Andries W. Coetzee
November 2008

# LEARNING LEXICAL INDEXATION*

*Abstract*. Morphological concatenation often triggers phonological processes. For instance, addition of the plural suffix /-ən/ to Dutch nouns causes vowel lengthening in some nouns due to the weight-to-stress principle ([xát] vs. [xáː.tən] 'hole'). These kinds of processes often apply only to a subset of words – not all Dutch nouns undergo this process ([kát] vs. [ká.tən] 'cat'). Nouns need to be lexically indexed as either undergoing this process or not. I investigate how phonological grammar and lexical indexation are learned when learners are confronted with data like these. Based on learnability considerations, I hypothesize that learners acquire a grammar with default non-alternation, so that novel items are treated as non-alternating. I report the results of artificial language learning experiments compatible with this hypothesis, and model these results in a version of the Biased Constraint Demotion Algorithm (Prince and Tesar 2004).

*This research has been presented at the Phonology Discussion group at the University of Michigan, and at the ExpOT conference held at the University of Michigan in 2007. My appreciation to these audiences for supporting, even if challenging, discussion. Particularly helpful have been discussions with Pam Beddor, Iris Berent, and Joe Pater. Also a word of appreciation to my fellow guest-editors for this volume (René Kager and Joe Pater), and to two anonymous reviewers for their thorough review of this research.

# 1. INTRODUCTION

## 1.1 Generalizing beyond learning data

The language learner faces a twofold task. He/she must acquire a grammar that would deem grammatical the data to which he/she is exposed. But this grammar must also be able to generalize beyond the learning data. This task is complicated by the fact that the learning data are sometimes compatible with multiple possible grammars. These different grammars, though they agree on their evaluation of the learning data, often differ in how they generalize beyond the data.

Language learners are often exposed to data in which a process applies only to a subset of the data. This is not an uncommon phenomenon – see, for instance, the differences between words of Latin and Greek origin in English, the Yamato, Sino-Japanese and borrowed lexicon of Japanese (Itô & Mester 1999), native and borrowed words in German (Féry 2003), etc. When confronted with data like these, the learner must memorize for each word to which lexical class it belongs. But this achieves only one part of the learning task. The grammar must also generalize beyond the learning data. The learner does not have access to information on the lexical class affiliation of words that were not in the learning data. Even so, he/she still needs to decide about the affiliation of such words.

In this paper, I discuss one example of this scenario. Imagine learning data that contain roots that are realized uniformly throughout their paradigms, and also roots that, though identical in all relevant aspects to the non-alternating roots, alternate between different parts of their paradigms. The learner must memorize for each root whether it belongs to the alternating or the non-alternating class. And when faced with a novel root, the learner must decide to which class to assign the root. Based on learnability considerations, I argue that the default choice is to treat novel roots as non-alternating. Only when evidence to the contrary exists, will novel roots be treated as alternating.

Section §1.2 discusses the realization of vowel length in Dutch nominal paradigms as an example of this scenario. Section §1.3 presents the learnability argument for treating non-alternation as the default. Results of artificial learning experiments consistent with default non-alternation are presented in §2. In §3, I show how this default can be accounted for in the Biased Constraint Demotion Algorithm of Prince and Tesar (2004).

## 1.2 Vowel length in Dutch nouns

Because of morphological concatenation, root morphemes can appear in different phonological environments, so that a phonological process can be conditioned to apply to a root in only some contexts in which it appears. For example, in Dutch, short vowels sometimes lengthen in open, stressed syllables, causing length alternations between noun singulars and plurals. In the plural of *gat* [xɑ́t] 'hole', addition of the suffix /-ən/ places the short /ɑ/ in an open, stressed, syllable and vowel lengthening results in [xáː.tən]. Some Dutch nouns undergo this process, while others do not. Examples are given in (1).

(1)

| | | sg. | | pl. | | |
|---|---|---|---|---|---|---|
| Alternating | *gat* | [xɑ́t] | *gaten* | [xáː.tən] | 'hole' |
| | *spel* | [spɛ́l] | *spelen* | [spéː.lən] | 'game' |
| | *schip* | [sxíp] | *schepen* | [sxéː.pən] | 'boat' |
| | *god* | [xɔ́t] | *goden* | [xóː.dən] | 'god' |
| Non-alternating | *kat* | [kɑ́t] | *katten* | [kɑ́.tən] | 'cat' |
| | *stel* | [stɛ́l] | *stellen* | [stɛ́.lən] | 'set' |
| | *wip* | [ʋíp] | *wippen* | [ʋí.pən] | 'trap' |
| | *vod* | [fɔ́t] | *vodden* | [fɔ́.dən] | 'rag' |

Vowel lengthening in alternating roots is due to the stress-to-weight principle (stressed syllables are heavy) (Prince 1990). However, in non-alternating roots, a short, stressed vowel is tolerated in an open syllable. The lexical class affiliation of nouns is not phonologically predictable. This illustrates the scenario discussed in §1.1. A learner exposed to these data needs to memorize for each root to which lexical class it belongs. But the learner must also be able to generalize beyond these data. Suppose that [nɔ́l] is a singular to which the learner was not exposed during learning. If the learner had to form the plural of [nɔ́l], he/she can choose the non-alternating [nɔ́.lən] or the alternating [nóː.lən] plural. There are at least three possible outcomes, stated in (2).[1]

(2) Possible lexical class assignment for words absent from learning data

a. *Random*. The learner randomly assigns novel items to a lexical class. Novel items are equally likely to be alternating or non-alternating.

b. *Default alternation*. Lexical items are assigned to the alternating class in the absence of evidence.

c. *Default non-alternation*. Lexical items are assigned to the non-alternating class in the absence of evidence.

Although Dutch has both alternating and non-alternating roots, the non-alternating pattern is more frequent.[2] A learner will therefore get more evidence for non-alternation, and this will likely influence how learners will treat novel forms. Similarly, learners could rely on analogy so that a novel item that accidentally shares more phonological properties with (non-)alternating words may be treated as (non-)alternating (Prasada & Pinker 1993). I return to these issues later (§2.1.1 and §2.3.1), but for now, I assume a learner without access to this kind of information. Imagine a learner that is exposed to exactly the same number of alternating and non-alternating forms, and that this learner is confronted with a novel item equally (dis)similar to the alternating and non-alternating roots in the learning data. The question asked here is what this learner will do. In §1.3, I use learnability considerations to argue that the option in (2c) is to be expected. In

[1] This is a simplification of the actual situation in Dutch. In addition to /-ən/, Dutch also has an /-s/ plural suffix. The distribution of these suffixes is largely overlapping, and a learner could choose to pluralize /nɔl/ as [nɔ́ls].

[2] I compiled a list of all monosyllabic Dutch singular nouns in CELEX (Baayen et al. 1995) to which vowel lengthening could apply in the plural. Only 5.6% of these nouns undergo vowel lengthening.

the absence of evidence about the lexical class affiliation of some word, it should be treated as non-alternating.

### 1.3 The subset learning problem

A classic problem in learnability research is the so-called "subset problem". When a learner is confronted with learning data, he/she can learn a restrictive grammar, allowing only patterns included in the data, or a permissive grammar, also allowing patterns not in the data. The restrictive grammar will generate a subset of the languages generated by the permissive grammar. If the learner opts for the subset grammar while the superset grammar is actually correct, exposure to further learning data will present him/her with evidence of this error so that the error can be corrected. However, if the learner incorrectly assumed the superset grammar, no additional data could give evidence of this error. Berwick (1985) proposed a general solution to this problem: learn the most restrictive grammar consistent with the learning data.

The subset problem is usually discussed in the input-output domain. However, it holds equally of surface relations in paradigms (McCarthy 1998; Hayes 2004; Tessier 2007). I first illustrate how this problem applies to the input-output relation, and then, following McCarthy (1998), do the same for intra-paradigmatic relations.

Imagine a learner exposed to data like those in (3). These data contain no tauto-syllabic consonant clusters. If the learner had to learn a syllable structure grammar based on these data, he/she has two options: (i) The *restrictive* grammar assumes that structures absent from the data are ungrammatical. The absence of consonant clusters from the learning data is thus taken as evidence of their ungrammaticality. (ii) The *permissive* grammar assumes that everything is grammatical unless the data contain evidence to the contrary. The absence of clusters is hence interpreted as accidental. The restrictive grammar allows only a subset of the forms allowed by the permissive grammar – shown graphically in *Figure 1*.

(3) The subset problem in input/output relations

| | | | |
|---|---|---|---|
| [bit] | [lap] | [guk] | [pef] |
| [sop] | [tip] | [sam] | [kus] |

Both grammars are consistent with the learning data – since both allow the structures in the learning data. However, they differ on whether additional learning is possible. Suppose that the language actually allowed consonant clusters, and that forms with clusters were accidentally absent from the learning data. A learner that has erroneously settled on the restrictive grammar will eventually encounter examples with clusters. Once this happens, this learner can correct his/her error. However, suppose that the language really does not allow clusters and that clusters were absent for this reason. A learner that acquired the permissive superset grammar will never encounter the information necessary to alert him/her to his/her error. This learner will only ever encounter data like those in (3), and since these are consistent with his/her grammar, the grammar will not be changed.

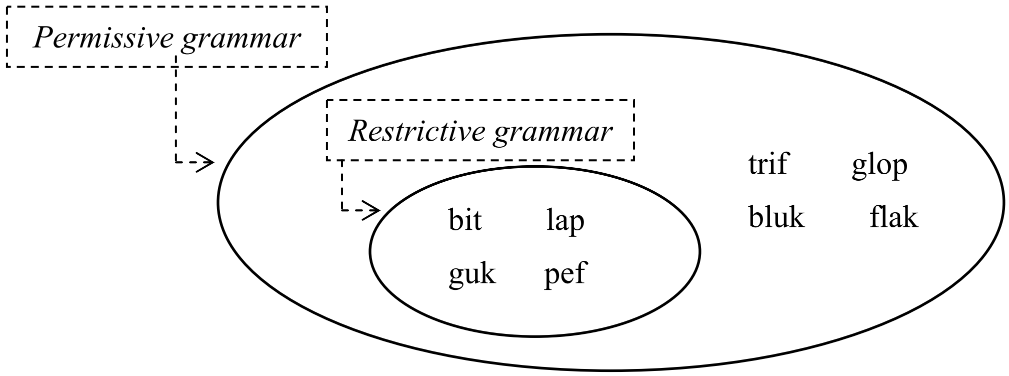


*Figure 1*: Restrictive and permissive grammars for (3).

From a learnability standpoint, it is better to learn the restrictive grammar in a situation like this. It is possible to recover from an error if the restrictive grammar is not actually correct. But recovery is impossible if the permissive grammar has incorrectly been learned. Based on this, Berwick (1985) proposes that learning algorithms should be conservative – they should learn the most restrictive grammar consistent with the data. This restrictive bias is also supported by first language acquisition. Children generally start with restrictive grammars, and gradually approach more permissive grammars based on evidence (Gnanadesikan 2004; Jakobson 1968; Smolensky 1996; etc.).

Although the subset problem is mostly discussed in the context of input/output relations, it applies equally to the relation between surface forms in a paradigm (Hayes 2004; McCarthy 1998; Tessier 2007). I illustrate this here with an example from McCarthy (1998). Imagine a learner acquiring a language with final obstruent devoicing. Suppose that this learner has already learned the final devoicing part of the grammar. At this point, he/she encounters data like those in (4a), where the plural is formed by /-i/ suffixation, and the root is realized identically in the singular and plural.

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| (4) | a. | In the learning data | | b. | Not in the learning data | |
| | | *Singular* | *Plural* | | *Singular* | *Plural* |
| | | gat | ga.ti | | bat | ba.di |
| | | pok | po.ki | | dok | do.gi |
| | | zep | ze.pi | | vep | ve.bi |

These data are consistent with more than one grammar. It might be that these nouns are derived from roots than end on voiced obstruents (/gad, pog, zeb/). In the singular, devoicing is motivated by ordinary phonological considerations, and the final obstruents devoice accordingly. In the plural, however, phonological motivation for devoicing is lacking, and it applies due to paradigm uniformity. The paradigms in (4b), where the noun roots alternate, are then absent from the learning data because such paradigms are ungrammatical.

Under an alternative interpretation of the data, there are no paradigm uniformity requirements. The nouns in (4a) are then derived from roots that end on voiceless obstruents (/gat, pok, zep/) – since intra-paradigmatic identity requirements cannot explain the voiceless realization of these obstruents in the plural. With this interpretation, the absence of the (4b) paradigms from the learning data becomes accidental. A root such as /bad/ will be realized as [bat] in the singular due to final devoicing. In the plural, the /d/ is no longer word-final, and hence not subject to devoicing. Intra-paradigmatic identity requirements can also not force the root to be identical in the singular and plural, and the plural is realized as [ba.di].

We again have a superset/subset relation, as shown in *Figure 2*. The permissive grammar allows non-alternating and alternating paradigms. The restrictive grammar allows only non-alternating paradigms. Since the learning data contain only non-alternating examples, both grammars are consistent with the learning data. However, the grammars again have different learnability implications.

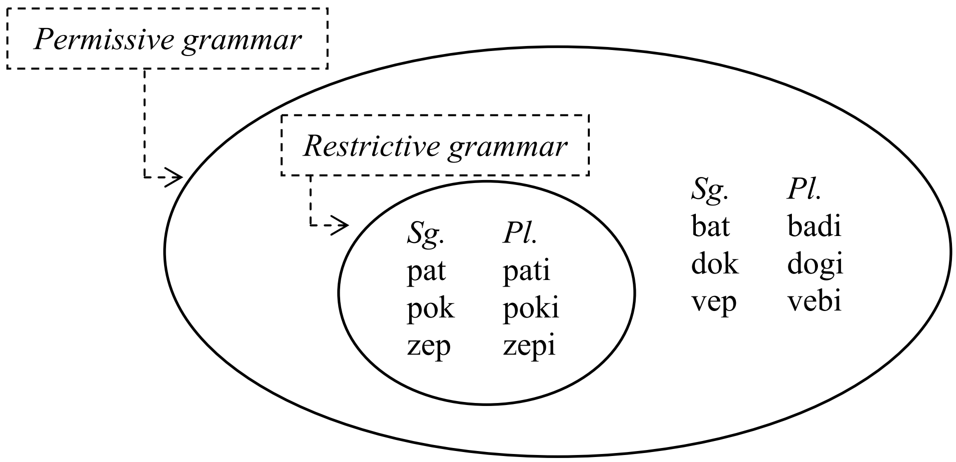


*Figure 2*: Restrictive and permissive grammars for (4).

Suppose that the language actually allowed alternating paradigms, and that these were accidentally absent from the initial data. A learner that has erroneously learned the restrictive grammar will eventually encounter examples of alternating paradigms, giving him/her evidence of his/her error. However, suppose that alternating paradigms are really ungrammatical, and the examples in (4b) were absent for this reason. A learner that acquired the permissive grammar will never encounter information to alert him/her to his/her error. He/she will only encounter paradigms like those in (4a) which are consistent with his/her grammar.

To prevent a learner from becoming stuck in an incorrect superset grammar, the learning algorithm used by the learner must be conservative in terms of both input/output relations and of relationships between morphologically related outputs.

(5) Restrictiveness in learning

a. If the learning data contain no *X*, assume that *X* is ungrammatical.

b. If the learning data contain no alternating paradigms, assume that alternating paradigms are ungrammatical.

The Dutch data in (1) are more complex. These data contain both alternating and non-alternating paradigms, and it cannot be predicted for any given word whether it will alternate or not. Learners have to acquire a hybrid between a restrictive and a permissive grammar, and memorize for each word whether the restrictive or permissive part of the grammar applies to it. However, the learner cannot rely on the learning data to determine how to treat a novel token. I propose that the learner will extend the restrictiveness principle in (5) to situations like these. The learner will assume that any new word is subject to the restrictive part of the grammar (no-alternation), unless if there is evidence to the contrary.

It is hard to test these predictions with real language data. There are many factors other than learnability that influence acquisition. In the Dutch lexicon, for instance, non-alternating plurals far outnumber alternating plurals (see footnote 2). Such frequency information could impact the acquisition. Similarly, learners might rely on analogy between novel forms and forms included in the learning data (Prasada & Pinker 1993). Due to considerations such as these, I opted to test the predictions relying on artificial grammar learning experiments, where it is easier to control for variables that are not of interest.

## 2. ARTIFICIAL LANGUAGE LEARNING EXPERIMENTS

Although the artificial language learning paradigm does not exactly mirror natural language learning exactly, there is a growing body of research that utilizes this paradigm to investigate problems in theoretical phonology (Wilson 2006; Pater & Tessier 2006; Carpenter 2006; etc.). While acknowledging that results from such experiments should be interpreted with caution, I follow this recent trend. I conducted three learning experiments in which participants were first exposed to data modeled after the Dutch paradigms from (1), and then tested on novel items to determine which of the patterns in the learning data they generalize to novel items.

### 2.1. Experiment 1

In Experiment 1, participants are exposed to learning data with an equal number of alternating and non-alternating forms. When presented with novel items, participants cannot rely on the learning data for evidence on which of patterns to extend to the novel items. Based on §1.3, I hypothesize that these learners will preferentially treat novel items as non-alternating.

#### *2.1.1. Methods*

*Participants*. Participants for all three experiments were native speakers of American English, recruited from the undergraduate population at the University of Michigan. All participants had normal hearing, and normal or corrected-to-normal eyesight. Participants were tested in groups of up to three. There was no overlap in participants between experiments. Thirteen participants participated in Experiment 1.

*Tokens*. The tokens were modeled after the Dutch paradigms in (1). I selected 32 English non-words ("singulars"), with the form [CəC.Cv́C] where [v́] is a lax vowel of English, [ɛ́, ɪ́, ɔ́, ʊ́]. Each vowel was used in eight tokens. For each singular, two "plurals" were created by addition of the suffix [-ən]. In one plural form the singular vowel did not change, while in the other it was replaced by its tense counterpart, [é, í, ó, ú]. In (6), I list a few examples – the full list is in Appendix A.

(6) Example of tokens for Experiment 1

| | *Plural* | |
|---|---|---|
| *Singular* | *Non-Alternating* | *Alternating* |
| [təl.sɛ́p] | [təl.sɛ́.pən] | [təl.sé.pən] |
| [kən.dɪ́p] | [kən.dɪ́.pən] | [kən.dí.pən] |
| [dəf.sɔ́k] | [dəf.sɔ́.kən] | [dəf.só.kən] |
| [pəm.bʊ́k] | [pəm.bʊ́.kən] | [pəm.bú.kən] |

Five repetitions of each token were recorded by a male native speaker of American English in a sound attenuated room. One example of each was selected. For the singular, I selected the repetition that I judged to be closest to the intended pronunciation. For the non-alternating plural, I selected a repetition that I judged to contain a vowel that is auditorily close to the stressed vowel in the singular. The alternating plural was selected so that its stressed vowel was a clear instantiation of the relevant tense vowel. All tokens were scaled to have a mean intensity of 70 dB, using the *Scale intensity…* command in *Praat* (Boersma & Weenink 2007).

*Learning*. Sixteen singulars were used in learning – four for each lax vowel. Two tokens for each vowel in the singular were assigned to the "non-alternating", and two to the "alternating" group. Non-alternating singulars were paired with plurals containing the same vowel as the singular, and alternating singulars with plurals with the corresponding tense vowel. The learning data consisted of 16 singular/plural pairs, 8 each that alternate or not.

Since there are only 8 examples of each kind, it is unlikely that participants could use analogy during learning (see §1.2). There just are not enough pairs in each group for reliable patterns to be identified. To minimize this possibility even more, singulars that were assigned to each group were selected so that the groups were phonologically very similar. To simulate what a learner will have to do if he/she was looking for patterns in the data, I ran a discriminant analysis (using *SPSS 16.0*) on the learning tokens. Each token was classified by the voicing, place, continuancy, and sonorancy of the consonants flanking the stressed vowel. The analysis finds the function that best classifies the tokens into two groups based on the features of these consonants. I then classified the novel tokens using the discriminant function. The function classified 10 of the novel forms as alternating and 6 as non-alternating. If the participants were to mimic the discriminant analysis, they should show a preference for alternation, the opposite of what is expected based on learnability considerations.

Each singular/plural pair was paired with a picture of a countable noun (baseball, flamingo, etc.). Table 1 shows the basic structure of the learning data. Participants were seated in a sound attenuated room in front of computer screens, and wore headphones. Token presentation and response collection were controlled with *SuperLab Pro 4.0*. Each participant had a response box, with buttons marked as "1" and "2".

| *Alternate?* | *Vowel* | *Singular* | *Plural* | *# of pairs* | *Picture* |
|---|---|---|---|---|---|
| Yes | /ɛ/ | [təl.sɛ́p] | [təl.sé.pən] | 2 | carrot |
| | /ɪ/ | [kən.dɪ́p] | [kən.dí.pən] | 2 | horse |
| | /ɔ/ | [dəf.sɔ́k] | [dəf.só.kən] | 2 | pencil |
| | /ʊ/ | [pəm.bʊ́k] | [pəm.bú.kən] | 2 | arm chair |
| No | /ɛ/ | [məl.tɛ́p] | [məl.tɛ́.pən] | 2 | nail |
| | /ɪ/ | [məs.pɪ́t] | [məs.pɪ́.tən] | 2 | truck |
| | /ɔ/ | [kən.tɔ́p] | [kən.tɔ́.pən] | 2 | bicycle |
| | /ʊ/ | [səŋ.gʊ́f] | [səŋ.gʊ́.fən] | 2 | cactus |

Table 1: Learning data for Experiment 1

A learning event proceeded as follows: Participants saw a picture of a singular object (a carrot), and heard the singular word associated with that picture. After 500 ms, the plural picture (a bunch of carrots) was presented, and the correct plural form (alternating or non-alternating) was played simultaneously. After 1000 ms, the same pair was presented again. Participants then pushed a button to move to the next item. Each singular/plural pair was presented five times like this (five repetitions, with two presentations in each, giving ten repetitions per pair). Presentations were randomized.

*Testing*. After the learning phase, testing was conducted. Participants were tested on the 16 words included in the learning data, but also on the 16 words to which they had no prior exposure. Learned and novel tokens were randomized together. A testing event proceeded as follows: Participants saw a picture of a single object (a carrot), and heard the singular word associated with that picture. After 500 ms, the corresponding plural picture was displayed (a bunch of carrots), and both possible pronunciations of the plural was presented, with 500 ms between the pronunciations. The display then changed to "1 or 2?". A participant indicated which plural he/she considered to be correct by pushing the corresponding button. Participants received no feedback. The next test event followed 2000 ms after a response was registered. Each singular/plural pair was presented twice. The order in which the two plural pronunciations was presented was changed the for the second presentation of a pair.

### *2.1.2. Results and discussion*

Participants performed well on the learned tokens, scoring 95.2% correct on non-alternating pairs and 90.9% on alternating pairs. Percent correct responses was significantly above chance for alternating and non-alternating pairs by participants (one-tailed *t*: alternating, $t(12) = 14.17$, $p < .001$; non-alternating, $t(12) = 22.36$, $p < .001$) and items (one-tailed *t*: alternating, $t(7) = 19.96$, $p < .001$; non-alternating, $t(7) = 32.11$, $p < .001$). These results are represented in Figure 3(a), and show that participants successfully learned which words alternate and which do not.

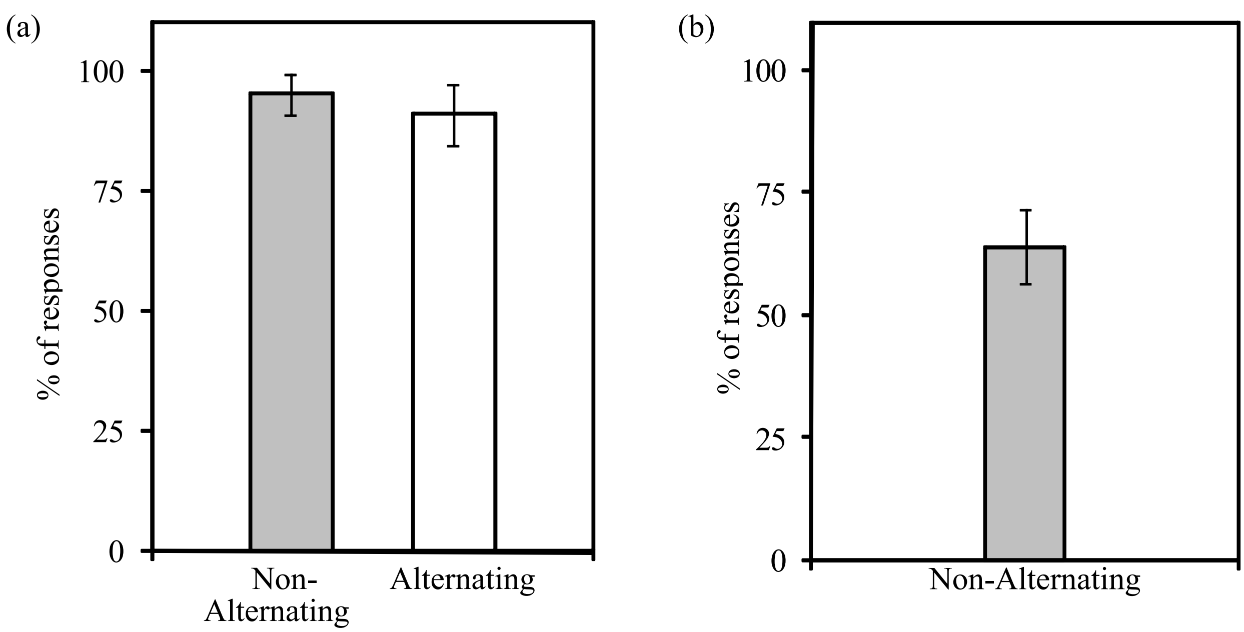


Figure 3: Experiment 1: (a) Percent correct responses to learned tokens. (b) Response pattern to novel items. Error bars show 95% confidence intervals.

For novel items, participants preferred non-alternation over alternation 64.2% of the time. These results are shown in Figure 3(b). Non-alternating responses were more likely than chance on one-tailed *t*-tests by participants ($t(12) = 2.01$, $p < .04$) and items ($t(15) = 5.58$, $p < .001$).

In Experiment 1, the alternating and non-alternating pattern had equal frequency in the learning data. Participants could not rely on the learning data to determine which pattern to treat as default. As shown in §2.1.1, there was also no phonological patterns that distinguished alternating and non-alternating pairs systematically, and participants could not rely on such patterns to decide the class membership of novel items. Under this learning condition, participants were more likely to extend the non-alternating pattern to novel items, showing that non-alternation is the default chosen in the absence of evidence. This agrees with the general principle from (5).

Although the participants preferred the non-alternating pattern for novel forms, this preference was not absolute – alternation was selected 35.8% of the time. This might be evidence of factors other than the default non-alternation preference also impacting the performance of the participants. For instance, it might be that participants calculated that alternation and non-alternation were equally likely in the learning data and that they are trying to maintain this ratio.[3] There would then be one force that wants alternation and non-alternation to be equally likely for novel forms, and another that wants all novel forms to be non-alternating. The participants' performance then reflects the outcome of these two conflicting forces.

### 2.2. Experiment 2

Although English does not tolerate stressed lax vowels in word-final open syllables, such vowels are tolerated non-word-finally – see (7). It is hence possible that the participants in Experiment 1 did not do grammar learning, but that they memorized the plurals in the learning data, and relied on their English grammar when presented with novel items. The non-alternation preference in the novel forms might then reflect transfer from their native grammar rather than the default option employed in learning. Experiment 2 was designed to test whether English speakers can learn a grammar that lengthens lax vowels in word-medial, stressed, open syllables. If they can, it would be support for the interpretation of the results of Experiment 1 as reflecting the outcome of grammar learning.

(7) Lax, stressed vowels in open syllables

| | |
|---|---|
| Non-finally: | *better* [bɛ́.ɾəɹ], *villain* [vɪ́.lən]. |
| But not finally: | *[im.plɛ́], *[ɹi.plɪ́] |

#### *2.2.1. Methods*

*Participants*. Eleven participants were recruited exactly as for Experiment 1.

*Tokens*. I selected 16 of the 32 singular/plural pairs from Experiment 1, four pairs with each of the lax vowels in the singular. These singulars had the structure [CəC.CV́C], and I refer to them here as the [VC#]-tokens. I also selected 16 additional non-words, all with the structure

3 Walter (2006) shows that gender assignment to Spanish loan words is subject to such a frequency constraint. A large number of Spanish nouns were borrowed from Arabic. Walter shows that the likelihood of these loans being assigned masculine or feminine gender reflects the frequency of the two genders in the native Spanish lexicon.

[CəC.Cv́CC], where [v́] was again one of the four lax vowels. Each of the lax vowels was represented by four non-words. These tokens are referred to as the [VCC#]-tokens. As before, two plurals were created for each non-word. The plurals were formed in the same manner as in Experiment 1, by the addition of [-ən], and with the same vowel as in the singular or with the corresponding tense vowel. Examples of the additional items are given in (8), with a full list in Appendix B. Tokens were recorded by the same speaker and prepared in the same way as in Experiment 1.

(8) Examples of additional tokens for Experiment 2

| | | *Plural* | |
|---|---|---|---|
| | *Singular* | *Non-Alternating* | *Alternating* |
| [ɛ́]: | [kən.dɛ́ft] | [kən.dɛ́f.tən] | [kən.déf.tən] |
| [ɪ́]: | [fən.sɪ́ŋk] | [fən.sɪ́ŋ.kən] | [fən.síŋ.kən] |
| [ɔ́]: | [vəl.dɔ́mf] | [vəl.dɔ́m.fən] | [vəl.dóm.fən] |
| [ʊ́] | [fək.sʊ́lt] | [fək.sʊ́l.tən] | [fək.súl.tən] |

*Learning*. Learning commenced as in Experiment 1. Half each of the [VC#] and [VCC#]-tokens were selected for learning – two for each of the vowels for the two token groups. All [VC#]-tokens were paired with their alternating plurals, and all [VCC#]-tokens with their non-alternating plurals. Although the learning data contained evidence of alternation and non-alternation, the alternating and non-alternating paradigms differ systematically. Every [VC#]-singular has an alternating plural, so that the learning data contain consistent evidence that stressed, lax vowels in open syllables are lengthened.

*Testing*. Testing was done as in Experiment 1.

For the novel [VCC#]-tokens, we expect participants to prefer the non-alternating plural, both because the learning data only contained evidence for non-alternation in this token type, and because the results of Experiment 1 showed a general preference for non-alternation. For novel [VC#]-tokens, it is less clear what to expect. The learning data contain only evidence of alternation in this context. If participants in Experiment 2 do learn a grammar based on the learning data, we would expect them to extend the alternating pattern to novel [VC#]-tokens. If they rely on their native English grammars (no alternation), they should extend the non-alternating pattern also to these tokens.

### *2.2.2. Results and discussion*

Participants performed well on the learned tokens, scoring 90.9% correct on the [VC#]-tokens and 85.3% on the [VCC#]-tokens. Percent correct responses was significantly above chance for both token types by participants (one-tailed $t$: [VC#], $t(10) = 22.04$, $p < .001$; [VCC#], $t(10) = 9.92$, $p < .001$) and items (one-tailed $t$: [VC#], $t(7) = 15.27$, $p < .001$; [VCC#], $t(7) = 7.00$, $p < .001$). These results are represented in Figure 4(a).

There are at least two explanations for this performance. It is possible that participants did not select the alternating plural for [VC#]-forms because they learned a grammar that lengthens lax vowels in open, stressed syllables, but that they simply memorized the correct plural for every singular. If this is what participants did, there is no reason to expect them to treat [VC#] and [VCC#] novel tokens differently. On the other hand, if they did a grammatical analysis of the learning tokens, and learned a vowel lengthening grammar consistent with the data, they should prefer non-alternating plurals for [VCC#]-forms and alternating plurals for [VC#]-forms. Participants were much more likely to prefer alternating plurals for novel [VC#]-tokens (46.5%) than for novel [VCC#]-tokens (18.1%), as shown in Figure 4(b). The difference in preference for alternating responses was significant on one-tailed *t*-tests by items ($t(14) = 4.32$, $p < .001$) and participants ($t(20) = 5.75$, $p < .001$).[4]

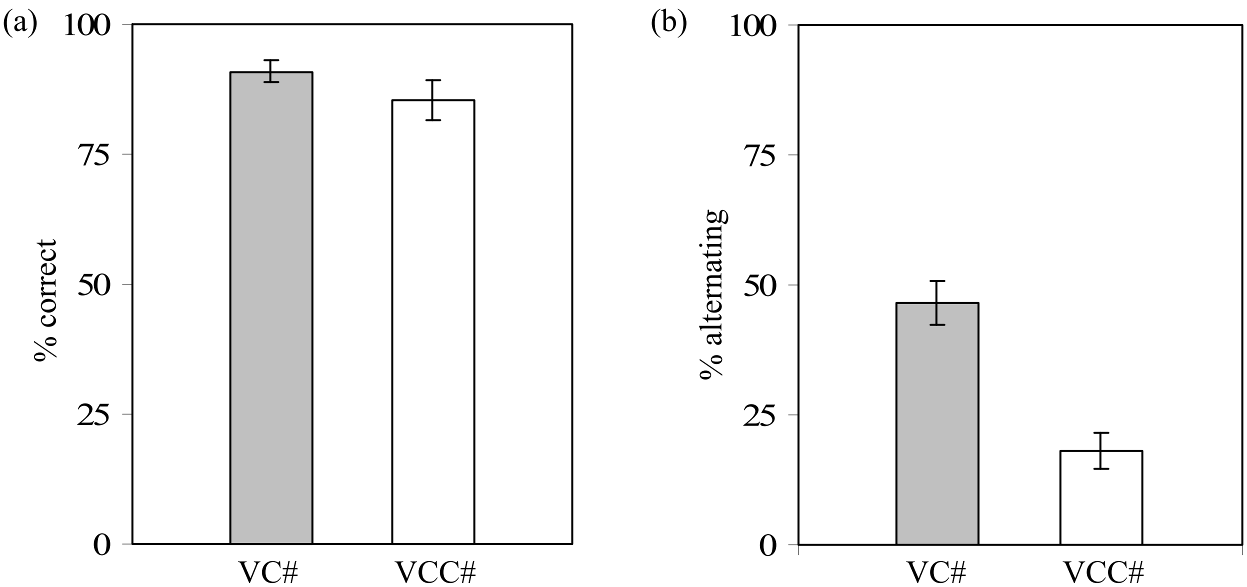


Figure 4: Experiment 2: (a) Percent correct responses to learned tokens. (b) Percent alternating response to novel items. Error bars show 95% confidence intervals.

I take the different response patterns to the two kinds of novel items as evidence that participants do differentiate grammatically between the two kinds. The first conclusion that can be drawn from this is that participants did perform a grammatical task rather than mere memorization of learned items and random guessing on novel items. Participants show a stronger aversion for the non-alternating plural in [VC#]-tokens. The non-alternating plurals of these tokens all contain stressed, lax vowels in open syllables. Although English tolerates such forms, participants avoided them, showing that English speakers can learn a grammar that differs from their native grammar. This supports the interpretation that participants in Experiment 1 performed a grammatical task rather than mere memorization.

Although participants in Experiment 2 were more likely to select alternation for novel [VC#]-tokens than [VCC#]-tokens, they did not overwhelmingly prefer alternation for [VC#]-forms. In this regard, participants in Experiment 1 and Experiment 2 performed similarly – both preferred non-alternation over alternation for novel [VC#]-forms (64.2% in Experiment 1, 53.5% in Experiment 2). Even so, there is evidence that Experiment 2 participants were less likely to select non-alternation than Experiment 1participants. The percent non-alternation selected for novel [VC#]-tokens is significantly lower in Experiment 2 than Experiment 1 by items (one tailed $t(22) = 1.84$, $p = .039$) and tended towards significance by participants (one-tailed $t(22)$ =

[4] Since every participant has a score for novel [VC#]- and [VCC#]-tokens, the participant *t*-test can also be performed as a paired samples *t*-test. This does not affect the result: $t(10) = 11.35$, $p < .001$.

1.27, $p$ = .109). The difference in the structure of the learning data between these experiments did result in a difference in performance on novel items.

## 2.3. Experiment 3

Based on learnability considerations discussed in §1.3, I predicted that novel items will by default be treated as non-alternating, and the results of Experiment 1 confirmed this. In Experiment 1, the learning data contained equal evidence for alternation and non-alternation, and participants could not rely on the learning data for guidance about how to treat novel items. These participants were more likely to extend the non-alternation pattern to novel items. However, it must be possible for learners to overcome the non-alternation bias. There are, after all, many examples of morphologically conditioned phonological alternations that apply without exception. Experiment 3 is designed to determine whether the frequency structure of the learning data can contribute to overcoming the non-alternation bias. In this experiment, there are two learning conditions. In one, the learning data contain only alternating plurals, giving evidence exclusively for alternation and hence against the default non-alternation preference. In the second condition, 75% of the learning data contain alternating plurals. The first condition is intended to test whether the bias can be overcome at all, and the second condition to determine whether a mere majority of alternating forms is sufficient.

### *2.3.1. Methods*

*Participants*. Twenty participants were recruited as described for Experiment 1. Participants were randomly assigned to the learning conditions so that there where ten participants in each condition.

*Tokens*. Exactly the same tokens as in Experiment 1 were used.

*Learning*. Half of the tokens were selected for learning – four with each of the lax vowels in the singular. In Condition 1, all tokens were assigned to the alternating class, and the singulars were all paired with plurals with the corresponding tense vowels. In Condition 2, three out of the four tokens for each lax vowel were assigned to the alternating group, and one to the non-alternating group. The singulars were paired with plurals with either the corresponding tense vowel, or with the lax vowel of the singular. There were hence three alternating plural pairs with the vowel [έ] in the singular, and [é:] in the plural, and one pair with [έ] in both the singular and plural, and similarly for the other three lax vowels. Learning commenced as in Experiment 1. Table 2 shows the structure of the learning data for the two conditions.

*Testing*. Testing was done in exactly the same manner as in Experiment 1.

### *2.3.2. Results and discussion*

Participants in both conditions did well on learned tokens. In Condition 1 (all alternation), the correct answer was selected 98.4% of the time, and in Condition 2, the correct response was given for alternating words 92.1% of the time and for non-alternating words 88.8%. All of these were significantly higher than chance on a one-tailed $t$-test by both participants (Condition 1: $t(9) = 69.32$, $p < .001$; Condition 2: alternating, $t(9) = 13.99$, $p < .001$; non-alternating, $t(9) = 8.91$, $p < .001$) and items (Condition 1: $t(15) = 48.85$, $p < .001$; Condition 2: alternating, $t(11) = 19.37$, $p < .001$; non-alternating, $t(3) = 10.33$, $p < .002$). These results are represented in Figure 5(a).

| | *Alternate?* | *Vowel* | *Singular* | *Plural* | *# of pairs* | *Picture* |
|---|---|---|---|---|---|---|
| Condition 1 | Yes | /ɛ/ | [təl.sɛ́p] | [təl.sé.pən] | 4 | carrot |
| | | /ɪ/ | [kən.dɪ́p] | [kən.dí.pən] | 4 | horse |
| | | /ɔ/ | [dəf.sɔ́k] | [dəf.só.kən] | 4 | pencil |
| | | /ʊ/ | [pəm.bʊ́k] | [pəm.bú.kən] | 4 | arm chair |
| Condition 2 | Yes | /ɛ/ | [təl.sɛ́p] | [təl.sé.pən] | 3 | carrot |
| | | /ɪ/ | [kən.dɪ́p] | [kən.dí.pən] | 3 | horse |
| | | /ɔ/ | [dəf.sɔ́k] | [dəf.só.kən] | 3 | pencil |
| | | /ʊ/ | [pəm.bʊ́k] | [pəm.bú.kən] | 3 | arm chair |
| | No | /ɛ/ | [məl.tɛ́p] | [məl.tɛ́.pən] | 1 | nail |
| | | /ɪ/ | [məs.pɪ́t] | [məs.pɪ́.tən] | 1 | truck |
| | | /ɔ/ | [kən.tɔ́p] | [kən.tɔ́.pən] | 1 | bicycle |
| | | /ʊ/ | [səŋ.gʊ́f] | [səŋ.gʊ́.fən] | 1 | cactus |

Table 2: Learning data for Experiment 3

In Condition 1, participants heard only alternating plurals during learning. These participants also overwhelmingly preferred alternation for the novel items (94.7%). This preference is significantly higher than chance on a one-tailed *t*-test by participants ($t(9) = 17.22$, $p < .001$) and by items ($t(15) = 29.88$, $p < .001$). In Condition 2, 75% of the learning data gave evidence for alternation. Participants in this condition selected alternation for novel items 52.5% of the time. This preference was not significantly higher than chance on a one-tailed *t*-test by participants ($t(9) = 0.31$, $p = .39$) or items ($t(15) = 1.10$, $p = .15$). It was, however, significantly different from 75%, showing that the response pattern does not match the frequency pattern of the learning data: by participants (two tailed $t(9) = -2.76$, $p < .03$), and by items (two tailed $t(15) = -9.86$, $p < .001$). These results are represented in Figure 5(b).

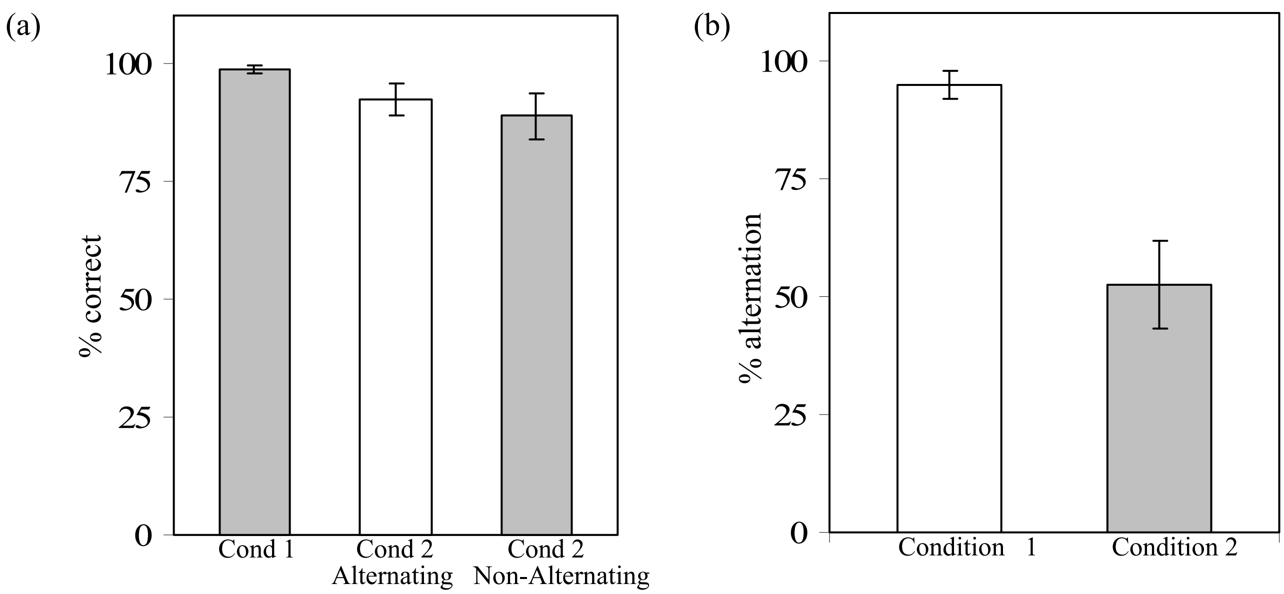


Figure 5: Experiment 3: (a) Percent correct responses to learned tokens. (b) Percent alternating response to novel items. Error bars show 95% confidence intervals.

The results on the novel items in Condition 1 show that the default non-alternation bias can be overcome. The results from Condition 2 give evidence for the robustness of the non-alternation bias. Although 75% of the evidence in the learning data was for alternation, participants did not prefer alternation more than non-alternation for novel items. However, the performance of Condition 2 participants does differ from that of participants in Experiment 1. In Experiment 1, participants received equal evidence for alternation and non-alternation, and they were significantly more likely to treat novel items as non-alternating (64.2%). In Condition 2 of Experiment 3, the learning data were skewed in favor of alternation. Although this skewedness was not sufficient to result in an overall preference for alternation, it was sufficient to wipe out the advantage of the non-alternating over the alternating pattern. Non-alternating responses were significantly more likely in Experiment 1 than in Condition 2 of Experiment 3 by items (one-tailed $t(30) = 3.79$, $p < .001$), and tended towards significance by participants (one-tailed $t(21) = 1.55$, $p < .07$). This gives some indication of the strength of the default non-alternation preference. When 75% of the data give evidence for alternation, the default non-alternation preference can be wiped out. However, this is not enough to result in a preference for alternation.

### 2.4. Summary discussion

Many languages have the type of pattern shown for Dutch in (1). In these data, there are some paradigms in which the root is realized the same throughout, and others in which the root alternates between different paradigm positions. Additionally, there is no consistent phonological difference between alternating and non-alternating roots. The question investigated here is what learners do in such a scenario. Which pattern are they more likely to extend to novel items? Many factors influence how novel items will be treated. For instance, if the majority of the learning data gives evidence for alternation (or non-alternation), learners might extend this pattern to novel items. Similarly, when presented with a novel item that is phonologically very similar to an actual alternating (or non-alternating) word in the learning data, the learner might extend this pattern to this novel item. But what will learners do in the absence of such frequency or analogical evidence? In §1.3, I argued, based on learnability considerations, that learners will assign novel items to the non-alternating class in the absence of evidence to the contrary. The results of the experiments discussed here are consistent with this claim.

In Experiment 1, learners were exposed to learning data with equal evidence for alternation and non-alternation. The results of this experiment show that learners were more likely to extend the non-alternating pattern to novel items under these conditions. In Experiment 3, learners were exposed either exclusively to evidence for alternation or to more evidence for alternation than non-alternation. In this experiment, learners can be guided by the learning data when faced with novel words. Learners exposed exclusively to evidence for alternation overwhelmingly extended the alternating pattern to novel items, showing that the default non-alternation preference can be overcome by the frequency structure of the learning data.

Learners in the second condition of Experiment 3 received three times as much evidence for alternation as for non-alternation. These learners were more likely to extend alternation to novel items than the learners in Experiment 1 (52.5% vs. 38.8%). However, they did not prefer alternation over non-alternation by a 3:1-ratio. This gives insight into the relationship between the limitations imposed on learning by learnability considerations and by frequency information in the data. Both factors contribute to the learners' performance. In this instance, learnability considerations favor non-alternation while the frequency structure of the data favors alternation. The performance of the participants shows the effects of the conflict between these factors.

## 3. LEARNING A DUTCH-LIKE VOWEL LENGTH GRAMMAR

The discussion in the preceding sections is not specific to any particular phonological theory. The subset problem follows from the structure of the data to which the learner is exposed, and these data are the same irrespective of the specific instantiation of phonological grammar assumed. In this section, I develop an Optimality Theoretic (OT) grammar for the Dutch data in (1), and show how this grammar can be learned with an expanded version of Prince and Tesar's (2004) Biased Constraint Demotion Algorithm (BCD). The participants in the experiments discussed in §2 were exposed to learning data modeled after the Dutch pattern, so that the grammar developed here for Dutch can also serve as a model for what these participants might have done. The purpose of this section is not to argue that OT is the only kind of theory in which the subset problem can be overcome, or even that it is the best theory for this purpose. The intention is to illustrate how this problem can be addressed in one specific instantiation of phonological theory. In §3.1, I first show how the difference between restrictive and permissive grammars can be encoded in OT. In §3.2, I develop the basic Dutch vowel length grammar, and §3.3 is dedicated to showing how this grammar can be learned.

### 3.1 Restrictive and permissive grammars in OT

In this section, I show how restrictive and permissive grammars can be encoded in OT. In (9), I repeat the data used above to illustrate the difference between restrictive and permissive grammars in the input/output domain. Let us assume that a child had to learn a grammar based on these data, and that his/her grammar contained only the constraints $\text{DEP}_{\text{IO}}$ and *COMPLEX. The child learning under these conditions can settle on either of the possible rankings between these constraints – both rankings will rate the forms in (9) as grammatical. However, the rankings differ in how they evaluate an input with a consonant cluster.

| | | | | |
|---|---|---|---|---|
| (9) | [bit] | [lap] | [guk] | [pef] |
| | [sop] | [tip] | [sam] | [kus] |

The tableaux in (10) show how each of these rankings evaluate an input with and without a cluster. Since the faithful candidate for an input without clusters (/bit/) does not violate *COMPLEX, the faithful candidate will be optimal irrespective of the ranking. However, when an input with a cluster (/blit/) is submitted to the grammar, the situation is different. The faithful candidate ([blit]) violates *COMPLEX. Under the ranking *COMPLEX ≫ DEP, as in (10a), the epenthetic candidate ([bəlit]) is hence optimal. Under the opposite ranking, as in (10b), the faithful candidate ([blit]) is optimal. The ranking with markedness ≫ IO-faithfulness therefore represents the restrictive grammar, where forms absent from the learning data (with clusters) are treated as ungrammatical. Since the restrictive grammar is to be preferred based on learnability arguments, it follows that in the absence of evidence to the contrary, the ranking markedness ≫ IO-faithfulness should be assumed (Itô & Mester 1999; Tesar & Smolensky 2000; Hayes 2004; Prince & Tesar 2004; Gnanadesikan 2004; etc.).

(10) a. Restrictive grammar

| | | *COMPLEX | $DEP_{IO}$ |
|---|---|---|---|
| In learning data | /bit/ ☞ bit | | |
| | bi.tə | *! | * |
| Not in learning data | /blit/ blit | *! | |
| | ☞ bə.lit | | * |

b. Permissive grammar

| | | $DEP_{IO}$ | *COMPLEX |
|---|---|---|---|
| In learning data | /bit/ ☞ bit | | |
| | bi.tə | *! | * |
| Not in learning data | /blit/ ☞ blit | | * |
| | bə.lit | *! | |

In (11), I repeat the data from (4) above. In (11a) are the data to which the learner is exposed, consisting of singular/plural paradigms, in which the noun root surfaces the same in the singular and the plural. In (11b) are examples not included in the learning data, with roots that alternate between the singular and plural. Assume that a child had to learn a grammar from these data, and that his/her grammar had only the constraints: *VOIOBS]$_\sigma$ (no syllable-final voiced obstruents), IDENT[voice]$_{IO}$ (no voicing change from input to output), and IDENT[voice]$_{OO}$ (voicing specification in derived forms of a paradigm agrees with the base form). There are two rankings between these constraints that are consistent with the data in (11a). Corresponding tableaux are given in (12).

(11) a. In learning data — b. Not in learning data

| a. *Singular* | *Plural* | b. *Singular* | *Plural* |
|---|---|---|---|
| gat | ga.ti | bat | ba.di |
| pok | po.ki | dok | do.gi |
| zep | ze.pi | vep | ve.bi |

These tableaux show that both grammars are consistent with the learning data. When presented with a singular/plural paradigm input that was contained in the learning data, like /gat/~/gat-i/, both grammars select [gat]~[ga.ti]. However, they differ on how they evaluate a paradigm input like /gad/~/gad-i/. The permissive grammar in (12a) selects [gat]~[ga.di] with the final root consonant alternating in voicing. The restrictive grammar in (12b) selects [gat]~[ga.ti] without this alternation. In the restrictive grammar, all inputs are mapped onto output patterns that were present in the learning data. In effect, only patterns included in the learning data will ever be generated by this grammar. The permissive grammar also generates patterns that were not in the learning data.

In both grammars, the ranking *VOIOBS]$_\sigma$ ≫ IDENT[voi]$_{IO}$ is assumed. This reflects the default ranking that was established just above in (9) and (10). In the absence of evidence to the contrary, the learner assumes that markedness constraints outrank IO-faithfulness constraints. Since the learning data in (11) contain no syllable-final voiced obstruents, the learner assumes that such outputs are ungrammatical, and settles on *VOIOBS]$_\sigma$ ≫ IDENT[voi]$_{IO}$. The grammars differ in where IDENT[voi]$_{OO}$ ranks relative to the other constraints. In the permissive grammar

(12a), IDENT[voi]$_{OO}$ ranks at the bottom. This is what causes this grammar to accept alternating paradigms as grammatical, as seen in the last two rows of (12a), where /gad-i/ is evaluated. The faithful candidate [ga.di] violates IDENT[voi]$_{OO}$ – since the root final consonant differs in voicing from its correspondent in the base [gat]. The unfaithful candidate [ga.ti], however, violates IDENT[voi]$_{IO}$. The ranking IDENT[voi]$_{IO}$ $\gg$ IDENT[voi]$_{OO}$ is directly responsible for the alternating paradigm being rated as grammatical.

(12) a. Permissive grammar: *VOIOBS]$_{\sigma}$ $\gg$ IDENT[voice]$_{IO}$ $\gg$ IDENT[voice]$_{OO}$

| | | | *VOIOBS]$_{\sigma}$ | IDENT[voi]$_{IO}$ | IDENT[voi]$_{OO}$ |
|---|---|---|---|---|---|
| In learning data | /gat/ | ☞ gat | | | |
| | | gad | *! | * | |
| | /gat-i/ | ☞ ga.ti | | | |
| | Base: [gat] | ga.di | | *! | * |
| Not in learning data | /gad/ | ☞ gat | | * | |
| | | gad | *! | | |
| | /gad-i/ | ga.ti | | *! | |
| | Base: [gat] | ☞ ga.di | | | * |

b. Restrictive grammar: IDENT[voice]$_{OO}$, *VOIOBS]$_{\sigma}$ $\gg$ IDENT[voice]$_{IO}$

| | | | IDENT[voi]$_{OO}$ | *VOIOBS]$_{\sigma}$ | IDENT[voi]$_{IO}$ |
|---|---|---|---|---|---|
| In learning data | /gat/ | ☞ gat | | | |
| | | gad | | *! | * |
| | /gat-i/ | ☞ ga.ti | | | |
| | Base: [gat] | ga.di | *! | | * |
| Not in learning data | /gad/ | ☞ gat | | | * |
| | | gad | | *! | |
| | /gad-i/ | ☞ ga.ti | | | * |
| | Base: [gat] | ga.di | *! | | |

In the restrictive grammar (12b), the ranking between IDENT[voi]$_{IO}$ and IDENT[voi]$_{OO}$ is inverted, so that a plural form that differs from its singular base is no longer tolerated. The ranking IDENT[voi]$_{OO}$ $\gg$ IDENT[voi]$_{IO}$ therefore assures that an alternating paradigm cannot be rated as grammatical in this grammar.

The ranking IDENT[voi]$_{IO}$ $\gg$ IDENT[voi]$_{OO}$ results in a permissive grammar, while the opposite ranking IDENT[voi]$_{OO}$ $\gg$ IDENT[voi]$_{IO}$ results in a restrictive grammar. Based on the discussion in §1.3, the restrictive grammar should be preferred all else being equal. From this follows the general principle that in the absence of evidence to the contrary, OO-faithfulness

≫ IO-faithfulness should be assumed. See also McCarthy (1988), Hayes (2004), and Tessier (2007). In (13), I give the two ranking schemas that are assumed in OT to ensure restrictive grammar learning. These schemas are built into the learning algorithms discussed in §3.3.

(13) Default rankings[5]
  a. Markedness ≫ IO-faithfulness
  b. OO-faithfulness ≫ IO-faithfulness

### 3.2 The basic Dutch vowel length grammar

The Dutch data in (1) contain examples of alternating ([xɑ́t]~[xáː.tən]) and non-alternating nouns ([kɑ́t]~[kɑ́.tən]). Since it cannot be predicted based on the phonological properties of a noun whether it will alternate or not, the class membership of every noun needs to be lexically encoded. This kind of idiosyncratic behavior of lexical items has been handled in several different ways in OT, including the cophonology approach where individual morphemes choose their own constraint rankings (Anttila 2002), underspecification approaches (Inkelas, Orgun & Zoll 1997), allomorphy approaches where alternating morphemes have two lexical entries (Burzio 1996; Kager to appear), and by using lexically indexed constraints (Pater 2000; Itô & Mester 1999, 2002). It is not the goal of this paper to evaluate different ways of accounting for lexical idiosyncracy in OT, and I do not discuss all of these approaches in detail. I opt to use the indexed constraints approach, primarily because this is the only one among these that has been studied from a learnability angle (Hayes 2004; Ota 2004; Pater 2005; Tessier 2007).

The constraints that I assume are listed in (14). The markedness constraint, SWP, is the stress-to-weight principle (Prince 1990). I use both IO and OO-faithfulness constraints, because the plurals of alternating nouns differ from their inputs, and from their singulars. The faithfulness constraints assume that vowel length is a bivalent feature [±long]. This interpretation of vowel length is not crucial to the point of the paper. If vowel length is interpreted in terms of moraic structure (short vowels are monomoraic, long vowels are bimoraic), vowel lengthening will be the result of mora addition. The faithfulness constraints will then have to be replaced with DEP-$\mu_{IO}$ and DEP-$\mu_{OO}$, but this will not affect the analysis or the learnability account presented here.

(14)

| | | |
|---|---|---|
| SWP | Stressed syllables are heavy. |
| IDENT[long]$_{IO}$ | Let $V_O$ be a vowel in some output form, and $V_I$ its input correspondent. If $V_I$ is [αlong], then $V_O$ is [αlong]. |
| IDENT[long]$_{OO}$ | Let $V_B$ be a vowel in the base form of a morphological paradigm, and $V_D$ its correspondent in some derived form of the paradigm. If $V_B$ is [αlong], then $V_D$ is [αlong]. |

Lexical items are stored as members of specific lexical classes – alternating (A) or non-alternating (NA) here. Constraints can then be indexed to lexical classes, so that there may be an IDENT[long]$_{IO}$-A and an IDENT[long]$_{IO}$-NA constraint. An indexed constraint evaluates only lexical items that share its index. There is disagreement in the literature on whether only correspondence constraints can be indexed (Itô & Mester 2002) or whether also markedness constraints can be indexed (Pater to appear). I develop grammars here under both of these assumptions.

[5] No default ranking between markedness and OO-faithfulness constraints is included here, and based on the discussion in this section it seems to be unnecessary. However, when dealing with indexed constraints, the default ranking OO-faithfulness ≫ markedness is also required. This is motivated in §3.3.

(15) Grammars compatible with Dutch plural patterns

Lexicon: A = {/xɑt/, /spɛl/, /sxɪp/, /xɔd/} NA = {/kɑt/, /stɛl/, /ʋɪp/, /fɔd/}

a. IO-Blocking grammar

| | | | Id[long]$_{IO}$-NA | SWP | Id[long]$_{IO}$-A | Id[long]$_{OO}$ |
|---|---|---|---|---|---|---|
| Alternating | /xɑt$_{A}$-ən/ | xá.tən | | *! | | |
| | OO-Base: [xát] | ☞ xáː.tən | | | * | * |
| Non-Alternating | /kɑt$_{NA}$-ən/ | ☞ ká.tən | | * | | |
| | OO-Base: [kát] | káː.tən | *! | | | * |

b. OO-Blocking grammar

| | | | Id[long]$_{OO}$-NA | SWP | Id[long]$_{OO}$-A | Id[long]$_{IO}$ |
|---|---|---|---|---|---|---|
| Alternating | /xɑt$_{A}$-ən/ | xá.tən | | *! | | |
| | OO-Base: [xát] | ☞ xáː.tən | | | * | * |
| Non-Alternating | /kɑt$_{NA}$-ən/ | ☞ ká.tən | | * | | |
| | OO-Base: [kát] | káː.tən | *! | | | * |

c. Indexed markedness grammar

| | | | SWP-A | Id[long]$_{OO}$ | Id[long]$_{IO}$ | SWP-NA |
|---|---|---|---|---|---|---|
| Alternating | /xɑt$_{A}$-ən/ | xá.tən | *! | | | |
| | OO-Base: [xát] | ☞ xáː.tən | | * | * | |
| Non-Alternating | /kɑt$_{NA}$-ən/ | ☞ ká.tən | | | | * |
| | OO-Base: [kát] | káː.tən | | *! | *! | |

The tableaux in (15) show the grammars that are consistent with the Dutch data. The tableaux only show the evaluation of plurals. All of the Dutch nouns considered here have /CVC/-roots, so that the faithful candidate of a singular can never violate SWP – assuming that /CVC/-

syllables are heavy – and the faithful candidate will always be selected as output for a singular. In each tableau, only one family of constraints (markedness, IO-faithfulness, OO-faithfulness) is indexed. In reality, it is possible that a single grammar can have indexed constraints from all three families (Pater to appear). These tableaux represent the simplest grammars (fewest constraints) consistent with the data, but they can be expanded to accommodate the other possible indexed constraints. Every unindexed constraint can be replaced by its two corresponding indexed versions. As long as the A and NA versions rank in the same place in the hierarchy, this will not impact output selection.

The grammar in (15a) has indexed IO-faithfulness constraints, that in (15b) indexed OO-faithfulness constraints, and that in (15c) indexed markedness constraints. Consider first the indexed IO-faithfulness grammar from (15a). Neither SWP nor IDENT[long]$_{OO}$ are indexed, and these constraints apply to all words. Both alternating and non-alternating nouns are hence subject to the ranking SWP ≫ IDENT[long]$_{OO}$, so that the grammar does not enforce identity between the singular base and the plural for any of the noun classes. IDENT[long]$_{IO}$, on the other hand, is indexed to the two noun classes, and this constraint treats the classes differently. The ranking SWP ≫ IDENT[long]$_{IO}$-A applies to the alternating class. These nouns are therefore also not required by the grammar to be realized identically to their inputs. The non-alternating nouns, on the other hand, are evaluated by the ranking IDENT[long]$_{IO}$-NA ≫ SWP. Although the plurals of these nouns are allowed to differ from their singulars, they are not allowed to differ from their inputs. The non-alternation of the non-alternation class is the result of these nouns being subject to a special requirement on input~output identity, and not because of any special requirement for uniform paradigms. The indexed OO-faithfulness grammar in (15b) is identical to the indexed IO-faithfulness grammar in (15a), except that the IO and OO-faithfulness constraints are flipped around. As a consequence, the non-alternating nouns are not subject to any special input~output identity requirement, but they are subject to a special requirement on paradigmatic consistency. Although both grammars result in selection of the same optimal forms, this selection is motivated differently in the two grammars.

The motivation for alternation or non-alternation is different in the indexed markedness grammar in (15c). For the alternating nouns, the markedness constraint outranks both faithfulness constraints – SWP-A ≫ {IDENT[long]$_{OO}$, IDENT[long]$_{IO}$}. To satisfy SWP, the plural of these nouns are hence allowed to differ both from their inputs and from their singular bases. To block vowel lengthening in the non-alternating nouns, the markedness constraint that applies to them, SWP-NA, has to be dominated by at least one of the faithfulness constraints.

All three of the grammars in (15) are consistent with the Dutch data from (1). To apply any of these grammars to some input form, the language user must know the lexical class of the input. This is not a problem for a noun that was contained in the learning data. The data contain the information the learner would need to determine the lexical class affiliation. These three grammars therefore perform identically with regard to nouns that were included in the learning data. The situation is different for nouns that were not in the learning data. Suppose that /nɔl/ were a Dutch noun that was accidentally absent from the learning data. What would the learner do when faced with an input like /nɔl-ən/? To evaluate the candidate outputs for this input, /nɔl/ must be assigned to either the NA or the A class.

The learner has two options: He/she can randomly assign new roots to lexical classes, or he/she can base class assignment on some principle. Based on the results of the experiments discussed above, I assume that the class assignment of novel forms is not done randomly. What is then the principle that determines class assignment? What is the default class affiliation assumed for lexical items in the absence of evidence about their class affiliation? I propose that class assignment is determined by the imperative to learn the most restrictive grammar consistent with the learning data. When presented with a new form, the learner assigns this form to the lexical class that will result in the form being subject to the default restrictive rankings from (13). The actual class assignment in different scenarios is done according to (16).

(16) Default lexical class assignment

a. Indexed IO-faithfulness constraints

Let $IO_X$ and $IO_Y$ be two indexed versions of the same IO-faithfulness constraint, indexed to lexical class *X* and *Y* respectively. In a grammar with $IO_X$ and $IO_Y$:

i. If there is some markedness constraint M such that M ≫ $IO_X$, and not M ≫ $IO_Y$, assign novel forms to lexical class *X*.

ii. If there is some OO-faithfulness constraint OO such that OO ≫ $IO_X$, and not OO ≫ $IO_Y$, assign novel forms to lexical class *X*.

b. Indexed OO-faithfulness constraints

Let $OO_X$ and $OO_Y$ be two indexed versions of the same OO-faithfulness constraint, indexed to lexical class *X* and *Y* respectively. In a grammar with $OO_X$ and $OO_Y$:

i. If there is some IO-faithfulness constraint IO such that $OO_X$ ≫ IO, and not $OO_Y$ ≫ IO, assign novel forms to lexical class *X*.

ii. (If there is some markedness constraint M such that $OO_X$ ≫ M, and not $OO_Y$ ≫ M, assign novel forms to lexical class *X*.)[6]

c. Indexed markedness constraints

Let $M_X$ and $M_Y$ be two indexed versions of the same markedness constraint, indexed to lexical class *X* and *Y* respectively. In a grammar with $M_X$ and $M_Y$:

i. If there is some IO-faithfulness constraint IO such that $M_X$ ≫ IO, and not $M_Y$ ≫ IO, assign novel forms to lexical class *X*.

ii. (If there is some OO-faithfulness constraint OO such that OO ≫ $M_X$, and not OO ≫ $M_Y$, assign novel forms to lexical class *X*.)[7]

We can now ask what each of the grammars in (15) will do when presented with a novel form. The grammar in (15a) has an indexed IO-faithfulness constraint, so that (16a) will apply. This grammar has the ranking IDENT[long]$_{IO}$-NA ≫ SWP ≫ IDENT[long]$_{IO}$-A. In accordance with (16a.i), novel words will therefore be assigned to class A. A learner who has acquired this grammar is expected to treat novel forms as alternating. The grammar in (15b) has indexed OO-faithfulness constraints, and (16b) applies. The ranking IDENT[long]$_{OO}$-NA ≫ IDENT[long]$_{IO}$ holds in this grammar, but not IDENT[long]$_{OO}$-A ≫ IDENT[long]$_{IO}$. In accordance with (16b.i), novel items will be assigned to NA, and will be treated as non-alternating.

As explained above, in the markedness grammar from (15c) it is not necessary that both IDENT[long]$_{OO}$ and IDENT[long]$_{IO}$ outrank SWP-NA. The tableau in (15c) therefore represents several different possible total rankings, all consistent with the data in (1). One of the possible total rankings represented by (15c) is SWP-A ≫ IDENT[long]$_{OO}$ ≫ SWP-NA ≫ IDENT[long]$_{IO}$. Clause (16c.ii) can only be satisfied by assigning a novel word to class NA. Clause (16c.i) can be

[6] The need for this clause is motivated in §3.3. It is included here for completeness.

[7] See previous footnote.

satisfied by either class assignment. Under this ranking, novel items will hence be assigned to class NA, and their plurals will not alternate.

Another total ranking consistent with (15c) is SWP-A ≫ IDENT[long]$_{IO}$ ≫ SWP-NA ≫ IDENT[long]$_{OO}$. No lexical class assignment of novel items is possible that will satisfy (16c.ii) – since both indexed markedness constraints outrank IDENT[long]$_{OO}$. Assigning novel items to class A does satisfy (16c.i) – since the ranking markedness ≫ IO-faithfulness does then apply to such forms. But assigning novel items to NA violates (16c.i). The best that can be achieved here is to assign novel items to class A, so that their plurals will then undergo vowel lengthening.

Although the grammars represented in (15) perform the same on words that were included in the learning data, they differ on how they treat novel forms. The indexed IO-faithfulness grammar extends the alternating pattern to novel words. Depending on the specific instantiation of the indexed markedness grammar, it extends either the alternating or the non-alternating pattern to novel items. The indexed OO-faithfulness grammar, however, always extends the non-alternating pattern. The participants in Experiment 1 consistently favored non-alternation over alternation for novel items. The OO-faithfulness grammar and some instantiations of the markedness grammar are therefore consistent with how these participants performed.

In the next section, I show that the indexed OO-faithfulness grammar is the grammar that is acquired when Prince and Tesar's (2004) BCD algorithm of is presented with data like that in (1). Applying this algorithm to the learning data therefore results in acquiring a grammar that will always be consistent with how the participants in Experiment 1 performed.

### 3.3 Learning the grammar

Tesar and Smolensky (2000) propose a learning algorithm, the Recursive Constraint Demotion algorithm (RCD), that learns an OT grammar by recursively submitting learning data to the current grammar, and changing the grammar every time it generates an error. However, the RCD does not incorporate the restrictiveness principle, and could result in the learner getting trapped in the superset grammar. Prince and Tesar (2004) augment the RCD by building a markedness ≫ IO-faithfulness bias into the algorithm. In the absence of evidence about the ranking between a markedness and an IO-faithfulness constraint, their algorithm ranks the markedness constraint higher. Because of this bias, they call their algorithm the "Biased Constraint Demotion" algorithm. Prince and Tesar do not consider grammars with OO-faithfulness constraints, and their algorithm therefore does not contain ranking biases for OO-faithfulness constraints. The default ranking OO-faithfulness ≫ IO-faithfulness motivated above can be incorporated into their algorithm in the same way that they incorporate the markedness ≫ IO-faithfulness bias. I give their algorithm in (17), with an additional step added – step ii in (17) ensures that OO-faithfulness constraints are ranked before IO-faithfulness and markedness constraints, all else being equal. Although I motivate the default ranking OO-faithfulness ≫ markedness only later in this section, I include it in the statement of the algorithm here. In the rest of this section, I show how application of (17) to data like those in (1) results in a grammar with indexed OO-faithfulness constraints, rather than any of the alternatives.

(17) Biased Constraint Demotion Algorithm: With markedness and an OO-faithfulness biases[8]

i. Identify constraints that prefer no Losers.
ii. If this set contains OO-faithfulness constraints, install them and return to i.
iii. If this set contains markedness constraints, install them and return to i.
iv. Else install IO-faithfulness constraints that prefer only Winners, and return to i.
v. Else install IO-faithfulness constraints that prefer no Losers, and return to i.

The BCD keeps a running index of data that have been encountered, called the "Support" by Prince and Tesar (2004). The Support is represented in a comparative tableau – i.e. each line contains a "Winner~Loser" pair, where Winner is the candidate that should be optimal, and Loser some competing candidate that wins under the current incorrect ranking. Cells are marked with "W" for constraints that prefer the Winner over the Loser, and with "L" for constraints that prefer the Loser. In (18), the Support that a Dutch learner will have formed after being exposed to one alternating and one non-alternating noun is represented.

(18) Initial Support[9]

| | Winner~Loser | SWP | IDENT[long]$_{IO}$ | IDENT[long]$_{OO}$ |
|---|---|---|---|---|
| Non-Alternating | kɑ́.tən~káː.tən | L | W | W |
| Alternating | xáː.tən~xɑ́.tən | W | L | L |

When faced with a Support like this, the BCD is stuck. The first step requires the collection of constraints that prefer no Losers. There are no such constraints here, and there is hence no ranking of just these constraints that are compatible with this Support. Pater (2005, to appear) proposes that this impasse is resolved by constraint cloning – creating lexically indexed versions of constraints. At this point, a decision has to be made about which constraint(s) to clone. One possibility is to clone all constraints. This would increase the size of the grammar considerably – minimally, it would double the number of constraints. It might be better if there was some way to determine which constraint(s) should be cloned to resolve the inconsistency in the most economic way. Pater (to appear) proposes the principle in (19).[10]

[8] The algorithm is stated here in a simplified version. Specifically on step iv, the algorithm will not install all IO-faithfulness constraints that prefer only Winners. Prince and Tesar (2004:257) propose that only the "smallest effective F-set" will be installed – that is, the smallest set of IO-faithfulness constraints that would enable a markedness constraint to become rankable on the next iteration of the algorithm. This is never an issue in the learning situations discussed here, and I abstract away from this complication.

[9] The singular forms are not included in the Support. The faithful candidate for a singular input violates no constraint and is optimal under all possible rankings of the constraints. Since the Support only includes items on which an error is generated, the singulars are therefore not included. It is clear that /kɑt/ → [kɑ́t] does not violate SWP or IDENT[long]$_{IO}$. I assume that the singular is the Base of the OO-faithfulness relation and hence that it cannot itself violate the OO-faithfulness constraint IDENT[long]$_{OO}$. See Kenstowicz (1997) for evidence that the morphosyntactically unmarked member of an inflectional paradigm (e.g. the singular) serves as the OO-base for the paradigm. See McCarthy (2005) for an alternative view.

[10] Clause (19b) is slightly different from Pater's statement, and in agreement with Becker's (2009:10, 167) interpretation of constraint cloning. Pater envisions cloning as addition of indexed constraints to the existing constraint set – if constraint *C* is cloned, the result is {*C*, $C_I$}. I, similarly to Becker (2009), assume that cloning

(19) Which constraint(s) to clone

a. Clone a constraint that prefers only Winners in all instances of some morpheme.

b. Index all morphemes that are evaluated similarly by the uncloned constraint to the same clone of the constraint.

SWP prefers only winners for the morpheme /xɑt/, while IDENT[long]$_{IO}$ and IDENT[long]$_{OO}$ prefer only winners for /kɑt/. All three constraints satisfy (19a), and all three are therefore cloned. Clause (19b) is responsible for creating lexical classes – morphemes that are treated the same by the original unindexed constraint are indexed to the same clone of the constraint. The Support after cloning is represented in (20).

(20) Support after cloning

| | Winner~Loser | SWP-A | SWP-NA | ID[long]$_{IO}$-A | ID[long]$_{IO}$-NA | ID[long]$_{OO}$-A | ID[long]$_{OO}$-NA |
|---|---|---|---|---|---|---|---|
| Non-Alternating | NA: kɑ́.tən~kɑ́ː.tən | | L | | W | | W |
| Alternating | A: xɑ́ː.tən~xɑ́.tən | W | | L | | L | |

In this post-cloning Support, the ranking inconsistency has been resolved. The first step of learning is to collect constraints that prefer no Losers: {IDENT[long]$_{IO}$-NA, IDENT[long]$_{OO}$-NA, SWP-A}. On step ii, the algorithm chooses to install the OO-faithfulness constraint. The Support is trimmed down by removing the lines where a W-preferring constraint has been ranked, and the columns corresponding to constraints that have been ranked. The top line of (20) is hence removed – since a W-preferring constraint from this line has been ranked. The final column is also removed, since the constraint represented by this column has been ranked. The grammar at this point and the trimmed down Support are shown in (21).

(21) Grammar: IDENT[long]$_{OO}$-NA ≫ …

Trimmed down Support

| | Winner~Loser | SWP-A | SWP-NA | ID[long]$_{IO}$-A | ID[long]$_{IO}$-NA | ID[long]$_{OO}$-A |
|---|---|---|---|---|---|---|
| Alternating | A: xɑ́ː.tən~xɑ́.tən | W | | L | | L |

The BCD first again collects the constraints that prefer no Losers: {SWP-A, SWP-NA, IDENT[long]$_{IO}$-NA}. Since this set does not contain an OO-faithfulness constraint, step ii is skipped. On step iii, the markedness constraints are installed in the current stratum. The Support is them trimmed down by removing the columns corresponding to the markedness constraints,

splits a constraint in two – if $C$ is cloned, the result is $\{C_1, C_2\}$. This difference influences how forms are assigned to lexical classes.

and the last line – since a W-preferring constraint from this line has been ranked. The Support now contains no more lines, but still have some constraints left. The new grammar and trimmed down Support are given in (22).

(22) Grammar thus far: IDENT[long]$_{OO}$-NA ≫ {SWP-A, SWP-NA} ≫ …

Trimmed down Support

| | Winner~Loser | ID[long]$_{IO}$-A | ID[long]$_{IO}$-NA | ID[long]$_{OO}$-A |
|---|---|---|---|---|

Learning could stop here. There are no lines left, and hence no more learning data on which an error could be generated. However, we can also apply the algorithm until all constraints have been ranked. This will result in unnecessary, but harmless, rankings. Should the algorithm apply again to the Support in (22), all remaining constraints will be selected in step i. Step ii favors ranking OO-faithfulness constraints, so that IDENT[long]$_{OO}$-A is installed, and the column corresponding to it is removed. The newly trimmed down Support will look exactly like that in (22), except for the absence of the final column. Applying the algorithm again will result first in the selection of both remaining constraints. Both of these are IO-faithfulness constraints, and steps ii and iii of the algorithm are not applicable. In step iv, both IDENT[long]$_{IO}$-A and IDENT[long]$_{IO}$-NA are installed. All remaining columns are removed, so that the Support is completely empty and no further learning is possible.

The grammar that has been learned is represented in (23a). This grammar contains two indexed versions of SWP ranked in the same stratum. Since these constraints are in the same stratum, there is no need for two separate constraints, and they can be merged into a single unindexed constraint that will evaluate all morphemes irrespective of their indexation. The same holds for the two indexed versions of IDENT[long]$_{IO}$. The final grammar in which the clones of these constraints have been merged is given in (23b). This merging is done for the sake of having a final grammar with fewer constraints. These two grammars will generate identical outputs. The grammar proposed for Dutch in (15b) is given in (23c) for comparison.

(23) a. Grammar at the end of learning

ID[long]$_{OO}$-NA ≫ {SWP-A, SWP-NA} ≫ ID[long]$_{OO}$-A ≫ {ID[long]$_{IO}$-A, ID[long]$_{IO}$-NA}

b. Final simplified grammar

ID[long]$_{OO}$-NA ≫ SWP ≫ ID[long]$_{OO}$-A ≫ ID[long]$_{IO}$

c. Dutch grammar from (15b)

ID[long]$_{OO}$-NA ≫ SWP ≫ {ID[long]$_{OO}$-A, ID[long]$_{IO}$}

The grammar in (23b) is identical to the OO-blocking grammar in (15b) with one exception – in (15b), IDENT[long]$_{OO}$-A and IDENT[long]$_{IO}$-NA are unranked relative to each other. As shown just above, the ranking between these constraints in (23b) is the result of continuing to learn until the Support is empty, even after the grammar has seized making errors. In spite of the additional ranking in (23b), these grammars will evaluate all known words the same. Looking back at (15b) will confirm that the competition is decided by either IDENT[long]$_{OO}$-NA or SWP for known words, so that the ranking between constraints below SWP cannot impact output selection.

The situation is not as simple for the evaluation of novel words. If the grammar in (23b) is presented with a novel word like /nɔl/, this word needs to be assigned to one of the lexical classes. Both indexed OO-faithfulness constraints outrank the IO-faithfulness constraint. Whether /nɔl/ is assigned to NA or N, it will be subject to the default restrictive OO-faithfulness ≫ IO-faithfulness ranking. However, these two options obviously make different predictions about how novel items will be treated. If /nɔl/ is assigned to NA, its plural will be non-alternating, i.e. [nɔ́.lən], while it will be alternating, [nɔ́ː.lən], if /nɔl/ is assigned to A. In (16), I stated the principles that determine the class assignment of novel forms. Clause (16b.ii) deals with markedness and indexed OO-faithfulness constraints: "If there is some markedness constraint M such that $OO_X \gg$ M, and not $OO_Y \gg$ M, assign novel forms to lexical class *X*." In (16), this clause was placed in parentheses, since the necessity for this clause has not yet been motivated at that point. The need for this clause is now clear. The participants in Experiment 1 were more likely to treat novel words as non-alternating. Assuming that they acquired a grammar similar to that in (23b), they were hence more likely to evaluate novel items with IDENT[long]$_{OO}$-NA than IDENT[long]$_{OO}$-A. (13b.ii) enforces this preference.

Principle (13) on the lexical class assignment of novel items is based on the default rankings – all else being equal, a novel item is assigned to a lexical class that results in it being subject to the default ranking. The default ranking Markedness ≫ IO-faithfulness is enforced by (16a.i) and (16.c.i). The default OO-faithfulness ≫ IO-faithfulness is enforced by (16a.ii) and (16b.i). The two remaining clauses, (16b.ii) and (16c.ii), deal with the ranking between OO-faithfulness and markedness constraints, and both of them result in the ranking OO-faithfulness ≫ markedness applying to novel items. If lexical class assignment of novel items is determined by default rankings, then OO-faithfulness ≫ markedness should be afforded the same default status as the other two default rankings. I therefore propose to add this as a third default ranking, so that the set of default rankings from (13) should be replaced by that in (24).

The reason for installing OO-faithfulness constraints before markedness constraints in the learning algorithm is also now evident. Default rankings are enforced by a ranking preference, and the default OO-faithfulness ≫ markedness ranking is enforced by step ii.

(24) Default rankings
   a. OO-faithfulness ≫IO-faithfulness
   b. Markedness ≫ IO-faithfulness
   c. OO-faithfulness ≫ Markedness

In this section, I have shown that with minimal augmentation, an existing learning algorithm (the BCD) automatically learns a grammar that is conservative in how it treats novel words. This grammar is also consistent with the response pattern observed in Experiment 1. Information from different sources converge here: learnability considerations favor conservative grammars, participants in an artificial language learning experiment respond as if they have acquired a conservative grammar, and applying an independently motivated learning algorithm also results in a conservative grammar.

## 4. CONCLUDING DISCUSSION: ON THE TREATMENT OF NOVEL ITEMS

The learning algorithm proposed in §3.3 will acquire exactly the same grammar when presented with learning data like those in Experiment 1 and those in Condition 2 of Experiment 3. The default ranking biases will be enforced no matter what the proportion of alternating to non-alternating plurals is in the learning data. The difference seen in the responses to novel items in these two experiments can therefore not follow from the grammar that is acquired. In this section, I first show that the algorithm will indeed acquire the same grammar irrespective of whether the majority of forms alternate or not. I then propose that the assignment of novel items to lexical classes is determined not solely based on grammatical considerations, but that the frequency structure of the lexicon also plays a role.

The learning data that were used in §3.3 are structurally identical to those to which participants in Experiment 1 were exposed (the data contained equal evidence for alternation and non-alternation). Given learning data like these, the algorithm settles on the OO-blocking grammar. However, even when the learning data contain more examples of alternating than non-alternating nouns, the algorithm will still acquire the same grammar. A representative sample of the learning data from Condition 2 in Experiment 3 is given in (25). These data contain three times as many examples of alternation as of non-alternation. As in §3.3, the Support constructed from these data would be inconsistent with any ranking, which would lead to constraint cloning. The post-cloning Support is represented in (26).

(25) Learning data for Condition 2, Experiment 3

| | *Singular* | *Plural* |
|---|---|---|
| *Non-Alternating:* | [məl.tɛ́p] | [məl.tɛ́.pən] |
| *Alternating:* | [təl.sɛ́p] | [təl.sé.pən] |
| | [kən.dɪ́p] | [kən.dí.pən] |
| | [dəf.sɔ́k] | [dəf.só.kən] |

(26) Support after cloning

| | Winner~Loser | SWP-A | SWP-NA | Id[long]$_{IO}$-A | Id[long]$_{IO}$-NA | Id[long]$_{OO}$-A | Id[long]$_{OO}$-NA |
|---|---|---|---|---|---|---|---|
| Non-Alternating | məl.tɛ́.pən~məl.té.pən | | L | | W | | W |
| Alternating | təl.sé.pən~təl.sɛ́.pən | W | | L | | L | |
| | kən.dí.pən~kən.dɪ́.pən | W | | L | | L | |
| | dəf.só.kən~dəf.sɔ́.kən | W | | L | | L | |

When the algorithm in (17) is applied to this Support, it first collects the constraints that prefer no Losers: {Ident[long]$_{OO}$-NA, Ident[long]$_{IO}$-NA, SWP-A}. On ii, the OO-faithfulness constraint Ident[long]$_{OO}$-NA is installed, and the first line and last column is trimmed from the Support. The trimmed down Support after this first learning cycle is the unshaded portion of (26). When (17) is applied again, the constraints {SWP-A, SWP-NA, Ident[long]$_{IO}$-NA} is selected on step i. Step ii does not apply, and on step iii, the markedness constraints are installed, and the columns corresponding to these constraints are trimmed from the Support. The three remaining lines are also trimmed, since a W-preferring constraint (SWP-A) for all three lines have been ranked. At this point, the grammar learned is Ident[long]$_{OO}$-NA $\gg$ {SWP-A, SWP-NA}, and the remaining Support contains only the constraints Ident[long]$_{IO}$-A, Ident[long]$_{IO}$-

NA and IDENT[long]$_{OO}$-A. This is exactly the scenario in (22) above, and it is hence clear that the same final grammar as in (23) will result here.

Although the grammars will look the same in these two conditions, the lexicons will differ. In (27), I show the final grammar in both conditions, and the two different lexicons.

(27) Final grammar and lexicons

a. Final grammar in both conditions
IDENT[long]$_{OO}$-NA ≫ SWP ≫ IDENT[long]$_{OO}$-A ≫ IDENT[long]$_{IO}$

b. Lexicon in §3.3
NA = {kɑt}, A = {xɑt}

c. Lexicon in this section
NA = {məltɛp}, A = {təlsɛp, kəndɪp, dəfsɔk}

The grammar and lexicon from section 3.3 are formally identical to that assumed in Experiment 1, and the grammar and lexicon learned just above are formally the same as that assumed in Condition 2 of Experiment 3. We must therefore be able to explain the results of these two experiments by using the grammar and lexicons from (27). The participants in these two experiments responded differently to novel items during testing. Since the same grammar is learned in both conditions, the difference in response pattern cannot originate directly in the grammar. It has to come from the structure of the lexicons.

In Experiment 1, the lexical classes contained the same number of entries. Participants could therefore not rely on the structure of the lexicon for information on which of the patterns is the preferred pattern. In the absence of lexical information, participants have to fall back on the default, as determined by the default ranking OO-faithfulness ≫ markedness ≫ IO-faithfulness. Novel items are therefore assigned to class NA so that this default ranking applies to them. However, in Condition 2 of Experiment 3, participants can consult their lexicons to determine which of the patterns is the preferred pattern. In this experiment, the class of alternating nouns is three times as large as the non-alternating class. The participants in this experiment were more likely to assign novel items to the alternating class, showing that in addition to the default rankings, the structure of the lexicon also influences how novel items are treated.

Novel items that are similar to many existing words are rated as more wordlike than novel items that are similar to fewer existing words (Bailey & Hahn 2001). We also know that humans can calculate frequency statistics generalized over unrelated words rapidly and based on only minimal exposure. Saffran, Newport and Aslin (1996), for instance, conducted an experiment in which they exposed listeners briefly to a speech stream consisting of nonsense syllables. The syllable sequence was designed so that some syllables appear together more often than others. Listeners were then tested on how they segment the syllable stream into word-sized chunks. They found that listeners rely on the frequency structure of the syllable sequence – syllables that appear together frequently are chunked together into one word, while word breaks are inserted between syllables that do not appear together often. It is not relevant to the current paper that language users use abstract frequency patterns to do word segmentation. What is relevant is that listeners can calculate abstract frequency statistics very rapidly and based on limited exposure. It is therefore reasonable to assume that the participants in Condition 2 of Experiment 3 were able to detect the frequency difference between alternating and non-alternating nouns.

However, the frequency structure of the lexicon is not the only determiner of how novel items are treated. When the lexicon contains an equal number of alternating and non-alternating examples so that the lexicon itself cannot cause a response bias, the influence of grammar on the treatment of novel items becomes evident. The default rankings, motivated on learnability grounds, then show their influence on how novel items are treated. For more evidence that nonce words with equal lexical frequency statistics are treated differently if they differ on grammatical grounds, see Berent *et al*. (2007), Coetzee (to appear, 2008), Moreton (2002), etc.

The results of both Experiment 1 and Experiment 3 also show that the influence of grammar and of the frequency structure of the lexicon are not exclusive of each other. In Experiment 1, participants did not categorically prefer non-alternation for novel items. Such a categorical preference would be expected if they relied solely on the drive towards grammatical restrictiveness in the treatment of novel items. The fact that they did sometimes (though less often) select the alternating plural for novel items could be interpreted as the result of the frequency structure of the lexicon. The lexicon to which they were exposed did contain some evidence for alternation.

Experiment 3 gives even more evidence that grammatical considerations and frequency patterns in the lexicon co-determine how novel items are treated. In Condition 2, the lexicon to which participants were exposed contained 75% evidence for alternation and only 25% for non-alternation. Had participants relied solely on the frequency structure of the lexicon, they should have preferred alternation for novel items 75% of the time. Had the participants ignored the frequency structure of the lexicon, they should have performed the same as the participants in Experiment 1. Neither of these two options was observed. Participants in Condition 2 of Experiment 3 preferred alternation 52.5% of the time, and hence less than the representation of alternation in the learning data. But they also preferred alternation more than the participants in Experiment 1 who opted for alternation only 35.8% of the time. The performance of the participants in Condition 2 of Experiment 3 shows evidence of being influenced by grammatical considerations and by the frequency structure of the lexicon.

The model presented above does not make any predictions about how these two influences interact: What determines how much influence the frequency structure of the lexicon has and how much influence grammatical considerations have? More research is needed before this question can be resolved.

## References

Anttila, Arto (2002). Morphologically conditioned phonological alternations. *NLLT* **20.** 1–42.

Baayen, R. Harold, Ricahrd Piepenbrock, & Leon Gulikers. (1995). *The CELEX Lexical Database (CD-ROM)*. Philadelphia, PA: Linguistic Data Consortium.

Bailey, Todd M. & Ulrike Hahn (2001). Determinants of wordlikeness: phonotactics or lexical neighborhoods. *Journal of Memory and Language* **44.** 568–591.

Becker, Michael. (2009). *Phonological Trends in the Lexicon: The Role of Constraints*. PhD Dissertation, University of Massachusetts, Amherst.

Berent, Iris, Donca Steriade, Tracy Lennertz, & Vered Vaknin (2007). What we know about what we have never heard: Evidence from perceptual illusions. *Cognition* **104.** 591–630.

Berwick, Robert (1985). *The Acquisition of Syntactic Knowledge*. Cambridge, MA: MIT Press.

Boersma, Paul & Weenink, David (2007). Praat: doing phonetics by computer (Version 5.0.02) [www.praat.org.]

Buzio, Luigi (1996). Surface constraints versus underlying representation. In Jacques Durand & Bernard Laks (eds.) *Current Trends in Phonology: Models and Methods*. Manchester: European Studies Research Institute, University of Salford,

Carpenter, Angela (2006). *Acquisition of a Natural Versus an Unnatural Stress System*. PhD Dissertation, University of Massachusetts, Amherst.

Coetzee, Andries W. (to appear). Grammar is both categorical and gradient. In Steven Parker (ed.) *Phonological Argumentation*. London: Equinox Publishers.

Coetzee, Andries W. (2008). Grammaticality and ungrammaticality in phonology. *Language* **84**. 218-257.

Féry, Caroline (2003). Final devoicing and the stratification of the lexicon in German. In Jeroen van de Weijer, Vincent Van Heuven, & Harry Van Der Hulst (eds.) *The Phonological Spectrum. Volume I: Segmental Structure*. Amsterdam: John Benjamins. 145-170.

Gnanadesikan, Amalia (2004). Markedness and faithfulness constraints in child phonology. In René Kager, Joe Pater & Wim Zonneveld (eds.) *Constraints in Phonological Acquisition*. Cambridge: Cambridge University Press.

Hayes, Bruce (2004). Phonological acquisition in Optimality Theory: the early stages. In René Kager, Wim Zonneveld, & Joe Pater (eds.) *Fixing Priorities: Constraints in Phonological Acquisition*. Cambridge: Cambridge University Press.

Inkelas, Sharon, Orhan Orgun & Cheryl Zoll (1997). The implications of lexical exceptions for the nature of grammar. In Iggy Roca (ed.) *Derivations and Constraints in Phonology*. New York: Oxford University Press. 393–418.

Itô, Junko & Armin Mester (1999). The structure of the phonological lexicon. In Natsuko Tsujimura (ed.) *The Handbook of Japanese Linguistics*. Oxford: Blackwell. 62–100.

Itô, Junko & Armin Mester (2002). One phonology or many? Issues in stratal faithfulness theory. In The Phonological Society of Japan (ed.) *Phonological Studies 5*. Tokyo: Kaitakusha. 121–126.

Jakobson, Roman (1968). *Child Language Aphasia and Phonological Universals: Aphasia and Phonological Universals*. The Hague: Mouton.

Kager, René (to appear). Lexical irregularity and the typology of contrast. In Kristin Hanson & Sharon Inkelas (eds.) *The Nature of the Word. Studies in Honor of Paul Kiparsky*. Cambridge, MA: MIT Press.

Kenstowicz, Michael (1997). Base identity and uniform exponence: alternative to cyclicity. In Jacques Durand and Bernard Laks (eds). *Current Trends in Phonology: Models and Methods*. Salford: University of Salford. 363–394.

McCarthy, John J. (1998). Morpheme structure constraints and paradigm occultation. In M. Catherine Gruber, Derrick Higgins, Kenneth Olson, & Tamra Wysocki (eds.) *Proceedings of the Chicago Linguistic Society 32, Vol. II: The Panels.* Chicago: Chicago Linguistic Society. 123–150.

McCarthy, John J. (2005). Optimal paradigms. In Laura Downing, Tracy Alan Hall, & Renate Raffelsiefen (eds.) *Paradigms in Phonological Theory*. New York: Oxford University Press.

Moreton, Elliot (2002). Structural constraints in the perception of English stop-sonorant clusters. *Cognition* **84.** 55–71.

Ota, Mitsuhiko (2004). The learnability of a stratified lexicon. *Journal of Japanese Linguistics* **20.** 19–40.

Pater, Joe (2000). Non-uniformity in English secondary stress: the role of ranked and lexically specific constraints. *Phonology* **17.** 237–274.

Pater, Joe (2005). Learning a stratified grammar. In Alejna Brugos, Manuella R. Clark-Cotton, & Seungwan Ha (eds.) *Proceedings of the 29th Boston University Conference on Language Development*. Somerville: Cascadilla Press. 482–492.

Pater, Joe (to appear). The locus of exceptionality: morpheme-specific phonology as constraint indexation. In Stephen Parker (ed.) *Phonological Argumentation*. London: Equinox.

Pater, Joe & Anne-Michelle Tessier (2006). L1 phonotactic knowledge and the L2 acquisition of alternations. In Roumyana Slabakova, Silvina A. Montrul, and Philippe Prévost (eds.) *Inquiries in Linguistic Development: Studies in Honor of Lydia White*. Amsterdam: Benjamins. 115-131.

Prasada, Sandeep & Steven Pinker (1993). Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes* **8.** 1–56.

Prince, Alan (1990). Quantitative consequences of rhythmic organization. *CLS* **26:2.** 355–398.

Prince, Alan & Bruce Tesar (2004). Learning phonotactic distributions. In René Kager, Joe Pater, & Wim Zonneveld (eds.) *Fixing Priorities: Constraints in Phonological Acquisition*. Cambridge: Cambridge University Press. 245–291.

Saffran, Jenny R., Elissa L. Newport, & Richard N. Aslin (1996). Word segmentation: The role of distributional cues. *Journal of Memory and Language* **35.** 606–621.

Smolensky, Paul. (1996). *The Initial State and "Richness of the Base" in Optimality Theory.* Ms., Department of Cognitive Science, Johns Hopkins University. Technical Report JHU-CogSci-96-4. [Rutgers Optimality Archive #154.]

Tesar, Bruce & Paul Smolensky (2000). *Learnability in Optimality Theory*. Cambridge, MA: MIT Press.

Tessier, Anne-Michelle (2007). *Biases and stages in Phonological Acquisition*. Ph.D. Dissertation, University of Massachusetts.

Walter, Mary Ann (2006). Grammatical gender via lexical statistics: the case of Arabic-to-Spanish loanwords. In F.F. Hsieh & Michael Kenstowicz (eds.) *Studies in Loanword Phonology, MIT Working Papers in Linguistics 52*.

Wilson, Colin (2006). Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science* **30**. 945–982.

## APPENDIX A: EXPERIMENTS 1 AND 3

| *Vowel* | *Singular* | *Plural* | |
|---|---|---|---|
| | | *Non-Alternating* | *Alternating* |
| /ɛ/ | [təl.sɛ́p] | [təl.sɛ́.pən] | [təl.sé.pən] |
| | [məl.tɛ́p] | [məl.tɛ́.pən] | [məl.té.pən] |
| | [bəm.pɛ́t] | [bəm.pɛ́.tən] | [bəm.pé.tən] |
| | [mək.tɛ́s] | [mək.tɛ́.sən] | [mək.té.sən] |
| | [fəs.mɛ́k] | [fəs.mɛ́.kən] | [fəs.mé.kən] |
| | [kəm.pɛ́s] | [kəm.pɛ́.sən] | [kəm.pé.sən] |
| | [kən.dɛ́f] | [kən.dɛ́.fən] | [kən.dé.fən] |
| | [vən.sɛ́p] | [vən.sɛ́.pən] | [vən.sé.pən] |
| /ɪ/ | [səm.fɪ́t] | [səm.fɪ́.tən] | [səm.fí.tən] |
| | [ləs.kɪ́f] | [ləs.kɪ́.fən] | [ləs.kí.fən] |
| | [fən.sɪ́k] | [fən.sɪ́.kən] | [fən.sí.kən] |
| | [fəm.bɪ́s] | [fəm.bɪ́.sən] | [fəm.bí.sən] |
| | [kən.dɪ́p] | [kən.dɪ́.pən] | [kən.dí.pən] |
| | [məs.pɪ́t] | [məs.pɪ́.tən] | [məs.pí.tən] |
| | [bəl.tɪ́k] | [bəl.tɪ́.kən] | [bəl.tí.kən] |
| | [təl.bɪ́f] | [təl.bɪ́.fən] | [təl.bí.fən] |
| /ɔ/ | [ləm.pɔ́s] | [ləm.pɔ́.sən] | [ləm.pó.sən] |
| | [məf.pɔ́t] | [məf.pɔ́.tən] | [məfó.tən] |
| | [təl.dɔ́n] | [təl.dɔ́.nən] | [təl.dó.nən] |
| | [vəl.dɔ́f] | [vəl.dɔ́.fən] | [vəl.dó.fən] |
| | [dəf.sɔ́k] | [dəf.sɔ́.kən] | [dəf.só.kən] |
| | [kən.tɔ́p] | [kən.tɔ́.pən] | [kən.tó.pən] |
| | [fəl.sɔ́t] | [fəl.sɔ́.tən] | [fəl.só.tən] |
| | [səm.pɔ́k] | [səm.pɔ́.kən] | [səm.pó.kən] |
| /ʊ/ | [zəl.tʊ́k] | [zəl.tʊ́.kən] | [zəl.tú.kən] |
| | [zəŋ.kʊ́p] | [zəŋ.kʊ́.pən] | [zəŋ.kú.pən] |
| | [fək.sʊ́t] | [fək.sʊ́.tən] | [fək.sú.tən] |
| | [bəs.fʊ́t] | [bəs.fʊ́.tən] | [bəs.fú.tən] |
| | [pəm.bʊ́k] | [pəm.bʊ́.kən] | [pəm.bú.kən] |
| | [səŋ.gʊ́f] | [səŋ.gʊ́.fən] | [səŋ.gú.fən] |
| | [bəl.fʊ́s] | [bəl.fʊ́.sən] | [bəl.fú.sən] |
| | [ges.tʊ́p] | [ges.tʊ́.pən] | [ges.tú.pən] |

## APPENDIX B: ADDITIONAL TOKENS FOR EXPERIMENT 2

| *Vowel* | *Singular* | *Plural* | |
|---|---|---|---|
| | | *Non-Alternating* | *Alternating* |
| /ɛ/ | [bəm.pɛ́ts] | [bəm.pɛ́t.sən] | [bəm.pét.sən] |
| | [məl.tɛ́sp] | [məl.tɛ́s.pən] | [məl.tés.pən] |
| | [kən.dɛ́ft] | [kən.dɛ́f.tən] | [kən.déf.tən] |
| | [vən.sɛ́mp] | [vən.sɛ́m.pən] | [vən.sém.pən] |
| /ɪ/ | [kən.dímp] | [kən.dím.pən] | [kən.dím.pən] |
| | [fən.síŋk] | [fən.síŋ.kən] | [fən.síŋ.kən] |
| | [məs.pílt] | [məs.píl.tən] | [məs.píl.tən] |
| | [səm.físt] | [səm.fís.tən] | [səm.fís.tən] |
| /ɔ/ | [kən.tɔ́mp] | [kən.tɔ́m.pən] | [kən.tóm.pən] |
| | [ləm.pɔ́st] | [ləm.pɔ́s.tən] | [ləm.pós.tən] |
| | [tən.kɔ́ps] | [tən.kɔ́p.sən] | [tən.kóp.sən] |
| | [vəl.dɔ́mf] | [vəl.dɔ́m.fən] | [vəl.dóm.fən] |
| /ʊ/ | [bəs.fʊ́nt] | [bəs.fʊ́n.tən] | [bəs.fún.tən] |
| | [pəm.bʊ́sk] | [pəm.bʊ́s.kən] | [pəm.bús.kən] |
| | [fək.sʊ́lt] | [fək.sʊ́l.tən] | [fək.súl.tən] |
| | [gəs.tʊ́mp] | [gəs.tʊ́m.pən] | [gəs.túm.pən] |