

The Role of Phonetic Naturalness in Phonological Rule Acquisition

Sharon Peperkamp^{1,2}, Katrin Skoruppa¹, and Emmanuel Dupoux¹

¹Laboratoire de Sciences Cognitives et Psycholinguistique,
EHESS-ENS-CNRS, Paris

²Université de Paris 8

1. Introduction

The speed and reliability of language acquisition in human infants remain puzzling facts, especially given the huge variability in the quality and quantity of the speech input. It has been argued that such a robust learning path is due to a specialized Language Acquisition Device (Chomsky 1965), which incorporates statistical mechanisms to extract regular patterns from noisy input (e.g. Saffran, Aslin & Newport 1996), as well as innate constraints that restrict learning to patterns that conform to universal grammar (Chomsky 1981). In the domain of phonological acquisition, discussion has centered around the role of *phonetic naturalness* as a possible learning constraint. That is, phonological rules and patterns tend to respect a number of constraints that are grounded in phonetic factors and that might guide the language learner (Archangeli & Pulleyblank 1994; Hayes 1999). Some experimental work with infants has been carried out that examines the role of phonetic naturalness in acquisition, in particular that of surface phonological regularities (Saffran & Thiessen 2003; Seidl & Buckley 2005). In this article, we focus on the acquisition of phonological rules. We begin by illustrating the presence of naturalness constraints with rules of the type shown in (1).

(1) /a/ → [b] / _[c₁,c₂ .. c_n]

First, the change affecting the phoneme /a/ tends to result in a segment [b] that is phonetically close to it, that is, the change only concerns a small number of distinctive features (*Phonetic Proximity*). Second, the set of triggers [c₁...c_n] is relevant to the change in that it tends to be homogeneous with respect to the feature(s) that undergo(es) the change (*Contextual Relevance*); that is, the target becomes either more similar or more dissimilar to all the targets. Third, the application of the rule reduces the markedness of the surface structure (*Markedness Reduction*).¹ Peperkamp et al. (in press) proposed a statistical

¹ This is not an exhaustive list of constraints. For instance, when rules apply to several target segments, these segments tend to form a natural class, sharing one or more articulatory features. Likewise, the set of context segments tends to

algorithm for the acquisition of allophonic rules of type (1) based on the distribution of segments in the phonetic stream. This algorithm works well with artificial languages but fails with natural languages, unless it is supplemented with at least some of the naturalness constraints described above. Yet, empirical evidence regarding the use of naturalness constraints during language acquisition is scant and inconsistent. We review three studies, all involving artificial language learning by adults.

First, Shane et al. (1974/75) compared the learning of a natural rule that deletes word-final consonants before consonant-initial words (resulting in the unmarked CV-CV pattern) to that of an unnatural one that deletes word-final consonants before *vowel*-initial words (resulting in the marked CVC-CV and CV-V patterns). Native speakers of English were trained to produce adjective-noun combinations in an artificial language, where the adjectives appeared in two forms, with or without a final consonant. Subjects were asked to produce combinations of trained adjectives followed by both trained and untrained nouns. It was found that learning was more rapid in the group who learned the natural deletion rule than in the group who learned the unnatural one. However, by the end of the experiment subjects had learned both the natural and the unnatural rule to the same extent.

Second, Pycha et al. (2003) taught English subjects to derive the plural form of singular nouns in an artificial language, where the singular was formed by a bare stem and the plural by the stem plus a suffix. This plural suffix appeared in two phonetic forms: In one test group (*Harmony*), it exhibited palatal vowel harmony with the preceding stem vowel, in another group (*Disharmony*), it exhibited palatal disharmony, and in a third group (*Arbitrary*), it exhibited harmony with three of the stem vowels and disharmony with the remaining three. After an initial exposure phase and a supervised learning phase, subjects performed a forced choice grammaticality judgment on novel items. Subjects in the Harmony and Disharmony groups performed well and did not differ from one another. Subjects in the Arbitrary group were at chance.²

Finally, Peperkamp & Dupoux (in press) (henceforth: P&D) exposed French subjects to an artificial language consisting of determiner + noun phrases which obey either a natural allophonic rule that voices a subclass of obstruents intervocalically, or an unnatural one that defines arbitrary relationships among certain obstruents intervocalically (for instance, [ʒ] is an allophone of /p/). All phrases were accompanied by a referential picture that allowed subjects to infer the meaning of the nouns. There were two determiners, one ending in a vowel and one in a consonant, and the crucial alternations concerned the initial

form a natural class.

² In a comparable study, Wilson (2003) used artificial languages with nasal assimilation, nasal dissimilation, or a combination of both. In these experiments, however, subjects were not exposed to the alternations but only to the resulting surface patterns; they were tested on novel words either obeying these patterns or not. The results were very similar to those of Pycha et al. (2003).

consonant of the nouns. After the exposure phase, a phrase-picture matching paradigm was used to test if subjects had learned the allophonic distributions and hence distinguished between phonemic and allophonic contrasts among obstruents for the purposes of word recognition. It was found that regardless of the language of exposure, subjects took a change in the initial obstruent of nouns to be lexically irrelevant more often for the obstruents that were subject to the alternation than for those that were not; this difference held for both old and novel nouns (see Figure 1). Hence, natural assimilatory rules and unnatural arbitrary rules were learned with equal ease.³

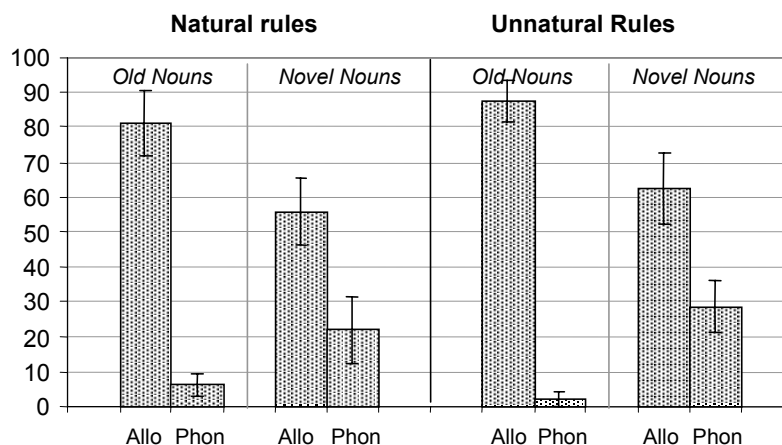


Figure 1. Mean percentages of matching of a phrase presented with a change in the initial obstruent of the noun to an image representing the unchanged noun. Data from P&D.

The divergence among these results could be due to differences in the naturalness constraints that were involved and/or in the experimental paradigms. First of all, the natural and unnatural rules tested by Shane et al. varied in Markedness Reduction only. The fact that Schane et al. found better learning for the natural rule shows the relevance of this constraint for rule learning. The harmony and disharmony rules of Pycha et al., however, likewise differed in Markedness Reduction only, but subjects learned these rules equally well. Intuitively, the difference in markedness between harmonic and disharmonic words seems relatively small, especially when compared to that between the syllable structures resulting from Schane et al.'s natural and unnatural rules. Note also that the unnatural rule of Schane et al. created two different marked

³ Note that subjects did not generalize the natural assimilation rule to a new untrained segment within the same natural class. That is, the learned allophonic pattern for /p/ and /k/ (or, alternatively for /f/ and /ʃ/) did not transfer to the untrained segment /t/ (or /s/).

structures, one with a consonant cluster, the other one with vowel hiatus. This might explain why Markedness Reduction had an effect in the experiment of Shane et al. only. Next, Pycha et al.'s arbitrary rule crucially violated Contextual Relevance. The fact that the arbitrary rule was not learned thus shows the importance of this constraint. Finally, the natural and unnatural rules of P&D differed on all three constraints. The absence of an effect of naturalness in this study is, therefore, rather surprising, but might be due to a ceiling effect. Recall that Shane et al. reported a significant difference only at the beginning of the experiment, suggesting that with enough exposure, subjects can learn rules that enhance markedness. Perhaps the exposure in P&D's experiment was long enough for both the natural and the unnatural rules to be learned, even though at different speeds. Alternatively, it might be the case that a forced choice phrase-picture matching task is in some way too easy or too explicit in this paradigm, in the sense that it allows subjects to use non-linguistic strategies and learn even unnatural rules. In the present experiment, we test the latter possibility. We thus use a more natural task, i.e. picture naming, while leaving the stimuli and exposure identical.

2. Experiment

As in P&D, we use four artificial languages, two containing a natural allophonic rule (Nat_A and Nat_B) and two an unnatural one (Unnat_A and Unnat_B). All languages share the same segmental repertoire but not the same set of underlying phonemes. In particular, there are 12 obstruents, of which only nine are phonemic in each language, the remaining three being contextual allophones (Table 1).

Table 1. Surface obstruent consonants and their underlying phonemes in the four languages.

	stops			fricatives		
	voiceless	voiced		voiceless	voiced	
surface segments	p t k	b d g	f s ʃ	v z ʒ		
underlying phonemes						
Nat_A	p t k	b d g	f s ʃ			
Nat_B	p t k		f s ʃ	v z ʒ		
Unnat_A	p k	b d g	s ʃ	v	ʒ	
Unnat_B	p t	d g	f ʃ	v z	ʒ	

In Language Nat_A, the allophones are introduced by a rule of intervocalic fricative voicing (2a), in Language Nat_B by a rule of intervocalic stop voicing (2b). In Languages Unnat_A and Unnat_B, the allophones also appear intervocalically, but they are the results of arbitrary mappings (2cd).

- (2) a. Nat_A: /f,s,ʃ/ → [v,z,ʒ] / V _ V
 b. Nat_B: /p,t,k/ → [b,d,g] / V _ V
 c. Unnat_A: /p,g,z/ → [ʒ,f,t] / V _ V
 d. Unnat_B: /ʃ,v,d/ → [b,k,s] / V _ V

Subjects are exposed to short phrases accompanied by referential pictures, in one of the four languages. These phrases are of the type determiner + noun, where the determiner is either *nel*, meaning ‘two’ or *ra*, meaning ‘three’, and the noun starts with either a stop or a fricative followed by a vowel. Crucially, the determiner *ra* but not *nel* creates the context for the allophonic rules. Certain nouns thus appear in two phonetic forms, depending on the determiner. After the exposure phase, a picture naming task is used to test if subjects have learned the allophonic distributions. Specifically, if learning is successful, we expect that they produce different noun-initial obstruents as a function of the preceding determiner if and only if these obstruents are involved in an alternation in the language of exposure.

2.1 Method

2.1.1 Stimuli

We use the same stimuli as the sets in Experiment 1 (natural rules) and Experiment 3 (unnatural rules) of P&D. We describe the stimuli for the languages with the natural rules only; for the languages with unnatural rules there is a matched set. There are 12 nouns for the exposure phase and 78 additional nouns for the test phase (of which 30 are used for filler trials). The nouns for the exposure phase that begin with a non-dental obstruent are embedded in three phrases, such that for both languages there is one grammatical phrase with *nel* and one with *ra* (e.g. Nat_A and Nat_B: [nel pemuʃ], Nat_A: [ra pemuʃ], Nat_B: [ra bemuʃ]); the nouns that begin with a dental obstruent, by contrast, are embedded in only one phrase, that is grammatical in both languages (e.g. [nel timu] and [ra daru]). The nouns for the test phase are likewise embedded in only one phrase, regardless of their initial consonant. All phrases are accompanied by a referential picture of two or three identical objects.

2.1.2 Procedure

At the beginning of the experiment, subjects were told that they would be studying short phrases of the type ‘two dolls’ or ‘three boats’ in an unknown language, in which *nel* means ‘two’ and *ra* means ‘three’, and that they were to memorize the nouns in this language.

The exposure phase was identical to that of Experiments 1 and 3 of P&D. That is, subjects listened to 20 phrases in one of the four languages,

accompanied by their referential pictures. Table 2 and Table 3 show the exposure for the natural and the unnatural languages, respectively.

Table 2. Phrases used for the exposure phase in the two languages with a natural rule.

	Nat_A: allophonic fricative voicing		Nat_B: allophonic stop voicing	
'rabbit'	nel pemuʃ	ka pemuʃ	nel pemuʃ	ka bemuʃ
'flower'	nel bovi	ka bovi	nel povi	ka bovi
'apple'	nel kelaf	ka kelaf	nel kelaf	ka gelaf
'fork'	nel ginɛl	ka ginɛl	nel kinɛl	ka ginɛl
'hat'	nel timu		nel timu	
'tree'		ka daʁu		ka daʁu
'cat'	nel foʒam	ka voʒam	nel foʒam	ka foʒam
'nose'	nel fulek	ka vulek	nel vulek	ka vulek
'bottle'	nel ʃagip	ka ʒagip	nel ʃagip	ka ʃagip
'house'	nel ʃubo	ka ʒubo	nel ʒubo	ka ʒubo
'balloon'	nel sano		nel sano	
'snail'		ka zelum		ka zelum

Table 3. Phrases used for the exposure phase in the two languages with an unnatural rule.

	Unnat_A: /p,g,z/ → [ʒ,f,t] / V_V		Unnat_B: /ʃ,v,d/ → [b,k,s] / V_V	
'rabbit'	nel pemuʃ	ka ʒemuʃ	nel pemuʃ	ka pemuʃ
'flower'	nel bomi	ka bomi	nel ʃomi	ka bomi
'apple'	nel kela	ka kela	nel vela	ka kela
'fork'	nel girel	ka firel	nel girel	ka girel
'balloon'	nel doba		nel doba	
'tree'		ka tiʁuk		ka tiʁuk
'cat'	nel goʒa	ka foʒa	nel foʒa	ka foʒa
'nose'	nel vusen	ka vusen	nel vusen	ka kusen
'bottle'	nel ʃanip	ka ʃanip	nel ʃanip	ka banip
'house'	nel puko	ka ʒuko	nel ʒuko	ka ʒuko
'hat'	nel zifu		nel zifu	
'snail'		ka setum		ka setum

Note that for nouns starting with a dental obstruent, only one of the two phrases occurred, such that subjects did not receive evidence as to whether these segments participated in an alternation. For the languages with a natural rule, this allowed us to test if subjects would generalize what they learned about the status of voicing in labial and velar obstruents to dental obstruents. By contrast, the presence of the four single phrases containing a dental-initial noun in the languages with the unnatural rules only ensures that the exposure is comparable to that for the languages with a natural rule. Indeed, the alternations /z/ → [t] (Unnat_A) and /d/ → [s] (Unnat_B) are impossible to infer, since they do not represent a generalization within a natural class of one of the alternations that are present overtly. Hence, we expect that subjects treat both the [z-t] and the [d-s] contrasts in the Untrained condition as phonemic; our classification of these contrasts as phonemic or allophonic in the two languages is arbitrary to begin with.

Phrases in which the noun started with a labial, palatal, or velar obstruent were presented 16 times each, and those in which the noun started with a dental obstruent were presented 8 times each. All phrases were presented in a semi-random order with an SOA of 3000ms., for a total of around twenty minutes.

The test phase consisted of two parts and was identical for all subjects. In the first part, subjects were tested on their lexical representations of the nouns they had learned during the exposure phase, and during the second part, they were tested on their lexical representations of novel nouns. The first part consisted of 12 test trials (one for each noun of the exposure) and 6 filler trials. The second part consisted of 48 different test trials and 24 different filler trials; during this part, all phrase-picture pairings were novel. Half of the nouns started with a stop (1/3 labial, 1/3 dental, 1/3 velar) and the other half with a fricative (1/3 labial, 1/3 dental, 1/3 palatal).

Subjects exposed to Languages Nat_A and Nat_B received the same test trials; likewise for subjects exposed to Languages Unnat_A and Unnat_B. We describe the test trials for the languages with the natural rules only, the structure of the ones for the languages with the unnatural rules being identical. A test trial started with the presentation of a probe for 2000ms, consisting of a phrase-picture pairing that was identical in both languages. After a pause of another 2000ms, a picture containing the same object as in the probe was presented, now appearing in a different number. The subjects' task was to produce the phrase describing the new picture. For instance, during the first test phase, they would be probed with the phrase [nel pemuʃ] and the picture of two rabbits, and then see the picture of three rabbits, for which they had to produce the corresponding phrase (i.e., according to the language of exposure, [ʁa pemuʃ] or [ʁa bemuʃ]). Likewise, there would be a probe [ʁa bovi] with the picture of three flowers, and subjects would have to produce the phrase describing the picture of two flowers (i.e. [nel bovi] or [nel povi]). Finally, filler trials differed from test trials in that the test picture was strictly identical to the probe picture.

According to the language of exposure and the initial consonant of the noun, a test trial belonged to either the Phonemic or the Allophonic condition. For instance, for subjects exposed to Language Nat_A (with allophonic fricative voicing), trials with a stop-initial noun belonged to the Phonemic and trials with a fricative-initial noun to the Allophonic condition; the reverse for subjects exposed to Language Nat_B. Trials in which the noun started with a labial, palatal, or velar obstruent belonged to the Trained condition; during exposure, these nouns had been presented both in the context of *nel* and in the context of *ra*. Trials in which the noun started with a dental obstruent belonged to the Untrained condition; during exposure, these nouns had only been presented either in the context of *nel* or in the context of *ra* (see Tables 2 and 3).

The entire experiment lasted around 35 minutes.

2.1.3 Participants

Thirty-two native French speakers with normal hearing and normal or corrected-to-normal vision, aged between 17 and 31 years (mean: 22), participated. They were randomly distributed to one of four groups, with each of the groups being exposed to a different language.

2.2 Results and discussion

The subjects' productions were transcribed by a native French speaker who was unaware of the purposes of the experiment. We discarded responses with one of the following characteristics: incorrect determiner, hesitation pause between determiner and noun, unintelligible, or missing noun (total: 0.8%).

We analyzed the results for the natural and the unnatural rules separately. For the natural rules, we first analyzed the responses in the Trained condition (see Figure 2, left part). The mean percentages of responses with a change (whether following the rule or not) were subject to two separate ANOVAs, one for the old and one for the novel nouns, with between-subject factor Language (Nat_A vs. Nat_B) and within-subject factor Contrast (Phonemic vs. Allophonic). For both the old and the novel nouns, the effect of Contrast was significant (old nouns: $F(1,14) = 18,7$ $p < .001$; novel nouns: $F(1,14) = 5,3$ $p < .038$). No other main effect or interaction reached significance.

In the Untrained condition, responses with a change in the noun-initial consonant were overall very rare: there was none for the old nouns, while for the novel nouns, they occurred in only 3,9% and 1,6% in the Allophonic and the Phonemic condition, respectively. The difference between the means in the Allophonic and the Phonemic conditions for the novel nouns was marginally significant ($F(1,14) = 3.32$, $p = .09$).

For the unnatural rules, the mean percentages of responses with a change in the Trained condition are shown in the right part of Figure 2. For the old nouns, an ANOVA yielded significant main effects of Language ($F(1,14) = 5.7$, $p < .031$, with Language Unnat_B giving rise to more responses with a change

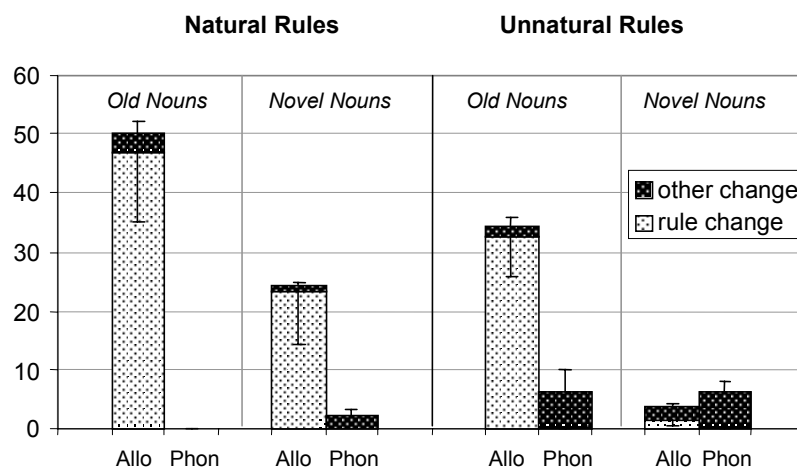


Figure 2: Mean percentages of productions with a segmental change in the Trained condition for natural and unnatural rules. The percentages of changes that follow the rule ('rule change') are shown in light gray, those that do not ('other change') in dark gray; the upward error bar concerns the former and the downward one the latter type of change. (In the absence of a rule in the Phoneme condition, all changes are considered as 'other change' in this condition.)

than Language Unnat_A) and Contrast ($F(1,14) = 10.2, p < .006$). For the novel nouns, there was a significant interaction between Language and Contrast ($F(1,14) = 16.7, p < .001$). No other main effect or interaction reached significance. Two restricted analyses showed that the interaction between Language and Contrast for the novel nouns was due to the fact that in Language Unnat_A, there were significant effects of Contrast ($F(1,7) = 6.4, p < .039$, with more responses with a change been given in the *Phonemic* condition) and Segment ($F(1,7) = 9.2, p < .019$), as well as a significant interaction between them ($F(1,7) = 10.6, p < .014$), whereas no main effect of interaction was significant in Language Unnat_B.

In the Untrained condition, responses with a change in the noun-initial consonant were rare. For the old nouns, they occurred in 3,1% of the cases in both the Allophonic and the Phonemic condition. For the novel nouns, they occurred in 2,5% and 1,6% of the cases in the Allophonic and the Phonemic condition, respectively, representing a non-significant difference.

Finally, we compared the results in the Trained conditions of the natural and the unnatural rules. We thus submitted the mean percentages of responses with a change to two ANOVAs, one for the old and one for the novel nouns, with the between-subject factor Naturalness (Nat_A/B vs. Unnat_A/B) and the within-

subject factor Contrast. For the old nouns, there was a significant effect of Contrast ($F(1,30) = 29.4, p < .0001$). For the novel nouns, there were marginally significant effects of Naturalness ($F(1,30) = 3.0, p < .094$) and Contrast ($F(1,30) = 3.9, p < .057$), and a significant interaction between these two factors ($F(1,30) = 6.4, p < .017$). Moreover, we submitted the mean percentages of ‘other’ changes to two one-way ANOVAs, one for the old and one for the novel nouns, with the factor Naturalness. There was a significant effect for the novel nouns only ($F(1,30) = 5.5, p < .026$).

For the natural rules, these results replicate with a picture naming task the results previously obtained by P&D with a phrase-picture matching task. That is, subjects are able to learn the context-dependant alteration in surface word form that is introduced by an allophonic rule in an artificial language; they generalize their responses to novel lexical items; and generalization does not extend to untrained segments, showing that the rule is learned on a segment-by-segment basis. For the unnatural rules, by contrast, the results are different. That is, whereas subjects learned variation in word forms arising from an unnatural rule for nouns that they have been trained on, this learning did *not* extend to novel nouns. Overall, the results of this experiment show a clear effect of naturalness with novel nouns: subjects learned to apply a natural but not an unnatural rule.

3. General discussion

Using a picture naming task, we obtained an unambiguous effect of naturalness for the acquisition of allophonic rules. This contrasts starkly with the P&D study, which failed to find any effect of naturalness with a phrase-picture matching task. What are the possible factors that could account for this difference? In order to answer this question, it is useful to look at a broader picture obtained by comparing these and the other studies reviewed in the introduction. Table 4 presents these studies, together with a characterization of the conditions that were tested, and the results.

Table 4. Summary of rule-learning studies.

Study (<i>task</i>)	Group / Condition	Phonetic Proximity	Contextual Relevance	Markedness Reduction	Learning outcome
Pycha et al. (<i>grammaticality judgment</i>)	<i>Harmony</i>	+	+	+	<i>good</i>
	<i>Disharmony</i>	+	+	-	<i>good</i>
	<i>Arbitrary</i>	+	-	0	<i>bad</i>
Schane et al. (<i>phrase production</i>)	<i>Natural</i>	+	+	+	<i>good</i>
	<i>Unnatural</i>	+	+	-	<i>1st part: bad</i> <i>2nd part: good</i>
P&D (<i>phrase-picture matching</i>)	<i>Natural</i>	+	+	+	<i>good</i>
	<i>Unnatural</i>	-	-	0	<i>good</i>
present study (<i>picture naming</i>)	<i>Natural</i>	+	+	+	<i>good</i>
	<i>Unnatural</i>	-	-	0	<i>bad</i>

The first part of the table presents a typology of the rules that were tested in the different studies, according to the three constraints listed in the introduction: Phonetic Proximity, Contextual Relevance, and Markedness Reduction. For the latter constraint, we use a three-way classification, with ‘+’ meaning that markedness is reduced, ‘-’ that it is enhanced, and ‘0’ that it is reduced in some cases while enhanced in others, depending upon the trigger. It should be noted that these three constraints neither form an exhaustive partition of naturalness (see footnote 1), nor are fully independent from one another. For instance, when Contextual Relevance is not respected, Markedness Reduction cannot be maximal. The second part of the table presents the results obtained in the various studies. Two aspects are worth noting. First, performances on unnatural rules are systematically worse than or equal to performances on natural rules. Second, a given contrast in naturalness does not always yield a significant difference in performance; different tasks seem to be more sensitive than others.

One possible global account of these different patterns is that given enough training, all rules can eventually be learned, be they natural or unnatural. Under this view, naturalness does not affect learnability per se, but rather the speed at which learning takes place: the more natural the rule, the faster it is learned. In a given experiment, then, participants might reach ceiling performance for some types of rules (the more natural ones), but not for others (the less natural ones), flattening part of the continuum. An alternative possibility is that the difference between natural and unnatural rules lies with the ease with which they are put to use once they have been learned. In this case, more demanding tasks would result in stronger effects of naturalness than less demanding tasks.⁴

Could speed of learning account for the contrasting results of P&D and the present experiments? Given that across the experiments, not only the stimuli and conditions were the same but so was the training phase, this could only be the case if subjects were still learning during the *test* phase. One might argue that learning especially continues during the first test phase, containing the old items. Indeed, old items may benefit from memory traces, which could be exploited by subjects to extract the rule as they are responding. Subjects would thus reach ceiling faster in the phrase-picture matching task than in the naming task, because the former is easier than the latter, leaving more resources for learning the rule on the fly. However, in an unpublished experiment with the picture-matching task, we simply removed the first test phase on old items and found that although the overall amount of learning was reduced, there was still no difference between the performances on the natural and the unnatural rules.

The alternative account, based on rule difficulty, would perhaps fare better in accounting for the difference between our phrase-picture matching task and our picture naming task. Since picture naming is a rather demanding task (subjects have to generate a response that is fully specified in terms of all the

⁴ Both accounts thus allow for the ultimate acquisition of unnatural rules, in accordance with the empirical observation that such rules exist (see, for instance, Anderson 1981).

segments), they might not have the extra processing capacity needed to simply formulate and execute the costly unnatural rules. In contrast, in phrase-picture matching, subjects only have to choose between two alternatives. This is not a complex task, and therefore subjects could use the available resources to implement the complicated unnatural rules as well as the natural ones.

These two accounts do not exhaust the theoretical possibilities and are rather post-hoc in nature. An understanding of the task parameters as well as a better control of the processing load is needed to fully account for the pattern of data found thus far. Over and above these methodological problems, however, it seems quite clear that naturalness plays a role in phonological rule acquisition. Further work is needed in two directions. First, it is important to do more theoretical work to define the decomposition of phonetic naturalness more precisely. How many different constraints are there? What are their relative strengths? Second, it is important to do more experimental work to understand how different tasks are influenced by these different linguistic constraints, as opposed to by performance factors involving task difficulty and response strategies.

References

- Anderson, S. (1981) Why phonology isn't "natural". *Linguistic Inquiry* 12, 493-539.
- Chomsky, N. (1965) *Aspects of the Theory of Syntax*. Cambridge, MA, MIT Press.
- Chomsky, N. (1981) *Lectures on Government and Binding*. Dordrecht: Foris.
- Hayes, B. (1999) Phonetically driven phonology: The role of Optimality Theory and inductive grounding. M. Darnell, E. Moravcsik, M. Noonan, F. Newmeyer, and K. Wheatly (eds). *Functionalism and Formalism in Linguistics 1: General Papers*. Amsterdam, John Benjamins: 243-285.
- Peperkamp, S. and E. Dupoux (in press) Learning the mapping from surface to underlying representations in an artificial language. J. Cole and J. Hualde (eds) *Laboratory Phonology IX*.
- Peperkamp, S., R. Le Calvez, J.-P. Nadal, et E. Dupoux (in press) The acquisition of allophonic rules: statistical learning with linguistic constraints. *Cognition*.
- Pycha, A., P. Nowak, E. Shin and R. Shosted (2003). Phonological rule-learning and its implications for a theory of vowel harmony. G. Garding and M. Tsujimura (eds) *Proceedings of WCCFL 22*. Somerville, Cascadilla Press: 423-435.
- Saffran, J. R., R. N. Aslin and E. L. Newport (1996). Statistical learning by 8-month-old infants. *Science* 274: 1926-1928.
- Saffran, J. and E. Thiessen (2003) Pattern induction by infant language learners. *Developmental Psychology* 39: 484-494.
- Schane, S., B. Tranel and H. Lane (1974/75). On the psychological reality of a natural rule of syllable structure. *Cognition* 3/4: 351-358.
- Seidl, A. and G. Buckley (2005). On the learning of arbitrary phonological rules. *Language Learning and Development* 1: 289-316.
- Wilson, C. (2003). Experimental investigation of phonological naturalness. G. Garding and M. Tsujimura (eds). *Proceedings of WCCFL 22*. Somerville, MA, Cascadilla Press: 533-546.