Lexical and morphological conditioning of paradigm gaps

Adam Albright
MIT

Draft: April 2006 (Comments welcome)

1 Introduction

Although paradigm gaps are an analytical puzzle for all linguistic theories, until recently they have received only sporadic discussion in the literature. Since the advent of Optimality Theory, however, gaps have gradually become the subject of more systematic attention. This is certainly due in part to the particular mechanical challenge that gaps pose for OT; hence, a major focus has been on providing a mechanism for all overt candidates to be eliminated, rather than the more usual outcome of one emerging as optimal (Prince and Smolensky 2004, p. 57; Orgun and Sprouse 1999; Fanselow and Féry 2002, Raffelsiefen 2004; Rice 2005, 2006; McCarthy and Wolf [this volume]). Equally important, though, is the fact that OT provides a natural way to formalize an intuition expressed in some earlier discussions, that gaps often appear when the expected form would violate a surface-true phonotactic constraint and that silence is simply one part of a larger conspiracy to avoid illegal configurations (Hetzron 1975; Iverson 1981). For this reason, discussions in the OT literature have focused primarily on cases that appear to have clear phonotactic motivation, in that the expected faithful candidates would involve illegal configurations like stress lapse, OCP violations, sonority sequencing violations, and so on (Orgun and Sprouse 1999; Raffelsiefen 2004; Rice 2005).

Not all cases of paradigm gaps involve such obvious phonotactic violations, however. The focus of this paper will be on cases that affect only certain words, while other, seemingly parallel words surface as expected without gaps. For example, many speakers of American English find the past participles of some irregular verbs to be problematic (Pinker 1999):

(1) Problematic past participles in American English

a. *dive* ∼ *dove* ∼ ???

> "He has diven, dived, dove, doved, doven or whatever . . . in every major and most minor bodies of water on this planet"[1]

b. *stride* ∼ *strode* ∼ ???

> "I have strode (stridden? strided?) into the backyard, boots strapped on. . . "[2]

c. *smite* ∼ *smote* ∼ ???

[1] 4/05/2006: http://www.scubaboard.com/archive/index.php/t-4550.html
[2] 4/05/2006: http://spanglemonkey.typepad.com/spanglemonkey/2006/01/the_what_the_fu.html

"It's like you're smitten by God. Smote? Smoten? Smoted? Smited? Well, whatever."[3]

"So I think I've been smote. Smited? What is the past tense of 'to smite'? Anyway, God got me!"[4]

"…the Arab planes are smote (smited? smut? smeet?) mysteriously from the sky to the bafflement of everyone"[5]

d. *strive* ∼ *strove* ∼ ???

"I've stroven/strove/striven to escape this insanity…"[6]

The problematic forms in (1) do not appear suffer from irreparable phonotactic violations—in fact, all of the possible options find parallels in other, non-problematic verbs: *arrived* [əɹaɪvd], *driven* [drɪvn̩], *ridden* [rɪdn̩], *woven* [woʊvn̩], *written* [rɪtn̩]. For this reason, such cases have sometimes been referred to as *lexically arbitrary paradigm gaps* (Hetzron 1975; Albright 2003). The aim of this paper is to understand why such gaps arise, and why they affect particular words in particular parts of the paradigm. I claim that they are neither arbitrary not lexical, but that their occurrence is in large part predictable: namely, gaps occur when speakers know that an inflected form must stand in a certain relation to another inflected form, but the language does not provide enough data to be certain of what that relation should be.

To see how the account works at an intuitive level, consider the case of **diven*/*doven*/*dove*. In most cases in English, the form of the past participle can be predicted fairly accurately by looking at the form of the simple past: if the simple past is suffixed with [t] or [d], the past participle is identical to it; otherwise, it is created using a set of vowel changes, with or without additional suffixation ([æ] → [ʌ] (*drank* → *drunk*); [ʊ] → [eɪ]+*en* (*shook* → *shaken*); [u] → [oʊ]+*n* (*grew* → *grown*); etc.). In the case of irregular pasts with the vowel [oʊ], however, there are several different competing patterns:

(2) Participle formation for verbs with [oʊ] pasts

a. No change: *shone*

b. [oʊ] → [ɪ], and suffix *-n*: *written*, *driven*, *risen*, *ridden*

c. Suffix *-n*: *worn*, *torn*, *sworn*, *born(e)*, *broken*, *woken*, *spoken*, *stolen*, *frozen*, *woven*, *chosen*

Not only do these patterns compete with one another, but they are also based on relatively small amounts of data—even the most robust pattern has just eleven examples, while the others involve a few verb roots each. The observation at the core of the analysis is that generalizations that cover so few forms (and suffer from exceptions at that) are not well enough supported to be trustable. This means that for verbs with irregular past tense forms in [oʊ], speakers do not have any usable generalizations that let them confidently predict the participle. In most cases, this is not a problem; the verbs in (2) are generally common enough that speakers can memorize their participles, and need not use their grammar to derive the output. For relatively uncommon verbs like *stride* or *strive*, however, speakers may not have sufficient exposure to be sure of their

[3] 4/12/2006: http://favorabledicta.blogspot.com/2004/12/my-name-is-espat-and-im-addicted-to.html
[4] 4/12/2006: http://heather.tadma.net/archives/00000661.html
[5] 4/12/2006: http://www.trashcity.org/ARTICLES/IBFS0007.HTM
[6] 4/05/2006: http://raven.utc.edu/cgi-bin/WA.EXE?A2=ind0202&L=scuba-se&P=29779

participles, and are forced to resort (unsuccessfully) to their grammar.[7] The grammar provides several possibilities (*stroven*, *striven*, *strove*), but all are poorly supported. In such cases, speakers have no way to produce a past participle: they lack a lexically listed form, and the grammar does not provide any way to derive one.

In order to formalize this intuition, we need two components. First, it is crucial for this account that forms are projected from particular other forms in the paradigm—e.g., that past participle in English are generated with reference to the simple past forms, and cannot be projected directly from present tense forms. To see why this is necessary, consider the mapping from presents to past participles: here, we find an overwhelmingly strong generalization, namely, that verbs of any shape can form past participles with *-ed*. In fact, this pattern is extremely likely and well-supported, because most verbs in English are regular. If speakers had access to this robust pattern, they could use it to confidently derive past participles like *strided*, *strived*, *dived*, and *smited*. The fact that they do not do this indicates the irregular past tense forms (*strode*, *strove*, *dove*, *smote*) play a crucial role in the generation of the participle. A major challenge in the analysis of lexically arbitrary paradigm gaps is to understand why certain statistically well-supported generalizations are unavailable to speakers. The approach taken here is to say that generalizations about the relation between the present form and the past participles cannot be used because the grammar of English does not derive past participles directly from present tense forms. Naturally, in order to give this claim substance, we need a theory of how relations are established between different parts of the paradigm. One goal of this paper, therefore, is to show how a model of paradigm organization can predict the occurrence of paradigm gaps in certain parts of the paradigm.

The second thing that is needed to develop this account is an understanding of what constitutes sufficient evidence for a trustable grammatical generalization. Looking at the participle forms of [oʊ] pasts in (2), we see that two factors seem to be important: the inconsistency of the data (three competing patterns), and the paucity of the data (less than a dozen forms for each pattern). In order to make use of this observation, we need a theory of how and why these factors influence grammar. Developing such a theory raises numerous questions: how much data is required? Must it be both sparse and inconsistent to create uncertainty, or is one factor alone enough? What is the tradeoff between sparseness and inconsistency? How is this information stored, and how does the grammar make use of it? A second goal of this paper is to provide a preliminary model of how learners decide which generalizations can be confidently extended, by evaluating the consistency and abundance of data.

The organization of the rest of this paper is as follows: in section 2, I present the details of a lexically arbitrary paradigm gap in the Spanish verbal system (Albright 2003). I show why this data cannot be satisfactorily explained using mechanisms that have been successful elsewhere, such as MPARSE relativized to particular morphological categories (McCarthy and Wolf, this volume) or lexical conservatism (Steriade 1997; Pertsova 2005), and why a model that relates members of the paradigm to particular other forms is needed. Next, I sketch a model of grammatical learning that attempts to discover an optimal set of relations between members of the paradigm, based on considerations of predictability and contrast maintenance (Albright 2002a). This model has two important properties for the analysis of paradigm gaps: first, it correctly pre-

[7] The verbs *smite* and *dive* are missing past participles for different reasons: *smitten* has undergone semantic drift and is no longer clearly associated with *smite*, while *dive* was historically regular, so did not inherit a participle that corresponds to irregular *dove*.

dicts which forms in the paradigm may potentially suffer from gaps. In addition, it learns grammars that take into account both the reliability of generalizations, and also the amount of data that supports them. Under this system, generalizations receive higher confidence to the extent that they are based on many forms, and do not vie with competing patterns. When competition is fierce, or when data is sparse, there may be no high-confidence generalizations, leaving the speaker to resort to word-specific knowledge. After showing how this model has the potential to explain key aspects of the cases discussed in section 2, I sketch the outline of an OT analysis of lexically arbitrary gaps. Finally, I consider the relation between these lexically arbitrary cases, and the more clearly phonotactically motivated gaps discussed by Orgun and Sprouse, Raffelsiefen, Rice, McCarthy and Wolf, and others.

2 Some challenges for a theory of paradigm gaps

In this section, I lay out a set of explicanda for a theory of lexically arbitrary paradigm gaps. As an illustration, I use data from two types of gaps in Spanish present tense paradigms, described in more detail in Albright (2003). After this backgound, I will then show in sections 2.5-2.4 why current mechanisms for licensing gaps or alternations in OT do not capture the full range of facts.

2.1 Gaps in Spanish present tense paradigms

Spanish verbs exhibit a wide variety of lexically idiosyncratic properties; the ones that will relevant for this discussion are those that are seen in the present indicative paradigm. The first major division concerns the vowel that shows up between the verb stem and the person/number endings: *-a-* (class 1), *-e-* (class 2) or *-i-* (class 3; /i/ reduces to [e] when stressless). The paradigms in (3) also show that stress (indicated here with V́) falls on the root in the 1sg, 2sg, 3sg and 3pl forms, and the suffix elsewhere. The first singular suffix is always *-o*, while the remaining suffixes reveal the conjugation class to varying degrees. Class 1 (*-ar*) is the productive class.

(3) Spanish conjugation classes

a. *habl**á**r* 'speak' (Class 1)		b. *com**é**r* 'eat' (Class 2)		c. *viv**í**r* 'live' (Class 3)	
hábl-o	*habl-**á**mos*	*cóm-o*	*com-**é**mos*	*vív-o*	*viv-**í**mos*
*hábl-**a**s*	*habl-**á**is*	*cóm-**e**s*	*com-**é**is*	*vív-**e**s*	*viv-**í**s*
*hábl-**a***	*hábl-**a**n*	*cóm-**e***	*cóm-**e**n*	*vív-**e***	*vív-**e**n*

Some verbs that contain mid vowels *e*, *o* in the final syllable of the stem undergo lexically conditioned changes in stressed forms. In some verbs, *e* and *o* diphthongize to *jé*, *wé*, while in others, *e* raises to *i*; in yet others, no change is observed.

(4) Lexically conditioned mid-vowel alternations

a. Diphthongization of *e*, *o*

i. *sentar* 'seat'

s[**jé**]*nt-o*	*s*[**e**]*nt-ámos*
s[**jé**]*nt-as*	*s*[**e**]*nt-áis*
s[**jé**]*nt-a*	*s*[**jé**]*nt-an*

ii. *contar* 'count'

c[**wé**]*nt-o*	*c*[**o**]*nt-ámos*
c[**wé**]*nt-as*	*c*[**o**]*nt-áis*
c[**wé**]*nt-a*	*c*[**wé**]*nt-an*

b. Raising of *e*

i. *pedir* 'request'

p[**í**]*d-o*	*p*[**e**]*d-ímos*
p[**í**]*d-es*	*p*[**e**]*d-ís*
p[**í**]*d-e*	*p*[**í**]*d-en*

c. Neither

i. *rentar* 'rent'

r[**é**]*nt-o*	*r*[**e**]*nt-ámos*
r[**é**]*nt-as*	*r*[**e**]*nt-áis*
r[**é**]*nt-a*	*r*[**é**]*nt-an*

ii. *montar* 'mount'

m[**ó**]*nt-o*	*m*[**o**]*nt-ámos*
m[**ó**]*nt-as*	*m*[**o**]*nt-áis*
m[**ó**]*nt-a*	*m*[**ó**]*nt-an*

There are also significant interactions between conjugation class and mid-vowel alternations. Diphthongizing verbs occur in all three classes, but they are a minority in class 1, somewhat more prevalent in class 2, and a majority in class 3. Raising occurs only in class 3, and in that class, almost every mid-vowel verb either diphthongizes or raises (Harris xxx). There is also an effect of vowel backness: the front mid vowel *e* undergoes both diphthongization and raising, while back *o* diphthongizes but does not raise. Furthermore, the segmental contexts that encourage diphthongization differ substantially between *e* and *o*, and the rate of *o* diphthongization is overall lower than that of *e* (Albright, Andrade, and Hayes 2001). In sum, it is difficult to form any language-wide generalizations about mid-vowel alternations, since they depend heavily on the particular vowel involved, the segmental context, and the inflectional class of the verb.

Another alternation is seen in the 1sg, in which some verbs in classes 2 and 3 show a process of velar insertion: [k] or [g] added between the stem and the suffix. Voiceless [k] is frequently added after stem-final [s]/[θ] (depending on the dialect), while [g] often occurs after stem-final [l] or [n].

(5) Velar insertion in the 1sg

a. *crecer* 'grow'

*cré[s-**k**]-o*[8]	*cre[s]-émos*
cré[s]-es	*cre[s]-éis*
cré[s]-e	*cré[s]-en*

b. *valer* 'be worth'

*vál-**g**-o*	*val-émos*
vál-es	*val-éis*
vál-e	*vál-en*

The are a number of other minor patterns affecting 1sg or stressed forms, which I will not discuss here. What matters for present purposes is that the major unpredictable properties of verbs include (1) conjugation class, (2) diphthongization and raising of mid vowels, and (3) velar insertion.

[8]The stem-final fricative in this stem, written <c>/<z>, is pronouced [θ] in some Iberian dialects.

With this in mind, we now turn to the two patterns of present tense paradigm gaps reported by Spanish grammars (de Gámez 1973; Butt 1997). The first involves verbs that lack all inflected forms in which stress would fall on the root:

(6) Missing stressed forms: *abolir* 'to abolish'

—	*abol-ímos*
—	*abol-ís*
—	—

The second pattern of gaps involves verbs that lack specifically the 1sg:

(7) Missing 1sg forms: *asir* 'to grasp'

—	*as-ímos*
ás-es	*as-ís*
ás-e	*ás-en*

Significantly, verbs with gaps also meet the conditions for undergoing stem alternations: *abolir* contains a potentially diphthongizing mid vowel, while *asir* contains a stem-final strident associated with velar insertion. More generally, verbs missing stressed forms typically belong to the *-ir* class, in which virtually all mid vowel verbs are irregular in some way, while verbs missing the 1sg belong to the *-er* and *-ir* classes, where velar insertion often occurs. Furthermore, affected verbs are missing just those forms where the alternation would apply. (Exactly parallel facts are also remarked on for Russian by Hetzron 1975, p. 861). Finally, when speakers are asked to inflect the verb *abolir*, they frequently utter both diphthongized and non-diphthongized possibilities (e.g., 3pl *abolen*, *abuelen*), before eventually settling on one or rejecting both. I take this as very strong evidence that the gaps in these cases involve uncertainty about whether to employ an alternation or not. The connection between irregularity and uncertainty in creating lexically arbitrary paradigm gaps is important, and must be accounted for.

What is the source of this uncertainty? Unlike cases such as Norwegian imperatives (Rice 2005, et seq.), where speakers are forced to choose between a faithful form that is unpronounceable in Norwegian (*[sikl]) and forms that employ fixes not otherwise seen in verbs ([sikl̥]? [sikkəl]? [siklə]?), in Spanish the competing choices both find at least some degree of support from parallel verbs: 3pl *abuelen* would parallel *cuentan* 'count', *suelen* 'are used to', *huelen* 'smell', *duelen* 'are in pain', *vuelan* 'fly', and so on, while 3pl *abolen* would parallel (somewhat more distantly) *controlan* 'control', *violan* 'violate', *tremolan* 'flutter', etc. A major analytical challenge of such cases, therefore, is to explain why certain words suffer from gaps, while others are permitted to surface in one form or another. I will call this property *lexical selectivity*.

A closely related problem is the *morphological selectivity* of gaps—namely, the fact that only certain parts of the paradigm are affected. In principle, speakers could be uncertain about other properties of verbs as well, such as the conjugation class, leading to gaps in other parts of the paradigm (e.g., infinitive and 1pl, 2pl). Even restricting the discussion to stem vowel alternations, the paradigms in (8) show that in the *-ir* class of verbs, the difference between underlying /i/ vs. "raised /e/" leads to ambiguity in those forms in which the root is stressed (*v**i**vir* ∼ *v**i**ve* vs. *p**e**dir* ∼ *p**i**de*).

(8) Ambiguity in stressed forms: invariant *i* vs. raising *e* ∼ *i*

a. *vivir* 'to live'	
vív-o	*viv-**í**mos*
*vív-**e**s*	*viv-**í**s*
*vív-**e***	*vív-**e**n*

b. *pedir* 'to request'	
p[**í**]*d-o*	*p*[**e**]*d-ímos*
p[**í**]*d-es*	*p*[**e**]*d-ís*
p[**í**]*d-e*	*p*[**í**]*d-en*

In principle, this could lead to uncertainty, if a speaker has only encountered a particular verb in stressed forms before—a not unlikely scenario, given that the 3sg, 1sg, and 3pl forms are generally among the most frequent parts of the paradigm, both for adults and children (Rodríguez Bou 1952; Juilland and Chang-Rodríguez 1964; Bybee 1985, p. 71). Yet strikingly, there are no *-ir* verbs with stem vowel [i] in stressed forms, but gaps in all stressless forms (where the vowel could be either [i] or [e]). This is reminiscent of the fact noted by both Rice (2005) and McCarthy and Wolf, that even in more phonotactically motivated cases, only certain morphemes or morphological categories are affected. This is handled in their analyses by stipulating that MPARSE or MAX{mcat} be relativized to different morphological categories, and ranked lower for some than for others. Something we would like to explain here is *why* certain parts of the paradigm are affected and not others.

Another feature of the Spanish data which is not evident from grammatical descriptions is the fact that uncertainty about gapped forms is *gradient*. Albright (2003) reports results of an elicitation study showing that ratings of potentially gapped forms fall along a continuum from very certain to not at all certain. This raises an important question about the object of analysis, since when speakers' intuitions are consulted, there is no obvious watershed of uncertainty that corresponds to a criterion for declaring the verb to be "gapped'. (A similar point is made by Sims 2005 in a study of genitive plural "gaps" in Greek.) One factor that plays an obvious role is lexical frequency—unsurprisingly, the greatest uncertainty surrounds low frequency items. (In fact, many of the verbs listed as gapped by grammars have come to be used only in the infinitive or participial forms, or have switched morphological classes, or have fallen out of the language completely.) Low frequency alone is typically not sufficient to cause inflection to break down, however. In the usual case, speakers are willing to inflect even nonce ("wug") words. Thus, we must explain why it is just in these particular cases that speakers are unwilling to synthesize unknown forms.

A fact that goes hand in hand with gradient uncertainty is the relation between gaps and variability. In Albright (2003), I showed that when speakers report uncertainty about inflected forms, they also tend to produce divergent forms (i.e., one speaker says *abolen*, another says *abuelen*, but both profess uncertainty). A plot demonstrating this correlation is reproduced in Figure 1. In more naturalistic settings, this seems to translate into considerable individual variability in whether speakers are willing to utter an odd-sounding form, or whether they seek a periphrastic. In some cases, such as Sims' Greek example, an obvious alternative is available and speakers can easily avoid the gap. In the Spanish case, a more serious maneuver is required, and apparently speakers sometimes simply fill the gap—as can be seen from the significant numbers of Google hits for both *abole* and *abuele*, sometimes even on the same page.[9] A similar point is made by Rice (2003) concerning Norwegian imperatives, for which many speakers have a gap, but some speakers settle on a "solution" to avoid periphrasis. Rice treats this difference as a grammatical difference between speakers (some have a grammar that provides a pronounceable solution,

[9] E.g., the chronicles listed on http://cenaculo.org/cenaculo/ius/cuarto.htm use both forms within just a few lines of each other.

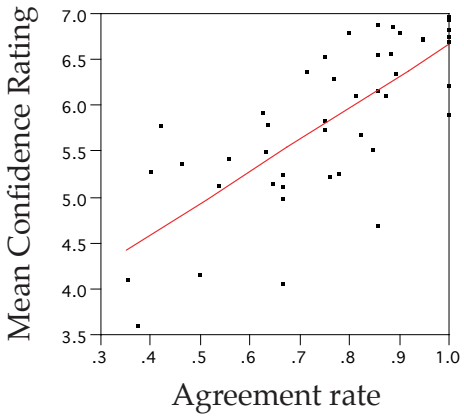

Figure 1: Relation between certainty and between-speaker agreement

others do not). The Spanish results suggest another possibility, though, which is that all speakers agree that the form is uncomfortable (i.e., there is a common grammar which does not derive any form with any degree of certainty), but speakers differ, for possibly non-grammatical reasons, about whether they are willing to venture one of the unappealing possibilities.

Finally, the fact that communal uncertainty manifests itself as variability raises perhaps the deepest and most difficult challenge in explaining lexically arbitrary paradigm gaps: frequently, uncertainty persists in spite of the fact that the "fix" is (at least somewhat) attested elsewhere in the language. This is true at several different levels. In cases where the gap affects all relevant lexical items for some speakers, others speakers are apparently uttering repaired forms (e.g., Rice's speakers who are willing to devoice to [sykl̥]) —yet these repairs are unable to propagate through the entire population. Across morphological categories, the fact that the fix is attested in one part of the grammar (e.g., for nouns) does not automatically carry over to other morphological categories. And at the level of the individual word, we see that gaps persist even though speakers are able to inflect other words, and are even occasionally willing to inflect that particular word. For some reason, neither the existence of *controla*, *tremola*, *viola*, nor the the 15,000 Google hits for *abole* itself (20 March 2006) are sufficient to guarantee that speakers are willing to accept it. We could add that the same issue is seen also from a historical point of view, since gapped forms were frequently attested as expected at earlier stages of the language. Evidently, this type of sporadic attestation is not enough to ensure the survival of a previously attested solution; gaps are *aggressive*. This property is puzzling, since the gaps themselves provide no positive evidence to "defend themselves" against the occasional data provided by speakers who are willing to utter [sykl̥] or [abole]—and, more generally, there is no possible overt evidence that such forms cannot be used. Given that there do exist small amounts of evidence in the environment *against* gaps, the challenge is to explain why these solutions do not take over. One possibility is that the level of attestation for the fixes is simply too low to be learned. An alternative possibility that I will pursue here, however, is that there is a more active force involved: gaps follow from a deeper issue, which sporadic pieces of data are not able to alleviate.

To summarize, then, lexically arbitrary paradigm gaps (and, to some extent, perhaps phonotactically motivated ones, as well) exhibit the following set of properties which must be explained: they are morphologically and even lexically selective, they are marked by gradience

and variability, they are modulated by lexical frequency, and they emerge in spite of a lack of direct, overt evidence for their existence. In the next few sections, I consider the extent to which existing proposals are able to explain these facts.

2.2 Semantic implausibility, or homophony avoidance?

Grammars and native speakers frequently rationalize non-occuring forms by declaring them to be semantic implausible, potentially homophonous with unrelated words. It is certainly true that individual cases may suffer from one or more of these problems: for example, the diphthongized candidate for the 1sg of *abolir* 'abolish' is *abuelo*, which also happens to mean 'grandfather', while the 1sg of *pacer* 'graze' would normally be said only by livestock. These factors alone are not enough to explain the gaps, however. Speakers are often comfortable forming semantically or pragmatically improbable inflected forms (*lluevo*'I am raining', *descaffeíno* 'I decaffeinate'), and they are generally not bothered by homophony (*creo* 'I believe/I create'). Many authors have commented on the fact that these types of explanations do not lend themselves to a viable and explanatory theory of paradigm gaps—see., e.g., Halle (1973) regarding Russian, and Albright (2003) for Spanish. I will not consider such explanations further here.

2.3 Blocking by null?

One approach to arbitrary gaps, put forward by Halle (1973), posits that gapped forms are individually blocked by lexically specific surface filters, which he accomplishes by marking the relevant forms as [−LEXICAL INSERTION]. This is equivalent to saying that speakers have learned that the affected form (irregularly) does not exist—that is, its phonological form is ∅—and that this listed irregular null form blocks the creation of the otherwise expected overt form. In OT, an equivalent account could be constructed by letting MPARSE constraints be relativized not only to different morphological contexts, but even to different lexical items: MPARSE$_{asir-1SG}$.

This approach runs into several problems with the Spanish data. First and most important, it raises a learnability issue. As noted above, the words most strongly affected by gaps are the lowest frequency words in the relevant classes. This is unexpected under Halle's account, since ordinarily, the words that provide the strongest evidence about their irregularity are *high* frequency words. How do speakers know that for these particular low frequency words, an unattested 1sg form is a true gap, and not an accidental gap? McCarthy and Wolf [this volume] suggest a conservative learning strategy for the ranking of MPARSE when it is relativized to specific morphological contexts: the learner assumes that all MPARSE constraints start low, and rerank to allow pronouceable fixes only in response to positive evidence about the context in question. That is, hearing the answer for one morphological context does not automatically permit the same fix in a different morphological context—each context is assumed to have gaps, until the learner receives overt evidence otherwise. The solution cannot be extended to word-specific MPARSE constraints, however, since it predicts that learners treat every inflected form of every word as a gap until hearing otherwise. This is plainly false; speakers are frequently comfortable constructing inflected forms that they have never heard before, even in cases involving less plausible forms and more frequent verbs (e.g., the 1sg of *llovar* 'rain' is widely agreed to be *lluevo*). It is not safe to assume that unattested forms reflect gaps, particularly for low frequency words, which are most susceptible to incomplete sampling. It is hard to see how the existence of gaps

could be inferred for these rare words. This is a serious problem for any account that accomplishes lexical selectivity by requiring something to be learned about each affected word.

A blocking account is also unable to explain the morphological selectivity of gaps, since in principle, the same mechanism could be used to prevent any part of any paradigm from surfacing. For example, there is no reason why just a 2sg or infinitive form could not be marked as non-occurring. The fact that there are only two gap patterns (stressed forms, and first singulars), and that these patterns apply only in the irregular conjugation classes, is left unexplained. Lexically-specific blocking is too powerful to predict where gaps actually occur. At the same time, blocking is too weak to capture the observation that gaps are gradient and lead to variability. The experimental results of Albright (2003) suggest that gaps are not an all-or-nothing phenomenon, affecting a limited class of lexical items that can be listed as exceptional.

A "blocking by null" account thus fails to capture the intuition that lexically arbitrary gaps affect potentially irregular forms of word that belong to small and highly irregular inflectional classes, and that the effect is stronger when *less* is known about the word. This is not just a peculiarity of the Spanish example; it also seems to be true for the Russian case discussed by Halle (more on this below).

2.4 A competition-based account?

As noted above, lexically arbitrary gaps frequently surround forms with irregular morphophonology. For that reason, it is tempting to suppose that they arise when the language provides two competing patterns, leaving speakers unable to decide between them. Concretely, perhaps the fact that Spanish has both diphthongizing and non-diphthongizing verbs (*vuelo* 'I fly' vs. *violo* 'I violate') creates a tie or a ranking paradox, and when faced with two equally good candidates, speakers are somehow frozen in indecision. There are several reasons to think that the failure is not as simple as "no unique winner", however.

The first obstacle to attributing gaps to indecision among multiple winners is that we would lose a plausible explanation for a different type of data, namely, grammatically licensed variation. Numerous works in OT have linked the availability of multiple surface variants to incomplete or indeterminate grammars (Boersma 1997; Boersma and Hayes 2001; Anttila 2002). In these cases, the response to multiple winning candidates is to allow *all* of them as possible outputs, rather than *none* of them. Even in cases of irregularity, variation between happily coexisting competing patterns can be observed (e.g., doublets like *shrunk* ∼ *shrank* or *pleaded* ∼ *pled* in American English; see also Zuraw (2000) for examples of competition leading to variation). It is natural to suppose that when the grammar produces two outputs, either one should be pronounceable, and the widely attested phenomenon of side-by-side variants lends support to this idea. There is a fundamental contradiction between the idea that variation should be modeling via multiple winners on the other hand, versus the idea that ties might cause paralyzing uncertainty, on the other. The use of multiple winners to model variation appears to have the upper hand, and we must look elsewhere to explain gaps.

There is yet another reason to think that irreconcilable competition between two equally good candidates is not the appropriate analysis of paradigm gaps. In all of the cases to be discussed here, gaps affect those morphological classes where the incidence of irregularity is highest, affecting not just some, but *nearly all* relevant lexical items. We see that gaps are strongest

in cases where irregularity nears 100%; this is unexpected, if gaps stem from inability to resolve competition between two evenly matched patterns.[10]

The fact that gaps can occur even the in absence of strong competition from a second pattern shows that evenly matched competition is not a necessary precondition for lexically arbitrary paradigm gaps. The fact that when evenly matched patterns occur, they sometimes lead to free variation shows that such competition is also not sufficient to cause gaps. This serves to strengthen the conclusion that there is no promising easy solution involving a contradiction between two insufficiently distinguished candidates.

Given the discussion in section 2.1 above, it must be acknowledged that empirically, variation and gaps can mimic each other to a certain extent. When there's a gap, people tend to arrive at different solutions (leading to variation in responses), and when variation is brought to people's attention, they sometimes begin to worry about which form is normatively correct (leading to reported uncertainty). I take it for granted here, however, that there truly is a distinction between peaceable free variation and uncomfortable gaps, and I conclude that the uncertainty that leads to gaps must be modeled not as competition between two equally *good* outputs, but as the attempt to pronounce one or more *bad* outputs. This conclusion is by no means novel; it is, in fact, assumed without argument in all current approaches to phonotactically motivated gaps in OT cited above, which distinguish "bad" outputs as those that lose to the null parse or are eliminated in the Control component.

2.5 Faithfulness to listed allomorphs?

A different intuition, which drives most current analyses of paradigm gaps in some fashion, is that the alternation needed to repair a markedness violation would require an illicit faithfulness violation—encoded, for example, by FAITH ≫ MPARSE (McCarthy and Wolf), or SEG FAITH ≫ MAX{mcat} (Rice). Applying this to the case of Spanish anti-stress verbs, we might try to claim that mid-vowel verbs are subject to a markedness constraint that bans stressed mid vowels (stated here, in ad hoc fashion, as *V́-MID), but that all possible fixes (diphthongization, raising, lowering) are ruled out because the relevant faithfulness constraints (UNIFORMITY, IDENT[±high], etc.) are all ranked above MPARSE:

(9) A ranking that would produce a 3pl gap

/abol-e-n/	*V́-MID	IDENT[±hi]	UNIFORMITY	MPARSE
a. abólen	*!			
b. abuélen			*!	
c. abúlen/abálen/...		*!		
☞ d. Ø				*

[10] Hansson (1999) points out that the irregular pattern may exert its influence through a different mechanism than the regular pattern (analogical vs. grammatical mechanisms), and that when irregularity nears 100%, the answers from the two mechanisms would be most at odds with each other. This is an intriguing idea, though I do not know how to reconcile it with results showing that even "irregular" patterns seem to be extended in grammar-like ways (Albright 2002b; Albright and Hayes 2003), and that words that fall under multiple strong patterns, regular and irregular, can happily take all of them: *spling* → {*splinged, splung, splang*}.

Unfortunately, the ranking in (9) predicts that no stressed forms of mid vowel verbs should ever occur. This is patently false: as we saw in (4) above, some Spanish verbs do diphthongize or raise, while others have stressed mid vowels. This leads to several apparent ranking paradoxes, created by the simultaneous existence of verbs that alternate (*V́-MID ≫ UNIFORMITY) and those that do not (UNIFORMITY ≫ *V́-MID), as well as verbs with gaps ($\mathcal{F}$ ≫ MPARSE) alongside verbs without (MPARSE ≫ $\mathcal{F}$).

One way to accommodate the side-by-side existence of alternating, non-alternating, and gapped words is to resort to lexically listed allomorphs. In particular, if we assume that the stressed and unstressed allomorphs of verbs are listed separately, then we do not need to predict mid vowel alternations at all. With the relevant faithfulness constraints (IDENT(height), IDENT(stress), UNIFORMITY) ranked above *V́-MID, the grammar simply uses whatever stressed allomorph it is given, regardless of whether it is alternating or non-alternating.[11] If a verb happens to be lacking a listed stressed allomorph, the decision falls to lower ranked constraints—and with MPARSE ranked just below $\mathcal{F}$, it is better to have a gap than to generate a novel stressed allomorph. This is essentially a lexical conservatism account (Steriade 1997), and it is parallel to one proposed by Pertsova (2005) for paradigm gaps in Russian nouns.

(10) A lexical conservatism approach

a. Listed non-alternating allomorph

/{viol,viól}-a-n/	$\mathcal{F}$	MPARSE	*V́-MID
☞ a. viólan			*!
b. viuélan	*!		
c. Θ		*!	

b. Listed diphthongizing allomorph

/{vol,vuél}-a-n/	$\mathcal{F}$	MPARSE	*V́-MID
a. vólan			*!
☞ b. vuélan			
c. Θ		*!	

c. No listed stressed allomorph

/{abol}-e-n/	$\mathcal{F}$	MPARSE	*V́-MID
a. abólen	*!		*
b. abuélen	*!		
☞ c. Θ		*	

The lexical conservatism approach provides a way to describe word-by-word differences, but the solution comes at a steep cost. The grammar in (10) prohibits speakers from projecting any alternations, essentially choking off the generative capacity of the grammar. This prediction is too extreme; in fact, Spanish speakers are often perfectly comfortable generating stressed allomorphs that they have never heard before, including both non-alternating and also diphthongizing ones (Albright, Andrade, and Hayes 2001). As Pertsova (2005) also notes, the analysis

[11] In theory, this has the potential to allow all sorts of idiosyncratic alternations beyond diphthongization and raising of mid vowels; these would have to be prevented with higher-ranked OO-Faith constraints regulating possible relations between related allomorphs.

does not give us any insight into why speakers behave so conservatively in some cases, when equivalent processes apply automatically in other cases.

In addition, this analysis does not actually do much to explain the lexical selectivity of gaps. There is no reason why verbs that are defectively lacking a stressed allomorph should be limited to the second and third conjugations—or indeed, why they should be restricted to mid vowel verbs, even. In principle, any verb could be missing a stressed allomorph, and the ranking in (10) would prevent a new one from being generated. This account also goes only partway in explaining the morphological selectivity of gaps, since there is nothing that would prevent a verb from defectively missing the *stressless* allomorph, predicting the possibility of verbs that occur only in stressed forms. It is also difficult to extend this analysis to the 1sg gaps, since this would require stipulating a separate listed allomorph for just this form, even if it is usually identical to the other stressed forms.

Although resorting to lexically listed allomorphs does not explain the particulars of the Spanish data, there are indeed reasons to think that a lexical conservatism may nonetheless play a role in shaping lexically arbitrary paradigm gaps. A possibly telling case from Modern Icelandic is documented by Hansson (1999). In Icelandic, both the weak (regular) past tense and the clipped imperative are formed by the addition of a dental -T suffix, which, depending on the preceding context, is realized either as unaspirated *-t* (further lenited to [ð] in some cases), or as aspirated *-t^h*. Abstracting away from some of the details, it can be said that the *-t^h* allomorph occurs when the preceding segment is underlyingly aspirated or when it is itself a /t/, and the unaspirated *-t*/ *-ð* allomorph occurs elsewhere (see Hansson's presentation for a fuller treatment of the distribution). When the *-t^h* allomorph is chosen, its aspiration is typically realized as preaspiration or regressive devoicing of a preceding sonorant.

(11) Icelandic past/clipped imperative formation

a. *-t^h* after aspirated and after /t/

Verb stem	Past/Imperative UR	Past stem	Imper. stem	Gloss
/t^{h}ak^h-/	/t^{h}ak^h-t^h/[12]	[t^{h}axt-]	[t^{h}axt-]	'take'
/sɪnt-/	/sɪnt-t^h-/	[sɪn̥t-]	[sɪn̥t-]	'swim'
/mɪrt-/	/mɪrt-t^h-/	[mɪr̥t-]	[mɪr̥t-]	'murder'

b. *-t* elsewhere

/sin-/	/sin-t-/	[sint-]	[sint-]	'show'
/mail-/	/mail-t-/	[mailt-]	[mailt-]	'measure'
/heir-/	/heir-t-/	[heirð-]	[heirð-]	'hear'
/hav-/	/hav-d-/	[havð-]	[havð-]	'have'

There are, however, some exceptional verbs that end in sonorants, but take the *-t^h* allomorph of the past tense (parallel to English irregular pasts like *burnt*, *dwelt*, etc.). Crucially, when the past tense of a verb takes the "wrong" allomorph of the suffix, *so does the clipped imperative*:

(12) Exceptional *-t^h* after sonorants

[12] As Hansson points out, there is, strictly speaking, no reason to call the suffix aspirated after aspirated stops (/t^{h}ak^h-t-/ and /t^{h}ak^h-t^h-/ both yield [t^{h}axt-]). I assume an "agreeing" distribution for expository ease.

/mail-/	/mail-t^{h}-/	[mail̥t-]	[mail̥t-]	‘speak’
/mein-/	/mein-t^{h}-/	[mein̥t-]	[mein̥t-]	‘mean’

The isomorphism of the past and the clipped imperative is not a historical accident; in fact, the clipped imperative is a relatively recent innovation in Icelandic, formed by reanalyzing the [θ] of the following 2sg pronoun as including a suffix ([sin θu] ‘show you’ ⇒ [sin**t** θu]). The only way for aspirated *-t^{h}* to have entered the imperative in these exceptional cases is for it to have been imported directly from the past stem. Stated informally, we could say that the past provides a -T suffixed allomorph, and the imperative parasitically employs whatever is found in the past.

The relevance of Icelandic is clear when we turn to strong verbs, which do not form their past by suffixing *-t(h)*. Here, the clipped imperative is on its own, and in most cases, it takes the contextually appropriate form of the *-t*/*-t^{h}* suffix.

(13) Clipped imperatives from strong verbs

/trak-/	trak-t-	(troː-)	[traɣð-]	‘drag’
/halt-/	/halt-t^{h}-/	(heːlt-)	[hal̥t-]	‘hold’

There is one systematic exception to this, however: verbs ending in /nn/ and /ll/ typically do not take either *-t* or *-t^{h}*, but have a gap: /vinn-/ ‘work’ (past /vann-/) → *[vɪnt-], *[vɪn̥t-]. For these verbs, a non-clipped imperative form must be used (periphrasis). What is special about verbs ending in /nn/ and /ll/, that would prevent them from taking the expected *-t* allomorph? Hansson argues that the problem stems from the fact that among weak verbs, verbs ending in these sequences *usually* take “exceptional” *-t^{h}*: [fil̥t], [hɛl̥t], [spɪl̥t], [prɪn̥t], [klɛn̥t], etc. In fact, fully 21/27 of the relevant verbs are exceptional in this way, making “regular” […lt] and […nt] the minority pattern. Hansson’s intuition is that speakers know that verbs with this shape are typically irregular, but they are unable or unwilling to assume the “expected irregular” *-t^{h}* form. Thus, they are left with an irresolvable clash between the grammatically expected regular form, and the analogically excepted irregular form. This is highly reminiscent of both the Spanish and Russian cases, in which gaps also apply specifically in those inflectional classes where irregularity is the norm.

Strikingly, there is also one exception to the exceptions in Icelandic: the verb /finn-/ ‘find’ does take the phonologically expected clipped imperative form [fint-] (with unaspirated *-t* after the sonorant). What makes /finn-/ unique, however, is that it *does* have a -t- in some past tense forms: sg. [fan], [fanst], [fan], pl. [fʏn**t**ʏm], [fʏn**t**ʏð], [fʏn**t**u]. Historically, the *-t-* of [fʏn**t**ʏm] is part of the verb stem (it is not the dental past suffix), but it appears that speakers nonetheless treat it as evidence that /finn-/ does not take -t^{h}(*[fʏn̥tʏm]). It is also important to note that the clipped imperative does not simply import wholesale the past plural allomorph /fʏnt-/; rather, it preserves the vowel of the present, and uses the past plural only for evidence about *-t*/*-t^{h}* allomorphy (a *split base* effect; Steriade 1997).

The Icelandic data show two lexical conservatism effects: first, the form of the clipped imperative depends on the past allomorph more closely than would be predicted by morpheme concatenation. In particular, there is no reason why the imperative should be bound to take exceptional *-t^{h}* when the past does ([mail̥t-], [mein̥t-]), yet speakers prefer to maintain past/imperative identity rather than creating the regularly expected imperatives [mailt], [meint]. Second, when a regular past allomorph is unavailable and there is uncertainty about whether to use *-t* or *-t^{h}*,

speakers look to whatever other allomorphs happen to be present to settle the uncertainty. In the case of [finn-], this yields a *-t* form, but for other verbs, no such help is available. Hansson (1999) claims that the analysis of the clipped imperative requires reference to past tense form, and I do not see any alternative.

Lexical conservatism alone is insufficient to provide a complete account of Icelandic, however. First, it is unable to explain why the form of the clipped imperative depends on an attested past stem, and why speakers do not have similar difficulties with the past itself. Hansson observes that *-nn* and *-ll* final verbs never have equivalent gaps in the past; presumably novel verbs take regular *-t* (or perhaps *-t^h*; Hansson does not say which). Lexical conservatism also cannot explain why the problem is confined to *-nn* and *-ll* stems, and why elsewhere, speakers are willing to create new suffixed forms that are unattested in the past. More generally, it cannot explain why speakers are sometimes perfectly happy to create new allomorphs (i.e., in cases of regular phonological alternations), and why they sometimes act conservatively. This issue is also commented on by Pertsova 2005.

What I conclude from the Icelandic example is that the ability to consult related forms as evidence for the shape of a needed allomorph is an important part of the puzzle, and that faithfulness to listed forms can serve a useful role in explaining why speakers are able to avoid gaps in cases like [fɪnt] 'find!'. Although the principle of lexical conservatism can't actually explain why in certain cases speakers are unable or unwilling to create new allomorphs, it *can* explain why speakers are sometimes able to find a solution for problems that the grammar can't adequately resolve. The overall picture that this suggests, then, is that if we can find an answer to why the grammar sometimes fails to produce a clear winner in particular morphological contexts (e.g., in producing 1sg forms in Spanish, or imperative forms in Icelandic), then faithfulness to listed allomorphs can help to explain why speakers are able to overcome the problem for certain words but not for others. In other words, the machinery of a lexical conservatism account will turn out to be useful in the final formalization, but it cannot not actually explain the core mystery of why gaps occur where they do.

The basic challenge remains, therefore, to explain why the grammar is unable to provide a default winner in cases of gaps.

3 Modeling grammatical confidence with stochastic generalizations

In section 2.1, I argued that in order to understand paradigm gaps, we must understand not only their lexical selectivity, but also their morphological selectivity—that is, why only certain parts of the paradigm are affected, while parallel structures are fixed or tolerated in other words or other morphological contexts. A key insight is that the forms that suffer paradigm gaps are the ones whose properties are determined with reference to another forms in the paradigm. This is seen very clearly in Hansson's discussion of Icelandic: the clipped imperative is identical to the past stem whenever possible, including even details such as irregular selection of *-t^h*. It is also seen, indirectly, for Spanish: stressed forms, and the 1sg in particular, are the present tense forms that have undergone the most historical restructuring (Penny 2002, p. 183), that show the most dialect variation, and that exhibit the greatest number of child errors (Clahsen, Aveledo, and Roca 2002). The fact that these forms suffer from gaps is not merely a fact about gaps; it is a consequence of the more general morphological organization of the language. If this is correct,

then the first step in understanding paradigm gaps is to understand why certain forms are built on other forms.

3.1 Asymmetric relations within the paradigm

In Albright (2002a), a model is proposed that incorporates directionality as an intrinsic part of how paradigms are learned and organized. In particular, it is hypothesized that learners seek to discover which forms in the paradigm are the most *informative*, or have the greatest predictive power for how the remainder of the paradigm is formed. These forms are then selected as base forms, from which the remaining forms are projected. In order to be predictive, a form must reveal all idiosyncratic properties of the lexical item; for example, in Spanish, a maximally predictive verb form would need to reveal the conjugation class, diphthongization and/or raising, velar insertion, and any other irregular properties that the verb may have. According to this hypothesis, grammars are organized in a way that makes use of predictability relations, by storing the less predictable (more idiosyncratic, or informative) forms, and deriving more predictable forms from them.[13]

In the case of Spanish, the most pervasive and unpredictable property of verbs is their conjugation class. This can be seen intuitively by inspecting the paradigms in (3)-(5), and observing that every single verb belongs to a conjugation class, while verbs in only certain classes and with certain phonological shapes display alternations like diphthongization or velar insertion. Therefore, the optimal form to memorize in the paradigm is one that clearly reveals conjugation class, such as the infinitive, 1pl, or 2pl, and use this as a base for deriving the remainder of the paradigm. Assuming for the moment that the infinitive is chosen, this makes it very easy to derive the 1pl and 2pl, since for all but a handful of extremely irregular verbs, they simply involve changing the endings. For the stressed forms, however, things are somewhat more difficult, since the forms that most clearly reveal conjugation class also systematically lack information about mid vowel alternations (which occur only under stress), or about velar insertion (which occurs only in the 1sg). For the most part, once a stressed form is known, however, it is straightforward to derive the others—for example, starting with the 3sg, the 2sg and 3pl can virtually always be derived by adding *-s* or *-n*. The only remaining ambiguity concerns the 1sg, where velar insertion sometimes creates non-identity with the other stressed forms. This makes the 1sg slightly unpredictable, even given the other stressed forms. In general, this type of reasoning establishes a series of unidirectional relations among paradigm members; a possible end result of this learning procedure for Spanish in particular is shown in Figure 2. The figure shows that stressed forms are derived from stressless ones (without mid vowel alternations), and the 1sg is derived from a form that shares vowel alternations, but lacks velar insertion. For details of the learning algorithm, and the computer simulations that were used to derive this result from a lexicon of Spanish verbs, the reader is referred to Albright 2002a, chap. 6.

What is crucial about this model is that it establishes asymmetrical dependencies between forms in the paradigm, on a language-particular basis. A model of this sort has the potential to explain the morphological selectivity of paradigm gaps, in the following way: building on

[13]This is clearly related to the traditional idea in generative phonology that all lexically arbitrary information is memorized as part of the underlying form, which serves as the input to grammar. The chief difference is that in this model, the input consists of a *surface* form, which cannot combine unpredictable information from multiple parts of the paradigm; see (Albright 2002a) for discussion.

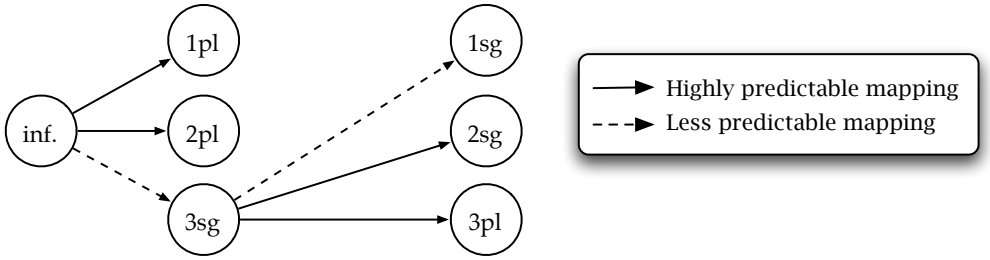

Figure 2: Directionality of mappings in the Spanish present tense paradigm

the observation that gaps occur in forms that are computed with reference to other parts of the paradigm, we can hypothesize that gaps occur when something about the mapping from a base to a derived form fails. This is an important first step in limiting the sets of forms that can potentially suffer from gaps: gaps may affect leaf nodes in a derivation independently (the 1pl, the 2pl, the 1sg, the 2sg, or the 3pl), or sets of forms that are all derived from the same non-terminal node (e.g., all of the 3sg, 1sg, 2sg, and 3pl). The next step, then, is to understand why mappings fail only for the 1sg and 3sg, and only for particular words in particular inflectional classes.

3.2 Confidence of mappings

The structure in Figure 2 posits six different mappings (infinitive → 1pl, infinitive → 2pl, etc.). Of these, four are nearly 100% predictable across all verbs and all classes:

(14) Virtually exceptionless mappings in Spanish

infinitive → 1pl	-r → -mos
infinitive → 2pl	-r → -is
3sg → 2sg	∅ → -s
3sg → 3pl	∅ → -n

Looking back to the data in (2.1), we see that the two remaining mappings (infinitive → 3sg and 3sg → 1sg) are also the ones that involve the greatest degree of unpredictability, because of the unpredictability of mid vowel alternations and velar insertion. (This is reflected in the dashed lines in Figure 2.) The challenge is to understand why unpredictability leads to uncertainty in just a certain subset of 1sg and 3sg forms.

Let us consider first the mid vowel alternations. Here, it is instructive to compare class 1 (*-ar*) with class 3 (*-ir*). In class 1, there are many verbs that diphthongize (e.g., *sentar*, *contar*; (4a)), and many verbs that do not (e.g., *rentar*, *montar*; (4c)). Table 1 shows counts of verbs from the LEXESP corpus (Sebastián, Cuetos, Martí, and Carreiras 2000), categorized according to their "dictionary" alternation pattern (de Gámez 1973; Butt 1997). The counts reveal that overall, mid vowels tend not to alternate in this class (91% of /e/, 89% of /o/). Verbs in this class never suffer from gaps, and when speakers need to project stressed forms, they generally select non-alternating mid vowels as a default. This can be seen in several ways: first, when speakers are given a wug test in which they hear novel class 1 verbs in the infinitive and must produce 1sg forms, they confidently and overwhelmingly choose [é] and [ó]—e.g., novel *rempar*

Table 1: Counts of mid vowel alternations in classes 1 and 3

Unstressed	/e/			/o/		Total
Stressed	é	jé	í	ó	wé	
Class 1 (*-ar*)	916	93	0	588	71	1668
Class 3 (*-ir*)	2	32	42	0[15]	3	79

→ 1sg *rempo* preferred over *riempo* (Albright, Andrade, and Hayes 2001). In addition, class 1 verbs have historically tended to lose diphthongization alternations: older *apuerta*, *entriega*, *confuerta* ⇒ newer *aporta*, *entrega*, *conforta* (Penny 2002, p. 183; Morris 2005). If we view such leveling as the replacement of older, irregular forms with newer, grammatically prefered forms (Paul 1920; Kiparsky 1978; Albright 2002b, and many others), then this fact indicates that in class 1, diphthongization is irregular, and non-alternation is generally the grammatically preferred option.[14] This seems unsurprising, given the statistical predominance of non-alternating forms, both in class 1, and in the language taken as a whole. The ability to productively produce non-alternating forms is significant, however, because it shows that in general, there is no reluctance to produce novel stressed forms (i.e., lexical conservatism does not hold here).

Of course, a grammar that always predicts non-alternation and never allows diphthongization cannot capture the full set of attested data, since many words do diphthongize. One possibility is that every case of diphthongization is simply memorized as exceptional irregularity; Bybee and Pardo (1981), Eddington (1998), and Clahsen, Aveledo, and Roca (2002) all this claim, in one way or another. In the model depicted in Figure 2, this could be captured by saying that speakers have stored an irregular diphthongized 3sg form, which simultaneously blocks the default non-alternating mapping, and at the same time provides a diphthongized input for further mappings to the 1sg and 2sg.

There is evidence that such a model is too simple, however. Albright, Andrade, and Hayes (2001) showed that in addition to the general preference for non-alternation, speakers also have knowledge about particular segmental contexts that tend to favor or disfavor diphthongization. For example, all class 1 verbs that contain *e* before *rr* diphthongize,[16] while class 1 roots containing *e* before *ch* never diphthongize. When speakers are presented with novel verbs containing these sequences, their likelihood to select diphthongization varies, depending on its attested rate for that context in the lexicon. Albright et al. interpret this result in terms of a grammar with multiple competing rules, at varying degrees of generality. In other words, the mapping from infinitive to 3sg is not solely an identity map, but rather, contains numerous sub-rules, including extremely general ones (unstressed *e* maps to stressed *ê*) and also very specific ones (*e* → *jé* / ___ *rr*, *e* → *é* / ___ *ch*). These rules compete according to the degree of support they get from the lexicon. A well-supported rule is one which applies with high accuracy (or *reliability*), meaning that it should be exceptionless, or close to exceptionless. A well-supported rule should also be well-attested; generalizations that are based on just a few forms (types) are not terribly impressive, even if perfect. Albright and Hayes (2002) propose a rule evaluation metric called *confidence*, which uses lower confidence limit statistics to combine the reliability and amount of relevant data into a single measure. As can be seen in Figure 3, reliability is the most important determinant of confidence. However, confidence also decreases considerably when there is very

[14] See also Clahsen, Aveledo, and Roca (2002) for the same claim, made on the basis of data from children's errors.

[16] This includes verbs built from several different roots: *cerr-* 'close', *err-* 'wander', *ferr-* 'grasp', *serr-* 'saw', *terr-* 'bury', *terr-* 'terrify'.

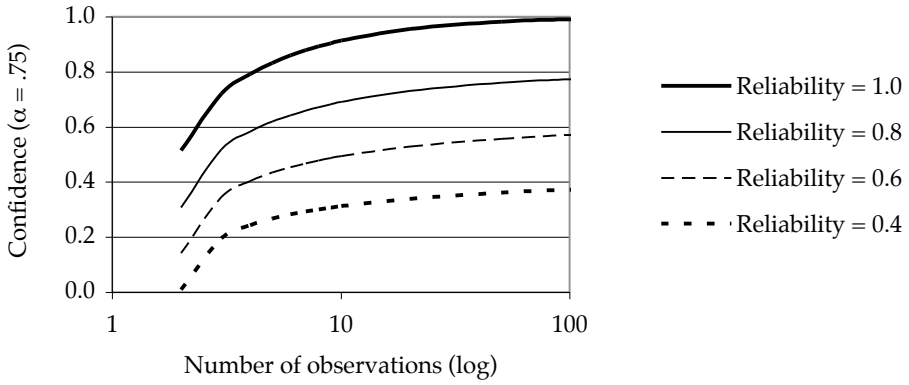

Figure 3: Confidence as a function of size and accuracy

little data to go on—especially when there are fewer than ten observations.

In most cases, phonological and morphological processes are seen in more than just a few forms, so the confidence adjustment in Figure 3 is felt only in very small, irregular classes. This turns out to be crucial in the comparison between Spanish class 1 and class 3. Looking back at Table 1, we see that unlike class 1, in class 3 virtually all verbs with mid vowels alternate. For the back vowel /o/, all existing verbs diphthongize, but in fact there are so few of them that it is difficult to form any real generalization. This is similar to the paucity of data seen in the English example in the introduction, and this observation forms the core of the current proposal: lexically arbitrary gaps occur in just those cases where there is too little data to be sure about any of the available generalizations. Concretely, there is no *o*→*ó* rule that would map *abolír* to 3sg *abóle* (no non-alternation in class 3), while the *o*→*ué* rule that would map *abolír* to *abuéle* is supported by really just two verbs (*dormir* 'sleep', *morir* 'die'; a prefixed form *premorir* 'predecease' also occurs in LEXESP). For the front vowel /e/, there are relatively more cases, but they are evenly split between diphthongization (32 verbs) and raising (24 verbs). If we look at the curves in Figure 3, we see that these cases are not as different from one another as they might seem: a generalization that works perfectly for 2/2 or 3/3 cases receives a confidence value of around 0.6–0.7 (left edge of the top line), while a generalization that works for 32/56 cases has a confidence value of just under 0.6 (third line down, towards the right of the graph). Thus, generalizations covering both /e/ and /o/ verbs in class 3 suffer from low confidence. Well attested forms like *duerme* 'sleeps' or *muere* 'dies' can survive by lexical listing of the stressed form, but for rarer words like *abolir*, a listed form is likely to be unavailable and the grammar is unable to synthesize a new form.

In order to make use of the observation that there are very few mid-vowel verbs in the third conjugation, it is crucial that the grammar of vowel alternations be subdivided into separate generalizations, depending on the particular vowel involved (*e* vs. *o*), and also the particular conjugation class (class 1 vs. that in class 3). This may seem counterintuitive, since there is no obvious formal reason why, for example, diphthongization in class 1 should be treated as a separate process from diphthongization in class 3. There are various reasons to think that native speakers do treat them separately, however. First, we can observe from Table 1 that the overall rate of alternation does differ substantially between conjugation classes in Spanish, due partly to the fact that many class 1 verbs have regularized (become non-alternating) over time. The fact that class 2 and 3 verbs have not also done this may indicate that the pressure to diph-

thongize is not uniform, but rather is assessed separately in different classes. Furthermore, the segmental contexts that encourage diphthongization differ not only from class to class, but also between /e/ and /o/; Albright, Andrade, and Hayes (2001) found that the best model of likelihood to diphthongize is one that considers diphthongization contexts independently for each vowel and each class. A similar finding is presented by Eddington (1996), who shows that the rate of diphthongization differs substantially across different derivational suffixes, and that speakers respect these differences when producing nonce forms. This fracturing of the grammar for separate segments and separate contexts also mirrors a conclusion of Rice and also McCarthy and Wolf (both this volume), in which rankings that license alternations in one morphological context are not automatically transferred to another morphological context.[17] These results provide converging evidence that the segmental environments and overall rates of alternations are at least sometimes calculated separately for different segments and different morphological contexts.

It is natural to wonder why generalizations should be so fragmented here, when in other cases, speakers *do* form generalizations that cover different morphological contexts, and multiple segments. For example, Kiparsky (1965) discusses the innovation of a new umlaut correspondence in the Kesswil dialect of Swiss German, which could not have arisen unless speakers treated umlaut as a unified process across different segments. I conjecture that in the case of Spanish mid vowel alternations, the division of the grammar into separate generalizations for different segments and different classes may be due to the fact that their statistics simply are so different across these different contexts, and, for accidental historical reasons, always have been. For a learner which is seeking the most accurate description of the conditions under which an alternation should occur, it would not be advantageous to throw away information about the particular vowel or conjugation class, since generalizations that ignore this information provide a worse fit to the data. This effect is likely to be greatest in inflectional classes that contain rather small numbers of words; if Spanish classes 2 and 3 had thousands of words each, then the chances of accidental statistical differences between them would be much smaller.

Another question that arises is how much data is enough to support a trustable generalization. The particular confidence limit function shown in Figure 3 suggests that confidence may fall perilously low under only the most extreme of circumstances—a common rule of thumb is a dozen observations or less. (The exact slope of the penalty depends on the particular value of alpha that is chosen—I have shown here a somewhat liberal tolerance of 25% uncertainty.) As the comparison between the class 3 front and back vowels above shows, the claim of this adjustment is that a perfect but poorly attested generalization is no better than a better-attested but exceptionful one. The rationale is that when we have so little data, there is a good chance that additional data might in fact, reveal exceptions. Low confidence can also arise when there is more data, but it is hopelessly conflicting (the lower lines towards the right of Figure 3. This, too, is only likely to happen in smallish inflectional classes, however, for the following reasons: first, it rarely happens that a large class of words is evenly split between two conflicting patterns (one almost always predominates, as we seen in class 1 in Spanish). Second, if a large class does happen to be split evenly between two patterns, then this would constitute a major source of unpredictability in the language, and, for reasons outlined above, the grammar would be be attempting to derive it in the first place. (A form that reveals it would be selected as basic.) Thus, by

[17] In fact, the Swedish example of neuter adjective gaps discussed at the end of McCarthy and Wolf's paper also demonstrates differential treatment of different segments; Iverson (1981) shows that the ban on /dd-t/ coalescence that blocks **rädd-t* [rɛtː] 'scared-NEUT.SG.' does not apply to /tt-t/: *rätt-t* 'right-NEUT.SG' surfaces fine as [rɛtː].

attributing lexically arbitrary gaps to low confidence, we derive a strong prediction: gaps should only occur in small and irregular inflectional classes, in which the number of relevant attested (non-gapped) forms is below some critical threshhold (roughly a dozen to twenty forms?), or the attested cases are extremely inconsistent. This correctly characterizes Spanish class 3, and captures the intuition that there is a deep connection between gaps and morphological irregularity.

3.3 Local summary

In this section, I have outlined a grammatical model that attempts to predict where and when gaps can occur. The features of the model which are essential to the analysis of gaps are: (1) morphological relations are directional, (2) the direction of relations is not universal, but is established by learners in an attempt to find the most reliable or accurate mappings from one form to another, (3) mappings are probabilistic, and are evaluated in terms of "confidence" values, (4) mappings are also established rather locally, such that generalizations are established independently across different segments and different inflectional classes, and (5) when there is too little data about how a particular segment should behave within a certain inflectional class, the resulting generalizations may have such low confidence as to be untrustable. Naturally, competition between different patterns can compound the problem of untrustable generalizations, but it is also not necessary; extreme paucity of data is by itself sufficient to create low confidence. When speakers lack a good enough grammatical mapping to derive a form confidently, they are forced to resort to an "irregular" listed form if it exists, and otherwise, they have a gap.

In the final section, I discuss similar data from Russian, sketch how the analysis of these cases could be recast in OT terms, and consider the relation between these cases and the phonotactically motivated cases that have been discussed in the literature.

4 Extensions, and discussion

4.1 Russian 1sg present forms

One of the most famous cases of lexically arbitrary paradigm gaps occurs in the 1sg present forms in Russian (Shvedova 1970, pp. 412-413; Halle 1973, pp. 7-8; Hetzron 1975, p. 861). In particular, 1sg forms are avoided in certain second conjugation verbs ending in *t*, *d*, *s*, *z*, where consonant alternations are expected to occur. Crucially, however, there exist parallel forms in which alternations occur as expected:

(15) Alternations in 1sg present forms in Russian

a. Existing (non-gapped) forms (inf., 1sg)

šutít'	šučú	'jest'
budít'	bužú	'waken'
gasít'	gašú	'strike'
podvozít'	podvožú	'haul'

b. Missing (gapped) forms

mutít'	*muču	'stir up'
pobedít'	*pobežu	'beat'
dubásit'	*dubašu	'batter'
lázit'	*lažu	'climb'

As argued by Halle, Hetzron, and others, there is no obvious semantic reason why the 1sg should be missing for verbs like *mutit'* but not *šutit'*. There is also no clear phonological reason why non-occurring **muču* would be any worse than occurring *šuču*. As above, the challenge is to explain why structures that are acceptable in some words are not acceptable in others.

The Russian 1sg gaps show several obvious commonalities with Spanish. First, they are morphologically selective: they affect only 1sg present forms. There is also a systematicity to the particular lexical items that are selected: they all belong to the second conjugation and end in coronals. It is surely not a coincidence that this class of verbs always undergoes alternations in the 1sg—that is, that gaps are associated with morphophonological alternations, and fall in a highly irregular inflectional class. Finally, although I do not know of any data concerning gradience or variability of gapped forms, some of the forms listed by Shvedova appear to be less problematic for speakers than others. (The form *lažu*, in particular, seems to occur not infrequently.)

As discussed above, Halle's analysis of these cases is to prevent gapped forms from surfacing by marking them as [−LEXICAL INSERTION]. Just as for Spanish, this approach is too powerful for the Russian data: it does not explain why gaps are limited to just those cases where they actually occur. If the analysis presented in the previous section is on the right track, it should predict the possibility of gaps in precisely these forms, and no others.

In order to understand why paradigm gaps target the 1sg in Russian, we must briefly consider the major sources of unpredictability in verbal inflection (see, e.g., Garde 1998, pp. 306-313 and 344-392 for a comprehensive overview). Russian verbs are traditionally classified into two conjugation classes, which, as in Spanish, differ in their theme vowel (class 1 = *e*, class 2 = *i*) (Jakobson 1948; Garde 1998). The different between these two classes is most clearly revealed in present tense forms other than the 1sg, but can usually also be inferred from the infinitive. The 1sg suffix of both conjugation classes is *-u*, reminiscent of the invariant 1sg suffix *-o* in Spanish.

In addition to theme vowels, there are also numerous alternations within the verb stems itself. Frequently, the form of the stem that occurs in the infinitive differs from that found in present tense forms:

(16) Russian infinitive ∼ present tense differences

Infinitive	3pl	Gloss
ž-ít'	*ži**v**-út*	'live'
ž-át'	*ž**m**-út*	'cut'
st-át'	*stá**n**-ut*	'become'

Within the present tense paradigm, various suffixes induce stem-final alternations. All person suffixes except the 1sg and 3pl induce palatalization: *n'es-ú*, *n'e**s'**-óš*, *n'e**s'**-ót*, *n'e**s'**-ót'e*, *n'es-út* 'carry'; this has the effect of adding secondary palatalization to non-velar consonants (e.g., *p*, *b*, *v*, *t*, *d*, *s*, *z*, *n*, *r*, *l*), and fronting velars (*k* → *č*, *g* → *ž*). In the second conjugation class, the 1sg also induces an additional set of changes, illustrated in (17):

(17) Palatalization in the 1sg

Alternation	Infinitive	1sg	2sg	Gloss
p, b, f, m ∼ pl', bl', fl', ml'	*l'ub'-ít*	*l'u**bl'**-ú*	*l'ub'-iš*	'love'
s, z ∼ š, ž	*pros'-ít'*	*pro**š**-ú*	*prós'-iš*	'ask'
t ∼ č	*trát'-it'*	*tra**č**-ú*	*trát'-iš*	'waste'
t ∼ šč	*zapr'et'-ít'*	*zapr'e**šč**-ú*	*zapr'et'-íš*	'prohibit'
d ∼ ž	*sl'ed-ít'*	*sl'e**ž**-ú*	*sl'ed-íš*	'follow'

Another alternation that can be observed in (17) is that in some verbs, stress differs between the 1sg (falling on the suffix) and the remainder of the present tense paradigm (falling on the root).

The important question for the purpose of predicting paradigm gaps is how these unpredictable properties influence the organization of the Russian paradigm. Intuitively, it is evident that the 1sg is the least predictive of all forms in the paradigm: it does not clearly reveal the conjugation class of the verb, it suffers from neutralizations caused by the alternations in (17), and it frequently differs from the remaining forms in the location of stress. Although computational simulations deriving the full predicted organization of Russian paradigms is beyond the scope of this paper, what is important for present purposes is that the 1sg has a very similar status to that in Spanish: it is predicted to be a derived form, and therefore it is generated based on some other form in the paradigm which does not show the palatalization alternations in (17). This establishes the correct directionality to predict 1sg gaps, and predicts that speakers may be uncertain about the stress and consonant alternations that apply in the 1sg.

The next question concerns the lexical selectivity of the gaps: why are second conjugation verbs ending in coronals specifically targeted? The restriction to the second conjugation is not mysterious; this is the class that exhibits the alternations in (17). What is puzzling is why there should ever be any doubt about whether the palatalization alternations in (17) should apply, since they always apply in this class. It is also mysterious why only the coronals are affected, since labials also alternate. This situation is in some ways reminiscent of Spanish, in which diphthongization does not confidently apply to class 3 verbs in *-o*, in spite of the fact that all attested expamples do undergo diphthongization. In that case, we hypothesized that although the data was consistent, there were too few examples to support a trustable diphthongization rule. Could it be that Russian second conjugation coronal-final verbs are equally sparse?

In order to answer this question, I started with a database of all of the second conjugation verbs that occur in a comprehensive grammatical dictionary of Russian (Zalizniak 1977[18]), and filtered out all duplicate entries and verbs that did not occur at a rate of at least one instance per million in a frequency corpus of 40M words (Sharoff 2002[19]). Overall, there are quite a few second conjugation verbs ending in all consonants—e.g., 66 ending in /d/, approximately 30 each of /t/ and /s/, and so on, casting doubt on a paucity of data account. Importantly, however, a large number of these verbs are derived from just a few verb roots. For example, among the 46 /d/-final verbs with alternating stress, only 7 unique verb roots are represented;[20] the rest are derived by prefixation, mostly from the two common verb roots *xod'-it'* 'go' and *vod'-it'* 'lead'. Additional prefixed forms of the same verb root arguably do not provide speakers with

[18]Available for download in electronic form at: http://starling.rinet.ru/download.htm

[19]Available for download at: http://www.artint.ru/projects/frqlist/frqlist-en.asp

[20]These are: *xod'-it'* 'go', *brod'-it'* 'work', *vod'-it'* 'lead', *bud'-it'* 'wake', *sad'-it'* 'put', *ud'-it'* 'fish', and *s'erd'-it'* 'anger'.

Table 2: Russian second conjugation verb roots, by final consonant and stress pattern

	Stress fixed	Stress alternating	Total	% alternating
p	12	2	14	14%
b	10	0	10	0%
m	9	2	11	18%
v	12	2	14	14%
t	19	10	29	34%
d	19	7	26	27%
s	7	3	10	30%
z	9	1	10	10%

additional information about how words of that shape and inflectional class behave in general, meaning that raw counts of verbs are likely to greatly overestimate the amount of data available to learners.

In order to get a truer estimate of the amount of data that learners have about each class, I removed from the dataset set all verbs that transparently consisted of a prefix + independently occurring verb root (e.g., *p'er'e-vod'it'* 'transfer' from *vod'it'* 'lead'). This eliminated 75 of the 199 verbs in the dataset that end in one of the mutable consonants in (17). We must also keep in mind the fact that Russian verbs also belong to different stress types, meaning that generalizations about final consonants may well be learned separately for each stress type. (This is parallel to the idea that Spanish speakers maintain separate generalizations mid-vowels for different vowels and conjugation classes, for which evidence was presented in the previous section.) The counts for the Russian second conjugation, divided up by consonant and by stress pattern (fixed vs. alternating), are given in Table 2.[21]

The table reveals two important facts: first, most cells have just a few examples (a dozen or less), making this a prime area for potential gaps. Interestingly, however, there is also a noticeable interaction between place of articulation and stress pattern, such that labial-final verbs consistently prefer not to alternate, whereas coronal-final verbs tend to be much more divided (in differing proportions from segment to segment). In the previous section, I suggested that one factor that may discourage learners from generalizing across multiple segments and multiple contexts may be differences in the rate of alternation; if this is on the right track, it would imply that labials form a coherent class (all show roughly the same degree of preference), while coronal-final verbs are less consistent (multiple stress patterns, different alternation types, and in some cases, multiple competing outputs). Therefore, although coronals are overall better represented than labials, they are also more fragmented. As with Spanish, this observation appears to provide the needed insight into why gaps occur specifically in this conjugation class, and to verbs with this specific phonological shape.

Clearly, these results for Russian would benefit from further empirical work quantifying the reluctance to produce forms of different shapes and across different conjugations, as well as modeling work confirming the directionality within the present tense paradigm and the effect of different stress patterns and alternation classes on predicted confidence. Nonetheless, the

[21] It should be noted that the Zalizniak dictionary does not indicate separate inflection patterns for verbs with gaps; therefore, counts from this data set are likely to be slightly inflated, since they include verbs that are claimed elsewhere to include gaps, and therefore could not provide speakers with evidence about alternations.

data does seem to support the basic predictions of the current approach to predicting when and where gaps will occur: namely, they should occur in derived forms (those with weak predictive power), and they should affect inflectional classes that are small and fragmented by various irregularities.

4.2 Sketch of an OT formalization

The focus of the preceding sections has been on identifying principles that predict when and where lexically arbitrary gaps may occur. The crucial elements appear to be the following:

(18) Components of a explanatory theory of lexically arbitrary gaps

1. A language-particular set of asymmetric relations between members of the paradigm (more predictable forms based on less predictable forms)
2. A conservative strategy for learning those relations, that generalizes across different segments and different morphological contexts only in case the evidence suggests that they do indeed pattern identically
3. An ability to assess strength of generalizations that relies not only on the accuracy of the generalization, but also the number of examples
4. Word-specific (lexicalized) forms can be used, even when the grammar itself is uncertain

The requirement that forms in the paradigm be computed with respect to particular other forms is reminiscent of base-prioritizing or transderivational correspondence (Benua 1997; Kenstowicz 1997; Kager 1999). In particular, in the case of Spanish, we need to assume that the 3sg is based on the infinitive; I will assume Benua's TCT formalism, in which the grammar first generates the infinitive, and then generates the 3sg with reference to the previously derived infinitive form.

In addition to asymmetric relations, we need a way of allowing word-by-word differences, while at the same time generating predictions for unknown or novel forms. I will follow a proposal by Zuraw (2000) for handling such blocking effects in phonology, by employing a highly ranked USELISTED constraint, which demands that if a lexically listed form exists, it must be used. (If no form is listed, then USELISTED is vacuously satisfied.) Zuraw actually proposes that USELISTED consists of a set of constraints that refer to the *strength* of listing, so better-instantiated forms have stronger blocking power; this would be crucial in modeling gradient gaps in medium-to-low frequency items, but I will abstract away from it in the current analysis.

The last thing we need is a way to state the relation between the infinitive and the 3sg. As the forms in (3) above show, the relation between the infinitive and the 3sg sometimes involves simply stressing the root, and sometimes involves diphthongization or raising. The data from Albright, Andrade, and Hayes (2001) suggest that this relation actually consists of numerous independ statements about these processes in different segmental contexts, such as diphthongization of *e* → *jé* before *rr*, retaining the monophthong before *ch*, etc. Thus, we need to be able to state a relation that is more than just an identity map, and which involves several different changes, with different probabilities in different contexts. These relations are very naturally stated in rule terms (*e* → *jé* / ___*rr*), and a straightforward way to accommodate them in OT is

by using constraints that essentially recode the relevant rules in the form of anticorrespondence constraints (Hayes 1999). (Note that the "anti-" part of anti-correspondence is not crucial; in some cases, the constraints may actually require identity, but this is just a special case of a defined relation between two surface forms.)

(19) (Anti)correspondence constraints relating Spanish infinitive and 3sg forms

a. *err*]$_{\text{Class 1 infinitive}}$ → *jérr*]$_{\text{3sg}}$

b. *et͡ʃ*]$_{\text{Class 1 infinitive}}$ → *ét͡ʃ*]$_{\text{3sg}}$

c. *e* X]$_{\text{Class 1 infinitive}}$ → *é* X]$_{\text{3sg}}$ (default non-alternation)

There are various ways to formalize such constraints, but for present purposes, the following definition is easiest to work with: the constraint *err*]$_{\text{Class 1 inf.}}$ → *jérr*]$_{\text{3sg}}$ means that if the infinitive of a class 1 verb ends in *-err*, the 3sg must end with *-jerr*. Such constraints are in many cases quite arbitrary, and language-specific. It seems likely that they are learned inductively from the data of Spanish, and ranked stochastically in order to produce the gradient and variable pattern that is observed for nonce words; see Albright and Hayes (in press) for a preliminary proposal for how this might be done.

So far, then, we have a grammar of the following form: high-ranking USELISTED requires that lexically listed forms be used when available; otherwise, the choice of the 3sg form falls to a set of stochastically ranked constraints describing the mapping between the infinitive form and the 3sg. This is shown for both diphthongizing and non-diphthongizing 1st conjugation forms in (20). For ease of exposition, I simply show the most general constraints demanding alternation and non-alternation for *o* in class 1. Mapping constraints are abbreviated by what class they refer to (class 1 or class 3), and they specify how the final vowel of the root in the infinitive is mapped to the corresponding vowel in the 3sg (e.g., Class 1: *o*→*ó*).

(20) Lexically arbitrary diphthongization

a. *vuela* 'fly-3sg'

/FLY-inf/ (listed: *volar*)	USELISTED	Class 1: *o*→*ó*	Class 1: *o*→*ué*
☞ a. volár			
b. vuelár	*!		

/FLY-3sg/ (listed: *vuéla*)	USELISTED	Class 1: *o*→*ó*	Class 1: *o*→*ué*
a. vóla	*!		*
☞ b. vuéla		*	

b. *viola* 'violate-3sg'

/VIOLATE-inf/ (listed: *violar*)	USELISTED	Class 1: *o*→*ó*	Class 1: *o*→*ué*
☞ a. violár			
b. viuelár	*!		

/VIOLATE-3sg/ (listed: none)	USELISTED	Class 1: *o*→*ó*	Class 1: *o*→*ué*
☞ a. vióla			*
b. viuéla		*!	

Next, we need a way of capturing the fact that there is sometimes no trustable mapping from the infinitive to 3sg. In the analysis sketched above, this was stated by giving the mapping a low confidence value. As the literature on gaps in OT has pointed out repeatedly, however, this is difficult to translate into a competition-based framework in which the best available candidate is chosen as optimal. As suggested already in section 2.5, I will follow Prince and Smolensky (2004), Raffelsiefen (2004), McCarthy and Wolf, and others, in formalizing gaps as selection of a null output (Ø). The real issue at hand is what motivates the selection of the null parse in the case of lexically arbitrary gaps. In general, the strategy for deriving phonotactically motivated paradigm gaps is to identify the constraints that remain above MPARSE and can eliminate the candidates that would otherwise be expected to win. As argued above, in the case of **abole*/**abuele*, this is unlikely to be a general markedness or faithfulness constraint of the language—could it be the infinitive → 3sg mapping constraints?

One way to accomplish this would be to introduce MPARSE, and to stipulate that the mapping constraints for class 1 are ranked lower than MPARSE, while the mapping constraints for class 3 are above it. This has the effect of requiring an overt output with the default vowel for class 1 verbs ((21)), but prefering the null parse for class 3 verbs ((22)).

(21) *viola* 'violate-3sg'

/VIOLATE-inf/ (listed: *violar*)	USELISTED	Class 3: *o*→*ó*	Class 3 *o*→*ué*	MPARSE	Class 1: *o*→*ó*	Class 1: *o*→*ué*
☞ a. violár						
b. viuelár	*!					

/VIOLATE-3sg/ (listed: none)	USELISTED	Class 3: *o*→*ó*	Class 3 *o*→*ué*	MPARSE	Class 1: *o*→*ó*	Class 1: *o*→*ué*
☞ a. vióla						*
b. viuéla					*!	
c. Ø				*!		

(22) **abole*/**abuele* 'abolish-3sg'

/ABOLISH-inf/ (listed: *abolir*)	USELISTED	Class 3: *o*→*ó*	Class 3 *o*→*ué*	MPARSE	Class 1: *o*→*ó*	Class 1: *o*→*ué*
☞ a. abolír						
b. abuelír	*!					

/ABOLISH-3sg/ (listed: none)	USELISTED	Class 3: *o*→*ó*	Class 3 *o*→*ué*	MPARSE	Class 1: *o*→*ó*	Class 1: *o*→*ué*
a. abóle			*!			
b. abuéle		*!				
☞ c. Ø				*		

This system correctly describes the basic pattern in a more restrictive fashion than the lexical conservatism analysis outlined in (10), but it raises a learnability issue. Given that the mapping constraints employed here must be learned inductively from the data and added to the constraint set gradually, is there a ranking procedure would be able to arrive at the ranking in (21)–(22)? Note that in this ranking, the constraints that refer to class 3 must be ranked above

MPARSE, in spite of the fact that this class provides very little positive data. This is consistent with the claim that MPARSE must be ranked low by default—class 1 provides sufficient evidence to rerank MPARSE ≫ Anticorrespondence, but class 3 does not, leaving gaps.

Although this scenario fits well with the general OT approach to poverty of the stimulus issues, there are also reasons to think that inductively learned language-particular constraints start low and rise only in response to abundant positive evidence. In brief, the quandary is that small-scale, inductively learned constraints are frequently accidentally exceptionless in the data, but this is not enough to guarantee their productivity in the grammar. A ranking procedure that is biased to put small-scale idiosyncratic constraints high in the ranking often overestimates the role that they should play in the final grammar. Boersma (1997) proposes that such constraints should actually start at the bottom of the grammar, and Albright and Hayes (in press) propose a ranking procedure using the Gradual Learning Algorithm that imposes a bias *against* small-scale inductively learned constraints. In the present case, this would favor an analysis in which all mapping constraints start out extremely low in the grammar, and class 1 constraints rise farther than class 3 constraints do.

Interestingly, it is, in fact, possible to flip the constraint ranking to be consistent with this "climb from the bottom" approach to ranking inductively learned constraints, to arrive at a ranking in which better-supported constraints are ranked highest and unsupported constraints stay at the bottom. The intuition beyond this approach is that better supported constraints clamor more strongly to have their mappings applied; small, barely attested classes do not pull much weight in enforcing their mappings. In order to implement this, however, we need a force that counteracts small-scale mappings and demands that their mappings *not* be used. That is, instead of MPARSE, we need a constraint that militates against generating previously unheard outputs: this could be a version of *STRUC (Prince and Smolensky 2004; Zoll 1996), or more simply, a constraint against novel formations: LEX (Steriade 1997).[22] In addition, we must now assume that the null parse violates the mapping constraints; this is actually more consistent with their definition in any event, since the null parse does not contain the corresponding material that the mapping constraint requires.

(23) *viola* 'violate-3sg'

/VIOLATE-inf/ (listed: *violar*)	USELISTED	Class 1: *o*→*ó*	Class 1 *o*→*ué*	LEX	Class 3: *o*→*ó*	Class 3: *o*→*ué*
☞ a. violár						
b. viuelár	*!					

/VIOLATE-3sg/ (listed: none)	USELISTED	Class 1: *o*→*ó*	Class 1 *o*→*ué*	LEX	Class 3: *o*→*ó*	Class 3: *o*→*ué*
☞ a. vióla			*	*		
b. viuéla		*!		*		
c. Θ		*!	*			

(24) **abole*/**abuele* 'abolish-3sg'

[22] Gouskova (2003) presents a number of compelling arguments that *STRUC is not the correct way to formalize economy constraints in OT. I am using it here simply as a cover term for the reluctance to derive any form that has not already been listed, not to derive specific economy effects. In this use, it could just as well be called LEX or *UNLISTED or *NOVEL; a better formalization is left for future work.

/ABOLISH-inf/ (listed: *abolir*)	USELISTED	Class 1: *o*→*ó*	Class 1 *o*→*ué*	LEX	Class 3: *o*→*ó*	Class 3: *o*→*ué*
☞ a. abolír						
b. abuelír	*!					

/ABOLISH-3sg/ (listed: none)	USELISTED	Class 1: *o*→*ó*	Class 1 *o*→*ué*	LEX	Class 3: *o*→*ó*	Class 3: *o*→*ué*
a. abóle				*!		*
b. abuéle				*!	*	
☞ c. Ø					*	*

The intuition behind this account, then, is that speakers are biased against constructing novel forms unless they have sufficient evidence that the mapping employed to generate the form is trustable enough to outrank LEX; this assumes also a constraint ranking procedure which is able to decide when enough data has amassed to motivate ranking a mapping above LEX (presumably mirroring the confidence function seen in Figure 3). Crucially, the mapping constraints shown in (23)–(24) are schematic stand-ins for a much larger set of stochastically ranked constraints, reflecting the relative likelihood of different mappings in different segmental contexts, and within different classes. Finally, it is worth reiterating that one of the most important aspects of this analysis for predicting the morphological selectivity of gaps is the directionality of the mapping constraints. This does not follow from anything in the OT formalism, but is assumed to have been established in the early stages of learning, by a procedure such as the one described in Albright (2002a).

4.3 Relation to phonotactically motivated gaps

In the analysis of lexically arbitrary gaps, directionality of mappings and strength of generalizations both play a crucial role in predicting where gaps will occur. In cases of phonotactically motivated gaps, on the other hand, neither of these factors has been invoked to date. This immediately raises the question of whether the two types of gaps are fundamentally different, as has sometimes been claimed (Albright 2003; McCarthy and Wolf, this volume), or whether directionality of mappings and scarcity of data might play a role in other cases, as well.

It seems entirely possible that directionality of mappings could help to shed light on why certain phonological alternations are learned as general processes that can extend across multiple morphological contexts, and why others are morphologically restricted. Consider, for example, the example of Norwegian imperative gaps discussed by Rice (this volume): **sykl*. Here, the mystery is framed as why epenthesis is not employed, even though it is attested in nouns: /sikl/ → [sikkəl] 'bike'. But how do speakers know that the noun *sykkel* involves epenthesis from underlying /sikl/? Rice points out that the major source of evidence is the plural [siklər]; the side-by-side occurrence of words without a schwa in the plural (*sykler*) and those that retain the schwa (*Mikk**e**ler* 'guys named Mikkel') motivates an underlying contrast between /kl/ words and /kəl/ words. Suppose, however, that the organization of Norwegian noun paradigms was such that plurals are obligatorily formed with reference to singulars, and not vice versa; in this case, the question of what the underlying form is of *sykkel* is moot, because speakers decide whether or not to put a schwa in the plural based directly on the singular form. This would effectively require speakers to analyze plurals like *sykler* as involving (irregular) syncope, rather

than epenthesis. If this scenario is right, then the only way to learn that there is an epenthesis process is from a part of speech than contains derived *kl#* clusters. In Norwegian, these appear to occur only in the imperatives, with no other possible source of evidence about the fate of such clusters. In Albright (2002a), I discuss additional cases from Latin and Lakhota in which the directionality imposed by the learning procedure determines which alternations can be learned as productive phonology, and which become morphologized. The implication for the analysis of gaps is that perhaps the restriction of phonological processes to particular morphological contexts is not a necessary fact of how phonology is learned, but rather, is a consequence of the direction in which the mappings are applied.

The amount of data that is available to motivate a trustable generalization also appears to be important in some cases of phonotactically motivated paradigm gaps. Rice (2003) notes, citing data from Iverson (1981), that the set of Swedish adjectives that lack neuter forms (chiefly, those ending in [dd]) happens to be a very small class, containing just a handful of words. Iverson (1981) also mentions another class of adjectives that lack gaps in Swedish, namely, those ending in *-id*: *rapid* does not have a neuter form **rapitt*/**rapidt*. Here, too, the number of relevant lexical items is small (under a half dozen[23]). Neither Iverson nor Rice makes use of this fact in their analyses, but it suggests that the amount of data may well play a role in creating phonotatically motivated gaps, as well.

It is also of interest to note that such cases involve not only small classes of words, but sometimes also a certain amount of exceptionality or competition: there is, in fact, one Swedish *-id* adjective that does have a widely used neuter form: *solid* 'solid' → *solitt*, and also at least one *-dd* adjective that is attested with a *-tt* neuter: *högljudd* → *högljutt* 'loud'. This shows that individual words may sometimes use lexical listing to avoid phonotactically motivated gaps, as well (i.e., USELISTED may apply in these cases, too). A fundamental prediction of the current account is that any case of gaps may have sporadic, "exceptionally filled" existing forms; as long as there are only a few of them, they will not reach critical mass to let a general pattern emerge and fill in all remaining cases of the gap. The origin of these exceptionally non-gapped forms seems to be somewhat haphazard. In some cases, they consist of inherited forms (Swedish adjectives in *-id* used to take *-itt* neuters more generally but lost them, perhaps precisely because the class was so small and the necessary generalization could not be learned). In other cases, one imagines that a particular innovative speaker was willing to take the plunge and create an otherwise uncomfortable form, which listeners were then able to lexicalize and use.[24] What I conclude from this is that further empirical work is also needed on phonotactically motivated paradigm gaps, to determine in what respects they truly differ from or are parallel to lexically arbitrary gaps.

5 Conclusion

In this paper, I have attempted to lay out some basic principles that predict which parts of the paradigm, and which particular lexical items, may be affected by paradigm gaps. In particular, it is hypothesized that gaps occur only in those forms that are computed with reference to another base form in the paradigm, and only in cases where the mapping between the base and the de-

[23] Namely: *rapid, morbid, gravid, stupid, rigid, solid.*

[24] Raffelsiefen (2004) points out that Turkish musicians have a lexicalized 1sg possessive form *dom* from the musical note *do* which is subminimal and would ordinarily be a gap; presumably they use it with high enough frequency to ensure that it remains lexicalized among a particular community of speakers.

rived form requires an inference over small amounts of possibly conflicting data. Although the principle is an extremely simple one, it makes strong predictions about what types of languages and inflectional classes gaps should occur in, and why gaps appear to be so restricted. Finally, although the primary purpose has been to provide an analysis of "lexically arbitrary" cases, further inquiry is needed to determine to what extent the mechanisms proposed here may also play a role in shaping other, phonotactically motivated gaps.

References

Albright, A. (2002a). *The Identification of Bases in Morphological Paradigms*. Ph. D. thesis, UCLA. http://www.mit.edu/~albright/papers/AlbrightDiss.html.

Albright, A. (2002b). Islands of reliability for regular morphology: Evidence from Italian. *Language 78*(4), 684–709.

Albright, A. (2003). A quantitative study of Spanish paradigm gaps. In G. Garding and M. Tsujimura (Eds.), *WCCFL 22 Proceedings*, Somerville, MA, pp. 1–14. Cascadilla Press.

Albright, A., A. E. Andrade, and B. Hayes (2001). Segmental environments of spanish diphthongization. In A. Albright and T. Cho (Eds.), *UCLA Working Papers in Linguistics, Number 7: Papers in Phonology 5*, pp. 117–151. http://www.linguistics.ucla.edu/people/hayes/Segenvspandiph/SegEnvSpanDiph.pdf.

Albright, A. and B. Hayes (2002). Modeling English past tense intuitions with minimal generalization. *SIGPHON 6: Proceedings of the Sixth Meeting of the ACL Special Interest Group in Computational Phonology*, 58–69.

Albright, A. and B. Hayes (2003). Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition 90*, 119–161.

Albright, A. and B. Hayes (in press). Modeling productivity with the Gradual Learning Algorithm: The problem of accidentally exceptionless generalizations. In G. Fanselow, C. Féry, M. Schlesewsky, and R. Vogel (Eds.), *Gradience in Grammar: Generative Perspectives*. Oxford University Press.

Anttila, A. (2002). Morphologically conditioned phonological alternations. *Natural Language and Linguistic Theory 20*(1), 1–42. Available for download at http://roa.rutgers.edu/view.php3?roa=425.

Benua, L. (1997). *Transderivational Identity: Phonological Relations Between Words*. Ph. D. thesis, University of Massachusetts, Amherst.

Boersma, P. (1997). How we learn variation, optionality, and probability. *21*, 43–58. http://fon.hum.uva.nl/paul/.

Boersma, P. and B. Hayes (2001). Empirical Tests of the Gradual Learning Algorithm. *Linguistic Inquiry 32*(1), 45–86.

Butt, J. (1997). *Spanish Verbs*. Oxford University Press.

Bybee, J. (1985). *Morphology: A study of the relation between meaning and form*. Amsterdam: John Benjamins Publishing Company.

Bybee, J. L. and E. Pardo (1981). On Lexical and morphological conditioning of alternations: a nonce-probe experiment with Spanish verbs. *Linguistics 19*, 937–968.

Clahsen, H., F. Aveledo, and I. Roca (2002). The development of regular and irregular verb inflection in Spanish child language. *Journal of Child Language 29*, 591–622.

de Gámez, T. (Ed.) (1973). *Simon and Schuster's International Dictionary: English/Spanish Spanish/English*. Simon and Schuster.

Eddington, D. (1996). Diphthongization in spanish derivational morphology: An empirical investigation. *Hispanic Linguistics 8*, 1–35.

Eddington, D. (1998). Spanish Diphthongization as a Non-derivational Phenomenon. *Rivista di Linguistica 10*, 335–354. URL: P 1 R58.

Fanselow, G. and C. Féry (2002). Ineffability in grammar. In G. Fanselow and C. Féry (Eds.), *Resolving Conflicts in Grammars: Optimality Theory in Syntax, Morphology, and Phonology*, Number 11 in Linguistische Berichte Sonderheft, pp. 265–307. Hamburg: Helmut Buske Verlag. http://www.ling.uni-potsdam.de/~fery/pdf/ineffability.pdf.

Garde, P. (1998). *Grammaire Russe: Phonologie et Morphologie*. Paris: Institut d'études slaves.

Gouskova, M. (2003). *Deriving Economy: Syncope in Optimality Theory*. Ph. D. thesis, University of Massachusetts Amherst.

Halle, M. (1973). Prolegomena to a theory of word formation. *Linguistic Inquiry 4*, @3–16.

Hansson, G. O. (1999). 'when in doubt…': Intraparadigmatic dependencies and gaps in Icelandic. In P. Tamanji, M. Hirotani, and N. Hall (Eds.), *NELS 29: Proceedings of the 29th meeting of the North Eastern Linguistic Society*, pp. 105–119. Amherst, MA: GLSA Publications.

Hayes, B. (1999). Phonological restructuring in Yidin^{y} and its theoretical consequences. In B. Hermans and M. van Oostendorp (Eds.), *The Derivational Residue in Phonological Optimality Theory*, pp. 175–205. Amsterdam: John Benjamins.

Hetzron, R. (1975). Where the grammar fails. *Language 51*, 859–872.

Iverson, G. (1981). Rules, constraints, and paradigm lacunae. *Glossa 15*, 136–144.

Jakobson, R. (1948). Russian conjugation. *Word 4*, 155–167.

Juilland, A. and E. Chang-Rodríguez (1964). *Frequency dictionary of Spanish words*. The Hague: Mouton.

Kager, R. (1999). Surface opacity of metrical structure in Optimality Theory. In B. Hermans and M. van Oostendorp (Eds.), *The derivational residue in phonology*. Amsterdam: Benjamins.

Kenstowicz, M. (1997). Base identity and uniform exponence: alternatives to cyclicity. In J. Durand and B. Laks (Eds.), *Current Trends in Phonology: Models and Methods*, pp. 363–394. Salford: University of Salford.

Kiparsky, P. (1965). *Phonological Change*. Ph. D. thesis, MIT.

Kiparsky, P. (1978). Analogical change as a problem for linguistic theory. In *Linguistics in the Seventies: Directions and Prospects*, Number 8:2 in Studies in the Linguistic Sciences. Dept. of Linguistics, University of Illinois, Urbana, Illinois. Reprinted in **?**), pp. 217-236.

Morris, R. E. (2005). Attraction to the unmarked in Old Spanish Leveling. In D. Eddington (Ed.), *Selected Proceedings of the 7th Hispanic Linguistics Symposium*, pp. 180–191. Somerville, MA: Cascadilla Proceedings Project. http://www.lingref.com/cpp/hls/7/paper1097.pdf.

Orgun, C. O. and R. Sprouse (1999). From MPARSE to CONTROL: deriving ungrammaticality. *Phonology 16*, 191–224.

Paul, H. (1920). *Prinzipien der Sprachgeschichte* (Fifth ed.). Tübingen.

Penny, R. (2002). *A History of the Spanish Language* (2nd ed.). Cambridge University Press.

Pertsova, K. (2005). How lexical conservatism can lead to paradigm gaps. In J. Heinz and K. Pertsova (Eds.), *UCLA Working Papers in Linguistics, Number 11: Papers in Phonology 6*, pp. 13–38.

Pinker, S. (1999). *Words and Rules: The Ingredients of Language*. New York: Basic Books.

Prince, A. and P. Smolensky (2004). *Optimality Theory: Constraint Interaction in Generative Grammar*. Blackwell Publishing.

Raffelsiefen, R. (2004). Absolute ill-formedness and other morphophonological effects. *Phonology 21*, 91–142.

Rice, C. (2003). Dialectal variation in norwegian imperatives. *Nordlyd 31*, 372–384.

Rice, C. (2005). Optimal gaps in optimal paradigms. *Catalan Journal of Linguistics 4*, 155–170.

Rice, C. (2006). When nothing wins: Gaps and repairs at the phonology-morphology interface. In C. Davis, A. R. Deal, and Y. Zabbal (Eds.), *NELS 36: Proceedings of the 36th meeting of the North Eastern Linguistic Society*. Amherst, MA: GLSA Publications. http://roa.rutgers.edu/view.php3?id=1076.

Rodríguez Bou, I. (1952). *Recuento de Vocabulario Español*. Río Piedras, Puerto Rico: Universidad de Puerto Rico.

Sebastián, N., F. Cuetos, M. A. Martí, and M. F. Carreiras (2000). *LEXESP: Léxico informatizado del español. Edición en CD-ROM*. Barcelona: Edicions de la Universitat de Barcelona (Colleccions Vàries, 14).

Sharoff, S. (2002). Meaning as use: exploitation of aligned corpora for the contrastive study of lexical semantics. In *Proc. of Language Resources and Evaluation Conference (LREC02). May, 2002, Las Palmas, Spain*.

Shvedova, N. J. (Ed.) (1970). *Grammatika sovremennogo russkovo literaturnogo jazyka*. Moscow: Nauka.

Steriade, D. (1997). Lexical conservatism. In *Linguistics in the Morning Calm, Selected Papers from SICOL 1997*, pp. 157–179. Linguistic Society of Korea, Hanshin Publishing House.

Zalizniak, A. A. (1977). *Grammaticheskij Slovarj Russkogo Jazyka*. Izdatel'stvo Russkij Jazyk.

Zoll, C. (1996). *Parsing below the Segment in a Constraint-based Framework*. Ph. D. thesis, University of California, Berkeley.

Zuraw, K. (2000). *Patterned Exceptions in Phonology*. Ph. D. thesis, UCLA.