Experimental Investigation of Phonological Naturalness

Colin Wilson
UCLA

1. Introduction

Phonological theories and textbook introductions to phonology typically focus on a relatively small set of process types. In the domain of segmental phonology, this set includes deletion and insertion of segments in a restricted range of environments, processes that make one segment more similar (assimilation) or less similar (dissimilation) to another segment with respect to a single feature or group of features, and a few others. What is the motivation for focusing on a limited set processes, to the exclusion of others that are attested in natural languages?

Consider, for example, the two hypothetical patterns in (1). Both involve the distribution of two allomorphs ([–la] and [–na]) of a single suffix. In pattern (1a), the suffix surfaces as [–la] except when the preceding stem consonant is [+nasal], in which case the suffix is realized as [–na]. In pattern (1b), the default suffix allomorph is again [–la], but the property of the stem that conditions the non-default allomorph is different: the [–na] allomorph appears iff the preceding stem consonant is [+dorsal].

(1) a. Consonant harmony

dume	dumena	tuko	tukola	suto	sutola
binu	binuna	dige	digela	dabu	dabula

b. “Random” alternation

dume	dumela	tuko	tukona	suto	sutola
binu	binula	dige	digena	dabu	dabula

---

* For comments and questions on parts of the material presented here, thanks to Adam Albright, Eric Bakovic, William Badecker, Jonathan Barnes, Jill Beckman, Mary Beckman, Luigi Burzio, Lisa Davidson, Matt Goldrick, Bruce Hayes, Junko Ito, Michael McCloskey, Dominic Massaro, Armin Mester, Jaye Padgett, Joe Pater, David Perlmutter, Janet Pierrehumbert, Anne Pycha, Brenda Rapp, Carson Schütze, Paul Smolensky, Donca Steriade, Andrew Wedel, Kie Zuraw, participants in the MIT Workshop on Markedness and the Lexicon, and audiences at Johns Hopkins, UC Santa Cruz, and WCCFL. This research was supported by UCLA Academic Senate Grants to the author.

 WCCFL 22 Proceedings, ed. G. Garding and M. Tsujimura, pp. 101-114. Somerville, MA: Cascadilla Press.

Pattern (1a) can be described in terms of an assimilation (more specifically, consonant harmony) process: the initial consonant of the suffix assimilates in [nasal] to the immediately preceding stem consonant. Assimilation in general has been intensely studied within generative phonology (e.g., Chomsky & Halle 1968, Clements 1985, Halle 1995, Ní Chiosáin & Padgett 2001), and several studies have been devoted to consonant harmony in particular (e.g., Gafos 1996, Hansson 2001, Rose & Walker 2002).

In contrast, pattern (1b) cannot be described in terms of any canonical process type, and is "random" in the sense that there is no apparent formal or substantive relation between the alternation in the suffix and the property of the stem that conditions the alternation. The characterization of this pattern as random is not meant to imply that it could not be attested. Indeed, there is substantial evidence that random processes can exist in languages, and are sometimes productive (e.g., Buckley 2000, Cena 1978, Pierrehumbert 2002). It therefore seems likely that the analyst, and the language learner, could be confronted with this pattern, or one that is equally random. What then could justify devoting theoretical and pedagogical attention to patterns such as (1a) — for example, developing special rule notations or constraint types to account for assimilation processes — and placing less emphases on patterns such as (1b)?

There is a way of understanding Universal Grammar ('UG') that could justify this asymmetric treatment and that is not contradicted by the observation that both "natural" patterns such as (1a) and "random" patterns such as (1b) are attested.[1] UG is often characterized as a perfectly rigid system, one that allows certain grammars and disallows others. An alternative view is that UG also contains a set of *cognitive biases*: aspects of the system of mental representation and computation that make processes with certain properties easier to learn and/or represent, but which do not exclude processes that fail to have those properties.[2]

Under this conception, the fact that attention has been devoted to a relatively small set of process types amounts to a claim that UG is biased in favor of these types, although it may allow others. This claim is not falsified by the existence of random processes. What evidence could support it?

Only a convergence of evidence from many sources (typology, acquisition, psycholinguistics, etc.) will provide a clear picture of the cognitive biases that are relevant for phonology. In this paper, I contribute to the larger research program by introducing an experimental paradigm

1. I use the term "naturalness" in this paper somewhat loosely, without committing to any specific theory; see section 2.2 and section 3 for some discussion.
2. At least with respect to phonology, this conception of UG seems close to the one developed in Chomsky & Halle (1968). See in particular the discussions of grammar "complexity" (402ff.) and the "freedom" of phonological changes (426ff.)

that can be used to address questions about cognitive bias, and by presenting some initial results from that paradigm. The specific experiments reported here provide evidence for the claim that there are cognitive biases that favor the "natural" assimilation process in (1a), as well as a corresponding dissimilation process, over formally-matched "random" processes such as the one in (1b). In sections 2 and 3, I describe the experiments on assimilation and dissimilation, respectively. Particular attention is devoted to the experimental methodology and the sense(s) in which the assimilation and dissimilation processes investigated here are "natural". I conclude the paper in section 4 with some remarks on whether the cognitive biases revealed here are language-specific.

2. Experiment 1: Nasal assimilation vs. random alternation

The goal of Experiment 1 was to test for a cognitive bias that favors the assimilation (consonant harmony) pattern in (1a) over the formally-matched random pattern in (1b). This experiment and Experiment 2 employed a slightly motified version of the Artificial Grammar paradigm; the method and background citations are provided in section 2.1. An overview of the materials and critical manipulation for the experiment is given in (2).

(2) Design of Experiment 1

a. Group 1A: Nasal assimilation

m/n-stems	k/g-stems	other stems
dumena binuna (N=8)	tukola digela (N=8)	sutola dabula (N=24)
Rule: Suffix is [–na] iff the final stem consonant is [+nasal]		

b. Group 1B: Random rule

m/n-stems	k/g-stems	other stems
dumela binula (N=8)	tukona digena (N=8)	sutola dabula (N=24)
Rule: Suffix is [–na] iff the final stem consonant is [+dorsal]		

Participants assigned to Group 1A were exposed to a pattern that can be described with an assimilation rule. According to the rule, the initial consonant of the suffix assimilates on the feature [nasal] with the immediately preceding stem consonant. This rule is attested (with some variations) in several natural languages, including Lamba (Hyman 1995), Nyangumarda (Hoard & O'Grady 1976), and Ulithian (Sohn & Bender 1973: 59); see Hansson (2001:111–124) for a more complete typology.

Participants in Group 1B were exposed to a pattern in which there is no formal or substantive relationship between the allomorphy of the suffix and the conditioning feature of the stem. This pattern can be described with a random rule, one that ensures that the initial consonant of the suffix is [+nasal] iff the final consonant of the stem is [+dorsal]. Although this specific alternation seems to be unattested, Hansson (2001:90ff.) discusses a case of consonant harmony in Tlachichilco Tepehua that is in some cases made phonetically arbitrary by the application of another process.

In the remainder of this section, I refer to the processes illustrated by these examples as the *assimilation rule* and the *random rule*, respectively.[3] These two rules are formally matched, in the sense that the [nasal] specification of the suffix-initial consonant is determined in both cases by exactly one feature specification of the last stem consonant. Neither rule is part of the phonology of English, which contains no rules that assimilate a consonant or make a random feature change under the influence of another, non-adjacent consonant. And the rules give rise to sequences that are of approximately equal frequency (by both type and token) in the English lexicon. Table (3) shows that, for example, sequences of the type nVn, which are grammatical according to the assimilation rule, are not substantially more or less frequent in the English lexicon than sequences of the type nVl, which are grammatical according to the random rule.[4]

(3) Relative frequencies (= freq(XVn)/[freq(XVn) + freq(XVl)])

Sequences		type	token	$\log_{10}$(token)
mVn	mVl	.80	.82	.81
nVn	nVl	.40	.33	.39
kVn	kVl	.58	.69	.60
gVn	gVl	.44	.66	.45

Because the rules are matched in these ways, we can be reasonably confident that the observed differences between Group 1A and Group 1B are due to a universal bias that favors assimilation over random alternation.

2.1 Method

3. This terminology is used for expository convenience only, and does not imply a commitment to a rule-based phonological framewoork.
4. The relative frequencies in this table are based on counts in the 125,000 word CMU Pronouncing Dictionary (http://www.speech.cs.cmu.edu/cgi-bin/cmudict). Similar numbers were obtained from counts in the 160,000 wordform CELEX English corpus (Baayen et al. 1993). The suffixes *–ment* and *–man* account for a large proportion (approx. 1/2) of the mVn tokens; eliminating suffixed words from the frequency counts would therefore substantially reduce the relative frequencies.

Participants. 24 adult native American English speakers were recruited from the UCLA community to participate in Experiment 1. Each participant was randomly assigned to Group 1A (N=12) or Group 1B (N=12), and received a nominal fee or course credit for participating.

Stimuli. A female native American English speaker, trained in phonetics but naive to the purposes of the experiment, recorded 40 [CVCV] stems with initial main stress and 2 suffix allomorphs ([–la] and [–na]). The stems and suffixes were recorded separately, and then spliced together to make a total of 80 items (40 stems × 2 suffix allomorphs). The resulting [CVCVC] stimuli were matched for overall amplitude and duration.

The stems were created by pseudo-randomly combining segments from a subset of the English inventory (consonants: /p b t d k g m n s l/, vowels: /i e ɑ o u/).[5] None of the stems was a real word of English. The pseudo-random combination ensured that, across the set of 40 [CVCV] stems, (i) each consonant and vowel in the experimental inventory had approximately the same frequency as other segments of the same type (C or V), and furthermore (ii) no CV, VC, C_C_, or _V_V combination occurred significantly more often than any other combination of the same type. This high degree of statistical control over the stimuli was necesssary to avoid unintended regularities that could be used by one of the experimental groups to (partially) predict the distribution of the suffix allomorphs.

Procedure. The experiment used a slightly modified version of the Artificial Grammar (AG) paradigm (e.g., Reber 1989, Gomez & Schvaneveldt 1994, Pothos & Bailey 2000, Saffran 2002). A typical AG experiment has two phases. In the exposure phase, participants are presented with example stimuli that have been generated with a grammar (usually a small finite-state machine). Then, in the testing phase, participants are tested on their ability to distinguish examples that are generated by the same grammar from examples that are not. Previous research in this paradigm has shown that participants can rapidly learn regularities over various stimulus domains (written consonant sequences, shape sequences, nested shapes, tone sequences, etc.) and that the knowledge acquired is generally implicit (in the sense that participants have little if any ability to verbalize what they learned during the experiment). Furthermore, experiments in the AG paradigm typically do not provide participants with any explicit negative evidence (i.e., ungrammatical examples labeled as such, or feedback that the participant's hypothesis about the grammar is incorrect). The absence of explicit negative evidence is one important respect in which learning in AG experiments resembles natural first-language acquisition.

5. The mid vowels were diphthongized on the surface, in accordance with the phonology of unreduced vowels in English.

The instructions for the testing phase in a typical AG experiment inform the participants that the examples in the exposure phase were created with a set of “rules”, and that their task during the testing phase is to judge whether each stimulus item “obeys the same rules”. In an attempt to bring the AG paradigm somewhat closer to the natural language setting, the instructions were modified as follows. At the beginning of the exposure phase, participants were told that they would be hearing a list of words from a made-up language — one that was like real languages, but simplified for the purposes of the experiment — and that their task was to remember the words that they head. Then, at the start of the testing phase, participants were told that they would be tested on how well they remembered the words from the first part of the experiment. The instructions did not include any reference to the existence of rules, patterns, regularities, or generalizations in the list of words.

The details of stimulus presentation and response collection in the present experiment are as follows. In the exposure phase, each participant heard a list of 20 [CVCVCV] items. All of the items initial consisted of a stem (initial [CVCV]) spliced together with one of the suffix allomorphs ([la] or [na]); the combination of stem and suffix allomorph was always grammatical according to the rule for the participant’s group. The list was presented twice, each time in a different random order.[6]

In the testing phase, each participant heard all 80 [CVCVCV] items. The 80 items were arranged in two blocks of 40 each; each block contained 10 *old* items (items that the participants had heard in the exposure phase) and 30 *new* items (items that had not been presented during the exposure phase). Participants hit one button to respond “yes” (they remembered hearing an item in the exposure phase) and another button to respond “no” (they did not remember hearing an item in the exposure phase). Responses and response times (RTs) were collected with the PsyScope button box; for reasons of space, the RT results are not presented in this paper.

Two aspects of stimulus presentation in both the exposure phase and the training phase should be emphasized. First, stimuli were presented auditorily without accompanying orthographic, semantic, or other information. Second, what have been referred to as the [CVCV] “stems” were never presented in isolation; the same is true for the [CV] syllables referred to as the “suffix allomorphs”. These terms are useful for describing

6. There was a short (1500ms) pause between each word in the list, and a short break between the two presentations of the list. The list of stimuli for the exposure phase was similar for every participant. Each member of the set /m n k g/ occured in the second consonant position of exactly 2 stems (i.e., there were 2 CVmVCV items, 2 CVnVCV items, etc.). The remaining 12 stems were selected pseudo-randomly for every participant. Stimulus presentation and response collection were controlled by PsyScope.

the structure of the experiment. But any morphological (or other) parse that the participants might have assigned to the stimuli was not encouraged by the instructions, and was not explicitly tested in the experiment.

2.2 Results and discussion

For the purposes of interpreting the results of the experiment, it is useful to distinguish four types of stimulus that occur during testing, as shown in (4) below. There was an equal number (N=20) of each stimulus type. The old-grammatical items are identical to the items that were heard during the exposure phase. The other items diverge from these along two dimensions. The old-ungrammatical items contain stems (initial [CVCV] sequences) that were heard during the exposure phase; but those stems are combined with suffix allomorphs (final [CV]s) that are ungrammatical according to the rule for the participant's group. The new-grammatical items contain stems that were not heard during the exposure phase, combined with suffix allomorphs that are correct according to the rule for the participant's group. Finally, the new-ungrammatical items consisted of stems that had not been heared during the first part of the experiment combined with the incorrect suffix allomorph. Recall that exactly the same 80 testing items were presented to all of the participants. The assignment of particular stimuli to the four categories in (4) of course varied by group.

(4) Stimulus categories in the testing phase

		Grammaticality	
		grammatical	ungrammatical
Stem type	old	N=20 ex. dumena, sutola	N=20 ex. *dumela, *sutona
	new	N=20 ex. kinena, tagola	N= 20 ex. *kinela, *tagona

(Examples correspond to stimulus types for a participant in Group 1A).

In interpreting the results of the experiment, I assume the following hypothesis relating error types and underlying knowledge:

(5) Linking hypothesis: If a participant responds "yes" significantly more often to **new-grammatical** items than to **new-ungrammatical** items, then there is evidence that he/she learned a rule in the exposure phase.

The items with new stems (i.e., [CVCV] sequences that were not heard during the exposure phase) provide the best test of rule-learning, because the new-grammatical and the new-ungrammatical items can be distinguished only on the basis of the relationship between the stem and the suffix allomorph. A participant that responds "yes" significantly more to the new-grammatical items cannot be doing so simply by recalling specific words (i.e., [CVCVCV] sequences) from the exposure phase. Instead, he/she must have some knowledge of the suffix allomorphy pattern that generalizes beyond those particular items.[7]

For every participant, the proportion of "yes" responses for each stimulus category in (4) was calculated. These proportions were then submitted to an arcsine transformation ($y = \sin^{-1}(\text{sqrt}(x))$, where x is the original proportion). The untransformed proportions are given in the tables below, but all statistical tests were performed on the transformed proportions (see Winer et al. (1991)).

An ANOVA with one between-subjects factor (Group: 1A vs. 1B) and two within-subjects factors (Stem type: old vs. new, Grammaticality: gramm. vs. ungramm.) showed that no significant main effect for Group $F(1,22)<1$).[8] In other words, neither group made significantly more errors than the other. Further analysis revealed that, though the overall error rate was similar, the *patterns* of errors for the two groups were quite distinct. The mean proportion of "yes" responses for each item type is given in (6). (Note that "yes" responses to all item types except old-gramm. are errors).

(6) Mean (SD) proportion "yes" by group and item type

		Group 1A		Group 1B	
		Grammaticality		Grammaticality	
		gramm.	ungramm.	gramm.	ungramm.
Stem type	old	.70 (.14)	.44 (.20)	.60 (.13)	.54 (.14)
	new	.53 (.19)	.34 (.21)	.46 (.12)	.38 (.13)

Participants in Group 1A, who were exposed to examples of the natural assimilation rule, accepted grammatical items significantly more often than ungrammatical items. A repeated-measures ANOVA with Stem type as the within-subjects factor showed no significant difference between items with old stems and items with new stems ($F(1,11)=4.12$, $p<.1$). In contrast, there is a significant difference when Grammaticality is taken as the within-subjects factor: participants responded "yes" more often to grammatical items than to ungrammatical items ($F(1,11)=10.46$, $p<.01$*). And the crucial comparison between new-grammatical and new-ungrammatical items was

7. An alternative analysis of the results presented below, namely that participants preferred certain testing items independently of their experience in the exposure phase, is ruled out by the results of Experiment 2; see discussion in section 3.2.
8. In all of the ANOVA and *t*-test analyses reported in this paper, Subject is the random factor. An □ level of .05 was used to determine significance in all cases.

also significant: participants accepted new-grammatical items more often than new-ungrammatical items (planned comparison $t(11)=2.45$, $p<.05$*).

This pattern of results suggests that participants in Group 1A acquired knowledge of the assimilation rule during the exposure phase. This knowledge allowed them to correctly reject items that were ungrammatical, even if they contained previously-heard stems. But it also led them to incorrectly accept the new-grammatical items, which obey the rule but were not heard during the first part of the experiment.

In contrast, there is no evidence that participants in Group 1B acquired knowledge of the random rule; instead, their responses seem to be based largely on familiarity with the stem. The same statistical tests that were performed on the results from Group 1A were also conducted for this group. Participants accepted items with old stems significantly more often than items with new stems ($F(1,11)=10.50$, $p<.01$*, with Stem type as the within-subjects factor). But they did not accept grammatical items significantly more often than ungrammatical items ($F(1,11)=4.05$, $p<.1$, with Grammaticality as the within-subjects factor). And the planned pairwise comparison was also non-significant: participants in Group 1B did not respond "yes" significantly more often to new-grammatical items than to new-ungrammatical items ($t(11)=1.81$, $p<.1$).

According to the logic behind the linking hypothesis in (5), a higher acceptance rate to grammatical items than to ungrammatical items, when familiarity with entire words is eliminated as an interfering factor, provides the best evidence of rule-learning in the AG paradigm. Such evidence was found for the participants in Group 1A but not for the participants in Group 1B. Because the rules for the two groups are formally matched (in the sense that the [nasal] feature of the suffix consonant is determined in both cases by one feature of the closest stem consonant), and because neither rule has an apparent advantage over the other with respect to English phonology or the statistical properties of the English lexicon, these results strongly suggest that the differences between the response patterns of the two groups are due to a cognitive bias that makes learning and/or representation of the assimilation rule more robust. This in turn supports phonological theories that give an explicit formalization of the cognitive bias.

There are several specific properties of the assimilation rule that distinguish it from the random rule and that could individually or collectively be the focus of a cognitive bias. Perhaps the most basic is that the assimilation rule involves a formal relation of *identity*. In all of the grammatical stem–suffix combinations, the [nasal] feature of the suffix consonant and that of the final stem consonant have identical specifications. Indeed, when the final stem consonant is [n], there is complete identity between the consonant in the suffix and its determinant in the stem. Many phonological theories treat relations of (partial) identity as special in ways that could underlie the results presented here (see, for example, the references on assimilation cited above and Reiss 2003).

A related, but more substantive property of the assimilation rule is that it applies to segments that are independently similar on particular features For example, based on their survey of consonant harmony, Rose & Walker (2002) claim that agreement on [sonorant], [continuant], and place features makes consonants particularly susceptible to agreement on other features. For the assimilation rule, the suffix consonant always agrees in [sonorant] and [continuant], and sometimes agrees in place, with the segment in the stem that conditions the non-default allomorph ([–na]). No comparable similarity relation obtains for the random rule. (See Hansson 2001:ch.4 for further discussion of the role that similarity plays in consonant harmony, and Burzio 1998 for a general theory of similarity effects in phonology).

3. Experiment 2: Nasal dissimilation vs. random alternation

Experiment 2 was designed to test for a cognitive bias that favors the opposite type of feature-changing process: dissimilation. Although assimilation and dissimilation are both well-attested in natural languages, they are not always considered to have the same status in synchronic phonology, and have been claimed to have distinct diachronic sources (Ohala 1981). The materials and critical manipulation for this experiment are given in (7).

(7) Design of Experiment 2

a. Group 2A: Nasal dissimilation

m/n-stems	k/g-stems	other stems
dumela binula (N=8)	tukona digena (N=8)	sutona dabuna (N=24)
Rule: Suffix is [–la] iff the final stem consonant is [+nasal]		

b. Group 2B: Random rule

m/n-stems	k/g-stems	other stems
dumena binuna (N=8)	tukola digela (N=8)	sutona dabuna (N=24)
Rule: Suffix is [–la] iff the final stem consonant is [+dorsal]		

Participants in Group 2A were exposed to a pattern that can be described with a dissimilation rule, one that forces the [nasal] specification of the suffix consonant to be distinct from that of the final stem consonant. Dissimilation processes similar to this one are attested in Takelma

(Goodman 1992) and several languages of Australia (e.g., Gurindji; McConvell 1988). (See Alderete 1997 and Suzuki 1998 for typological surveys and constraint-based theories of dissimilation). Participants in Group 2B were exposed to a pattern that involves an arbitrary, or random, relation. As in Experiment 1, the natural class of sounds that triggered the non-default suffix allomorph (here, [–la]) was the [+dorsal] consonants.

These rules are matched along the same dimensions as the two rules of Experiment 1. Both rules depend on one feature of the same stem consonant to determine the consonant of the suffix. Neither rule exists as part of the phonology of English. And the frequencies in table (3) show that neither rule produces surface structures that are substantially more frequent, in the English lexicon, than the structures produced by the other rule.

3.1 Method

24 adult native English speakers were recruited from the UCLA community to participate in Experiment 2. Each participant was randomly assigned to Group 2A (N=12) or Group 2B (N=12), and received a nominal fee or course credit for participating. The stimuli and procedure for this experiment were identical to those of Experiment 1.

3.2 Results and discussion

The stimulus categories and linking hypothesis developed in section 2.2 were also used to analyze the results of this experiment. A three-way ANOVA with one between-subjects factor (Group: 2A vs. 2B) and two within-subjects factors (Stem type: old vs. new, Grammaticality: grammatical vs. ungrammatical) was performed on the arcsine transformed proportion "yes" responses. There was no main effect of Group ($F(1,22)<1$), just as in Experiment 1. The analysis then proceeded to investigate the patterns within each group, which are summarized in (8).

(8) Mean (SD) proportion "yes" by group and item type

		Group 2A		Group 2B	
		Grammaticality		Grammaticality	
		gramm.	ungramm.	gramm.	ungramm.
Stem type	old	.73 (.17)	.47 (.17)	.68 (.12)	.52 (.19)
	new	.50 (.13)	.35 (.12)	.47 (.15)	.41 (.15)

The acceptance rates the cells of this table are very similar to those in the corresponding cells of table (6) for Experiment 1, suggesting a cognitive bias for dissimilation that is at least as strong as the one for assimilation.

Participants in Group 2A, who were exposed to an attested type of dissimilation, accepted items with old stems significantly more often than items with new stems (main effect for Stem type, F(1,11)=14.90, $p<.01$*) and accepted grammatical items significantly more often than ungrammatical items (main effect for Grammaticality, F(1,11)=21.21, $p<.01$*). The planned pairwise comparison between new-grammatical and new-ungrammatical items — the comparison that is critical according to the linking hypothesis in (5) — was also significant for Group 2A ($t(11)=3.54$, $p<.01$*). This pattern of significance suggests that participants in Group 2A acquired knowledge of the dissimilation rule.

The responses of participants in Group 2B, like those in Group 1B, provided no evidence of rule-learning. Participants in this group, who were exposed to a random rule of suffix allomorphy, accepted items with old stems significantly more often than items with new stems (main effect for Stem type, F(1,11)=11.26, $p<.01$*), but they did not respond significantly differently to grammatical and ungrammatical items (main effect for Grammaticality, F(1,11)=3.38, $p<.1$). And the crucial planned pairwise comparison was non-significant ($t(11)=.94$, n.s.).

These results make two contributions to the goals of the paper. First, they provide evidence for a cognitive bias that favors dissimilation, one that appears to be as strong as the bias favoring assimilation. This in turn supports theoretical frameworks that give a privileged status (e.g., with special classes of rules or constraints) to both process types. Second, they rule out the possibility that the results of Experiment 1 are due to a preference that the participants had, independently of their experience in the exposure phase, for items that obey the assimilation rule. The response patterns for Group 1A and Group 2A are almost identical, despite the fact that all of the items that are grammatical for Group 1A are ungrammatical for Group 2A (and vice versa). The entire set of results from both experiments can therefore only be explained in terms of *learning* during the exposure phase: learning that is biased, and therefore more successful with exposure to natural rules than with exposure to random ones.

4. Conclusion

The results presented in this paper are among the first to provide experimental support for the claim, widely held in theoretical phonology, that certain process types have a privileged cognitive status.[9] There is insufficient space in this paper to address many of the questions that these results raise (e.g., how directly results from AG experiments bear on claims

---

9. Related results appear in Bailey et al. (1999) and Nowak et al. (this volume); see also Pater (2002), which addresses different issues with a similar paradigm.

about first- and second- language acquisition, and what specific formalizations of the cognitive biases are supported by the results). I will conclude the paper with a brief remark on one important issue: namely, whether the cognitive biases investigated here are specific to language.

I think it is plausible that the biases for assimilation and dissimilation, to the extent that they refer to relations of identity and similarity, are not entirely language-specific, but rather more domain-general aspects of cognition. This has consequences not only for the interpretation of the present experiments, but also for the theories that motivated them. The experiments could reveal less about the linguistic system *per se* than has been assumed up to this point in the paper. But if so, then the possibility arises that certain components of phonological theories, in particular those that refer to featural identity and similarity, could be reduced to more domain-general biases and mechanisms. The same possibility exists, of course, for other principles and components of phonological theory that have been assumed to be language-particular. Regardless of whether any particular reduction is realized, the AG paradigm, which has been applied with a wide variety of stimulus types and in all modalities, promises to shed light on the nature of cognitive biases, and ultimately on the central theoretical question of which biases are specifically linguistic.

References

Alderete, John. 1997. Dissimilation as Local Conjunction. In Kiyomi Kusumoto (ed.), *Proceedings of North East Linguistic Society 27*, 17–31. Amherst: Graduate Linguistic Student Association.

Baayen, R. Harald, Richard Piepenbrock & H. van Rijn (eds). 1993. The CELEX Lexical Database (CD-ROM). Linguistic Data Consortium, University of Pennsylvania, Philadelphia (PA).

Bailey, Todd M., Kim Plunkett & Ester Scarpa. 1999. A Cross-Linguistic Study in Learning Prosodic Rhythms: Rules, Constraints, and Similarity. *Language & Speech* 42.1, 1–38.

Buckley, Eugene. 2000. What should phonology explain? Handout from Linguistics colloquium, State University of New York at Buffalo, March 17.

Burzio, Luigi. 1998. What's Hebb's Learning Rule doing in Phonology? Ms., Johns Hopkins University, Baltimore, MD.

Cena, Richard. M. 1978. *When is a Phonological Generalization Psychologically Real?* Bloomington, Indiana: Indiana University Linguistics Club.

Chomsky, Noam & Morris Halle. 1968. *The Sound Pattern of English*. New York: Harper & Row.

Clements, G.N. 1985. The geometry of phonological features. *Phonology Yearbook* 2, 225–252.

Gafos, Adamantios. 1996. *The Articulatory Basis of Locality in Phonology*. Doctoral dissertation. Baltimore, MD: Johns Hopkins University.

Gomez, Rebecca L. & Roger W. Schvaneveldt. 1994. What Is Learned from

Artificial Grammars? Transfer Tests of Simple Association. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20.2, 396–410.
Goodman, Beverley. 1992. Takelma Dissimilation and the Form of the OCP. *Working Papers of the Cornell Phonetics Laboratory*, vol.7, 41–63.
Halle, Morris. 1995. Feature geometry and feature spreading. *Linguistic Inquiry* 26.1, 1–46.
Hansson, Gunnar. 2001. *Theoretical and Typological Issues in Consonant Harmony*. Doctoral dissertation. Berkeley, CA: University of California, Berkeley.
Hoard, James E. & Geoffrey N. O'Grady. 1976. Nyangumarda phonology: a preliminary report. In R.M.W. Dixon, ed., *Grammatical categories in Australian languages*, 51-77. Canberra: AIAS.
Hyman, Larry. 1995. Nasal Consonant Harmony at a Distance: The Case of Yaka. *Studies in African Linguisics* 24.1, 5–30.
McConvell, Patrick. 1988. Nasal Cluster Dissimilation and Constraints on Phonological Variables in Gurindji and Related Languages. *Aboriginal Linguistics* 1, 135–165.
Ní Chiosáin, Máire & Jaye Padgett. 2001. Markedness, segment realization, and locality in spreading. In Linda Lombardi, ed., *Constraints and representations: segmental phonology in Optimality Theory*, 118–156. Cambridge: Cambridge University Press.
Nowak, Pawel, Anne Pycha, Eurie Shin & Ryan Shosted. 2003. Phonological rule-learning and its implications for a theory of vowel harmony. Handout of talk presented at WCCFL 22, UC San Diego, March 22, 2003.
Ohala, John. 1991. The Listener as a Source of Sound Change. In C.S. Masek, R.A. Hendrick & M. Miller (eds.), *Papers from the Parasession on Langauge and Behavior*, 178–203. Chicago: Chicago Linguistic Society.
Pater, Joe. 2002. Does phonotactic knowledge aid rule learning? Talk presented at the Boston University Conference on Language Development, April 24, 2002.
Pierrehumbert, Janet B. 2002. An Unnatural Process. Under review for *Laboratory Phonology* VIII.
Pothos, Emmanuel M. & Todd M. Bailey. 2000. The Role of Similarity in Artificial Grammar Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26.4, 847–862.
Reber, Arthur S. 1989. Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General* 118, 219–235.
Reiss, Charles. 2003. Quantification in Structural Descriptions: Attested and Unattested Patterns. [In press in *The Linguistic Review* 20.]
Rose, Sharon & Rachel Walker. 2002. A typology of consonant agreement as correspondence. Submitted. [ROA-458-0801]
Saffran, Jenny. 2002. Constraints on statistical language learning. *Journal of Memory and Language* 47, 172–196.
Sohn, Ho & Byron W. Bender. 1973. *A Ulithian grammar*. Canberra: Australian National University, Research School of Pacific Studies, Dept. of Linguistics.
Suzuki, Keiichiro. 1998. *A Typological Investigation of Dissimilation*. Doctoral dissertation. Tucson, AZ: University of Arizona, Tucson.
Winer, B.J., Donald R. Brown & Kenneth M. Michels. 1991. *Statistical principles in experimental design,* 3rd ed. New York: McGraw-Hill.