. (The p-values for the restricted-unrestricted difference as determined by Fisher’s exact tests for the four participants were .046, .001, 1.0, and .042) These results suggest that phonotactic-like learning is possible in this paradigm and that the null result of Experiment 2 was not due to the difficulty of the paradigm.

General Discussion

Together, the results from these two experiments suggest that there are limits to the types of experimental constraints we can implicitly learn. Experiment 1 found that participants learned a phonotactic-like nonadjacent constraint, but this learning did not appear until the second day of testing. This delayed learning is consistent with other studies of second-order constraint learning in production (Warker & Dell, 2006). Warker and Dell proposed a three-layer feed forward connectionist model of phonotactic constraint learning (see Figure 2). Like the experimental participants, the model takes longer (almost 10 times as long) to learn a second-order constraint compared to a first-order constraint. In second-order constraint conditions, the model receives conflicting input – sometimes /f/ is an onset and sometimes a coda. It requires hidden units and more trials to overcome this conflict to learn artificial contingencies between phonological elements. Critically, the model learns nonadjacent contingencies as easily as adjacent ones because its input *simultaneously* represents all the segments of each string, and its output is the binding of these to positions. Hence it does not matter what is next to what. This simultaneity assumption is clearly extreme. We note that the mean restricted-unrestricted difference over Days 2–4 for second-order *adjacent* effects is .14 (Warker & Dell, 2006; Exp 1a-b), whereas the mean of the nonadjacent ones is .07 (Experiment 1 & Warker & Dell, 2006, Exp 2a-b). Thus, the learning of nonadjacent constraints is less powerfully expressed in speech errors, suggesting they are more weakly learned. The model currently does not explain these differences.

Experiment 2 found that over the course of four days, participants were unable to learn a phonotactic-like constraint that hinged on speech rate. Given the learning patterns from previous studies, we expected learning to appear on the second day of testing. However, this was not the case; in fact, there was no evidence of learning on any testing days. This supports previous findings on perceptual constraint learning where participants were unable to learn a second-order constraint that depended on whether a male or female voice was speaking (Onishi et al., 2002). The lack of learning in both studies implies that information about some extralinguistic features (speech rate and speaker voice) was not coindexed with syllable content in the representations that underlie the implicit learning effect in these experiments.

Like Onishi et al. (2002), we propose that the inability to implicitly learn constraints dependent on extralinguistic features stems from these features being represented outside of the phonological system. Specifically, we claim that the weight-change mechanism that is the basis for implicit sequence learning in our studies has, by virtue of the modularity of the system, no conjunction detectors (e.g. hidden units) for co-occurrences of linguistic and nonlinguistic features (see Figure 3). The system is perfectly capable of detecting a variety of phonologically-internal co-occurrences such as absolute first-order effects (Dell et al., 2000; Taylor & Houghton, 2005; Onishi et al., 2002), graded first-order effects (Goldrick & Larson, in press), associations of phonological features to syllable positions (Goldrick, 2004), and second-order adjacent (e.g. Warker & Dell, 2006; Onishi et al., 2002) and nonadjacent (Experiment 1) effects between segments. However, this powerful learning mechanism is relatively impotent when it must link linguistic and nonlinguistic properties, and particularly nonlinguistic properties that are globally present for large chunks of the speech stream, such as speaker identity or speech rate. (Onishi et al., 2002: Experiment 2).

We have said that the learning in our study is implicit. The learning results from performance in a speech-production task, and is expressed without conscious retrieval of prior study episodes or awareness of the constraints. The implicit nature of the learning is supported by the fact that it is expressed in errors (and we do not intend to err) and that, when participants are explicitly told what the constraints are, it has no effect on their errors (Dell et al., 2000; Warker & Dell, 2006). Thus, the limitations that we demonstrate here are specific to this implicit learning

mechanism. We have no doubt that an extralinguistic dependency concerning phoneme sequencing could be learned through an explicit mechanism when attention is drawn to the dependency. For example, if we ask participants to look for a relationship between where the “f” is and whether they are speaking quickly or slowly, they should be able to “learn” the constraint, at least to the extent that they can state and remember it. However, because explicit knowledge of constraints does not affect slips, we would claim that such “learning” is irrelevant to the adaptability of the language processing system.

To conclude, two experiments tested how the speech-production system changes with experience. For one kind of constraint—a dependency between nonadjacent consonants—the system exhibited sensitivity. For another—a dependency between speech rate and consonant position—there was no evidence of learning. When it comes to implicitly acquiring phonotactic constraints from experience, dependencies between elements within the phonological system are readily learned.

Acknowledgments

This research was supported by HD-44458 and DC-00191. We thank Cynthia Fisher, Jennifer Cole, and Hahn Koo for helpful discussions. We also thank Soondo Baek for performing acoustic analyses.

References

- Chambers KE, Onishi KH, Fisher C. Infants learn phonotactic regularities from brief auditory experience. *Cognition* 2003;87:B69–B77. [PubMed: 12590043]
- Creel SC, Aslin RN, Tanenhaus MK. Heeding the voice of experience: The role of talker variation in lexical access. *Cognition* 2008;106:633–664. [PubMed: 17507006]
- Dell GS. A spreading activation theory of retrieval in language production. *Psychological Review* 1986;93:283–321. [PubMed: 3749399]
- Dell GS, Reed KD, Adams DR, Meyer AS. Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2000;26:1355–1367.
- Fromkin VA. The non-anomalous nature of anomalous utterances. *Language* 1971;47:27–52.
- Goldrick M. Phonological features and phonotactic constraints in speech production. *Journal of Memory and Language* 2004;51(4):583–604.
- Goldrick, M.; Larson, M. Constraints on the acquisition of variation. In: Fougeron, C.; D’Imperio, M., editors. *Papers in Laboratory Phonology 10: Variation, Detail and Representation*. Berlin: Mouton de Gruyter; (in press)
- Goldinger SD. Echoes of echoes? An episodic theory of lexical access. *Psychological Review* 1998;105:251–279. [PubMed: 9577239]
- Goldstein L, Pouplier M, Chen L, Saltzman E, Byrd D. Dynamic action units slip in speech production errors. *Cognition* 2007;103:386–412. [PubMed: 16822494]
- Hansson, GO. Theoretical and typological issues in consonant harmony. Department of linguistics, University of California; Berkeley: 2001.
- Hay J, Warren P, Drager K. Factors influencing speech perception in the context of a merger-in-progress. *Journal of Phonetics* 2006;34:458–484.
- Jusczyk PW, Friederici AD, Wessels JMI, Svenkerud VY, Jusczyk AM. Infants’ sensitivity to the sound patterns of native language words. *Journal of Memory and Language* 1993;32(3):402–420.
- Koo H, Cole J. On learnability and naturalness as constraints on phonological grammar. *ExLing-2006* 2006:165–168.
- MacKay DG. The problems of flexibility, fluency, and speed-accuracy trade-off in skilled behavior. *Psychological Review* 1982;89:483–506.
- Newman RS, Sawusch JR. Perceptual normalization for speaking rate: Effects of temporal distance. *Perception and Psychophysics* 1996;58:540–560. [PubMed: 8934686]

- Newport E, Aslin R. Learning at a distance: I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology* 2004;48:127–162. [PubMed: 14732409]
- Onishi KH, Chambers KE, Fisher C. Learning phonotactic constraints from brief auditory experience. *Cognition* 2002;83:B13–B23. [PubMed: 11814489]
- Rose S, Walker R. A typology of consonant agreement as correspondence. *Language* 2004;80:475–531.
- Sommers MS, Barcroft J. Stimulus variability and the phonetic relevance hypothesis: Effects of variability in speaking style, fundamental frequency, and speaking rate on spoken word identification. *Journal of Acoustical Society of America* 2006;119:2406–2416.
- Taylor CF, Houghton G. Learning artificial phonotactic constraints – time course, durability, and relationship to natural constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2005;31:1398–1416.
- Warker JA, Dell GS. Speech errors reflect newly learned phonotactic constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2006;32:387–398.

Figure 1. Differences in legality percentages between slips of restricted and unrestricted consonants from four experiments in Warker and Dell (2006). * indicates a significant difference.

Syllabification Model

Figure 2.

Figure 3.
No conjunction detectors (i.e. hidden units) between phonological system and extralinguistic features.

Table 1

Percentage of legal errors involving unrestricted and restricted consonants in Experiment 1.

	Unrestricted	SE	Restricted	SE
Day 1	81.3	1.82	77.9	3.03
Days 2–4	86.1	1.52	92.9	1.4
Total Errors	2699		1433	

Table 2

Percentage of legal errors involving unrestricted and restricted consonants in Experiment 2.

	Unrestricted	SE	Restricted	SE
Day 1	81.7	3.47	72.6	5.21
Days 2–4	84.6	1.83	84.1	3.28
Total Errors	982		537	