CHAPTER V

TESTING FOR THE HIGH-RANKING OO-FAITH BIAS

1. Introduction to the chapter

This chapter returns to the ranking bias proposed in Hayes (2004) and McCarthy (1998): the preference for high-ranking paradigm uniformity constraints, formalized here using Output-Output Faithfulness (OO-Faith; Benua, 2000). I will expand on the arguments discussed in chapter 2 that this bias provides a mechanism for preventing the acquisition of superset grammars, and also suggest following Hayes that it finds independent support in some children's innovative application of paradigm uniformity in the literature (§1.3).

One interesting aspect of this approach are its predictions (exemplified in §2.2 - §2.3) about the behavior of the phonological grammar at the point of *morphological* acquisition: that is, the point at which derived words are first decomposed into multiple morphemes and phonological patterns become attributable in principle to the demands of OO-Faith.

In sections 3 through 5 of this chapter, I report on an experiment which tested these predictions, using an artificial language 'wug test' (Berko, 1958) to test 4-year-old childrens' production of marked consonant clusters in two different morphological environments. The results (§5) support the claim that children prefer to repair clusters in ways that satisfy OO-Faith at the expense of other Markedness and Faithfulness pressures, and I provide some analyses (§6) of individual participant's productions and the associated rankings that match theoretical predictions.

Tessier, Anne-Michelle (2007). *Biases and Stages in Phonological Acquisition.*
Ph.D. dissertation, UMass Amherst

In section 7, I step back from the experimental details to consider the theoretical implications of these findings, including their positive contribution to the artificial learning literature, and some independent support for the OO-faith bias. In section 8, I discuss a possible experimental confound and its entailed questions for future research.

2. The OO-faith ranking bias and phonotactic learning

2.1 The role of OO-faith in enforcing restrictiveness

Chapter 2 (§3.2.3) made the point from McCarthy (1998) that a high-ranking OO-faith bias will make sure that paradigms are kept uniform and non-alternating unless evidence points to the contrary. Hayes (1999/2004) makes the same restrictiveness point for OO-faithful allophony, in languages where the normal distribution of some allophone is overriden just to keep a morphological paradigm uniform. The famous example he discusses is the interaction of flapping and Canadian raising (CR) in some dialects of English, already introduced in chapter 2 §3.2. To recall the facts: CR is purely allophonic in monomorphemic words: raised [ʌɪ] appears before voiceless obstruents, as in 'write' [ɹʌit], while [aɪ] appears elsewhere as in 'ride' [ɹaɪd]. However, derived forms with a base vowel [ʌɪ] exceptionally retain their raised quality even before a voiced flap, as in 'writer' [ɹʌiɾəɹ], *[ɹaiɾəɹ]. In the OO-faith analysis summarized in chapter 2, 'writer' contains a raised diphthong because it is faithful to the vowel quality of its morphological base 'write', whose raised vowel is allophonically-conditioned.

Hayes' learning argument is this: at the point that a learner encounters the word 'writer', how does he or she account for the presence of marked [ʌi] before a voiced flap? The child who does not yet know that 'writer' is derived from 'write' cannot explain ‘writer’s raised vowel with OO-faithfulness – instead, the error below (using the constraints from chapter 1) simply suggests that the [ʌiɾ] sequence is licit:

1) *The morphologically-naïve error caused by ‘writer’*

/ɹʌiɾəɹ/	Ident[lo]–OO	*ait	*ʌi	Ident[lo]–IO
(a) ☞ ɹaiɾəɹ				*!
(b) ɹʌiɾəɹ			*	

This error tells the BCD learner that IO faithfulness to any properties of raised vowels (here, Ident-[lo]) must rank above the general markedness constraint that disprefers raised vowels in the pre-flap context (see 2ii below).

2) *Initial Support during phonotactic learning, with no morphological relations*

input	winner ~ loser	Ident[lo]–OO	*ait	*ʌi	Ident[lo]–IO
(i) /ɹʌit/	[ɹʌit] ~ [ɹait]	e	W	L	W
(ii)/ɹʌiɾəɹ/	[ɹʌiɾəɹ] ~ [ɹaiɾəɹ]	e	e	L	W

Since OO-faith to vowel quality does not prefer any losers, BCD can rank it at the top of its ranking – however, it also does not prefer any winners, so it doesn’t resolve either of our errors. And while contextual markedness *does* explain the raised vowel in roots like ‘write’ in (2i), the ERC row in (2ii) can only be explained by IO-faithfulness, and thus BCD comes up with the ranking in 3):

3) *The resulting ranking that BCD learns: a superset grammar*
Ident[lo]-OO >> *ait >> **Ident[lo]-IO** >> >> ***ʌi**

As the bold emphasizes, this ranking is a superset grammar because it preserves input raised diphthongs in *all* contexts.

Hayes points out using data about the empirical timeline of acquisition that children do not know enough morphology early enough to avoid the superset trap illustrated in 2) and 3). And he therefore suggests that an OO-Faith bias offers part of a solution to this learning trap, if it is construed as a persistent bias that can return the grammar to a more restrictive state once morphological learning has occurred. When the learner realizes that 'writer' includes the base 'write', OO-faith constraints now prefer the raised vowel in the derived form. Imported into the present framework, this realization will mean that the Support’s ERC row resulting from 1) will be updated to look like 4):

4) *Revised Support, post-morphological learning (bases underlined)*

input	winner ~ loser	Ident[lo]–OO	*ait	*ʌi	Ident[lo]–IO
(i) /ɹʌit/	[ɹʌit] ~ [ɹait]	e	W	L	W
(ii) /ɹʌit/ + əɹ/	[ɹʌiɾəɹ] ~ [ɹaiɾəɹ]	**W**	e	L	W

Giving this revised Support to my BCD learner gets us the right ranking. In its first stratum, the Main Routine ranks OO-faith and resolves the error in (4ii); in the next stratum it ranks *ait and so resolves (4i). With all errors resolved, the algorithm ranks the rest of its constraints according to the M >> F bias, and the right ranking has been found:

5) *The ranking that BCD learns from 4): the correct grammar*
Ident[lo]-OO >> *ait >> ***ʌi** >> **Ident[lo]-IO**

I will return to the acquisition timeline of phonotactics vs. morphology, and the consequences for the nature of our ranking biases, in section 7.

Now that we have reviewed the end-state argument for a BCD-style ranking bias for OO-faith: what does this bias predict about stages of acquisition in the present system? The rest of this chapter considers this question in some detail.

2.2. Predictions for stages of acquisition

It was shown in (1) through (3) above that during pre-morphological learning, the OO-faith ranking bias causes BCD to continually re-install OO-faith constraints at the top of the hierarchy, even though they have no effect in resolving errors or driving mappings. However – once some morphological bases have been learned, these high-ranking OO-faith constraints will suddenly kick in and begin to assign violations. At this point, they should therefore begin to drive a new kind of error: enforcing paradigm uniformity on paradigms that are *not* OO-faithful in the target.

The following sections illustrate this, using an example that foreshadows the experimental test of the prediction to come.

2.2.1 The target: an OO-unfaithful language

Consider a language where a particular Markedness constraint is freely violated; to use an example relevant to the experiment to come, imagine this constraint is Agree(Voice), (Lombardi, 1996, 1999), which requires obstruent clusters to agree for voicing. In our toy grammar, the faithfulness constraints that conflict with Agree[voice] are:

6) Id[vce]-IO — Output segments must match their input correspondents for voicing

Id[vce]-OO — Output segments *in a derived form* must match their *base* correspondents for voicing[1]

Id[vce]-Ons-IO — Output segments *which are syllabified as onsets* must match their input correspondents for voicing

Before processing any errors, our BCD learner ranks these constraints as in 7):

7) Id[vce]-OO >> Agree[voice] >> Id[vce]-Ons-IO >> Id[vce]-IO

To see the effects of OO-Faith at different stages, we will consider the optimal outputs for two lexical items with different morphology:

8) *The toy lexicon*
(a) /zɪtʃdɪn/ — simple word
(b) /wʌtʃ + dəl/ — derived word – morphological base /wʌtʃ/

2.2.2 OO-faith kicks in at the initial state

How would the initial stage grammar treat these two words is this base ‘wutch’ had already been identified? Ignoring the implausibility of this order of acquisition, let us describe the grammars effect on our two-word lexicon in (8). With the ranking in (7), voicing mismatches would be repaired, but repaired differently depending on whether the cluster spans a base/affix boundary (in /wʌtʃ + dəl /) or not (in /zɪtʃdɪn/). In simple forms, the repair would remain the same as above: voicing is protected in onset, so it is changed in coda:

9) *Initial state, once post-morphology: simple word*

/zɪtʃdɪn/	Id-[vce]-OO	Agree[vce]	Id-Ons[vce]-IO	Id-[vce]-IO
zɪtʃdɪn		*!		
zɪtʃtɪn			*!	*
☞zɪdʒdɪn				*

In derived forms, however, coda voicing would now be protected by OO-Faith as part of a morphological base, so voicing would change in *onset*:

10) *Initial state,* ***post****-morphology: correctly analyzed /wʌtʃdəl/*

/wʌtʃ + dəl/	Id-[vce]-OO	Agree[vce]	Id[vce]-Ons-IO	Id-[vce]-IO
wʌtʃdəl		*!		
☞ wʌtʃtəl			*	*
wʌdʒdəl	*!			*

2.2.3 OO-faith kicks in at an intermediate stage

Returning to Hayes' point about learning timelines – it is probable by the time OO-faith begins to assign violations, much phonotactic learning will already have occurred. However this does not necessarily mean that the learner has reached the target ranking.

Imagine, for example, that the base 'wutch' were to be identified an intermediate stage between initial and final. Since the constraint set we are considering here includes just one Markedness constraint, the only possible intermediate stage is the Specific-F ranking in 11) below. This learner has promoted just a specific faithfulness constraint,

[1] This is a rather simplified definition of OO-Faith – see footnote 12 of chapter 1 – but the simplification does not affect the analysis here in any substantive way.

Ident(voice)-Onset, above Agree(voice), on the basis of some error that ESL has added to the Support (see §6.1 for what that error would look like.)

11) Id[vce]-OO >> Id[vce]-Ons-IO >> **Agree[voice]** >> Id[vce]-IO

Note, though, that this re-ranking has not been sufficient to tolerate voicing mismatches in general, because Agree(voice) still ranks above general Ident(voice).

What if morphology was learned at this point – what outputs would now be optimal for our two-word lexicon? Like the previous stage, the ranking in 11) chooses different outputs for this cluster, dependent on the morphology. In simple forms, the cluster is still repaired by a change in coda voicing:

12) *Intermediate stage, simple word*

/zɪtʃdɪn/	Id-[vce]-OO	Id-Ons[vce]-IO	Agree[vce]	Id-[vce]-IO
zɪtʃdɪn			*!	
zɪtʃtɪn		*!		*
☞zɪdʒdɪn				*

For the form /wʌtʃ + del/, however, OO-Faith blocks voicing change in coda, and onset-specific IO-Faith blocks a repair in onset. As a result, the optimal candidate is the faithful one, in which mid-ranking Markedness is violated and the voicing mismatch survives:

13) *Intermediate stage, derived word*

/wʌtʃ + del/	Id-[vce]-OO	Id-Ons[vce]-IO	Agree[vce]	Id-[vce]-IO
☞wʌtʃdel			*	
wʌtʃtel		*!		*
wʌdʒdel	*!			*

3. The experimental methodology: artificial language learning

3.1. The difficulties in testing for OO-faith in L1 acquisition

The previous chapter laid out the learnability case for high-ranking OO-faith and made predictions about stages of innovative OO-faithfulness in development. However, testing these predictions in natural L1 learning poses some problems.

The largest problem is catching children at the right stage – and being sure that it *is* the right stage. Unlike purely phonological analyses of children's production, the predicted effects of OO-faith at any stage are also tied to the learner's representational assumptions about morphology – its bases, relations, paradigms, and the like. Thus, the claim that a particular pattern results from the OO-faith influence is also a claim that children have learned enough morpho-semantics to calculate the right OO-faith relations and enforce their violation profiles in the ERCs added to the Support. In this respect, phonologically-transcribed data from sources like CHILDES will be insufficient in most cases to be sure of a child's state of morphological awareness.

Furthermore, English morphology does not provide many good testing grounds for such investigation. To see the innovative effects of OO-faith in a developing grammar, a child must have learned a morphological base that undergoes some phonological process or change in its target derived form, so that the child can block that process in their own derived word productions. One of the few good cases in English is flapping, which in the target causes an alternation between base-final [t]s and [d]s and their flapped correspondents in derived forms with vowel-initial suffixes:

14) *One potential English case of innovative OO-faith: base-final flapping*

a) The non-OO-faithful pattern in the target (bases underlined):

bases		*target derived forms*		*predicted OO-faithful forms*	
'wait'	[wai**t**]	'waiter'	[waî**ɾ**ɹ]	'waiter'	[waî**t**ɹ]
'sit'	[sɪ**t**]	'sitting'	[sí**ɾ**ɪŋ]	'sitting'	[sí**t**ɪŋ]
'need'	[ni**d**]	'needed'	[ní**ɾ**əd]	'needed'	[ní**d**əd]

In fact, Bernhardt and Stemberger (1998) do report the case of a child who went through this stage (see also the cases discussed in section 7.3.1) However, the biggest potential English cases come the shift of stress in fairly complicated derived words (cf. cýcle ~ cyclíty) that the average four year old has probably not heard very often (and whose morphological decomposition they have therefore probably not learned.)

For these two reasons, this chapter investigates the predictions of OO-faith in phonotactic learning using a somewhat novel experimental methodology: artificial language learning. In the next section I introduce this kind of testing method, and then present the child-directed version of the approach that I combined with a standard wug-test to tackle the OO-faith question.

3.2. The artificial language learning paradigm

Artificial language learning is an experimental technique where subjects are trained and tested on novel language material – words, sounds or the like – with the explicit understanding that they are from a language unknown to them. Training is implicit and quick, usually involving brief exposure to the novel language stimuli with little instruction other than to remember as much of the materials as possible. Testing can take a variety of forms, but its overall goal is to investigate what subjects have internalized about the novel language data – or, put differently, what effects the novel data have had on the subject's pre-existing linguistic knowledge. Usually, the materials are constructed specifically for the purposes of the experiment; although they are frequently constructed to resemble patterns attested in natural language learning, it is not necessary that they be the actual words or inventories of any *one* attested language. In addition, the paradigm allow us to compare the acquisition of *unattested* language patterns to attested ones.

The rationale behind artificial language learning is that it can be used ask questions or test predictions that are hard to test in natural settings, and also provides a way to control for as many confounds as possible. In the recent literature, research in this paradigm has compare adults' acquisition of minimally-different languages to test theories of how particular factors affect the ease or speed of phonological learning: properties such as phonetic naturalness, natural language attestedness, statistical patterns and probabilities, as well as the connection between static phonotactic and productive alternations. Research in this vein includes Esper (1925); Saffran, Aslin and Newport (1996); Saffran, Newport and Aslin (1996); Pater and Tessier (2003), (2005); Pycha et al (2003); Wilson (2003), (2006); Carpenter (2005), (2006); Peperkamp and Dupoux (2006); Morrison (2005). After discussing my own experimental results and their interpretation, I will return in section 8.1 to broader questions about this methodology and its applications.

3.3 The present application

The first crucial difference in this artificial language learning study is that it was with children: English-speaking four year olds who, (presumably) unlike adult subjects,

were still in the process of learning their L 1. To my knowledge, no previous work has used an artificial language learning paradigm of this sort with children. (It should be noted, however, that some infant speech perception research has asked related questions, in testing infants' abilities to acquire novel phonetic categories – e.g. Maye, 2000 – as well as novel lexical representations: Saffran, Aslin and Newport, 1996; Stager and Werker, 1997; Werker et al, 1998; Chambers et al, 2003.)

In the present context, the experimental goal was to test predictions about the effect of morphological discovery on phonological production. To induce morphological learning, I taught them a novel bound morpheme – a plural suffix – through direct comparison of singular and plural forms. To test the effects of newly-learned morphology on production I wanted to induce phonotactic errors, so the materials included marked coda-onset clusters, which were either very low-frequency or absent in English. This design was an attempt to simulate a morphologically-informed but phonologically-novel state, where morphological relations were known (and so OO-faith could be active), but high-ranking markedness constraints were still high enough to induce phonotactic repairs. By virtue of using novel language stimuli, we can be sure that participants were encountering the language's bases and plural suffix for the first time.

One question about this methodology is what grammar children are using when producing novel forms in this experiment; given that the participants were four years old, they had clearly learned most of the basic English phonological system. In most regards, the materials of this experiment were designed to allow for some agnosticism about the extent of English-specific grammar that participants brought to this task – the materials provided children with marked structures they were unfamiliar with, and so would probably have not yet acquired even outside the experimental context.[2]

4. Experimental Design

4.1 Experimental predictions

The words of the artificial language taught in this experiment were designed to compare faithfulness to bisyllabic forms in the two morphological conditions exemplified in section 2.2: within a morphologically-simple word, and across the base-affix boundary of a complex form. In my experiment, the Count Plural condition contained words like [wʌtʃ.del], where the first syllable was the singular base and the second syllable was the novel plural suffix; the Mass Noun condition contained words like [zɪtʃdɪn], where both syllables were neither base nor affix. As we saw above, grammars that have not completely demoted markedness constraints against a particular marked cluster should repair it differently in these two morphological contexts. This is summarized below:

15) *Two predicted asymmetric patterns of faithfulness, using OO-faith:*

	simple word	*derived word*
a) initial ranking in (7):	/zɪtʃdɪn/ → [zɪ**dʒd**ɪn]	/<u>wʌtʃ</u>$_{base}$+ dəl/ → [<u>wʌ**tʃ**</u>**t**əl]
b) intermediate ranking in (11):	/zɪtʃdɪn/ → [zɪ**dʒd**ɪn]	/<u>wʌtʃ</u>$_{base}$+dəl/ → [<u>wʌ**tʃ**</u>**d**əl]

Thus, the experiment was aimed at answering two specific questions about children's cluster productions in the two morphological conditions:

[2] One important related issue, however, is whether four-year old children have learned anything about English that drives them to be OO-faithful to clusters in the plural context that my experiment relied on. This is an issue I must look into further, and which will affect the design of subsequent experiments.

16) *Prediction 1, with respect to initial syllable codas*
Coda segments should be produced more faithfully in the first syllable of plural nouns than in the first syllable of mass nouns – e.g., more faithfulness to /tʃ/ in /wʌtʃ/ + /dəl/ than in /zɪtʃdɪn/

This asymmetry should manifest itself when children have made the morphological connection between base and affix and incorporated it into their phonology. Under both the initial and intermediate rankings, base codas should be protected – either at the expense of onset segments at the initial state, or Markedness violations at the intermediate state – whereas mass noun first-syllable codas should not.

17) *Prediction 2, with respect to unfaithful medial clusters:*
Among those tokens whose medial clusters are produced unfaithfully in some way, more of the unfaithfulness should be seen in onset position in the count nouns than the mass nouns.

This prediction is a specific test for the two-repair asymmetry of initial state. In those tokens where high-ranking Markedness has driven an unfaithful repair, the ranking of OO >> IO faith predicts onset repairs for count nouns, where OO-Faith protects the coda, but not for mass nouns where OO-Faith has no effect.

4.2 Materials

Table 23 below shows the representative properties of all the novel words children learned (details of how they learned them in the next section). Count noun singulars were all mono-syllables, of the shape CVC(C); each singular was suffixed with [dəl] to form a bisyllabic plural with the shape CVC(C).[dəl]. Mass nouns were all bisyllabic, of the shape CVC(C).dVC. All bisyllabic forms were initially stressed; all vowels in the second syllable were lax (of the set ɪ, ɛ, ə) and pronounced as unstressed but not completely reduced. Every effort was made by the experimenter to produce the clusters and their segments similarly in all tokens and contexts: in particular, coda stops were somewhat released (as they might be in very careful English speech). These measures were taken to attempt to make the two morphological types of bisyllabic words forms as prosodically-similar as possible, but to maximize the perceptability between the second syllables of plurals (always [dəl]) and mass nouns (always of the form [dVC].) Every cluster occurred both within a mass noun and across the count noun-plural suffix boundary.

The full set of clusters and items used:

18)

Coda Segment(s)	*Cluster*	*Count*	*Mass*
Stop(s)	b.d	pob (+ dəl)	gɪbdɪt
	g.d	wʌg (+ dəl)	mogdəm
	kt.d	lʌkt (+ dəl)	pæktdɪm
Fricative-Stop	ft.d	faft (+ dəl)	lʌftdɛk
Affricate	tʃ.d	wʌtʃ (+ dəl)	zɪtʃdɪn
	dʒ.d	bɪdʒ (+ dəl)	fodʒdɛt
	bʒ.d	gɪbʒ (+ dəl)	mæbʒdɪt
Nasal-Fricative	nf.d	nanf (+ dəl)	gʌnfdɛp
	mf.d	namf (+ dəl)	gʌmfdɛp

4.3 Methodology

The experiment was presented to the children as a novel language-learning game, in which their task was to learn some "alien" words, spoken by an alien puppet named Bozdim (operated by the experimenter.) Children were taught the alien words by association with pictures of familiar objects: count nouns in singular and plural contexts,

and mass nouns in two different containers. The children therefore learned a series of new nouns, as well as a novel plural suffix [dəl].

In training, children heard the puppet produce the alien words each paired with an object, and were encouraged to both imitate and spontaneously reproduce them in conversation. In the second part, the children played a picture-matching game with the puppet, in which the puppet named a picture and asked children to name a matching picture. In this game, children were tested for their pronunciation of the same clusters in bases (the two-syllable mass nouns with medial clusters) vs. derived forms (the plural, suffixed count nouns with clusters created at the morphological boundary.)

4.3.1 Participants

Twelve 4-year old children in the Amherst and Northampton, Massachusetts areas participated in the study. This age group was chosen because four-year-olds typically have fairly adult-like phonological grammars, but have not completely mastered difficult segments and clusters. With respect to morphology, four year olds are also reported to be at the stage of over-generalizing regular morphological patterns – e.g. *foots*, *mouses; ranned, bringed* – and thus presumably in the throes of productive morphological acquisition.[3]

4.3.2 Training

Initially, the experimenter (and puppet) presented children with picture one at a time. First, children were asked for the English name for the object, and engaged in short

[3] In fact, three of the children whose data is reported here produced English morphological errors of this sort during spontaneous conversation, including "He bringed the chair" and "I runned to the table".

discussion of the object and its properties – its colour, size, prototypicality, etc. – to get the child focused on the object. Then, the puppet was asked to give the name of the object in his language.[4] Once the puppet had given the object's name, the child was encouraged to repeat the name ("Can you say what Bozdim just said?"), and to use it in similar discussion as with the English name before ("Is this a blue wug? Or is it a yellow wug?" or "Does it look like this cup of zitchdin tastes good? Do you think the zitchdin is hot or cold?")

In this phase, children first learned three words of one noun class – count or mass – and then three of the other. Within a count noun block, participants first learned three singular nouns, and then their corresponding three plurals. Within a mass noun block, participants learned three mass nouns, and then heard the same three again in a different container – a glass of juice, and later a bottle of juice. Half of the participants saw count nouns before mass nouns, and the other half saw mass before count.

The set of materials that children learned in this initial block are given below:

[4] Ideally, the child would ask the puppet, but in the face of shyness the experimenter would do so instead.

19) *Materials*

Noun class	*Morphology*	*Prosodic shape*	*Sample words*	*Sample matching pictures*
Count	Three singulars (base)	CVC(C)	[pob]	one armchair
			[wʌtʃ]	one pick-up truck
			[nænf]	one flower
	Three plurals (base + suffix)	CVC(C).dəl	[pobdəl]	many armchairs
			[wʌtʃdəl]	a fleet of pick-up trucks
			[nænfdəl]	a garden of various flowers
Mass	Three mass nouns	CVC(C).dVC	[gɪbdɛt]	a glass of juice
			[zɪtʃdɪn]	a cup of hot chocolate
			[gʌnfdɛp]	a mug of milk
	Same three mass nouns	(same 3 as above 3)	[gɪbdɛt]	a bottle of juice
			[zɪtʃdɪn]	many cups of hot chocolate
			[gʌnfdɛp]	a carton of milk

4.3.3 Testing

Once children had learned three of each words, the experimenter asked the children to play a matching game with the puppet. All twelve pictures seen so far were laid out in front of the child; to play the game, the puppet pointed to one picture and named it for the child. The child would then find the matching picture, and name it for the puppet. In this game, the puppet pointed to one of each of the mass nouns, for the child to match by naming the other, and to each of the *singular* count nouns, for the child to match by naming the *plural*. Thus, testing asked the child to provide six words, all with difficult coda-onset clusters: three underived mass nouns, and three derived plural nouns from singular bases.

After the first game, children were presented another training block as in 19), with six more words, and another testing game was played. (For some children, only two more words of each category were taught, so that the second round of testing included only 4 words.)

5. Experimental Results

5.1 The data reported

To make any claims about the morphologically-sensitive phonological patterns in the data, we must be able to claim that participants had in fact learned the artificial language's morphology – that is, learned its plural suffix "del". In order to prove sufficient mastery of /dəl/, I required that participants provide at least one spontaneous token of more than one plural noun, associated with the right plural picture. This criterion eliminated 2 participants, leaving 10 children.

Of the 9 clusters tested, only 5 are included in the final results. Two criteria excluded the others: first, the cluster had to be pronounced unfaithfully in more than 2 tokens; second, it had to have been produced by more than 1 child in both the mass and plural contexts. The first criterion eliminated two clusters (gd, kd) and the second another three (ftd, ktd, bʒd), which leaves the following clusters:

20)	*obstruent/d*	*nasal-fricative/d*
	b.d	mf.d
	tʃ.d	nf.d
	dʒ.d	

All results reported are for these 5 clusters and 10 children. All tokens were reported, from what was referred to above as both training and testing – this is simply to get enough tokens.

In addition: plural tokens were only included when the participant produced a second syllable of type dV(C). In other words, the results do not include tokens with English plural affixes ("wutʃez, pobdelz") or zero morphology ("wutʃ, pob").

5.2 Testing the predictions

The majority of children's pronunciations were of two types: either faithful, or with reduction in the coda of the first syllable. To first give an impression of the data, table 21) summarizes the general results (variances across the 10 subjects given in parentheses):

21) *Results, across subject and by condition*

	total tokens	*faithful codas out of total tokens*		*unfaithful medial clusters out of total*		*faithful codas out of total unfaithful clusters*	
	#	#	%	#	%	#	%
plural nouns	87	69/87	**0.793** (0.035)	25/87	**0.287**	9/25	**0.36** (0.1711)
mass nouns	112	56/112	**0.50** (0.06)	52/112	**0.464**	1/52	**0.019** (0.006)
totals	199	125		77		10	

5.2.1 Testing prediction 1

The data in table 22) below allows us to test prediction 1 – that codas in initial syllables should be more faithful in count nouns than mass ones. The table shows the raw number of tokens of faithful coda productions that each subject produced in the two morphological conditions, and also the proportion of all tokens that were coda-faithful in each condition:

22) *Proportion of faithful σ1 codas by subject and condition*

	Mass Nouns			Plural Nouns		
Subject	faithful codas	total codas	% coda-faith	faithful codas	total codas	% coda-faith
C	17	20	**0.85**	11	11	**1**
E	4	9	**0.444**	9	9	**1**
A2	2	7	**0.286**	6	6	**1**
I	4	15	**0.267**	5	8	**0.625**
N2	8	12	**0.667**	9	12	**0.75**
A3	1	10	**0.1**	7	9	**0.778**
A1	3	9	**0.333**	10	15	**0.667**
D1	13	17	**0.765**	3	4	**0.75**
D2	1	3	**0.333**	7	10	**0.7**
N1	3	10	**0.30**	3	3	**1**
totals	**56**	**112**		**70**	**87**	
means			**0.5**			**0.805**
variance			0.06			0.024

Summing across all 10 subjects, a one-tailed t-test showed that codas were produced faithfully significantly less often in mass noun clusters, namely in 50% of tokens, than in the plural count noun clusters, where they were faithful 79.3% of the time ($p < 0.01$). Further, a pair-wise t-test, comparing the proportion means for each subject, also shows a significant difference between the lower proportions of faithful first-syllable codas produced in mass nouns compared to the higher proportions in plural nouns ($p < 0.01$).[5]

Thus, prediction 1 seems nicely borne out. This result provides some evidence for intermediate stage rankings, where plural nouns are faithful to both members of medial clusters, but mass nouns are still unfaithful in coda position.

[5] Statistics were calculated both for all 10 subjects, but also using just the first 8, since the low number of total items for the last two subjects, D2 and N1, might have skewed the proportions. Either way, however, the result is significant at $p < 0.01$.

5.2.2 Testing prediction 2

Table 23) below tests prediction 2 – that derived forms will show not only more faithfulness in their cluster codas, but also less faithfulness in cluster onsets. This result would correspond to an initial state ranking, where clusters are repaired in the *codas* of mass nouns, but the *onsets* of plural ones. To test for such a possibility, I consider again the proportion of faithful codas, but only among those that were unfaithful somewhere in the medial cluster:

23) *Proportion of faithful codas in unfaithful medial clusters by subject and condition*

	Mass Nouns				Plural Nouns			
Subject	faithful codas	faithful onsets	totals	% coda-faith	faithful codas	faithful onsets	totals	% coda-faith
C	1	3	4	**0.25**	0	0	0	**0**
E	0	5	5	**0**	1	0	1	**1**
A2	0	3	3	**0**	0	0	0	**0**
I	0	8	8	**0**	0	3	3	**0**
N2	0	4	4	**0**	1	3	4	**0.25**
A3	0	9	9	**0**	3	2	5	**0.6**
A1	0	6	6	**0**	1	5	6	**0.1667**
D1	0	4	4	**0**	0	1	**1**	**0**
D2	0	2	2	**0**	0	3	3	**0**
N1	0	7	7	**0**	0	0	**0**	**0**
totals	**1**	**51**	**52**		**6**	**17**	**23**	
means				**0.0192**				**0.261**
variance				0.0063				0.1711

The above table shows that unfaithful clusters were overwhelmingly unfaithful in coda position: only 1/52 mass nouns and 9/25 plural nouns with unfaithful clusters had faithful codas. This is not too surprising, given the privileged faithful status of onsets over codas (as encoded in the positional faithfulness Ident-Onset constraints used through this dissertation.)

Despite this clear tendency to apply repairs in coda, table 23) does provide some support for prediction 2. Across all 10 subjects, a one-tailed t-test assuming unequal variances shows that the mean proportion of coda faithfulness in these clusters is lower for mass than for plural nouns ($t = -2.077$; $p = 0.0322$.)[6]

5.3 Summary of results

The results of the previous section were positive: the data from our 10 children and 5 clusters does show the influence of morphological structure on phonological faithfulness, in the ways predicted by high-ranking OO-faith and the two developmental stages discussed.

While the results in the previous section were described in terms of faithful vs. unfaithful segments, it is clear that children differ in their patterns of repair, and for different clusters. Using table 23) as a guide, it seems that A3 and N1 are both at a stage where the effects of both the initial state and intermediate stages can be seen. The next section considers these two subjects' treatments of particular clusters in the two conditions in-depth, which as we will see provide ranking arguments that match the schematic stages already seen.

6. Rankings in the results

In this section, I focus on a few examples in the data collected with illustrate the particular rankings predicted in section 2.2.

[6] Across just the first 8 subjects, the difference between the means remains significant with a one-tailed t-test ($t = -2.2131$; $p = 0.0289$).

6.1 A3's cluster voicing: an intermediate ranking

Two nice examples of the predicted stages come from A3's treatment of the cluster /tʃ.d/. First, consider her pronunciation of two words below, with respect to cluster *voicing.*

24) /zɪtʃdɪn/ [zɪdʒdɪn]

25) /wʌtʃ + dəl/ [wʌtʃdʒəl]
c.f. /wʌtʃ/ [wʌtʃ]

The pronunciation [zidʒ.din] in 24) suggests that A3 does not tolerate affricate-stop clusters that disagree for voice, and that she can repair this disagreement by voicing the offending coda. The ranking in 26) derives this pattern: Agree(voice) >> Ident(voice) requires that voicing be repaired somewhere, and coda changes in voicing rather than the onset due to Ident(voice)-Onset.

26)

/zɪtʃdɪn/	Agree(voice)	Id(voice)-Onset	Id(voice)
zɪtʃ.dɪn	*!		
zɪtʃ.tɪn		*!	*
→ zɪdʒ.dɪn			*

Comparing this tableau to the token in 24) however, we can see that A3's grammar chooses different winners when this cluster is created at a base-suffix boundary: here, voicing surfaces faithfully, violating Agree(voice). The morphological explanation offered here is that the winners preserve the *base's* input voicing in this derived form.

Note that in the plural noun /wʌtʃ + dəl/, the coda affricate is part of the word's morphological base /wutʃ/ – and that A3 produces /wutʃ/ faithfully, with a voiceless coda. As 27) again shows, ranking the OO version of Id-voice above Agree(voice) predicts this pattern. Id(voice)-OO prevents coda voicing and positional faithfulness prevents onset devoicing in, so the (voicing) faithful candidate wins:

27)

/wʌtʃ + dəl/	Id(vce)-OO	Id(vce)-Ons-IO	Agree(vce)	Id(vce)-IO
(a) wʌtʃtəl		*!		*
(b) wʌdʒdəl	*!			*
→ (c) wʌtʃdʒəl			*	

This result is positive in two respects. First, A3's asymmetrical treatment of cluster voicing mismatches, despite the language's equal tolerance of them, is an intermediate stage ranking of the sort predicted by this gradual learner.

Second, the nature of this asymmetry in A3's grammar lends support to the hypothesized OO-faith bias as well. Without such a pressure, it is not clear why A3 should protect coda voicing – at the expense of the markedness of voicing disagreement – only in forms whose codas form part of a morphological base.

6.2 A3's treatment of coda affricates

Now we will consider two of A3's other tokens, with respect to the affricate [tʃ]:

[7] A3 also produced the form [wʌtʃʒəl]: while different in its treatment of cluster continuancy, this form also shows the progressive voicing of the winner in 31c) and so is compatible with 31)'s ranking.

28) /zɪtʃdɪn/ (a) [zɪʔdʒɪn]

29) /wʌtʃ + dəl/ (a) [wʌtʃdʒəl]
c.f. /wʌtʃ/ (b) [wʌtʃ]

In both morphological contexts, here A3 repairs some marked aspect of the cluster /tʃ.d/, although she produces coda /tʃ/ on its own faithfully in 29)b).[8]

For present purposes, I assume that the repair in the mass case, [zɪʔdʒɪn], represents the effect of a simple ban on coda affricates. The mass form is repaired by moving the continuant over to onset position, violating low-ranking IO faithfulness to continuancy.[9]

30)

/zɪtʃdɪn/	*CodaAffricate	Ident[+cont]-IO
(a) zɪ**tʃ**dɪn	*!	
(b) → zɪʔ**dʒ**ɪn		**

In the count noun case, the singular base also violates *CodaAffricate, and yet it is faithfully produced. This can be explained as another specific faithfulness effect: i.e. to word-final position.

31)

/wʌtʃ/	MaxSeg-WdFinal	*CodaAffricate	Ident[+cont]-IO
(a) → wʌtʃ		*	
(b) wʌʔ	*!		*

[8] I will leave aside the spreading of continuancy onto the plural affix in 29)a) although this effect seems ripe for analysis using a Markedness pressure for clusters to agree in continuancy, plus OO-faithfulness to the base.
[9] I assume that simply deleting the continuant feature is not an option, due to other high-ranking faithfulness constraints.

The crucial test case is A3's treatment of the count noun /wʌtʃ + dəl/: while in this form the coda affricate is no longer at the word edge, it is still produced faithfully. The explanation here is that its deletion is blocked by OO-faith, i.e. Ident[+cont]-OO.

32)

/wʌtʃ$_{base}$ + dəl/	Id[+cont]-OO	*CodaAffricate	Id[+cont]-IO
(a) wʌʔdʒəl	*!		*
(b) → wʌtʃdʒəl		*	*

This ranking in 32) simply demonstrates one preference of OO-faith over IO-faith in A3's grammar:

33) *OO-Faith* >> *Markedness* >> (*General) Faith*
Id[+cont]-OO >> *Coda Affricate >> Id[+cont]-IO

Again, the choice of two different repairs seems inexplicable without a morphologically-sensitive pressure like OO-faith.

6.3 N's treatment of [mf.d] clusters: an initial state ranking

A similar example comes from N's unfaithful realizations of /mfd/ clusters. N does sometimes produce this cluster faithfully, but looking just at the unfaithful ones reveals another two-repair asymmetry:

	Input	*Output*
34)	a) /gumfdin/	[gunfdin]
35)	a) /næmf + dəl/	[næmfəl]
	b) c.f. /næmf/	[næmf]

(Note that I am abstracting away from [p], which N inserts variably for all three of these forms.)

Whatever the markedness problem with a coda [mf] sequence – described here with the purely ad-hoc constraint *Coda[mf] -- N's grammar offers two different ways of solving it. In the mass noun, he retains all the input segments, but changes the place of articulation of the coda nasal from labial to coronal:

36)

/gumfdin/	*Coda[mf]	Max(Seg)-IO	Max(Labial)-IO
(a) gumfdin	*!		
(b) gumfin		*!	*
(d) → gunfdin			**

However – this repair is not applied in 35b) in the plural /næmf + dəl/ *[nænfdəl]. Instead, N deletes a following *onset* segment. On the current story, this is understood as another effect of high-ranking OO-Faith: the labial coda [m] in /næmf + dəl/ is protected by its membership to the base by Max(Labial)-OO, so a different IO constraint must be violated to satisfy markedness:

37)

/næmf+ dəl/	Max(Lab)-OO	*Coda[mf]	Max(Seg)-IO	Max(Lab)-IO
(a) næmfdəl		*!		
(b) → næmfəl			*	
(d) nænfdəl	*!			*

6.4 Summary of analyses

This section has highlighted three cases where the experiment results match the theoretical model. In all three cases, the child produced the same coda-onset cluster differently according to its morphological contexts, remaining preferentially faithful to base material in a derived context just as the OO-faithfulness account predicts.

7. Theoretical discussion

7.1 The intermediate stage, and Error-Selective Learning

The predictions that lead to my experiment were about the kinds of errors OO-faith could induce in a phonotactic learner – both at the initial stage in which markedness is still all-powerful (the ranking in 7), and at an intermediate stage where specific faithfulness alone has overcome markedness (the ranking in 11). This first stage, where Markedness is always obeyed, was a starting assumption of the learning theory I've adopted throughout this dissertation. But we have not yet seen why the Error-Selective BCD learner should go through this particular intermediate stage, whose ranking from 11) I repeat below:

38) *(repeated from 11)*
Id[vce]-OO >> Id[vce]-Ons-IO >> **Agree[voice]** >> Id[vce]-IO

To build such a ranking, the Error-Selective learner will have to have added an error to the Support where positional Ident-Ons[voice] prefers the winner so that BCD can install it high. With just the two errors and four constraints of section 2.2, however, this wouldn't happen – in that Error Cache, the only faith constraint doing any work is the general IO-Ident[voice]:

39) *Initial Support during phonotactic learning, with no morphological relations (repeated from 2)*

input	winner ~ loser	Ident[lo]–OO	*ait	*ʌi	Ident[lo]–IO
(ii)/ɹʌirəɹ/	[ɹʌirəɹ] ~ [ɹairəɹ]	e	e	L	W

If, however, we had a slightly more complex CON that included more markedness constraints, it could well be the case that some errors would optimally satisfy Agree-Voice by changing onset rather than coda voice. Such is the case in the hypothetical Error Cache below:

40) *An Error Cache*

winner ~ loser	Id[vce] -OO	*Vcd Velar (= *[g])	*Affricate	Agree [vce]	Id-Ons[vce] -IO	Id-[vce] -IO
(i) zɪtʃdɪn ~ zɪ**ʒ**dɪn	e	e	L	L	e	W
(ii) zɪtgɪn ~ zɪt**k**ɪn	e	L	e	L	W	W

If Agree[voice] were now to trigger Error-Selective learning on this Cache: the Markedness criterion of the ESA wouldn't choose between the two errors, because they each violate Agree[voice] and one other Markedness constraint. Thus the Faithfulness criterion would be allowed to choose, and this would mean choosing (ii) as the best error because of its additional faithfulness W for the specific faithfulness constraint:

41) *The resulting Support entry, chosen by the ESA*

winner ~ loser	Id[vce] -OO	*[g]	*Affricate	Agree [vce]	Id-Ons[vce] -IO	Id-[vce] -IO
(ii) zɪtgɪn ~ zɪt**k**ɪn	e	L	e	L	W	W

And adding just this error to the Support will make BCD choose the ranking in (42) below – which is the intermediate voicing grammar from above:

42) Id[vce]-OO >>*Affricate >> **Id-Ons[voice]-IO** >> **Agree[vce]**, *[g] >> **Id[vce]-IO**

The goal of the Error-Selective Learning proposal is not to insist that *all* learning goes through the same set of stages. For example, I have not claimed that learners of a language where Agree-Voice is freely violated will *always* go through an intermediate stage where onset voicing determines the clusters' voice properties. To recall the real-life example of french complex onsets acquisition in section 2, it appears that children impose positional faithfulness in different amounts, on different marked structures. This approach simply provides a mechanism by which the frequency of different markedness violations *could* produce intermediate stages like 42) above.

7.2 Independent evidence of innovative OO-faith

Hayes (2004) cites two cases in the literature where children appear to have innovated an OO-faithful pattern, in which their outputs fail to undergo a target-language process (i.e. satisfy Markedness) just in the higher interests of paradigm uniformity. As

mentioned in chapter 2's discussion of how to treat variation in children's developing grammars (§5.1): these innovations provide a serious challenge to any claim that production errors are a result of performance-only factors. That is: the evidence below suggests that children have innovated an OO-faith constraint *for which there is no evidence in the target grammar,* and then systematically ranked it high enough to make their child-specific patterns contingent on the bases that this OO-faith constraint protects.

One case that Hayes analyzes comes from data described by Kazazis (1969) from Marina, a Greek-learning child. Modern Greek has an allophonic alternation between velar and palatal fricatives which can be predicted by a following vowel. This alternation disrupts paradigm uniformity, because some verbal stems alternate between velar and palatal final consonants:

43) *Greek imperfective present paradigms: allophonic* [x,y] ~ [ç,j]

"to have "		"to leave"	
1sg. [e**x**o]	1pl. [e**x**ume]	1sg. [fev**ɣ**o]	1pl. [fev**ɣ**ume]
2sg. [e**ç**is]	2pl. [e**ç**ete]	2sg. [fev**j**is]	2pl. [fev**j**ete]
3sg. [e**ç**i]	3pl. [e**x**ume]	3sg. [fev**j**i]	3pl. [fev**ɣ**une]

This means that in the target Greek grammar, the markedness constraints that control the distribution of palatal and velar fricatives outrank both IO *and OO* faithfulness to fricative place:

44) *Target Greek ranking*
*[xi, ɣi] >> *[ç, j] >> **OO-Max[place],** IO-Max[place]

The child Marina had been correctly producing Greek velar and palatal fricative allophones – up until 4;7, when for a few weeks she began defaulting to the velar fricative just where the palatal disrupts the paradigm uniformity of verbal stems:

45) *Marina at 4:7:fricative palatalization blocked by OO-Faith*

		Target	*Child*
2pl.	'to have'	[e**ç**ete]	[e**x**ete]
2pl.	'to leave'	[fev**j**ete]	[fev**ɣ**ete]

46) *Marina's grammar*
OO-Max[velar] >> *[xi, ɣi] >> *[ç, j] >> IO-Max[velar]

This data fits neatly with the present theory. The BCD algorithm predicts that Marina should indeed have been installing OO-Max[velar] at the top of her ranking all along, for no reason other than her biases. Below that, the distributional evidence from her phonotactic learning ranked her markedness constraints in the right allophonic order. With the assumption that she has now realized that words like [e**ç**ete] have a verbal base /ex-/, these new errors have emerged under new pressure from undominated OO-faith:

47)a) *Simple words: fricative place decided by *[xi]*

/exete/ (hypothetical)	OO-Max [velar]	***[xi, ɣi]**	***[ç, j]**	IO-Max [velar]
(i) exete		*!		
(ii) → eçete			*	*

47)b) *Derived words: palatalization can be blocked by OO-Max[velar]*

/e**x-**/$_{base}$ + /ete/	OO-Max [velar]	*[xi, ɣi]	*[ç, j]	IO-Max [velar]
(i) → exete		*		
(ii) eçete	*!		*	*

Hayes also cites Bernhardt and Stemberger's (1998) case mentioned in §3.1, of an English-learning child who, from 2:0 – 3:8, consistently flapped in simplex words (e.g. 'water' [wɑɾər]), but only produced base-faithful voiced and voiceless stops in derived words even where the adult phonology required a flap (e.g. 'sitting' [sɪ**t**ɪŋ] from base [sɪ**t**] and 'needed' [niɾəd] from base [ni**d**])[10].

Another case of OO-faith innovation comes from Smith (1973)'s seminal diary study of English acquisition by Amahl. Both Macken (1980) and Jesney (2005) point out a paradigm uniformity quirk in Amahl's puzzle-puddle-pickle chain shift. This shift in Amahl's grammar between 2;2 and 2;11 caused coronal stops to become velar before laterals, in both simple and derived words:

48) *Amahl's velarization before laterals:* /t,d/ → [k,g]

'pu**dd**le'	[pʌgəl]	'gen**t**le'	[dɛŋkəl]	'pa**dd**ling'	[pægəlin]
'tur**t**le'	[tʰəːkəl][11]	'gen**t**ly'	[dɛŋkliː]	'pe**d**aling'	[pɛgəlin]

This pattern is a more extreme version of the English ban on tl and dl onset clusters (cf. the syllabifications of ma.**tr**ess vs. a**t.l**as) – in Amahl's grammar, these sequences are ruled out regardless of syllabic position. For present purposes I simply

[10] See Bernhardt and Stemberger (1998) for their alternative analysis of this data.
[11] Amahl was learning a British English dialect that lacks this post-vocalic [ɹ].

adopt the constraint *tl, intended as an OCP constraint that disallows sequences of coronal stops and laterals, and rank this constraint above faithfulness to consonantal place (see Jesney, 2005 for a full treatment of this chain shift in Amahl's grammar.)

49) *Amahl's pattern of pre-lateral velarization*

/pʌdəl/	*tl	Max [place]
(a) pʌdəl		*
☞ (b) pʌgəl	*!	

There is one class of words in which pre-lateral coronals do not become velarized: derived words whose *base* had only the stop and not the following lateral. As shown in 21) below, words like 'tight' surface with their normal coronal stop, and this stop is retained in derived words like 'tightly' even though provide the phonological context for velarization:

50) Velarization blocked *when the base has no velarized segment:*

'hard'	[haːd]	'har**d**ly'	[haːdli]
'soft'	[sɔft]	'sof**t**ly'	[sɔftliː]
'tight'	[tait]	'tigh**t**ly'	[taitliː]

The explanation adopted here is that coronal stop in 'tightly' is required to be OO-faithful to the stop in its base, 'tight'.[12]

[12] An alternative account of this data is that Amahl's pre-coronal velarization was the result of misperception (that he was hearing tl as [kl] in words like 'gentle' or 'gently'), and that those words whose bases had no l allowed him to perceive the coronal correctly – see Macken (1980).

51) *Adding OO-faith to Amahl's ranking*

/tait/	OO-Max [place]	*tl	IO-Max [place]
☞ (a) tait			
(b) taik			*!

/tait + ly/	OO-Max [place]	*tl	IO-Max [place]
☞ (a) taitly		*	
(b) taikly	*!		*

7.3 The persistent OO-faith bias and the GLA

7.3.1 The empirical need for persistent OO-faith

Section 2.1 alluded to the argument in Hayes (2004) that the OO-faith bias must be persistent, given the facts of order of acquisition. The facts that Hayes refers to come from the growing body of experimental work about *receptive* learning in very young children. On the one hand, this work has demonstrated that children have internalized native phonotactic distributions in some sense roughly by the age of 8-10 months (e.g. Werker & Tees 1983, 2002; Jusczyk et al 1993, 1994; Frederici & Wessels 1993; Johnson and Jusczyk, 2001). Second, what experimental evidence we have about the early receptive acquisition of morphology suggests the beginning of such learning occurs somewhere in the second year of life (e.g. Shady, 1996, Santelman and Jusczyk 1998, as well as a brief survey of references in Hayes, 2004).

Of course, this kind of evidence from perception experimentation does not translate directly into a claim about the relative order of acquisition of *productive* grammars or rankings. The entire survey of production data discussed throughout chapter 2 makes it obvious that the production grammar still has much to learn about phonology for several years after the perception revolution of 8-10 months. The anecdotal OO-faith data from section 1.3 also suggests that morphological basehood and its effects on phonological patterns and paradigms is also still very much under construction throughout early childhood.

7.3.2 The GLA problem with persistent biases and OO-faith

The lag between the acquisition of surface phonotactics and morphological bases provides another kind of winner misparse of the type described in chapter 4. As we've already seen, ERC rowss that are added to the Error Cache or Support before any relevant morphological bases have been identified are missing the input structure that allow them to be assigned OO-faith violations.

52) *(repeated again from 2)*

input	winner ~ loser	Ident[lo]–OO	*ait	*ʌi	Ident[lo]–IO
(ii)/ɹʌiɾəɹ/	[ɹʌiɾəɹ] ~ [ɹaiɾəɹ]	**e**	e	L	W

And so the same argument can be made as in previous sections. The GLA is in danger of end-state overgeneration if the morphological information missing in (52) stays missing long enough – because once this misparsed winner has promoted IO-faith above Markedness, the GLA learner has no way to reverse that ranking.

Imagine that we equip the GLA with an *a priori* bias for OO-faith, which ensures that regardless of input data, the OO version of any faithfulness constraint must always remain 20 points above the IO version. Now let us try to learn the distribution of Canadian Raising from section 2. Before morphological discoveries have been made, the learner makes errors that are parsed as in 53) below:

53) *Errors during phonotactic learning, with no morphological relations*

input	winner ~ loser	Ident[lo]–OO	*ait	*ʌi	Ident[lo]–IO
(i)/ɹʌiɾəɹ/	[ɹʌiɾəɹ] ~ [ɹaiɾəɹ]	e	e	L	W

These provide evidence for demoting *ʌi and promoting IO-Ident – and by virtue of this a priori ranking bias, OO-Ident will be promoted as well when IO-Ident gets too close, as shown in 54):

54)a) *The GLA learning effect of 51), early in acquisition*

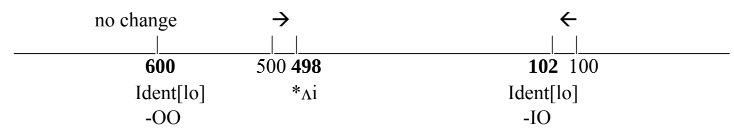

54)b) *The potential role of the* a priori *OO-faith bias, later in acquisition*

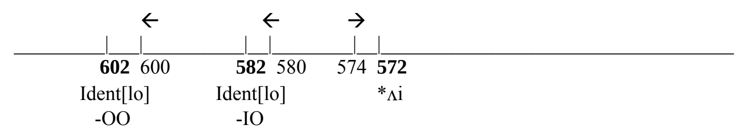

In 54)b), IO-faith has gotten close enough to OO-faith that the apriori bias pushes *both* constraints up the hierarchy, even though 53) only provides evidence to move the IO version.[13]

[13] For the reader who wonders how *ʌi can have gotten to a ranking value of 582, considerably higher than its starting value: recall that the grammar could be making other independent errors in which *ʌi prefered the winner. This wouldn't be the case in our Canadian Raising example -- but my real point is not the OO >> IO ranking but rather the reversal of IO-faith and Markedness.

By the later stage ranking above, the GLA has learned the same grammar that BCD initially does – but since it's stopped making errors, there is no more learning to be done. Fixing a lexical entry for 'writer' to include a base won't cause any errors, because the grammar in 54)b) has high-ranking Faith along *both* dimensions, even though either would do:

55) *No error, post-morphology*

/ɹʌit + əɹ/	Ident[lo]–OO	Ident[lo]–IO	*ʌi
(a) ɹairəɹ	*!	*!	
(b) ☞ ɹʌirəɹ			*

Kie Zuraw (p.c.) makes the interesting suggestion that the GLA's general problem with faithfulness and stringency could be addressed by building in a persistent bias for demoting IO-faithfulness independent of errors. For example: the GLA could demote every IO-faithfulness constraints by some small amount every time the grammar is used (or every day, or at some other frequent interval.) This would seem to be the best GLA version of a persistent low-faithfulness bias: it would have the effect of demoting all IO-faith constraints as far as they can go without causing errors – and if OO-faith is now doing the work of preventing errors, IO-faith will be allowed to sink to the bottom. Whether such an approach would provide an adequate answer to all the superset problems raised here is a question open for further investigation.

8. Experimental discussion

8.1 The connection between natural and artificial language learning

The overall results of this experiment matched those that we expect if learners initially rank OO-faith constraints at the top of their grammars. Recall that the learning theory that predicted high-ranking OO-faith was one built in response to superset grammar traps in natural language learning – e.g. McCarthy (1998)'s example in chapter 1 section 3.2.3 of learning static generalizations about non-alternating paradigms. Thus, the fact that these predictions were confirmed here supports the idea that artificial language experiments tap the same kind of phonological knowledge, of constraint rankings and biases, used in natural language acquisition.

From just this study, one might conclude that this connection is only a property of the behaviour of children – that is, that four year olds are still sufficiently engaged in the L1 learning task that their acquisition of artificial forms and paradigms can be influenced by true phonological learning mechanisms. But a number of artificial language learning studies have drawn similar conclusions – even in experiments with adult speakers.

For example, Carpenter (2005) taught native English speakers two patterns of sonority-influenced stress, and found that speakers learning the attested pattern, in which stress is attracted to low vowels, were better at predicting the stress of unfamiliar words than the learners of the opposite pattern of high-vowel stress attraction, which is unattested in the natural language typology. In experiments by Wilson (to appear), English-speaking adults learning patterns of velar palatalization were found to generalize the process from mid vowels to high vowels, but not vice versa, in keeping with the typological fact that natural languages whose velars palatalize next to mid vowels also do so next to high vowels, but not vice versa. And in Pater and Tessier (2003, 2005), adults were better at learning a phonological alternation that served to meet a static phonotactic generalization of their native language (the English minimal word requirement) than a comparable one that had no such L1 justification.[14] In sum, these results demonstrate that artificial langauge learning can indeed produce results that accord with a range of assumptions about natural language knowledge and its acquisition.

In the present experiment, one unanswered question is why children are *ever* unfaithful to initial syllable codas in plurals. According to the theory outlined here: once children have learned the suffix 'del', the predicted rankings protect base material at the expense of something else (namely affix material or markedness), and therefore a plural noun's initial syllable coda should remain untouched. This prediction of OO-faithfulness is clearly too strong for my results: section 5 showed that base codas were more faithful than other codas, but that base codas were still much less faithful than onsets in general.

One answer may lie with the mental resources required to implement an OO-faithful grammar: setting up a lexical entry for a closed class affix like 'plural', constructing a morphologically-complex input to the phonological grammar and the like. It is remains unknown, at least in this methodology, how and when the morphological knowledge that [dəl] is an affix was used online, either to prompt learning via constraint re-ranking or to rule out suboptimal candidates that violated OO-faithfulness. But it seems reasonable to suggest that all of this required a certain amount of concentration and effort, and thus explains part of this variability.

[14] I will also cite the results of Peperkamp and Dupoux (2006) here, but I confess I do not quite understand them yet.

8.2 A potential perceptual confound, and the next step

One alternative reading of this experiment's results is that the different morphological conditions did not induce different cluster repairs, but rather different percepts. Recall the three morphological contexts in which subjects heard coda segments, e.g. [tʃ]:

56)	(a) *count singulars*	(b) *count plurals*	(c) *mass nouns*
	[wʌtʃ]	[wʌtʃdəl]	[zɪtʃdɪn]

In the mass nouns like 56)c), codas affricates were only ever heard in a cluster, before a following [d]. In the count nouns, however, codas were heard in both the same pre-consonantal context of 56)b), but also word-finally in the related count singular of 56)a). So, it could be that children produced more accurate coda segments in count singulars only because those were the segments that they'd *heard* more accurately. If a subject misheard a coda consonant in the plural form of (56b), their accurate perception of that coda in its related singular could still let them choose the right input form as the base. For the mass noun, however, there is no related form with a word-final version of the coda to suggest an alternative input:

57) *Potential perceptual asymmetry*

	count singulars	(b) *count plurals*	(c) *mass nouns*
(a) sound:	[wʌtʃ]	[wʌtʃdəl]	[zɪtʃdɪn]
(b) subject perceived:	[wʌtʃ]	[wʌ**ts**dəl], or [wʌ**ʃ**dəl]	[zɪ**ts**dɪn], or [zɪ**ʃ**dɪn]
(c) inferred input:	/wʌtʃ/	/wʌ**tʃ** + dəl/ (from the sing.)	/zɪ**ts**dɪn/, or /zɪ**ʃ**dɪn/

Knowing precisely how much each participant perceived in each morphological condition is crucial to making claims about the grammars being used or acquired by subjects in the course of experiment.[15] And given its design, this methodology cannot tell us what was perceived.

To eliminate this experimental confound, the best next step is probably to use a similar training methodology, but to test subjects' resulting knowledge using a receptive task -- one that would tap learner's acceptability judgements about *new* forms.

9. Chapter Summary

The experiment reported here provides novel experimental evidence of high-ranking OO-faithfulness constraints in phonological acquisition. When four-year-old children were faced with marked consonant clusters in a novel language, their repairs to those clusters demonstrated preferences for OO-faithfulness over both Markedness and IO faithfulness.

These results provide novel empirical support for the present view of learning, in which learners come to their task with a bias for uniform paradigms independent of data triggers from the target. Further, this experiment also provides novel evidence that children are both willing and able to engage in artificial language learning of this type, particularly in learning new functional material like a plural suffix. Such results may pave the way for a fruitful new brand of experimental work on children's phonologies – because they suggest that artificial language experiments that use more novel materials than traditional wug tests can be used to tap the state of learner's phonological knowledge throughout the course of development.

[15] Thanks to Adam Albright for early discussion of these issues.