Structure and substance in artificial-phonology learning

Elliott Moreton — UNC-Chapel Hill
Joe Pater — UMass-Amherst [*]

Version 4.1, March 26, 2012

Contents

1 Introduction **2**

2 Formal complexity **3**
2.1 Number of relevant features . . . 3
2.2 Relations between features . . . 9
2.2.1 Featural agreement . . . 9
2.2.2 Contiguity-similarity tradeoff . . . 11
2.3 Summary: Complexity . . . 13

3 Phonetic substance **15**
3.1 Consonants versus vowels . . . 16
3.2 Segmental harmony . . . 17
3.3 Consonant-vowel interactions . . . 19
3.4 Prosodic rules . . . 20
3.5 Summary: Substance . . . 22

4 General discussion **22**
4.1 An analytic artifact? . . . 23
4.2 Is artificial phonology phonology at all? . . . 24
4.3 Exposure and robustness . . . 25
4.4 Structurally biased phonology . . . 26

[*]The authors gratefully acknowledge the advice and suggestions already contributed by two anonymous reviewers, Bruce Hayes, Anne Pycha, Jen Smith, Kie Zuraw, and the participants in Linguistics 751 at the University of Massachusetts in Spring 2011, especially John Kingston, Claire Moore-Cantwell, Wendell Kimper, and Robert Staubs.

1 Introduction

Laboratory study of the acquisition of artificial phonological patterns is potentially very significant as a window on the inductive biases involved in acquiring natural-language phonology.[1] Some such biases must exist, since generalizations cannot be learned without them (e.g., Pinker, 1979; Mitchell, 1990; Gallistel et al., 1991); what is at issue is rather their content and their causal relationship to phenomena of natural language acquisition, change, and typology. Artificial phonology may offer a way to study these biases in near-isolation from each other and from other confounding factors. The present article reviews the empirical literature on artificial-phonology learning in the context of this program.

What factors make phonological patterns harder to learn? And do these highly artificial tasks reveal anything about natural-language phonology? Two hypothesized biasing factors have been studied the most intensively, *formal complexity* and *phonetic substance*. A learner with formal complexity bias would acquire simpler patterns faster or better than complex ones, whereas a learner with substantive bias would acquire phonetically-motivated patterns better than phonetically-arbitrary ones, assuming in both cases that the training data instantiates the patterns equally well, that the learner perceives the training data correctly, and that other factors are controlled for. By skewing the direction of language change, these biases could cause simple or phonetically-motivated patterns to accumulate, thus producing a corresponding skew in natural-language typology. (Of course, these inductive biases are not the only factors that could be affecting typological frequencies, and in particular substantive bias is not the only way that phonetic substance could be influencing typology. More will be said of this below, §4.)

Natural-language typology is consistent with the effects of both kinds of bias, but the same is not true of phonological learning in the lab. Studies which directly compare simple patterns with complex ones nearly always find an advantage for the simple one, whereas studies which directly compare similar patterns instantiated by different features usually find no significant advantage for the phonetically-motivated or typologically-frequent pattern. Since similar methods and participant populations are used in both kinds of study, the systematic success with complexity and failure with substance corroborate the early conclusions of Pycha et al. (2003): If substantive biases exist at all, they are considerably weaker than complexity biases.[2]

[1] By *inductive bias* (also called "analytic bias" or "learning bias"), we mean any tendency of a pattern-learning algorithm to acquire one pattern faster or better than another from training sets that instantiate both patterns equally well. This definition is deliberately broad. It includes absolute distinctions between learnable and utterly unlearnable patterns, as well as relative distinctions between easier and harder ones. It is indifferent to details of implementation, applying to explicit penalties against specific patterns as well as to emergent consequences of the learner's architecture, representational system, or similarity metric. The term as we use it includes, but is not limited to, anything that would qualify as "Universal Grammar".

[2] The validity and necessity of considering null results (i.e., failures to find a statistically-significant difference) is discussed further in §4. For now, we note that studies which found null results *must* be reviewed because excluding them would falsely inflate the apparent robustness of the effects. Five percent

The scope of this review is the effects of formal complexity and phonetic substance on the acquisition of artificial analogues of categorical phonology i.e., patterns which partition a discrete stimulus space into positive ("legal", "pattern-conforming") versus negative ("illegal", "non-conforming") instances on the basis of phonological features.[3] In the non-linguistic psychological literature, such partitions are often referred to as "concepts". We are not concerned here with learning how to partition a continuous stimulus space into phonetic categories (e.g., Maye et al., 2002; Goudbeek et al., 2008), nor with analogues of lexical (e.g., Peña et al., 2002; Perruchet et al., 2004; Newport and Aslin, 2004) or syntactic (e.g., Gómez, 2002) dependencies.

We have tried to be as inclusive as possible, but have necessarily omitted some studies, either because they do not fall within the purview of the article, or because their results were not reported in a way that we could use, or because a discussion would have consisted of criticizing the experiments rather than interpreting their results, or because we were simply unaware of them at the time when the article went to press. Our interpretations of the experiments which we do review are not necessarily those of the original authors.

2 Formal complexity

Many formal theories of natural-language phonology are designed to favor patterns which have a simple expression in phonetic terms, and some proposals expressly impute this bias to human learners (e.g., Chomsky and Halle 1968, 330–334; Kiparsky 1971, 623; Bach and Harms 1972; Sagey 1990, 1; Hayes 1999; Gordon 2004, 304). Two main formal complexity factors have been studied in phonology: the number of features relevant to the pattern, and the relations between them.

2.1 Number of relevant features

For an independently motivated standard of formal complexity, we turn to the psychology literature. A stimulus space described by three binary-valued features can be divided into two equal-sized categories in only six ways, if we ignore trivial variants obtained by permuting features or inverting feature values. Examples are shown in (1) for stimuli that are geometric figures varying in color (black vs. white), shape (circle vs. triangle), and size (large vs. small). Only color matters for the Type I distinction. Type II requires attention

of experiments testing a false hypothesis will find statistically significant support for it at the conventional $p < 0.05$ level. If the inconclusive or contradictory results which constitute the other 95% are suppressed on the grounds that they are "null", the hypothesis will appear to be true.

[3] By *feature*, we mean any discrete-valued variable created by partitioning a continuous or discrete physical dimension. A *phonetic feature* is one for which the physical dimension is phonetic, and a *phonological feature* is a phonetic feature that is used in a model of human phonology. Experimental patterns are implemented in terms of phonetic features but interpreted in terms of phonological ones, with the result that the formal complexity and phonetic substance of a pattern depend on the choice of model. We will return to this point in §4.1 once individual examples have been presented.

to color and shape, but size can be ignored. Types III through V involve all three features, but some subsets can be decided with fewer (e.g., white triangles). For Type VI, not even this is possible; even a subset requires all three features.

(1) *Representatives of the six possible equal partitions of a stimulus space defined by binary features of color, shape, and size. Boxes enclose the (arbitrary) positive class. Concepts are arranged in increasing order of difficulty, with III, IV, and V being about equal. (After Shepard et al. 1961.)*

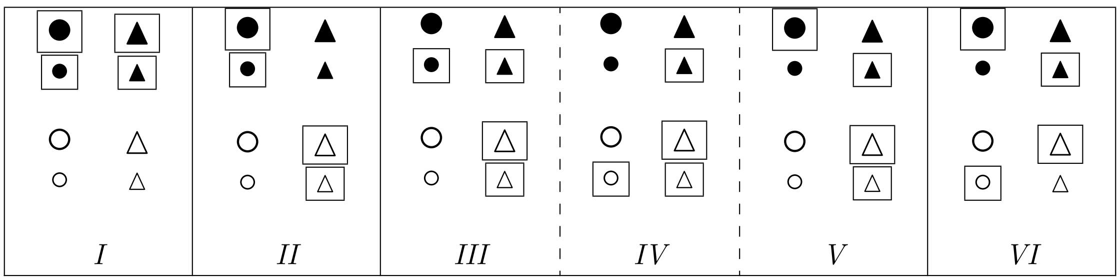

These six concepts have been extensively studied in connection with supervised learning of non-linguistic categories.[4] In a typical experiment, the participant is shown a randomly-selected stimulus, judges whether it belongs to the target concept, and is then told the correct response. This cycle repeats until some performance criterion is met. The main finding is that difficulty increases along with the number of relevant features: Type I is easier than Type II, which is easier than Types III, IV, and V, which in turn are easier than Type VI (Shepard et al., 1961; Neisser and Weene, 1962; Nosofsky et al., 1994; Feldman, 2000; Love, 2002; Smith et al., 2004).

Analogous phonological stimuli have been used in both supervised and unsupervised learning experiments. Participants are either trained with feedback to divide stimuli into “legal” and “illegal” categories, or are familiarized without feedback on “legal” stimuli only. They are then asked to categorize stimuli as legal or illegal, or to decide which of two stimuli is the legal one. The phonological experiments have replicated the non-linguistic difficulty hierarchy for Types I, II, and VI. We know of no published studies on the other three types.

A phonological pattern that depends on a single stimulus feature (Type I) has often proven easier, and never harder, than one that requires more. Saffran and Thiessen (2003,

[4]The established terminology here invites confusion. The psychology literature uses “category” or “concept” to mean a partition of a stimulus space into disjoint labelled subsets. For example, “black triangles” would be a category or concept in the three-dimensional space shown in (1). In phonology and phonetics, “category” often has the more specific meaning of a phoneme, a region in phonetic space whose elements are phonologically equivalent (as in the phrase “categorical perception”. We use the terms here in their broader sense; e.g., “voiced stops” would be called a category or concept.

Exps. 2, 3) familiarized English-learning 9-month-olds with isolated positive nonword instances, exposed them to a continuous stream of two positive and two negative nonwords, and then compared listening times to these four nonwords using headturn preference. When the pattern restricted [p t k] to some positions and [b d g] to others, the negative stimuli were preferred, but when the pattern distinguished [p d k] vs. [b t g], there was no difference in means (see schematic in Figure 2a). Cristiá and Seidl (2008, Exp. 1) familiarized English-learning 7-month-olds on positive C_1VC_2 nonwords. When the pattern was "C_1 is a nasal or oral stop" (i.e., [−continuant]), the infants preferred novel negative instances over positive ones; when it was "C_1 is a nasal or fricative", for which there is no standard feature, they showed no preference (Figure 2b). LaRiviere et al. (1974, 1977) trained English-L1 adults to categorize a set of six or eight syllables into two equal-sized classes defined either by a single feature or in an unspecified "random" way that needed more relevant features. Performance was significantly better for the single-feature condition than the random condition in three out of ten experiments, and was numerically better in the other seven.[5]

(2) *Examples of phonological patterns in experiments comparing featural complexity. Some are defective representatives of their category types, owing to the impossibility of some feature combinations. In each case, the left-hand pattern proved the easier. (Compare Table (1).)*

(a) Saffran and Thiessen (2003, Exps. 2, 3). Features were voiced vs. voiceless, coronal vs. non-coronal, and labial vs. velar. (Defective.)

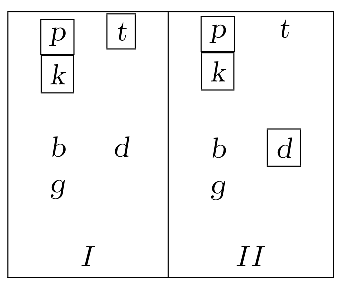

[5] If there were in fact no difference between the simple and complex conditions in any of the experiments, then the chance that all ten experiments would favor the simple condition is 1 in 1024, or 0.000977.

(b) *Cristiá and Seidl (2008): Features were oral vs. nasal, continuant vs. non-continuant, labial vs. non-labial. (Defective.)*

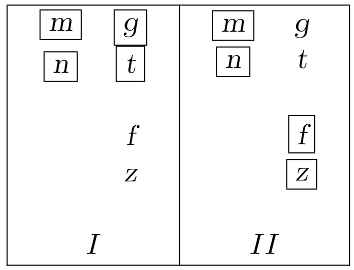

(c) *Kuo (2009): Features were plain initial stop vs. aspirated, labial initial stop vs. coronal, and palatal glide vs. labiovelar. Corresponding conditions (not shown) inverted the legal/illegal categories.*

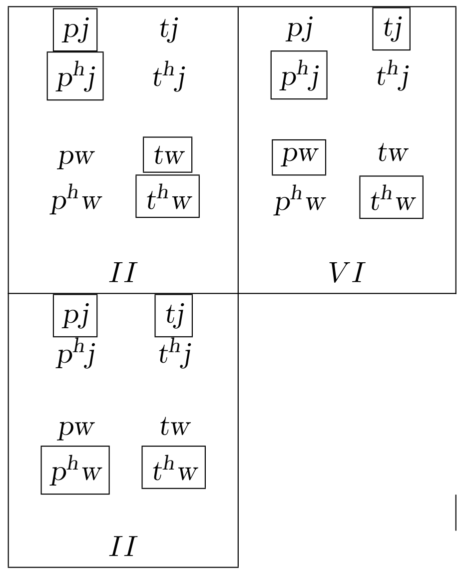

(d) *Pycha et al. (2003): Features are front first vowel vs. back, front last vowel vs. back, and high-lax first vowel vs. other first vowel.*

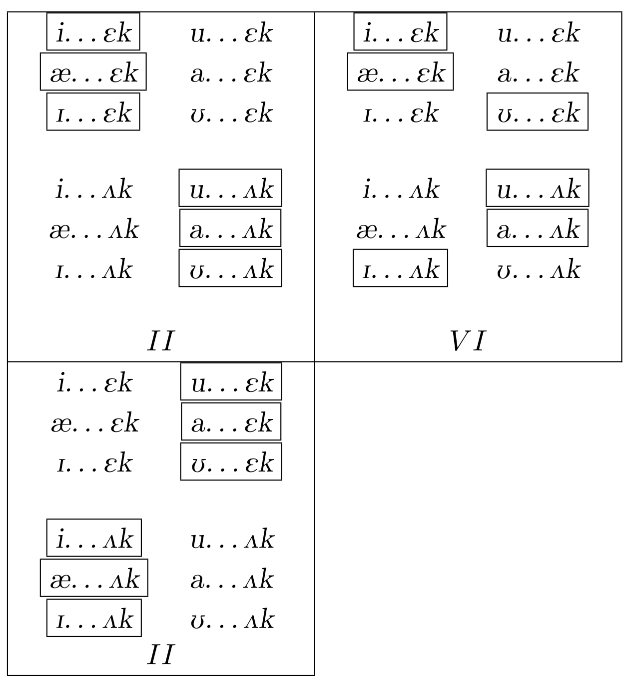

Two-feature Type II patterns enjoy a similar advantage over three-feature Type VI ones. Kuo (2009) familiarized L1 Mandarin speakers on syllables with two-consonant onsets in which the initial stop perfectly predicted whether the following glide would be [j] or [w] (Figure 2c). In two patterns, a single stop feature, aspiration or place, was relevant (Type II); in the third, both stop features were needed (Type VI). Both Type II conditions elicited a significantly greater preference for novel positive stimuli over non-conforming foils than did the Type VI condition. A similar result was found by Pycha et al. (2003) in an experiment in which adult English speakers were trained with feedback to make binary grammaticality judgments of $X \ldots XY$ stimuli, where Y was either [-ɛk] or [-ʌk] depending on the vowel of X (Figure 2d). Classification at test was more accurate for two Type II patterns (backness agreement and backness disagreement) than for a Type VI pattern. Skoruppa and Peperkamp (2011) exposed French speakers to spoken passages in their own language which had been modified so that a front vowel either agreed in rounding with the preceding vowel (Type II), disagreed, (Type II), or agreed if mid and disagreed if high (Type VI). Participants in the Type II conditions were better at recognizing new pattern-conforming stimuli than those in the Type VI condition.

The advantage for patterns with fewer relevant features extends to patterns which are in part phonetically arbitrary. Using a speeded-repetition paradigm, Chambers et al. (2010) familiarized English speakers with a pattern in which the unsystematic sets [b f k m t] and [p s g n tʃ] were restricted to opposite ends of a C_1VC_2 syllable when the nucleus was one of two vowels, but unrestricted when it was a third vowel. In four experiments with different vowel sets, novel probe syllables were repeated faster when their consonants obeyed the positional restriction, regardless of what the vowel was. Participants evidently did not detect the dependency between V and the presence of positional restrictions on the Cs. Their performance followed an inaccurate moderately-complex generalization rather than an accurate more-complex one.

A complexity disadvantage has also been reported for learned alternations in production. Peperkamp et al. (2006) exposed French-speaking adults to stimuli of the form XY paired with pictures of two or three of the same object. The number of objects determined X ([nɛl] or [ʀa]), and the identity of the object determined Y. The initial consonant C_Y of each Y varied depending on X. In two conditions, a phonetically-defined set of consonants switched voicing ([f ʃ]↔[v ʒ], or [p k]↔[b g]). In two others, both the sets and the change were phonetically unsystematic ([p z]↔[ʒ f], or [ʃ v]↔[b k]). When tested on XY phrases with novel Ys, participants in the voicing conditions changed C_Y in the pattern-conforming way about 25% of the time, whereas those in the unsystematic conditions did so only about 5% of the time (most responses left C_Y unchanged). The relevant complexity here seems to be that of the change rather than that of the segment classes undergoing it: Participants in the voicing conditions did not generalize the rule to new segments in the old classes. They must have induced two single-segment rules rather than a class-based one (e.g., [f]↔[v] and [ʃ]↔[ʒ], rather than "voiceless fricatives alternate with voiced ones"). Evidently a rule like [f]↔[v] is learned better than one like [ʃ]↔[b]. Using similar stimuli with a similar population, Skoruppa et al. (2009) found that an alternation in which only place of articulation changed was learned better than one in which place and manner changed together, and also better than another in which place, manner, and voice changed together.

On the other hand, an experiment by Finley and Badecker (2010) did not find a preference for a two-feature change over a one-feature change. They familiarized English speakers using stimuli consisting of three syllables X, Y, z, where X and Y agreed in backness but disagreed with Z, followed by their concatenation XYZ with the final syllable harmonized to the first two (i.e., every trial changed only the third syllable).[6] At test, they strongly preferred the familiar A, B, c, ABC, with one change, over A, B, c, abc, with two, but the one-change preference was reduced or even slightly reversed when the choice was between the unfamiliar stimulus types A, b, c, ABC (two changes) and A, b, c, abc (one). However, any bias towards fewer changes would have to compete with an opposing familiarity difference: A, b, c, ABC has an unfamiliar change in the second syllable; A, b, c, abc, an unfamiliar change in the first syllable and an unfamiliar lack of change in the third syllable. (Unfamiliarized participants had no significant preference.)

Thus, phonological-learning results from a wide variety of paradigms converge to show that patterns become harder to learn as the number of relevant features increases. These results are in agreement with what has been found for the learning of non-linguistic patterns.

2.2 Relations between features

A separate question is whether certain syntagmatic relations between features *within* a stimulus facilitate pattern learing when the number of relevant features is controlled. Two main relations have been studied in artificial phonology, featural agreement and the contiguity-similarity tradeoff.

2.2.1 Featural agreement

Dependencies between instances of the same feature within a word are conspicuously common cross-linguistically in the form of assimilation and dissimilation patterns (Archangeli and Pulleyblank, 2011; Baković, 2011; Bye, 2011; Rose and Walker, 2011). Many phonological theories make special provision for representing agreement or disagreement of features within some part of an utterance (e.g., Chomsky and Halle, 1968; Goldsmith, 1976; Alderete and Frisch, 2008). Domain-general theories of category learning have not addressed this possibility, and lack the means to recognize two features in the same stimulus as instances of the same abstract feature (Gluck and Bower, 1988; Kruschke, 1992; Nosofsky et al., 1994; Love et al., 2004; Feldman, 2006). However, there is non-linguistic evidence that patterns are easier to learn when they relate, e.g., color to color, or shape to shape, than when

[6] A mirror-image condition reversed the role of the first and third syllables, without major changes in the outcome.

they relate values on two different dimensions (Hunt and Hovland 1960; Ciborowski and Cole 1973; Ciborowski and Price-Williams 1974; not found by Shepard et al. 1961). This is true even for intra-dimensional patterns other than agreement (Rogers and Johnson, 1973). Several studies of phonological learning have found a similar phenomenon: Patterns relating two instances of the same feature produce larger familiarity effects than those relating instances of two different features.

Wilson (2003) familiarized English speakers to stimuli of the form $C_1V_1C_2V_2C_3a$. The identity of C_2 determined whether C_3 was [n] or [l], as shown in (3).

(3) *Conditions of Experiments 1 and 2 of Wilson (2003).*

	Familiarization			*% judged familiar*			
	Nasal C_2	*Dorsal* C_2	*Other* C_2	*Old*		*Novel*	
Exp.	4×	4×	12×	*Conf.*	*Non.*	*Conf.*	*Non.*
1A	$C_1V_1[m/n]V_2na$	$C_1V_1[k/g]V_2la$	$C_1V_1C_2V_2la$	*70*	*44*	*53*	*34*
1B	$C_1V_1[m/n]V_2la$	$C_1V_1[k/g]V_2na$	$C_1V_1C_2V_2la$	*60*	*54*	*46*	*38*
2A	$C_1V_1[m/n]V_2la$	$C_1V_1[k/g]V_2na$	$C_1V_1C_2V_2na$	*73*	*47*	*50*	*35*
2B	$C_1V_1[m/n]V_2na$	$C_1V_1[k/g]V_2la$	$C_1V_1C_2V_2na$	*68*	*52*	*47*	*41*

The patterns in Conditions 1A and 2A can be stated as agreement or disagreement in [nasal] between C_2 and C_3, whereas those in 1B and 2B crucially involve a relation between two features [Dorsal] and [nasal], or [Dorsal] and [lateral]. Concept membership significantly increased judged familiarity in the single-feature Conditions 1A and 2A, but not in the two-feature Conditions 1B and 2B.[7,8]

Healy and Levitt (1980, Experiment 3) found that a voicing-conditioned pattern was acquired better by English speakers than a phonetically arbitrary one when the pattern was voicing assimilation, but not when it was a correlation between voicing and vowel quality ([a] vs. [o]).[9] Moreton (2008, 2012) familiarized English speakers on $C_1V_1C_2V_2$ stimuli and tested discrimination between novel positive and negative instances. Performance was better when the pattern was height agreement between the vowels, or voice agreement between the consonants, than when it was correlation between the height of V_1 and voicing of C_2. This phenomenon is not peculiar to English, as Lin (2009) found the same result with speakers of Mandarin and speakers of Southern Min using a similar paradigm and

[7] Conditions 1A and 2A partially confound featural identity with segmental identity in the case of [n]; however, the difference between the C_2 = [m] and C_2 = [n] sub-conditions was not significant (Colin Wilson, p.c., 2010).

[8] This interpretation hinges on the traditional assumption that English post-tonic intervocalic [l] is phonologically [Coronal] despite its phonetic dorsal component (Sproat and Fujimura, 1993, 304). If [l] is phonologically [Dorsal] as well (Walsh Dickey, 1997, Ch. 2), then all four conditions can be stated as single-feature agreement or disagreement, and this study may exemplify a substantive rather than a complexity bias.

[9] The "arbitrary" patterns were constructed in the same way as those of Saffran and Thiessen (2003), by replacing the classes [p t k] vs. [b d g] with [p t g] vs. [b d k].

the same stimuli as Moreton (2008); nor is it confined to vowel height, since performance was better for backness agreement between the vowels than for correlation between the backness of V_1 and the voicing of C_2 (Moreton, 2012). Height-backness and voice-place dependencies showed no such advantage over height-voice, indicating a specific advantage for dependencies that involve two instances of the same feature, over and above any more general advantage for consonant-to-consonant or vowel-to-vowel dependencies.

Not all experiments that could have detected a single-feature advantage actually did so. Kuo (2009) found no difference between a place-place correlation (labial glide iff labial stop) and a place-aspiration correlation (labial glide iff aspirated stop; see the two Type II patterns in Figure 2c). Jen Smith points out to us that in some proposed phonological-feature systems, consonant labiality and vowel labiality are two different features (Odden, 1991), making both patterns two-feature patterns. Another null result was found by Seidl and Buckley (2005, Experiment 2), who familiarized 9-month-old infants on $C_1V_1C_2V_2(C_3)$ stimuli and tested listening preference for novel positive vs. negative stimuli. A novelty preference was obtained for an agreement pattern in which C_1 and V_1 agreed in labiality, but also for one in which C_1 was labial if and only if V_1 was high.

The overall tenor of the evidence is that dependencies between two instances of the same feature produce larger familiarity effects than dependencies between instances of two different features. These aspects of phonological learning have analogues in non-linguistic learning (reviewed in Moreton 2012).

2.2.2 Contiguity-similarity tradeoff

Phonological theory typically treats dependencies between adjacent elements as the normal case, excluding long-distance interactions unless the interacting segments share some property which is absent from intervening material (Jensen, 1974; McCarthy, 1981; Cole and Trigo, 1988; Pierrehumbert, 1993; Odden, 1995; Gafos, 1996; Hansson, 2001a; Frisch et al., 2004; Rose and Walker, 2004; Heinz, 2010). As with the other formal complexity biases reviewed above, there are parallels in non-linguistic learning: Two stimuli, or two elements of a compound stimulus, are more likely to cohere in perception and become associated in memory if they are contiguous in time or space, or are perceptually similar (Köhler 1941; Prentice and Asch 1958; Asch 1969; Arnold and Bower 1972; Rescorla 1980; Rescorla and Gillan 1980; Creel et al. 2004; Rescorla 2008; but see Pacton and Perruchet 2008.) It would therefore be surprising if contiguity and similarity did not facilitate acquisition within-stimulus dependencies in the lab. However, the relevant evidence is scanty, and what there is of it does not indicate a strong effect of either factor.[10]

Shorter-range dependencies have proven little, if any, easier to learn than longer-range ones. Majerus et al. (2004) familiarized French speakers on a continuous stream of CV

[10] We omit here studies of artificial long-distance lexical (e.g. Peña et al., 2002; Perruchet et al., 2004; Newport and Aslin, 2004) and syntactic (e.g. Gómez, 2002) dependencies, as outside the scope of this review.

syllables which contained phonetically unsystematic CV and $C \ldots C$ dependencies, then tested immediate recall of novel isolated nonword probes. Probes which belonged to both patterns simultaneously were recalled better than those which belonged to neither or only one, but there was no evidence that the two patterns differed in effect. Using a tongue-twister paradigm, Warker and Dell (2006, 2008) tested English speakers on stimuli in which two consonants ([f] and [s], or [k] and [g]) were constrained to appear at opposite ends of the stimulus. The positional restrictions were reversed depending on a third segment which was either adjacent to the marginal consonants ($C_1\mathbf{V}C_2$) or remote from them ($C_1V\mathbf{C}VC_2$). Exchange errors followed the positional restrictions to almost the same extent in the adjacent and remote conditions, with a numerical but nonsignificant advantage for the adjacent condition. Koo and Callahan (2011) used $C_1V_1C_2V_2C_3V_3$ stimuli to familiarize English speakers on a phonetically arbitrary $C_1 - C_2$ or $C_1 - C_3$ dependency. In both pattern conditions, novel conforming test items were judged as familiar significantly more often than were novel nonconforming items, and there were no significant differences between the closer and more-distant dependencies. Moreton (2008, 2012) familiarized English speakers on $C_1V_1C_2V_2$ stimuli and then tested their ability to discriminate novel positive versus negative stimuli. Conformity preference for phonetically-systematic continguous $V_1 - C_2$ dependencies did not differ significantly from that for phonetically-systematic non-contiguous $C_1 - V_2$, $C_1 - C_2$, or $V_1 - V_2$ dependencies, except when the latter involved two instances of the same feature (e.g., height-height)—in which case the stronger effect was produced by the *non*-contiguous dependency.

A similar result was obtained by Finley (2011a), but with an interesting twist. English-speaking participants were familiarized on stimuli of the form $X \ldots XY$, where X always contained exactly one [s] or [ʃ], and Y was [su] or [ʃu] to match. They were then tested on their ability to choose the pattern-conforming member of an X[su]–X[ʃu] pair. When all the familiarization stimuli had the form CV[s/ʃ]V[su/ʃu], with the critical consonants in adjacent syllables, participants in the test phase preferred new stimuli that fit the pattern over those that did not. Familiarization on [s/ʃ]VCV[su/ʃu], with the critical consonants far apart, likewise led to a preference for conforming test stimuli. As in other studies, there was no significant difference between the shorter- and the longer-range dependency in the rate of pattern-conforming responses. However, participants familiarized on [s/ʃ]VCV[su/ʃu] also significantly preferred CV[s/ʃ]V[su/ʃu] over CV[s/ʃ]V[ʃu/su], whereas those familiarized on the shorter-range dependency did not show an analogous preference for [s/ʃ]VCV[su/ʃu]-conforming test stimuli. Thus, familiarization on the longer-range pattern generalized to the shorter-range one, but not vice versa.[11] Subsequent experiments using $C_1V_1C_2V_2C_3V_3$[su/ʃu] stimuli found that a familiarized identity relation between C_1 and C_4 was learned as well as one between C_2 and C_4, and that each generalized equally well to the other (Finley, 2011b).

[11]This result is somewhat qualified by the fact that, when the rates of pattern-conforming response in the two generalization conditions were compared directly, the difference was only marginally significant.

In none of these studies have local dependencies produced a greater conformity preference than remote ones. Although no experiment directly tested for an interaction between contiguity and similarity, results have been comparable across experiments regardless of whether the segments involved are very similar, like the [s] and [ʃ] of Finley (2011a), or very different, like the [s] and [m] of Koo and Callahan (2011). The resemblance to natural-language typology, where local dependencies are the norm, is rather weak, raising the possibility that the typological bias may not be wholly due to inductive bias.

2.3 Summary: Complexity

There is abundant converging evidence that formal complexity impedes acquisition of artificial phonological patterns, in the sense that performance drops as the number of relevant features increases. The strongest result is that Type I patterns are easier than Type II, which are easier than Type VI (see §2.1 and Table 1). There is also evidence that stimulus-response mappings which change fewer features are easier than those which change more, and that within-stimulus dependencies are easier when they involve two instances of the same feature (i.e., assimilation or dissimilation) than instances of two different features. If the same inductive biases affect natural-language phonology, they should leave visible marks on cross-linguistic typology and within-language productivity.

A complexity bias in natural-language phonology would make more complex patterns harder to learn, hence harder to innovate and more likely to be changed (simplified) in transmission from one generation to the next (Bach and Harms, 1972). That in turn could lead to low long-term steady-state frequencies for the corresponding patterns (Bell, 1970, 1971; Greenberg, 1978). Phonologists have in fact noted informally that the patterns they discover tend to be featurally simple (e.g., Chomsky and Halle 1968, 401, Hayes 1999, Pierrehumbert 2001). The available quantitative evidence tends to confirm this observation.

Clements (2003) found that inventories tend to avoid both "holes" and "bumps": A given segment is more likely if all of its feature values are shared by other segments, and less likely if some of them are not. He proposes that inventories tend to maximize *feature economy*, the ratio of the number of segments in an inventory to the number of features required to distinguish among inventory members. Feature economy favors Shepard Type I inventories over those of Types III—VI. An example is shown in Figure (4). All three inventories contain four sounds, but the Type I inventories use only two contrastive features (feature economy index = 4/2 = 2), while the Type V inventory uses three (feature economy index = 4/3).

(4) *The probability that a segment will occur in an inventory increases if the inventory contains other segments minimally different from it. (Extrapolated from Clements, 2003, Figure 11).*

Favored				*Disfavored*	
$\boxed{p}$	$\boxed{t}$	$\boxed{p}$	$\boxed{t}$	$\boxed{p}$	$\boxed{t}$
$^*b$	$^*d$	$\boxed{b}$	$\boxed{d}$	$\boxed{b}$	$^*d$
$\boxed{f}$	$\boxed{s}$	$^*f$	$^*s$	$^*f$	$^*s$
$^*v$	$^*z$	$^*v$	$^*z$	$^*v$	$\boxed{z}$
Type I		*Type I*		*Type V*	

In a survey of 561 languages, (Mielke, 2004, 2008) studied "phonologically active classes", sets of sounds that pattern together by undergoing an alternation, triggering an alternation, or respecting a phonotactic restriction. One finding was that typologically common sound classes can usually be stated as a single feature value or a conjunction of a small number of feature values (.e.g, [–continuant –sonorant]), with typological frequency falling as the feature count rises. Of the non-conjunctive classes, most can be stated as disjunctions of conjunctions (e.g., [–sonorant Labial]∨[–sonorant Dorsal]); frequency falls as the number of disjuncts rises. Thus, featurally-complex patterns are attested but rarer, in the same way in which their artificial analogues are learnable but harder.[12]

An alternative hypothesis attributes the prevalence of simple classes to sampling error Pierrehumbert (2001, 2003). Suppose the learner decides whether to postulate a constraint by observing which of two classes, A or B, is more frequent. Classes defined by more features are rarer ($[+F+G]$ cannot outnumber $[+F]$), so if A and B both involve many features, the corpus of relevant examples will be small. For example, the learner can make a more reliable frequency comparison between $A = \{$ plosives $\}$ and $B = \{$ fricatives $\}$ than between $A' = \{$ labial plosives $\}$ and $B' = \{$ labial fricatives $\}$. Since there is more variability between smaller samples, learners will disagree more in their judgment of the relative frequency of A' and B', and hence also in the constraints they acquire. That makes highly specific ("fine-grained") constraints less likely to survive traditional transmission than very general ones. This idea may explain why conjunctive categories with more features are typologically rarer. However, it does not explain why *dis*junctive classes become rarer as the number of disjuncts goes up, since more disjuncts mean a larger class: The learner can make an even more reliable frequency comparison between $A^* = \{$ plosives and nasals $\}$ and $B^* = \{$ fricatives and laterals $\}$. In contrast, a bias towards

[12] Many of the natural-language classes in the Mielke (2004) study involve more relevant features than the artificial-phonology experiments. However, they agree where they overlap. A Type I problem uses a single-feature class. A Type II problem requires a disjunction of two two-feature conjunctions, e.g., "(black and triangle) or (not-black and non-triangle)", while Type VI needs four three-feature disjuncts (Feldman, 2000, 2006; Lafond et al., 2007).

featural simplicity can account equally well for the rarity of both complex conjunctions and complex disjunctions.

Another alternative is that the simplicity of natural-language phonology is inherited from the simplicity of natural-language phonetics. The phonological form of an utterance perceived by a listener is sometimes different from that intended by the speaker, owing to systematic distortions introduced by the articulation-transmission-perception channel. Such channel biases may serve as phonetic precursors for phonological innovations, e.g., if phonetic coarticulation by speakers is interpreted as phonological assimilation by listeners ("phonologization", Hyman 1976; Ohala 1993). If the precursors tend to be simple in phonetic terms, then their phonologizations will tend to be simple in phonological terms.

This hypothesis is not yet testable because there is no Mielke (2008)-like quantitative data on the typical complexity of phonetic precursors. In the meantime, it is more feasible to test a stronger hypothesis, namely, that phonological patterns tend to be simple *only* because precursors do (and not because inductive bias also favors innovation or preservation of simpler phonological patterns). Under this hypothesis, precursors which are equal in phonetic magnitude and differ only in complexity should be phonologized at the same rate, yield equally durable phonological patterns, and hence have the same typological frequency. On the other hand, if typology is also shaped by an inductive bias towards simplicity, the simpler precursor should be phonologized more often, yield a more-durable phonological pattern, and thus accumulate greater typological frequency.

Two such pairs of precursors have been examined in connection with a hypothesized inductive bias favoring single-feature dependencies over two-feature ones. (1) Vowel-height harmony and disharmony are typologically more common than phonological patterns relating vowel height to consonant voicing. However, the apparent phonetic precursors of these two patterns do not differ in their effect on first-formant frequency (Moreton, 2008). (2) Phonological patterns relating the height of tones in adjacent syllables are typologically more common than those relating tone height to consonant voicing, yet phonetic tonal coarticulation and tone-voice interaction do not differ in their effect on f_0 (Moreton, 2010). Both of these results suggest that the typological skew towards simplicity is at least partly due to inductive bias. They are on the one hand weakened by the use of acoustic rather than perceptual measures of phonetic-precursor magnitude (Yu, 2010, 2012), but on the other hand strengthened by evidence of an analogous inductive bias in learning experiments with English, Mandarin, and Southern Min speakers (see above, §2.2.1, p. 10).

3 Phonetic substance

Patterns of equal formal complexity can differ widely in typological frequency. For example, coronal-stop assibilation is asymmetrically triggered by following rather than preceding vocoids, and by high rather than low vocoids (Kim, 2001); vowel-height harmony is more common than consonant continuancy harmony (Hansson 2001b, 137–149, Rose and Walker 2004). Very often, the more-frequent phonological pattern is a stylized version of some kind of phonetic covariation (e.g., phonological vowel harmony resembles phonetic vowel-to-vowel coarticulation). One family of explanations proposes inductive bias favoring patterns which have phonetic motivations over those which do not (e.g., Stampe, 1973; Archangeli and Pulleyblank, 1994; Steriade, 1997; Hayes, 1999; Steriade, 2001; Wilson, 2006). The experiments that test this hypothesis typically compare a typologically-frequent, phonetically-motivated pattern to a rare, unmotivated (or even counter-motivated) pattern of the same complexity. The available studies can be divided into three categories: consonant/vowel asymmetries, segmental rules, and prosodic rules.

3.1 Consonants versus vowels

In natural language, phonological dependencies between non-adjacent vowels is thought to be much more common than dependencies between non-adjacent consonants (e.g., Gafos (1996, 7–8), Baković (2000, 5–6), Hansson (2001b, 1–2)), Attention in the lab has focused on a specific kind of non-adjacent dependency, namely repetition.

There have been several reports that a within-stimulus vowel-repetition dependency is learned faster than the analogous consonant pattern. Toro, Nespor, Mehler, and Bonatti (2008) familiarized native Italian speakers on $CVCVCV$ stimuli. In one condition, the first and last vowel were identical; in the other, the first and last consonant. Participants in the vowel condition preferred novel positive stimuli over non-conforming foils; those in the consonant condition did not (the two conditions were not directly compared to each other). The finding was replicated by Toro, Shukla, Nespor, and Endress (2008), where it was also shown to be robust against manipulation of the relative amplitudes of the consonants and vowels (again without direct comparison). It was replicated by Pons and Toro (2010) with Spanish-learning 16-month-olds; however, a direct comparison between the two conditions found no significant difference. Nevins and Toro (Nevins, 2010), in an experiment with Italian-speaking adults, directly compared consonant- with vowel-repetition patterns, and found a stronger preference for positive instances with the latter.

On the other hand, some studies have found no advantage for vowel repetition over consonant repetition. Koo (2007, Ch. 2), using a speeded-repetition paradigm with English speakers, found a conformity advantage for an [l. . . l]/[r. . . r] pattern and an [l. . . r]/[r. . . l] pattern, but not for analogous patterns with [i] and [u]. Since the responses in this experiment are utterances rather than judgments, the result may be due to the articulatory difficulty of co-occurring liquids, rather than to differences in learnability of the patterns. Two subsequent experiments in the same series with [l. . . l]/[r. . . r] and [i. . . i]/[u. . . u] found no difference in their effects on pattern-membership judgments of new stimuli.

On the whole, the experimental evidence is consistent with the hypothesis that learners are more sensitive to syntagmatic repetition of vowels than of consonants. If there is such an inductive bias, and if that bias shapes typology, we should find that patterns of non-adjacent vowel repetition outnumber the analogous consonant patterns in natural languages. Many languages have a phonological pattern in which one vowel is required to be identical to another, notably total vowel harmony (Aoki, 1968) and copy-vowel epenthesis (Kitto and de Lacy, 1999), whereas the consonantal analogues seem to us to be much rarer. However, we know of no quantitative test of this hypothesis.

Aside from the identity dependency, isomorphic consonant and vowel patterns have seldom been directly compared. Zaba (2008) found no differences between vowel backness agreement and two consonant featural-agreement patterns (see §3.2). Moreton (2012), using a paradigm described above (p. 10), found a modest but significant advantage for two phonetically-systematic non-adjacent vowel dependencies over two phonetically-systematic non-adjacent consonant dependencies. It is not known whether this difference would persist if other, possibly more salient, consonant features were used instead of place and voicing (LaRiviere et al., 1974, 1977).

3.2 Segmental harmony

In natural language, dissimilation is thought to be rarer than assimilation both synchronically (Bye, 2011) and diachronically (Campbell, 2004, 30). Some studies have tested corresponding agreement and disagreement patterns side by side, but there is hardly any evidence that agreement is easier to learn. As described above (§2.2), Pycha et al. (2003) found no difference in learnability between analogues of backness harmony and backness disharmony; Wilson (2003) found no difference between nasal agreement and disagreement; and Kuo 2009, 139 found no difference between place agreement and place disagreement (labial glide iff labial stop, vs. labial glide iff coronal stop). Similar results have been found in other studies as well.

Koo (2007, §2.1; Koo and Cole, 2006) used a speeded-repetition paradigm to familiarize English speakers with long-distance agreement and disagreement patterns for liquid laterality ([l] vs. [ɹ]) and vowel backness ([i e] vs. [a u]). New pattern-conforming test items were repeated quicker than non-conforming foils in both of the liquid conditions, and in neither of the backness conditions. The only sign of a harmonic-disharmonic difference was a higher error rate on non-conforming items in the liquid-agreement condition, alongside no difference in the liquid-disagreemnt condition. (The interpretation of these experiments is complicated somewhat by the absence of a baseline condition in which participants are familiarized on neither pattern. Since we do not know what performance is like for speakers unfamiliar with both pattern, we cannot be sure how much of the performance in the critical conditions is due to learning.)

Finally, Skoruppa and Peperkamp (2011) modified the non-low front vowels in Standard French words to create "Harmonic French", in which the vowels of a word agree in rounding (e.g., *pudeur* [pydœʁ] → [pydɛʁ]), and "Disharmonic French", in which vowels in odd- and even-numbered positions disagree in rounding (e.g., *ordinaire* [ɔʁdinɛʁ] → [ɔʁdinœʁ]). Participants were familiarized by hearing stories in one or the other accent. then tested by hearing paired Harmonic and Disharmonic French versions of words and trying to choose the one in the familiar "accent". This they did better than chance, but it made no significant difference whether the familiarized accent was Harmonic or Disharmonic. Peperkamp and Skoruppa conclude (p. 356) that there are "no abstract linguistic biases favoring harmony over disharmony in perceptual phonological learning"; the preponderance of assimilation over dissimilation in natural language must have other causes.

Other studies have compared harmony patterns differing in typological frequency. Usually, no significant advantage is found for the analogue of the more-common pattern.

In natural languages, rounding harmony in mid vowels asymmetrically implies rounding harmony in high vowels (Kaun, 2004). In an experimental analogue, Finley and Badecker (2009, Exp. 3) familiarized English speakers, by passive listening, with stimuli of the form $X \ldots XY$ (where X was a $CVCV$ nonsense word and Y a CV suffix), and then asked them to choose the positive member of a pair of new stimuli XZ_1, XZ_2. The vowel of Y agreed in backness and rounding with those of X. For one group of participants, the vowels of Y were high and those of the Z's were mid; the reverse was true for the other group. Although familiarized participants chose the positive stimuli much more often than did unfamiliarized control listeners, the rate did not differ between the two familiarization groups. In other words, no analogue of the asymmetrical implication in natural language typology was found.

Zaba (2008, Ch. 2) used a paradigm similar to that of Pycha et al. (2003) to compare artificial analogues of three patterns differing in natural-language typological frequency: backness agreement between non-adjacent vowels (common), nasality agreement between non-adjacent consonants (rare), and labiality ageement between non-adjacent consonants (unattested in adult language). English-speaking participants were familiarized by listening to stimuli of the form $X \ldots XY$, where X ended in $\ldots V_1C_1$, and Y could be either of two V_2C_2 sequences. The pattern deterimined the choice of Y as a function of V_1 (for backness agreement) or C_1 (for nasality and labiality agreement). Learning was then tested by asking participants to judge whether $X \ldots XY$ sequences conformed to the trained pattern. A control group of participants was likewise tested after being familiarized with $X \ldots XY$ stimuli in which Y was always the same. The results were analyzed in several different ways, none of which found significant differences between the three pattern conditions.

Sibilant harmony in natural languages can be triggered by only [–anterior] segments, only [+anterior] segments, or both. The first of these possibilities is much rarer typologically than the other two. Kosa (2010) compared two artificial analogues of rightward-spreading sibilant-harmony patterns, one triggered by [–anterior] and the other by [+anterior]. Words had the form XY, where X was one of many $CVCV$ sequences. When X contained no sibilants, Y was one of [næ], [sʌt], or [ʃɑ]. For participants in the [–anterior] condition, [sɑ] replaced [ʃɑ] when X contained a [–anterior] sibilant; for those in the [+anterior] condition, [ʃʌt] replaced [sʌt] when X contained a [+anterior] sibilant. In Experiment A, English-speaking participants listened to pattern-conforming XY words and then gave familiar/unfamiliar judgments to conforming and non-conforming test items. Novel items were more likely to be judged familiar when they fit the pattern, but the difference reached significance only in the [+anterior] condition (the analogue of the common natural-language pattern). However, the analysis did not directly compare the two conditions, so we do not know whether they differed significantly from each other. In Experiment B, participants listened to and repeated pattern-conforming XY words, then heard conforming-nonconforming pairs and tried to choose the conforming word. Proportion correct was significantly above chance in both conditions, but did not differ between them.

3.3 Consonant-vowel interactions

Wilson (2006) focused on two typological asymmetries in rules changing velars [k g] to palatoalveolars [tʃ dʒ] as a function of vowel context. One is that palatalization before more-back vowels implies palatalization before less-back ones; the other, that palatalization of voiced velars implies that of voiceless ones. English speakers were trained in a language game to respond to a subset of [ki gi ke ge] with [tʃi dʒi tʃe dʒe], and to both of [ka ga] with [ka ga] (the critical syllables occurred initially in disyllabic nonsense words). They were then tested on a mix of old and new stimuli to measure their velar-palatalization rate in different conditions. Experiment 1 focused on the effect of vowel context; Experiment 2, that of consonant voicing. A synopsis is given in Table (5).

(5) *Critical experimental conditions of Wilson (2006). n, number of stimuli; p, probability of velar palatalization (in the training stimuli or in the test responses).*

			ki		*ke*		*ka*		*gi*		*ge*		*ga*	
Exp	*Cond*	*Phase*	*n*	*p*	*n*	*p*	*n*	*p*	*n*	*p*	*n*	*p*	*n*	*p*
1	*i*	*Train*	*4*	*1.00*	–	–	*3*	*0.00*	*4*	*1.00*	–	–	*3*	*0.00*
		Test	*8*	*0.44*	*8*	*0.13*	*6*	*0.05*	*8*	*0.52*	*8*	*0.14*	*6*	*0.14*
1	*e*	*Train*	–	–	*4*	*1.00*	*3*	*0.00*	–	–	*4*	*1.00*	*3*	*0.00*
		Test	*8*	*0.20*	*8*	*0.19*	*6*	*0.15*	*8*	*0.48*	*8*	*0.49*	*6*	*0.39*
2	*k*	*Train*	*4*	*1.00*	*4*	*1.00*	*3*	*0.00*	*1*	*1.00*	*1*	*1.00*	*3*	*0.00*
		Test	*8*	*0.39*	*8*	*0.36*	*6*	*0.12*	*8*	*0.14*	*8*	*0.11*	*6*	*0.09*
2	*g*	*Train*	*1*	*1.00*	*1*	*1.00*	*3*	*0.00*	*4*	*1.00*	*4*	*1.00*	*3*	*0.00*
		Test	*8*	*0.26*	*8*	*0.20*	*6*	*0.00*	*8*	*0.50*	*8*	*0.44*	*6*	*0.23*

The rate of palatoalveolar responses to velar-initial stimuli was bimodal, with clusters around 15% (Low) and 45% (Medium), rates comparable to those found by Peperkamp et al. (2006). Changing the features made a clear difference in performance in Experiment 1. Participants trained to respond [tʃe dʒe ka ga] (1e) had the Medium rate of palatoalveolar responses to [gi ge ga], but the Low rate to [ki ke ka]. Participants trained on [tʃi dʒi ka ga] (1i) had the Medium rate only on [ki] and [gi] themselves. These results are to some extent consistent with typology, since velar palatalization before [e] asymmetrically implies velar palatalization before [i].

However, there are also differences. Participants in Condition 1e, unlike natural-language palatalization rules, disregarded vowel context entirely. Their palatalization rates were indistinguishable before [e], where they had been trained to palatalize, [a], where they had been trained *not* to palatalize, and [i], where they had been given no training (Wilson, *op. cit.*, Fig. 2). In natural language, [g]-palatalization asymmetrically implies [k]-palatalization. If this is a result of inductive bias, we expect the learner to interpret observed [g]-palatalization as evidence for [k]-palatalization, but not the other way around. This was not borne out in the experiments. In Experiment 1, [g] was palatalized significantly more often than [k], despite equal training on both. In Experiment 2, there was no significant difference between the effects of [k]-training on [g]-palatalization and that of [g]-training on [k]-palatalization.

3.4 Prosodic rules

Carpenter (2005, 2006, 2010) investigated acquisition of artificial stress patterns by native speakers of American English and Quebec French. Participants were trained to choose between correctly- and incorrectly-stressed versions of the same word, then tested with novel words. One set of experiments used sonority-sensitive stress, comparing the typologically uncommon "leftmost low vowel ([æ a]), else leftmost vowel" with the completely unattested "leftmost high vowel ([i u]), else leftmost vowel". Participants in both native-language groups preferred positive items more strongly in the leftmost-low condition. Here the lab results are aligned with typology.[13]

A second set of experiments in this series used quantity-sensitive stress, comparing the typologically frequent "leftmost heavy (CVC), else leftmost" pattern with the unattested "leftmost light (CV), else leftmost". Since the typology is much more strongly skewed here than in the case of sonority-sensitve stress, one would expect the same of learning performance. However, although preference for positive stimuli was significantly above chance in all conditions, it did not differ significantly between the two artificial patterns in either native-language group.

Stress in natural language is influenced much more often by codas than by onsets. In an experiment with 9-month-old English-learning infants, Gerken and Bollt (2008, Experiment 2) used a familiarization set of polysyllabic pseudowords in which stress was determined in part by the form of the syllables and in part by other principles. The initial $CV(:)$ syllable was always stressed.[14] In one condition, stress preferentially fell on syllables of the form $CV:C$ (e.g., /boʊm/) rather than those of the form $CV:$ (e.g., /mi/) or CV (e.g., /fa/).

[13] Although the vowels were equalized for duration and peak intensity to remove phonetic stress cues, the author cautions us that participants may have restored low vowels' inherently greater duration and intensity by subvocalizing the stimuli (Carpenter, 2010). However, the unedited high vowels were tense, while the unedited low ones were lax. Editing lengthened high vowels and shortened low ones (Carpenter, 2006, 72), exaggerating the tense/lax contrast in a way that would encourage English speakers to stress the high vowels, not the low ones.

[14] In another version of the experiment, initial and final syllables traded roles.

In the other, stress preferentially fell on syllables that began with /t/ rather than any of /d m f s l ɹ/. Test items always followed both of these form-based generalizations by stressing the unique [t]$V:C$ syllable, but, unlike familiarization items, never stressed the initial CV(:). In half of the test items, these two critical syllables were adjacent and could not both be stressed without clashing. For the other half, both critical syllables could have been stressed without clash (but were not). The rationale is that if participants learn the form-based generalization, they will, in effect, realize that test items of the first type have a valid excuse for not stressing the initial syllable, and be more likely to accept them as pattern-conforming. Infants in both conditions listened longer to the test items of the first kind than the second, but the difference only reached significance in the condition where stress was coda-sensitive.

We are reluctant to accept these results as evidence of an inductive bias in favor of coda-sensitive over onset-sensitive stress. First, a direct comparison between the onset and coda conditions (p. 241) found no significant differences between them (i.e., performance in one condition differed significantly from chance, while that in the other did not, but the two conditions did not differ significantly from each other). Second, the two patterns were not directly comparable. The coda-sensitive pattern depended on the presence vs. absence of a consonant, while the onset-sensitive one depended on the identity of an ever-present consonant. The former kind of difference may be more salient to infant pattern learners (Saffran and Thiessen, 2003).

Schane et al. (1974) trained participants to translate English adjective-plus-noun phrases word for word into an artificial language which had a context-sensitive rule deleting the final consonant of the first word (Figure 6). In one condition, deletion applied when the second word began with a consonant, simulating a cross-linguistically common process of cluster simplification (Wilson, 2001). Deletion in the other condition occurred when the second word was vowel-initial. Such intervocalic deletion in nature rarely applies to the voiceless obstruents used in this experiment (Picard, 2003).

(6) *Artificial-language conditions used by Schane et al. (1974).*

Condition	$/C_1\#C_2/$ /ˈtupak ˈsipu/	$/C_1\#V_2/$ /ˈtupak ˈoga/
Cluster simplification	$\rightarrow [\#C_2]$ [ˈtupa ˈsipu]	[ˈtupak ˈoga]
Intervocalic C deletion	[ˈtupak ˈsipu]	$\rightarrow [\#V_2]$ [ˈtupa ˈoga]

Participants were trained to a performance criterion. The cluster-simplification groups reached criterion before the intervocalic-deletion ones, and were less likely to erroneously give responses appropriate to the other group.

In both conditions, the presence of a consonant at the end of Word 1 was correlated with the presence of a consonant at the beginning of Word 2. Why would the negative correlation be easier to learn than the positive one? One possibility is substantive bias against marked phonological structures such as hiatus or syllable codas. Alternatively, participants in the cluster-simplification condition could have learned the pattern "exactly one intervocalic consonant", while those in the intervocalic-deletion condition would have had to learn the disjunction "zero or two intervocalic consonants". The results could then be accounted for by a complexity bias, the relative difficulty of disjunctive categories (Bruner et al., 1956; Conant and Trabasso, 1964; Ciborowski and Cole, 1972).

3.5 Summary: Substance

The best-supported finding, replicated with different stimuli and with speakers of different ages and languages, that it is easier to induce the pattern "contains the same vowel twice" than "contains the same consonant twice". The corresponding natural-language phenomena are typologically small potatoes. Meanwhile, the clearest non-finding is that, despite several attempts motivated by the typological importance of the patterns, no study has yet found positive evidence that vowel or consonant harmony is easier to learn than disharmony.

For other segmentally- and prosodically-defined patterns, results have been mixed at best, even within the same study. Even when positive (statistically significant) results are found, they may fail to match typology. For example, [g] was palatalized more than [k] in Wilson (2006, Experiment 1), whereas in natural language, [g]-palatalization asymmetrically implies [k]-palatalization. Such mismatches are hard to account for. When complexity bias mismatches typological frequency, the disparity may be due to the competing effects of channel bias (Moreton, 2008). This explanation is not available when substantive bias mismatches typology, since channel bias is predicted to *reinforce* substantive bias in shaping typology.

4 General discussion

An isolated "null result" (lack of significant difference between conditions) is hardly fatal to a hypothesis. It could be due to sampling error, or a flawed experiment. However, the null results in substantive-bias experiments are not isolated; they contrast with the positive results routinely obtained for complexity differences using similar experimental paradigms and participant populations. In §2.1, for example, we saw that significant advantages for featurally-simpler patterns were found by Pycha et al. (2003); Saffran and Thiessen (2003); Peperkamp et al. (2006); Cristiá and Seidl (2008); Kuo (2009); Skoruppa et al. (2009); Chambers et al. (2010); Skoruppa and Peperkamp (2011). Only one study yielded ambiguous results: LaRiviere et al. (1974, 1977) found a numerical advantage for the simpler pattern in ten experiments, but it only reached statistical significance in three of them. (The complexity affect vanishes in younger infants; Cristià et al. 2011,?). On the other hand, when experimenters compared phonetically-motived harmony patterns with corresponding disharmony patterns (§3.2), the opposite occurred: Not one study found an advantage for the harmony pattern (Pycha et al., 2003; Wilson, 2003; Koo and Cole, 2006; Kuo, 2009; Skoruppa and Peperkamp, 2011). Each individual null result could be attributed to bad luck, but the concentration of null results on the substance side of the ledger cannot be.[15]

The lab results so far thus add up to substantial and, in our view, convincing evidence that formal complexity impedes artificial-phonology learning in adults and infants, but the effect of phonetic substance is weaker if it exists at all. These findings corroborate the suggestion of Pycha et al. (2003) that complexity bias is stronger than substantive bias in artificial phonology. The leap from the existing artificial-phonology evidence to firm conclusions about natural-language phonology is a big one. It is decidedly better than a leap from no evidence at all, but much work needs to be done before the laboratory study of inductive biases approaches the level of reliability and sophistication found in that of channel biases. In the following sections, we consider some of the theoretical, empirical, and practical issues involved, and offer some suggestions for narrowing the gap.

4.1 An analytic artifact?

Calculations of both complexity and substance are affected by the analyst's choice of featural primitives, and phonological-feature models are deliberately engineered to encode substantive bias as complexity bias (McCarthy, 1988; Clements and Hume, 1995). This theoretical prejudice could inflate the apparent frequency of complexity biases at the expense of substantive ones. For example, when Cristiá and Seidl (2008) found that [m n b k] vs. [f z] was easier to learn than [m n f z] vs. [b k], we interpreted it as evidence that Type I patterns are easier than Type II (see 2b above). That interpretation depended on the absence of a phonological feature distinguishing [m n f z] from [b k]. If the learner in fact has such a feature (e.g., [continuous airflow] or [stable spectrum]), then reluctance to use it would constitute a substantive bias against a Type I pattern, rather than a structural bias against a Type II pattern.

Other experiments, however, resist such reanalysis. To convert the Saffran and Thiessen (2003) experiment from Type I vs. Type II into Type I vs. Type I, we would have to credit the learner with a phonological feature that separates [p d k] from [b t g]. To convert Kuo (2009)'s Type II vs. Type VI into Type II vs. Type II would require one that separates [p^h t] from [p t^h], and Pycha et al. (2003)'s II/VI comparison would need one that separated [i æ ʊ] from [u a ɪ]. Such phonetically arbitrary, *post hoc* features cannot be universal, nor learned from prior experience; they could only have been learned in the experiment itself. In that case, the unifying principle behind the experimental results would be that a pattern

[15] If we count the studies listed in this paragraph as 8 successes and 1 failure among the complexity studies, and no successes and 5 failures among the substance studies, then the difference in success rate is significant by Fisher's exact test ($p < 0.003$). A more sophisticated statistical meta-analysis would compare individual experiments, taking into account properties shared between experiments.

is harder when a phonetically complex feature has to be induced — a structural bias, not a substantive one.

The substantive-bias hypothesis is not just about classes of sounds, but about their behavior. A pattern does not qualify as phonetically motivated merely by employing a phonetically-defined segment class; that class has to *do* something for which there is a phonetic explanation. For example, final-obstruent devoicing and final-obstruent voicing both involve obstruents, but only the former is phonetically motivated (Steriade, 1997; Kiparsky, 2008). Although the Type I patterns used by LaRiviere et al. (1974, 1977); Saffran and Thiessen (2003), and Cristiá and Seidl (2008) were built on typologically common sound classes, the patterns themselves were phonetically unmotivated and typologically rare (e.g., "onsets must be nasal or oral stops"). Phonetic motivation therefore does not explain why they were learned better than equally unmotivated Type II patterns. Pycha et al. (2003) found no statistically significant difference in difficulty between a phonetically motivated, typologically common Type II pattern of vowel backness agreement and a phonetically less-motivated, typologically rare Type II pattern of backness disagreement, though both proved easier than a Type VI pattern. The three Type II patterns compared by Zaba (2008, Ch. 2) differed widely in phonetic motivation and typological frequency, but not in learning outcome.

4.2 Is artificial phonology phonology at all?

Perhaps participants are treating the artificial-phonology task as if it were a non-linguistic concept-learning task about red triangles or fictitious diseases. Use of the same domain-general cognitive processes would predict shared complexity biases and lack of substantive bias. It may be that such domain-general mechanisms are involved in natural-language phonology as well (Hume and Johnson, 2001). If so, then artificial phonology is informative about natural-language phonology after all. However, it is also possible that natural-language phonology is learned using a separate set of dedicated processes, and hence that artificial phonology is irrelevant to it. Several strands of evidence are germane to this hypothesis.

The sharing of structural biases between artificial phonology and non-linguistic category learning does not prove that they share a common mechanism to the exclusion of natural-language phonology. The currently known shared biases are few, and most are so generic that even radically different learning algorithms can share them (Gluck and Bower, 1988; Kruschke, 1992; Nosofsky et al., 1994; Love et al., 2004; Feldman, 2006). Nature is not guaranteed to be parsimonious, sometimes building two physically distinct circuits where behavioral similarity would lead us to expect one (e.g., Simons-Wiedenmaier et al., 2006). Furthermore, artificial phonology also shares structural biases with natural-language phonological typology (§2.3).

Since substantive biases are by nature domain-specific, they could provide very strong evidence that artificial phonology is learned using cognitive processes peculiar to phonology

— which could only be the same processes used for learning natural-language phonology. The more arbitrary and stipulative the bias, the stronger would be the evidence that artificial phonology experiments are directly relevant to natural language. However, as we have seen, the evidence for substantive biases is weak.

The separate-mechanism hypothesis would also be strengthened if there were an age at which infants show no knowledge of first-language phonology, but can rapidly learn artificial phonology. Like adults, infants receive much more exposure to the ambient language than to any experimental one. If artificial and natural phonology are acquired by the same mechanisms, then anyone old enough to be influenced by a few minutes' exposure to an artificial pattern should also be old enough to be influenced by comparable patterns in the phonology of the ambient language (where "comparable" means that the patterns have similar complexity and are instantiated with equal consistency in the respective training-data sets). For example, by six months of age, Turkish-learning infants, unlike German-learning ones, spontaneously prefer disyllabic pseudowords which have backness harmony over those which lack it (van Kampen et al., 2008). If Turkish-learning 4.5-month-olds do not have this preference spontaneously, but do acquire it from brief exposure to Turkish-like stimuli in an artificial-language experiment, that would support the separate-mechanism hypothesis.

Infants start showing sensitivity to first-language phonotactic patterns between 6 and 9 months of age (Jusczyk et al., 1993; Friederici and Wessels, 1993; Jusczyk et al., 1994; Mattys and Jusczyk, 2001). Cristià and colleagues (Cristià, Seidl, and Gerken, 2011; Cristià, Seidl, and Francis), using a paradigm described above (p. 5), found that infants as young as 4 months acquired a novelty preference after a few minutes' experience with a positional restriction on consonant occurrence, succeeding on both a Type I and a Type II pattern where 7-month-olds failed on the Type II pattern. Reviewing a range of studies, they conclude that complexity effects begin to emerge at around 7 to 9 months.

Interactions between natural and artificial phonology are evidence that the two are not wholly separate processes. Schane et al. (1974)'s cluster-simplification pattern was, by design, similar to French liaison, and was acquired better by those who had had more exposure to French. Healy and Levitt (1980, Experiment 2) found better performance with a pattern that resembled English voicing agreement (in -*ed*, -*'s*, etc.) than for one that did not. Pater and Tessier (2006), again with English speakers, found better performance on an artificial epenthesis rule when it was triggered by violation of the English minimal-word constraint than when it was triggered by vowel backness. It is not clear whether natural language phonology had these effects by facilitating rule learning in training, or by biasing participants' responses against non-conforming responses in testing, but it is clear that some aspect of the artificial-phonology task are ecologically valid enough to engage natural-language phonological preferences.

4.3 Exposure and robustness

An obvious difference between artificial-phonology experiments and natural first- or second-language phonology is amount and duration of exposure. Most of the studies reviewed above used less than 30 minutes of familiarization or training, and some used much less. Taylor and Houghton (2005, Exp. 5) used a tongue-twister paradigm to familiarize English speakers with a pattern that restricted different consonants to different syllabic positions (onset vs. coda). They then reversed the pattern unannounced in the middle of a block. Before the switch, transposition errors adhered to the original constraints 98% of the time. After it, the new constraints were adhered to 65% of the time, with the violations concentrated in the first 9 trials after the switch. In contrast, first-language phonology is notoriously persistent, even in highly motivated speakers with ample exposure to second-language input (Cutler et al., 1989; Darcy, 2006; Kager et al., 2008).[16]

Artificial phonology thus appears easier to lose than first-language phonology and easier to acquire than second-language phonology, which seems to support the different-mechanism hypothesis. However, the difference may not be between a linguistic and a non-linguistic mechanism, but between a short-term and a long-term one. In most artificial-phonology experiments, participants are tested immediately after training, which gives only those learning processes that are active during training an opportunity to apply. Learning actually continues long after the stimuli are experienced, as memories are consolidated and reprocessed. Several studies have found increased sensitivity to patterns when the interval between training and test includes sleep compared to when it does not (Wagner et al., 2004; Gómez et al., 2006; St. Clair and Monaghan, 2008). Integration of new words into the lexicon likewise continues after exposure (see Lindsay and Gaskell 2010 for a recent review), perhaps affecting the time course of their availability to pattern-finding processes. Artificial-phonology experiments using longer time scales may be therefore able to detect subtler effects than the ones reviewed above.

4.4 Structurally biased phonology

In short-term artificial-phonology learning, the effects of phonetic substance, if they exist at all, are overshadowed by those of formal complexity. If natural-language phonological learning works the same way, then the broader picture is essentially that proposed by Bach and Harms (1972): Inductive bias, a property of the learner's pattern-detection processes, facilitates faithful acquisition of simple patterns and rejection or innovative simplification of more complex ones, but is (relatively) insensitive to their phonetic motivation. Channel

[16] *Some* first-language patterns can be overcome with a small amount of lab exposure. Carpenter (2005, 2006, 2010) succeeded in teaching English speakers the artificial stress rule "stress leftmost low vowel, else leftmost vowel". Since the low vowels were also lax and the high ones also tense, the artificial pattern overcame the English L1 preference for stressing heavy syllables. Kuo (2009) found that Mandarin speakers' L1 phonotactics did not affect their ability to acquire L1-conforming vs. L1-nonconforming artificial patterns from brief exposure.

bias, a property of the articulatory-acoustic-perceptual channel, systematically distorts the phonological form of utterances in transmission, introducing new patterns into the learner's input (Hyman, 1976; Ohala, 1992, 1993). Together, the complexity-based inductive biases and the phonetic channel biases cause systematic asymmetries in the direction of language change, and hence in the long-term steady-state typological frequencies of different kinds of pattern (Bell, 1970, 1971; Greenberg, 1978).[17] Over time, phonological patterns are predicted to become featurally simpler and less phonetically-motivated (Bach and Harms, 1972).

This view could be called *Structurally-Biased Phonology*. Its main feature is that it separates phonetic from structural factors, identifying the former with channel bias and the latter with inductive bias, while abolishing the distinction between "synchronic" and "diachronic" factors — both kinds of bias are observable synchronically, in the lab and in their natural setting, and both shape typology only by skewing the outcome of diachronic change.[18] Elaborating and testing this hypothesis raises several research questions, a few of which we will touch on here (see our companion paper, Moreton and Pater (in preparation), for fuller discussion).

What are the inductive biases? What do they reveal about the learner? The work reviewed in this paper is a good start, but much still needs doing to perfect methodologies for studying inductive biases, to discover what kinds of pattern they favor, and to determine what properties a learner would have to have in order to exhibit them. Since a major source of inductive bias is the architecture of the learner, this mission will necessarily involve contact with models of non-linguistic pattern learning in psychology and computer science. Another major source of inductive bias is the phonological-representation schema used by the learner, since complexity is only meaningful in terms of the representational primitives available for expressing generalizations. This was the central insight behind Feature Geometry, a framework in which the feature system was engineered to allow all and only typologically common, phonetically transparent processes to be expressed by elementary operations of spreading and delinking (Clements 1985; McCarthy 1988; Clements and Hume 1995; for a critical review, see Padgett 2002a, 2002b), while more complex processes required unlikely rule conspiracies (Pulleyblank, 1988, 299). Structurally-Biased Phonology requires a renewed focus on feature systems, and provides a new tool for assessing them: their ability to predict complexity effects in the lab and in typology. (For example, the Cristiá and Seidl (2008) experiment, as discussed in §4.1 is evidence that human infants do not analyze patterns in terms of [continuous airflow] or [stable spectrum].)

[17] Social factors are crucial to the initiation, continuation, and spread of change (Weinreich et al., 1968), but we expect them to be neutral as to the phonetic or structural content of the patterns. The typological preponderance of vowel harmony over vowel disharmony, for example, does not have a social explanation. The nonexistence of substantive inductive bias has of course been proposed before, e.g., Ohala (1997); Hale and Reiss (2000); Hume and Johnson (2001).

[18] Structurally-Biased Phonology differs from Kiparsky's (2006, 2008) Amphichronic Program in this separation of phonetic from structural factors.

What are the channel biases? The richness of the phonetic literature and the concreteness of the variables creates the impression of solidity: In using channel bias to explain phonological typology, we seem to be using the known to illuminate the unknown. In fact, we still know next to nothing about the relative magnitudes of different channel biases across languages. A channel bias is the probability of a particular kind of misunderstanding between speaker and hearer; e.g., the probability that the listener will hear as harmonic a vowel sequence intended by the listener to be disharmonic. In order to make a prediction about phonological typology, such as whether vowel height harmony should be more frequent across languages than consonant voicing harmony, we need to know at the very least which of them is more strongly favored by channel bias across languages. The data needed to do this does not yet exist. The most convenient point of access to the speaker-hearer channel is its acoustic stage, and in a few cases there are pairs of phonetic dimensions whose interaction has been measured acoustically in a relatively large number of languages; for instance, vowel height and vowel f_0 Whalen and Levitt (1995), or consonant voicing and vowel f_0 (reviewed in Moreton (2010)[19] However, these studies are few, and, as Yu (2012) notes, they only measure acoustic effects, whereas channel bias is hypothesized to affect typology via mis*perception*. Nor do acoustic measurements allow us to compare the magnitude of channel biases affecting different phonetic dimensions, such as height harmony and voicing harmony (Moreton, 2008). This lack of knowledge should not be a deterrent, but a stimulus, to the rigorous testing of hypotheses about channel biases through the study of phonetic typology.

How complex are channel biases? Much previous research has asked whether inductive biases favor phonetically-motivated patterns. Once we turn our attention to explaining the featural complexity of natural-language patterns, we have to ask the complementary question: Do channel biases favor featurally simple patterns? To illustrate the point that inductive bias simplifies complex phonetic relationships, Hayes (1999) demonstrates the complex relationship between the difficulty of stop-closure voicing on the one hand and place of articulation, closure duration, phrasal position, and adjacent nasality on the other.[20] Is that complexity typical of phonetic interactions in general, or can they also be a source of simpicity?

What is the predicted typology? One straightforward prediction is that, other things being equal, if two phonological patterns have the same structure, the typologically more-frequent one should be favored by channel bias, whereas if they are equally favored by channel bias, the more frequent one should be favored by inductive bias (Moreton,

[19] Even these studies did not really measure the interaction between two phonetic dimensions. Instead, the independent variable was actually phonological, e.g., Whalen and Levitt (1995) looked at how vowel f_0 was affected by vowel height, not by vowel F_1.

[20] Production difficulty is not itself a channel bias, since it is not a probability of an error in transmission between speaker and hearer. However, difficult productions are likely to be produced wrong (that is what it means to say that they are difficult), and hence likely to lead to transmission errors. Consequently, production difficulty is probably positively correlated with channel bias.

2008). However, since phonological typology is hypothesized to be the long-term steady-state outcome of a Markov process, its relationship to the magnitudes of channel and inductive biases may not be transparent. Something which is unlikely to be innovated can become very frequent if it is even less likely to be extinguished, for example. We expect that in many cases iterated-learning simulations will be needed to derive predictions of the steady-state frequencies. These may also yield testable predictions about the frequency and direction of historical change.

What is the actual typology of structural complexity? Phonological typology is usually stated in substantive terms, classifying phonological patterns as vowel harmony, nasal place assimilation, etc. Structurally Biased Phonology makes predictions about the typology of structure and complexity, which has been little studied (Mielke, 2004, e.g.,). How frequent are, for instance, the different Shepard types in the phonological patterns of the world? And do those frequencies differ from those we would expect in the absence of specific complexity biases in the learner? This question leads to the theoretical problem of what the appropriate chance model is (Mackie and Mielke, 2011).

How does structure affect knowledge of language? In natural language, Structurally-Biased Phonology predicts that less-complex patterns will be learned better than more-complex ones, and that structurally ambiguous patterns will tend to be generalized in simple ways rather than complex ones. In "poverty-of-the-stimulus" experiments, the experimenter chooses an aspect of the native language phonology which is consistent with two different generalizations, then tests participants' reactions to novel stimulus types to discover which pattern they have actually acquired and infer the inductive bias that guided that acquisition. Complementary "surfeit-of-the-stimulus" (Becker et al., 2007, 2011) experiments compare the productivity of two patterns in the language. Experiments of this sort have been focused on substantive bias (Pertz and Bever, 1975; Davidson et al., 2004; Davidson, 2006; Kawahara, 2006; Zhang et al., 2006; Becker et al., 2007; Zuraw, 2007; Berent et al., 2007, 2008; Moreton et al., 2008; Berent et al., 2009; Hayes et al., 2009; Zhang and Lai, 2010; Becker et al., 2011; Daland et al., 2011; Zhang et al., 2011), but seeing how complexity affects the acquisition of artificial phonology, one would expect strong effects in first- and second-language acquisition as well.

Theories of phonological typology tend to rely primarily on one explanatory factor, channel bias or inductive bias, despite outside evidence that both kinds of bias exist. This reliance may be motivated by regard for theoretical parsimony, since the number of two-factor hypotheses seems unmanageably larger than the already large number of single-factor hypotheses. The main theoretical significance of the evidence we have reviewed in this article is that this dilemma may be avoidable. What Structurally-Biased Phonology offers is a way to control the combinatorial explosion while recognizing the existence and typological effectiveness of both factors.

References

Alderete, J. and S. A. Frisch (2008). Dissimilation in grammar and the lexicon. In P. de Lacy (Ed.), *The Cambridge handbook of phonology*, Chapter 16, pp. 379–398. Cambridge, England: Cambridge University Press.

Aoki, H. (1968). Towards a typology of vowel harmony. *International Journal of American Linguistics 34*(2), 142–145.

Archangeli, D. and D. Pulleyblank (1994). *Grounded Phonology.* Cambridge, Massachusetts: MIT Press.

Archangeli, D. and D. Pulleyblank (2011). Harmony. In P. de Lacy (Ed.), *The Cambridge handbook of phonology*, Chapter 15, pp. 353–378. Oxford, England: Cambridge University Press.

Arnold, P. G. and G. H. Bower (1972). Perceptual conditions affecting ease of association. *Journal of Experimental Psychology 93*(1), 176–180.

Asch, S. E. (1969). A reformulation of the problem of associations. *American Psychologist 24*(2), 92–101.

Bach, E. and R. T. Harms (1972). How do languages get crazy rules? In R. P. Stockwell and R. K. S. Macaulay (Eds.), *Linguistic change and generative theory*, Chapter 1, pp. 1–21. Bloomington: Indiana University Press.

Bakovć, E. (2011). Local assimilation and constraint interaction. In P. de Lacy (Ed.), *The Cambridge handbook of phonology*, Chapter 14, pp. 335–352. Oxford, England: Cambridge University Press.

Baković, E. (2000). *Harmony, dominance, and control.* Ph. D. thesis, Rutgers University, New Brunswick, New Jersey.

Becker, M., N. Ketrez, and A. Nevins (2007). Where and why to ignore lexical patterns in Turkish obstruent alternations. Handout, 81st annual meeting of the Linguistic Society of America, Anaheim, California.

Becker, M., N. Ketrez, and A. Nevins (2011). The surfeit of the stimulus: analytic biases filter lexical statistics in Turkish laryngeal alternations. *Language 87*(1), 84–125.

Bell, A. (1970). *A state-process approach to syllabicity and syllabic structure.* Ph. D. thesis, Stanford University.

Bell, A. (1971). Some patterns of the occurrence and formation of syllabic structure. *Working Papers on Language Universals 6*, 23–138.

Berent, I., T. Lennertz, J. Jun, M. A. Moreno, and P. Smolensky (2008). Language universals in human brains. *Proceedings of the National Academy of Sciences, U.S.A. 105*(14), 5321–5325.

Berent, I., T. Lennertz, P. Smolensky, and V. Vaknin-Nusbaum (2009). Listeners' knowledge of phonologial universals: evidence from nasal clusters. *Phonology 26*(1), 75–108.

Berent, I., D. Steriade, T. Lennertz, and V. Vaknin (2007). What we know about what we have never heard: evidence from perceptual illusions. *Cognition 104*, 591–630.

Bruner, J. S., J. J. Goodnow, and G. A. Austin (1956). *A study of thinking.* New York:

John Wiley and Sons.

Bye, P. (2011). Dissimilation. In M. van Oostendorp, C. Ewen, and E. Hume (Eds.), *The Blackwell companion to phonology*, Chapter 63. Wiley-Blackwell.

Campbell, L. (2004). *Historical linguistics: an introduction* (2nd ed.). Cambridge, Massachusetts: MIT Press.

Carpenter, A. C. (2005). Acquisition of a natural vs. an unnatural stress system. In A. Brugos, M. R. Clark-Cotton, and S. Han (Eds.), *Papers from the 29th Boston University Conference on Language Development (BUCLD 29)*, Somerville, pp. 134–143. Cascadilla Press.

Carpenter, A. C. (2006). *Acquisition of a natural versus an unnatural stress system.* Ph. D. thesis, University of Massachusetts, Amherst.

Carpenter, A. C. (2010). A naturalness bias in learning stress. *Phonology 27*(3), 345–392.

Chambers, K. E., K. H. Onishi, and C. Fisher (2010). A vowel is a vowel: generalizing newly learned phonotactic constraints to new contexts. *Journal of Experimental Psychology: Learning, Memory, and Cognition 36*(3), 821–828.

Chomsky, N. and M. A. Halle (1968). *The sound pattern of English.* Cambridge, Massachusetts: MIT Press.

Ciborowski, T. and M. Cole (1972). A cross-cultural study of conjunctive and disjunctive concept learning. *Child Development 43*, 774–789.

Ciborowski, T. and M. Cole (1973). A developmental and cross-cultural study of the influences of rule structure and problem composition on the learning of conceptual classifications. *Journal of Experimental Child Psychology 15*(2), 193–215.

Ciborowski, T. and D. Price-Williams (1974). The influence of rule structure and problem composition on conceptual learning among rural Hawaiian children. Technical Report 75, Kamehameha Early Education Program, Kamehamea Schools and Bernice P. Bishop Estate, 1850 Makuakane Street, Honolulu, Hawaii 96817.

Clements, G. N. (1985). The geometry of phonological features. In J. A. Goldsmith (Ed.), *Phonological theory: the essential readings*, pp. 201–223. Malden: Blackwell.

Clements, G. N. (2003). Feature economy in sound systems. *Phonology 20*(3), 287–333.

Clements, G. N. and E. V. Hume (1995). The internal organization of speech sounds. In J. A. Goldsmith (Ed.), *The handbook of phonological theory*, Chapter 7, pp. 245–306. Boston: Blackwell.

Cole, J. and L. Trigo (1988). Parasitic harmony. In H. van der Hulst and N. Smith (Eds.), *Features, segmental structure, and harmony processes*, pp. 19–38. Foris.

Conant, M. B. and T. Trabasso (1964). Conjunctive and disjunctive concept formation under equal-information conditions. *Journal of Experimental Psychology 67*(3), 250–255.

Creel, S. C., E. L. Newport, and R. N. Aslin (2004). Distant melodies: statistical learning of non-adjacent dependencies in tone sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition 30*, 1119–1130.

Cristiá, A. and A. Seidl (2008). Is infants' learning of sound patterns constrained by phonological features? *Language Learning and Development 4*(3), 203–227.

Cristià, A., A. Seidl, and A. Francis (2011). Phonological features in infancy. In G. N. Clements and R. Ridouane (Eds.), *Where do phonological contrasts come from? Cognitive, physical, and developmental bases of phonological features*, pp. 303–326. Amsterdam and Philadelphia: John Benjamins.

Cristià, A., A. Seidl, and L. Gerken (2011). Learning classes of sounds in infancy. *University of Pennsylvania Working Papers in Linguistics 17*(1).

Cutler, A., J. Mehler, D. Norris, and J. Segui (1989). Limits on bilingualism. *Nature 340*, 229–230.

Daland, R., B. Hayes, J. White, M. Garellek, A. Davis, and I. Norrmann (2011). Explaining sonority projection effects. *Phonology*, 197–234.

Darcy, I. (2006). *Assimilation phonologique et reconnaissance des mots.* Publications Universitaires Européennes, Série XXI: Linguistique. Bern: Peter Lang.

Davidson, L. (2006). Phonology, phonetics, or frequency: influences on the production of non-native sequences. *Journal of Phonetics 34*(1), 104–137.

Davidson, L., P. W. Jusczyk, and P. Smolensky (2004). The initial and final states: theoretical implications and experimental explorations of Richness of the Base. In R. Kager, J. Pater, and W. Zonneveld (Eds.), *Constraints in phonological acquisition*, Chapter 10, pp. 321–368. Cambridge, England: Cambridge University Press.

Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature 407*, 630–633.

Feldman, J. (2006). An algebra of human concept learning. *Journal of mathematical psychology 50*, 339–368.

Finley, S. (2011a). The privileged status of locality in consonant harmony. *Journal of Memory and Language 65*(74–83).

Finley, S. (2011b). Testing the limits of long-distance learning: learning beyond the three-segment window. To appear in *C*ognitive Science.

Finley, S. and W. Badecker (2009). Artificial language learning and feature-based generalization. *Journal of Memory and Language 61*, 423–437.

Finley, S. and W. Badecker (2010). Linguistic and non-linguistic influences on learning biases for vowel harmony. In S. Ohlsson and R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Sciety*, Austin, Texas, pp. 706–711. Cognitive Science Society.

Friederici, A. D. and J. M. I. Wessels (1993). Phonotactic knowledge of word boundaries and its use in infant speech perception. *Perception and Psychophysics 54*, 287–295.

Frisch, S., J. B. Pierrehumbert, and M. B. Broe (2004). Similarity avoidance and the OCP. *Natural Language and Linguistic Theory 22*(1), 179–228.

Gafos, A. (1996). *The articulatory basis of locality in phonology.* Ph. D. thesis, Johns Hopkins University, Baltimore.

Gallistel, C. R., A. L. Brown, S. Carey, R. Gelman, and F. C. Keil (1991). Lessons from animal learning for the study of cognitive development. In S. G. Carey and R. G. Gelman

(Eds.), *The epigenesis of mind*, Chapter 1, pp. 3–36. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Gerken, L. and A. Bollt (2008). Three exemplars allow at least some linguistic generalizations: implications for generalization mechanisms and constraints. *Language Learning and Development 4*(3), 228–248.

Gluck, M. A. and G. H. Bower (1988). Evaluating an adaptive network model of human learning. *Journal of Memory and Language 27*, 166–195.

Goldsmith, J. A. (1976). *Autosegmental phonology.* Ph. D. thesis, Massachusetts Institute of Technology.

Gómez, R. (2002). Variability and detection of invariant structure. *Psychological Science 13*(5), 431–436.

Gómez, R. L., R. R. Bootzin, and L. Nadel (2006). Naps promote abstraction in language-learning infants. *Psychological Science 17*(8), 670–674.

Gordon, M. (2004). Syllable weight. In B. Hayes, R. Kirchner, and D. Steriade (Eds.), *Phonetically-based phonology*, pp. 277–312. Cambridge, England: Cambridge University Press.

Goudbeek, M., A. Cutler, and R. Smits (2008). Supervised and unsupervised learning of multidimensionally varying non-native speech categories. *Speech communication 50*, 109–125.

Greenberg, J. H. (1978). Diachrony, synchrony, and language universals. In J. H. Greenberg, C. A. Ferguson, and E. A. Moravcsik (Eds.), *Universals of human language, volume 1, method and theory*, pp. 61–91. Stanford, California: Stanford University Press.

Hale, M. and C. A. Reiss (2000). 'Substance abuse' and 'dysfunctionalism': current trends in phonology. *Linguistic Inquiry 31*(1), 157–169.

Hansson, G. Ó. (2001a). The phonologization of production constraints: Evidence from consonant harmony. In M. Andronis, C. Ball, H. Elston, and S. Neuvel (Eds.), *Papers from the 37th meeting of the Chicago Linguistics Society (CLS 37)*, Volume 1, Chicago, pp. 187–200. Chicago Linguistic Society.

Hansson, G. Ó. (2001b). *Theoretical and typological issues in consonant harmony.* Ph. D. thesis, University of California, Berkeley.

Hayes, B. (1999). Phonetically driven phonology: the role of optimality in inductive grounding. In M. Darnell, E. Moravcsik, M. Noonan, F. Newmeyer, and K. Wheatly (Eds.), *Functionalism and Formalism in Linguistics*, Volume 1: General Papers, pp. 243–285. Amsterdam: John Benjamins.

Hayes, B., K. Zuraw, P. Siptár, and Z. Londe (2009). Natural and unnatural constraints in Hungarian vowel harmony. *Language 85*(4), 822–863.

Healy, A. F. and A. G. Levitt (1980). Accessibility of the voicing distinction for learning phonological rules. *Memory and Cognition 8*(2), 107–114.

Heinz, J. (2010). Learning long-distance phonotactics. *Linguistic Inquiry 41*(4), 623–661.

Hume, E. and K. Johnson (2001). A model of the interplay of speech perception and phonology. In E. Hume and K. Johnson (Eds.), *The role of speech perception in phonology*, pp. 3–26. San Diego: Academic Press.

Hunt, E. B. and C. L. Hovland (1960). Order of consideration of different types of concepts. *Journal of Experimental Psychology 59*(4), 220–225.

Hyman, L. M. (1976). Phonologization. In A. Juilland (Ed.), *Linguistic studies offered to Joseph Greenberg: second volume: phonology*, pp. 407–418. Saratoga, California: Anma Libri.

Jensen, J. T. (1974). A constraint on variables in phonology. *Language 50*(4, Part 1), 675–686.

Jusczyk, P. W., A. D. Friederici, J. M. I. Wessels, V. Y. Svenkerud, and A. M. Jusczyk (1993). Infants' sensitivity to the sound pattern of native language words. *Journal of Memory and Language 32*, 402–420.

Jusczyk, P. W., P. A. Luce, and J. Charles-Luce (1994). Infants sensitivity to phonotactic patterns in the native language. *Journal of Memory and Language 33*, 630–645.

Kager, R., N. Boll-Avetisyan, and C. Ao (2008, November). Gradient phonotactic constraints for speech segmentation in a second language. Poster presented at the Boston University Conference on Language Development.

Kaun, A. (2004). The typology of rounding harmony. In B. Hayes, R. Kirchner, and D. Steriade (Eds.), *Phonetically based phonology*, Chapter 4, pp. 87–116. Cambridge, England: Cambridge University Press.

Kawahara, S. (2006). Mimetic gemination in Japanese: a challenge for Evolutionary Phonology. *Theoretical Linguistics 32*(3), 411–424.

Kim, H. (2001). A phonetically based account of phonological stop assibilation. *Phonology 18*(1), 81–108.

Kiparsky, P. (1971). Historical linguistics. In W. Dingwall (Ed.), *A survey of linguistic science*, pp. 576–653. College Park: University of Maryland.

Kiparsky, P. (2006). The amphichronic program vs. Evolutionary Phonology. *Theoretical Linguistics 32*(2), 217–236.

Kiparsky, P. (2008). Universals constrain change, change results in typological generalizations. In J. Good (Ed.), *Linguistic universals and language change*, Chapter 2, pp. 23–53. Oxford, England: Oxford University Press.

Kitto, C. and P. de Lacy (1999). A Correspondence theory of epenthetic quality. In C. Kitto and C. Smallwood (Eds.), *Proceedings of AFLA (Austronesian Formal Linguistics Association VI*, Toronto Working Papers in Linguistics, pp. 181–200. Toronto: Unspecified publisher.

Köhler, W. (1941). On the nature of associations. *Proceedings of the American Philosophical Society 84*, 489–502.

Koo, H. (2007). *Change in the adult phonological processing system by learning non-adjacent phonotactic constraints from brief experience: an experimental and computational study.* Ph. D. thesis, University of Illinois, Urbana-Champaign.

Koo, H. and L. Callahan (2011). Tier adjacency is not a necessary condition for learning non-adjacent dependencies. *Language and Cognitive Processes*.

Koo, H. and J. S. Cole (2006). On learnability and naturalness as constraints on phonological grammar. In *Proceedings of ISCA Tutorial and Research Workshop on Experimental Linguistics, Athens, Greece, August 28-30*.

Kosa, L. A. (2010). Sibilant harmony: investigating the connection between typology and learnability. In *Proceedings of the 2010 annual conference of the Canadian Linguistic Association*.

Kruschke, J. K. (1992). ALCOVE: an exemplar-based connectionist model of category learning. *Psychological Review 99*, 22–44.

Kuo, L. (2009). The role of natural class features in the acquisition of phonotactic regularities. *Journal of psycholinguistic research 38*(2), 129–150.

Lafond, D., Y. Lacouture, and G. Mineau (2007). Complexity minimization in rule-based category learning: revising the catalog of Boolean concepts and evidence for non-minimal rules. *Journal of Mathematical Psychology 51*, 57–75.

LaRiviere, C., H. Winitz, J. Reeds, and E. Herriman (1974). The conceptual reality of selected distinctive features. *Journal of Speech and Hearing Research 17*(1), 122–133.

LaRiviere, C., H. Winitz, J. Reeds, and E. Herriman (1977). Erratum: The conceptual reality of selected distinctive features. *Journal of Speech and Hearing Research 20*(4), 817.

Lin, Y. (2009). Tests of analytic bias in native Mandarin speakers and native Southern Min speakers. In Y. Xiao (Ed.), *21st North American Conference on Chinese Linguistics*, Smithfield, Rhode Island, pp. 81–92. Bryant University.

Lindsay, S. and M. G. Gaskell (2010). A complementary systems account of word learning in L1 and L2. *Language Learning 60*(Suppl. 2), 45–63.

Love, B. C. (2002). Comparing supervised and unsupervised category learning. *Psychonomic Bulletin and Review 9*(4), 829–835.

Love, B. C., D. L. Medin, and T. M. Gureckis (2004). SUSTAIN: a network model of category learning. *Psychological Review 111*(2), 309–332.

Mackie, S. and J. Mielke (2011). Feature economy in natural, random, and synthetic inventories. In G. N. Clements and R. Ridouane (Eds.), *Where do phonological features come from? Cognitive, physical, and developmental bases of distinctive speech categories*, pp. 43–63. John Benjamins.

Majerus, S., L. Mulder, T. Meulmans, and F. Peters (2004). Verbal short-term memory reflects the sublexical organization of the phonological language network: evidence from an incidental phonotactic learning paradigm. *Journal of Memory and Language 51*, 297–306.

Mattys, S. L. and P. W. Jusczyk (2001). Phonotactic cues for segmentation of fleunt speech by infants. *Cognition 78*(2), 91–121.

Maye, J., J. F. Werker, and L. Gerken (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition 82*, B101–B111.

McCarthy, J. J. (1981). A prosodic theory of nonconcatenative morphology. *Linguistic Inquiry 12*, 373–418.

McCarthy, J. J. (1988). Feature Geometry and dependency: a review. *Phonetica 45*, 84–108.

Mielke, J. (2004). *The emergence of distinctive features.* Ph. D. thesis, Ohio State University.

Mielke, J. (2008). *The emergence of distinctive features.* Oxford, England: Oxford University Press.

Mitchell, T. M. (1980 [1990]). The need for biases in learning generalizations. In J. W. Shavlik and T. G. Detterich (Eds.), *Readings in machine learning*, Chapter 2.4.1, pp. 184–191. San Mateo, California: Morgan Kaufman.

Moreton, E. (2008). Analytic bias and phonological typology. *Phonology 25*(1), 83–127.

Moreton, E. (2010). Underphonologization and modularity bias. In S. Parker (Ed.), *Phonological argumentation: essays on evidence and motivation*, Chapter 3, pp. 79–101. London: Equinox.

Moreton, E. (2012). Inter- and intra-dimensional dependencies in implicit phonotactic learning. *Journal of Memory and Language*.

Moreton, E., G. Feng, and J. L. Smith (2008). Syllabification, sonority, and perception: new evidence from a language game. In R. L. Edwards, P. J. Midtlyng, C. L. Sprague, and K. G. Stensrud (Eds.), *Papers from the 41st Regional Meeting of the Chicago Linguistic Society*, pp. 341–355.

Neisser, U. and P. Weene (1962). Hierarchies in concept attainment. *Journal of Experimental Psychology 64*(6), 640–645.

Nevins, A. (2010). Two case studies in phonological universals: a view from artificial grammars. *Biolinguistics 4*(2), 218–233.

Newport, E. and R. N. Aslin (2004). Learning at a distance i: statistical learning of non-adjacent dependencies. *Cognitive Psychology 48*, 127–162.

Nosofsky, R. M., M. A. Gluck, T. J. Palmeri, S. C. McKinley, and P. Gauthier (1994). Comparing models of rule-based classification learning: a replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory and Cognition 22*(3), 352–369.

Nosofsky, R. M., T. J. Palmeri, and S. C. McKinley (1994). Rule-plus-exception model of classification learning. *Psychological Review 101*(1), 53–79.

Odden, D. (1991). Vowel geometry. *Phonology 8*, 261–289.

Odden, D. (1995). Adjacency parameters in phonology. *Language 70*(2), 289–330.

Ohala, J. J. (1992). What's cognitive, what's not, in sound change. In G. Kellerman and M. D. Morissey (Eds.), *Diachrony within synchrony: language history and cognition*, pp. 309–355. Frankfurt am Main: Peter Lang.

Ohala, J. J. (1993). The phonetics of sound change. In C. Jones (Ed.), *Historical linguistics: problems and perspectives*, pp. 237–278. Harlow: Longman.

Ohala, J. J. (1997). Phonetics in phonology. In *Proceedings of the 4th Seoul International Conference on Linguistics (SICOL), 11–15 August*, pp. 45–50.

Pacton, S. and P. Perruchet (2008). An attention-based associative account of adjacent and nonadjacent dependency learning. *Journal of Experimental Psychology: Learning,*

Memory, and Cognition 34(1), 80–96.
Padgett, J. (2002a). Feature classes in phonology. *Language 78*(1), 81–110.
Padgett, J. (2002b). The unabridged "feature classes in phonology". MS, Department of Linguistics, University of California, Santa Cruz.
Pater, J. and A.-M. Tessier (2006). L1 phonotactic knowledge and the L2 acquisition of alternations. In R. Slabakova, S. Motrul, P. Prévost, and L. White (Eds.), *Inquiries in linguistic development: in honor of Lydia White*, pp. 115–131. Benjamins.
Peña, M., L. L. Bonatti, M. Nespor, and J. Mehler (2002). Signal-driven computations in speech processing. *Science 298*, 604–607.
Peperkamp, S., K. Skoruppa, and E. Dupoux (2006). The role of phonetic naturalness in phonological rule acquisition. In D. Bamman, T. Magnitskaia, and C. Zoller (Eds.), *Papers from the 30th Boston University Conference on Language Development (BUCLD 30)*, pp. 464–475.
Perruchet, P., M. D. Tyler, N. Galland, and R. Peereman (2004). Learning nonadjacent dependencies: no need for algebraic-like computations. *Journal of Experimental Psychology: General 133*(4), 573–583.
Pertz, D. L. and T. G. Bever (1975). Sensitivity to phonological universals in children and adults. *Language 51*, 149–162.
Picard, M. (2003). On the emergence and resolution of hiatus. *Folia Linguistica Historica 24*(1–2), 47–57.
Pierrehumbert, J. B. (1993). Dissimilarity in the Arabic verbal roots. In A. Schafer (Ed.), *Proceedings of the Northeast Linguistics Society*, Volume 23, Amherst, Mass., pp. 367–381. Graduate Linguistics Students Association.
Pierrehumbert, J. B. (2001). Why phonological constraints are so coarse-grained. *Language and Cognitive Processes 16*(5–6), 691–698.
Pierrehumbert, J. B. (2003). Probabilistic phonology: discrimination and robustness. In R. Bod, J. Hay, and S. Jannedy (Eds.), *Probabilistic phonology*, Chapter 6, pp. 177–228. Cambridge, Massachusetts: MIT Press.
Pinker, S. (1979). Formal models of language learning. *Cognition 7*, 217–283.
Pons, F. and J. M. Toro (2010). Structural generalizations over consonants and vowels in 11-month-old infants. *Cognition 116*, 361–367.
Prentice, W. C. H. and S. E. Asch (1958). Paired associations with related and unrelated pairs of nonsense-figures. *American Journal of Psychology 71*, 247–254.
Pulleyblank, D. (1988). Underspecification, the feature hierarchy, and Tiv vowels. *Phonology 5*(2), 299–326.
Pycha, A., P. Nowak, E. Shin, and R. Shosted (2003). Phonological rule-learing and its implications for a theory of vowel harmony. In M. Tsujimura and G. Garding (Eds.), *Proceedings of the 22nd West Coast Conference on Formal Linguistics (WCCFL 22)*, pp. 101–114.
Rescorla, R. A. (1980). Simultaneous and successive associations in sensory preconditioning. *Journal of Experimental Psychology: Animal Behavior Processes 6*(3), 207–216.

Rescorla, R. A. (2008). Evaluating conditioning of related and unrelated stimuli using a compound test. *Learning and Behavior 36*(2), 67–74.

Rescorla, R. A. and D. J. Gillan (1980). An analysis of the facilitative effect of similarity on second-order conditioning. *Journal of Experimental Psychology: Animal Behavior Processes 6*(4), 339–351.

Rogers, C. J. and P. J. Johnson (1973). Attribute identification in children as a function of stimulus dimensionality. *Journal of Experimental Child Psychology 15*(2), 216–221.

Rose, S. and R. Walker (2004). A typology of consonant agreement as correspondence. *Language 80*(3), 475–531.

Rose, S. and R. Walker (2011). Harmony systems. In J. Goldsmith, J. Riggle, and A. C. L. Yu (Eds.), *The handbook of phonological theory* (2nd ed.)., Chapter 8, pp. 240–290. Malden, Massachusetts: Blackwell.

Saffran, J. R. and E. D. Thiessen (2003). Pattern induction by infant language learners. *Developmental Psychology 39*(3), 484–494.

Sagey, E. (1990). *The representation of features in non-linear phonology: the Articulator Node Hierarchy*. New York: Garland.

Schane, S. A., B. Tranel, and H. Lane (1974). On the psychological reality of a natural rule of syllable structure. *Cognition 3*(4), 351–358.

Seidl, A. and E. Buckley (2005). On the learning of arbitrary phonological rules. *Language Learning and Development 1*(3 & 4), 289–316.

Shepard, R. N., C. L. Hovland, and H. M. Jenkins (1961). Learning and memorization of classifications. *Psychological Monographs 75*(13, Whole No. 517).

Simons-Wiedenmaier, N. S., M. Weber, C. F. Plappert, P. K. D. Pilz, and S. Schmid (2006). Synaptic depression and short-term habituation are located in the sensory part of the mammalian startle pathway. *BMC Neuroscience 7*(38).

Skoruppa, K., A. Lambrechts, and S. Peperkamp (2009). The role of phonetic distance in the acquisition of phonological alternations. In S. Lima, K. Mullin, and B. Smith (Eds.), *Proceedings of the 39th Meeting of the North-East Linguistics Society (NELS)*, Amherst, Massachusetts. Graduate Linguistics Students Association.

Skoruppa, K. and S. Peperkamp (2011). Adaptation to novel accents: feature-based learning of context-sensitive phonological regularities. *Cognitive Science 35*, 348–366.

Smith, J. D., J. P. Minda, and D. A. Washburn (2004). Category learning in rhesus monkeys: a study of the Shepard, Hovland, and Jenkins (1961) tasks. *Journal of Experimental Psychology: General 133*(3), 398–404.

Sproat, R. and O. Fujimura (1993). Allophonic variation in English /l/ and its implications for phonetic implementation. *Journal of Phonetics 21*, 291–311.

St. Clair, M. C. and P. Monaghan (2008). Language abstraction: consolidation of language structure during sleep. In B. C. Love, K. McRae, and V. M. Sloutsky (Eds.), *Proceedings of the 30th annual conference of the Cognitive Science Society*, Austin, Texas, pp. 727–732. Cognitive Science Society.

Stampe, D. (1973). *A dissertation on Natural Phonology*. Ph. D. thesis, University of

Chicago.
Steriade, D. (1997). Phonetics in phonology: the case of laryngeal neutralization. MS, Department of Linguistics, University of California, Los Angeles.
Steriade, D. (2001). The phonology of perceptibility effects: the P-map and its consequences for constraint organization. MS, Department of Linguistics, University of California, Los Angeles.
Taylor, C. F. and G. Houghton (2005). Learning artificial phonotactic constraints: time course, durability, and relationship to natural constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition 31*, 1398–1416.
Toro, J. M., M. Nespor, J. Mehler, and L. L. Bonatti (2008). Finding words and rules in a speech stream. *Psychological Science 19*(2), 137–144.
Toro, J. M., M. Shukla, M. Nespor, and A. D. Endress (2008). The quest for generalizations over consonants: asymmetries between consonants and vowels are not the by-product of acoustic differences. *Perception and Psychophysics 70*(8), 1515–1525.
van Kampen, A., G. Parmaksiz, R. van de Vijver, and B. Höhle (2008). Metrical and statistical cues for word segmentation: vowel harmony and word stress as cues to word boundaries by 6- and 9-month-old Turkish learners. In A. Gavarroó and M. J. Freitas (Eds.), *Language Acquisition and Development*, pp. 313–324. Cambridge Scholars Publishing.
Wagner, U., S. Gais, H. Haider, R. Verleger, and J. Born (2004). Sleep inspires insight. *Nature 427*, 352–355.
Walsh Dickey, L. (1997). *The phonology of liquids*. Ph. D. thesis, University of Massachusetts, Amherst, Massachusetts.
Warker, J. A. and G. S. Dell (2006). Speech errors reflect newly learned phonotactic constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition 32*, 387–398.
Warker, J. A., G. S. Dell, C. A. Whalen, and S. Gereg (2008). Limits on learning phonotactic constraints from recent production experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition 34*(5), 1209–1295.
Weinreich, U., W. Labov, and M. Herzog (1968). Empirical foundations for a theory of language change. In W. P. Lehmann and Y. Malkiel (Eds.), *Directions for historical linguistics: a symposium*, pp. 97–195. Austin: University of Texas Press.
Whalen, D. H. and A. G. Levitt (1995). The universality of intrinsic F0 of vowels. *Journal of Phonetics 23*, 349–366.
Wilson, C. (2001). Consonant cluster neutralisation and targeted constraints. *Phonology 18*, 147–197.
Wilson, C. (2003). Experimental investigation of phonological naturalness. In G. Garding and M. Tsujimura (Eds.), *Proceedings of the 22nd West Coast Conference on Formal Linguistics (WCCFL 22)*, Somerville, pp. 533–546. Cascadilla Press.
Wilson, C. (2006). Learning phonology with substantive bias: an experimental and computational study of velar palatalization. *Cognitive Science 30*(5), 945–982.

Yu, A. C. L. (2010). Tonal effects on perceived vowel duration. In C. Fougeron, B. Kühnert, M. D'Imperio, and N. Vallée (Eds.), *Papers in Laboratory Phonology 10: phonology and phonetics*, pp. 151–168. Mouton De Gruyter.

Yu, A. C. L. (2012). On measuring phonetic precursor robustness: a response to Moreton. *Phonology 28*(3), 491–518.

Zaba, A. M. (2008). *Relative frequency of patterns and learnability: the case of phonological harmony*. Ph. D. thesis, University of Utah.

Zhang, J. and Y. Lai (2010). Testing the role of phonetic knowledge in Mandarin tone sandhi. *Phonology 27*(1), 153–201.

Zhang, J., Y. Lai, and C. Sailor (2011). Modeling Taiwanese speakers' knowledge of tone sandhi in reduplication. *Lingua 121*(2), 181–206.

Zhang, J., Y. Lai, and C. Turnbull-Sailor (2006). Wug-testing the "tone circle" in Taiwanese. In D. Baumer, D. Montero, and M. Scanlon (Eds.), *Proceedings of the 25th meeting of the West Coast Conference on Formal Linguistics (WCCFL)*, Somerville, Mass., pp. 453–461. Cascadilla Proceedings Project.

Zuraw, K. (2007). The role of phonetic knowledge in phonological patterning: corpus and survey evidence from Tagalog infixation. *Language 83*(2), 277–316.