EGG School, Banja Luka
July/August 2018

Does word frequency affect phonology? Reasons to be cautious... 2

Patrick Honeybone
University of Edinburgh
patrick.honeybone@ed.ac.uk

The contents of this session

1. What's really at issue?
2. Reasons to be cautious...
3. Is variation an issue?
4. What is the real balance of predictions?
5. Low frequency effects

What is really at issue?

Let's start with a brief history of phonology...

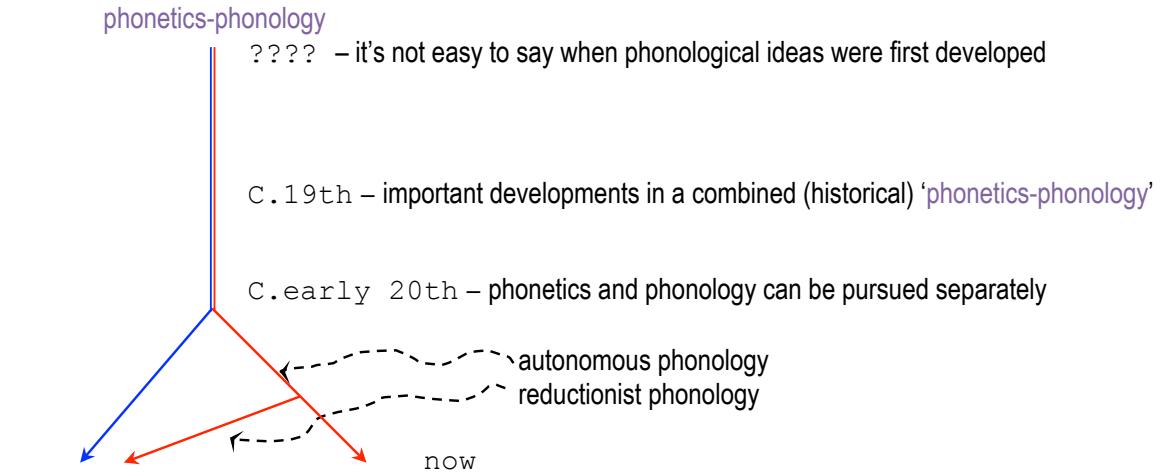

'phonetics'
reductionist
E-focused
concrete
empiricist
functionalist
gradient

'phonology'
autonomous
I-focused
abstract
rationalist
formal
categorial

We can recognise two fundamental currents in phonology:
'reductionist' = phonology can be reduced to phonetics
'autonomous' = there are some purely phonological primitives

There have been a range of theoretical currents within both of these fundamental approaches

- autonomous phonology = 'formal phonology' = emerged out of structuralist phonology, developing phonemic, morphophonemic and systemic approaches

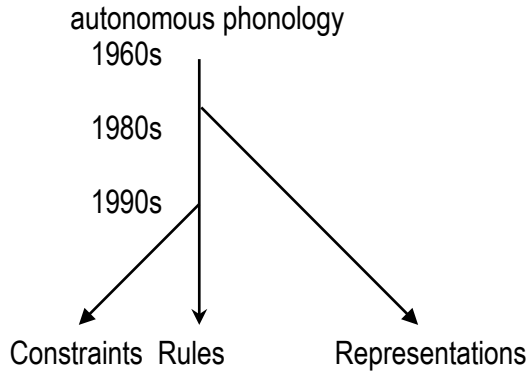

We can now recognise four main currents in contemporary phonological theory:

Rule-Based Phonology	= RBP	}	commensurable, autonomous, formalist
Representation-Based Phonology	= RepBP		
Constraint-Based Phonology	= CBP		
Usage-Based-Phonology	= UBP	–	iconoclastic, reductionist, functionalist

Autonomous phonology = Formal Phonology
Reductionist phonology = Functionalist = Usage-Based Phonology

While the range of theories within the broad current of Formal Phonology have considerably different approaches and assumptions

- they are all commensurable, assuming the same kind of 'atoms'
  - segments, features, syllabic structure

When taken at its word, Usage-Based Phonology may well not be commensurable with these approaches, although all approaches argue they are 'doing phonology'.

What *is* the difference between Formal Phonology and Usage-Based Phonology?

- Newmeyer (1998), writing on the basics of the approaches, says:

> I've been using the terms 'formal linguistics' and 'functional linguistics' as if they have unique well-understood referents. Unfortunately, they do not.
>
> There are, however, two broad orientations in the field. [...] One orientation sees as a central task for linguists characterizing the formal relationships among grammatical elements independently of any characterization of the semantic and pragmatic [and physiological and 'external'] properties of those elements. The other orientation rejects that task [...].
>
> It should be obvious why the former orientation is called 'formalist': it focuses centrally on linguistic form. The problem is the ambiguity of the word 'formal' and its derivatives. The term is ambiguous between the sense of 'pertaining to (grammatical) form', as opposed to meanings and uses, and the sense of 'formalized', i.e., stated in a mathematically precise vocabulary. This ambiguity has the danger of leading to confusion. When Pullum (1989), Chomsky (1990), and Ludlow (1992), for example, debate whether the 'principles and parameter' approach is a species of 'formal linguistics', they have the latter sense of the term in mind; functionalists' criticisms of 'formal linguistics' invariably refer to the former.

Hinskens, Hermans & van Oostendorp (2014)

> Although usage-based and 'rule-based' approaches to natural language share a number of constituent properties as a result of the fact that both look at language as a cognitive object, they differ in many respects. In essence, usage-based approaches contrast with 'rule-based', formal theory in that they do not assume language users to have abstract grammatical knowledge at their disposal. Instead, they postulate a close, organic connection between linguistic structure and language usage.
>
> With respect to the phonetic/phonological part of language, usage-based models assume that language users store detailed phonetic information about the words of their language each time that they are exposed to them. These models stipulate redundant mental storage of bundles of maximally concrete articulatory, acoustic, grammatical, semantic and pragmatic information concerning single occurrences ('tokens' or 'exemplars') of lexical items, along with characteristics of both the speaker and the situation, organized in 'clouds'. [...] Formal theory, on the other hand, conceives linguistic competence as a computational capacity based on internalized representations, rules, processes, constraints, principles and related abstract devices, which are usually categorical and generalize across many cases. [...] The designation 'formal' to refer to these various approaches goes back to De Saussure's definition of langue as "une forme, non une substance" (1916, Cours, Ch. III); where De Saussure's notion of 'forme' referred to the structure of the relations holding between linguistic elements.

So... *really*... this is a course about UBP *vs* FP

Why is frequency such an issue in this argument?

- the argument is typically that:
  - frequency effects are predicted to exist by UBP and are predicted not to exist by FP
  - they *do* exist
- ∴ UBP is right and FP is wrong [for my take on this: see the course's title...]

This can be seen as part of an 'assault on Formal Phonology', as in work like:

- Bybee, J. (2001) *Phonology and Language Use*. Cambridge: Cambridge University Press.
- Port, R. & Leary, A. (2005) 'Against formal phonology.' *Language* 81, 927-964.

How/why does UBP predict that frequency effects exist?

Within the field of UBP, there are multiple ways in which this happens

- all are tied to the existence of exemplars – it could be called Exemplar Phonology

For high-frequency effects, proponents have proposed two types of explanation
(i) articulation-based – speakers store lots of 'practised' exemplars for some words
(ii) listener-based – speakers know that listeners have lots of exemplars for some words

For low-frequency effects, there is more agreement in terms of explanation
(i) and (ii) – less frequent forms store less exemplars

The basic thrust of the UBP position is that the very idea that phonology involves underlying and surface representations is mistaken

- there are several flavours of exemplar theory (some of which are more conciliatory with standard phonological models than others)
  - I focus largely on 'strong positions' here, as they are testable, as in Goldinger (1998), Hawkins (2003), Johnson (2006), Wade & Möbius (2010)
- in exemplar-based UBP, whole words are the focus of phonology, rather than the type of units which most phonological theory works with (segments, features etc.)

The crux of exemplar theory (as in Johnson, 1997) is that the lexicon is a vast repository of highly-detailed memories of phonetic episodes experienced by the speaker

- these are the 'exemplars' – they replace the derivations of standard phonological models
  - they are stored in the lexicon on the basis of usage: on the basis of speakers' experience of production and perception (hence 'Usage-Based Phonology')
  - the strong position is that this is essentially all that speakers need and have in terms of phonological knowledge
  - as Bermudez-Otero (2007) puts it, the strong position assumes that "phonological categories do not exist independently of the exemplars"
  - as Pierrehumbert (2006) points out, in this form "the phonological principle [is] not be in force ... instead, each word [is] an individual point somewhere in phonetic hyperspace"

Van de Weijer (2009) writes the following, in part citing Bybee:

> (2) In exemplar theory, every token of experience is classified and placed in a vast organizational network as part of the decoding process. New tokens of experience are not decoded and discarded, but rather they impact memory representations. In particular, a token of linguistic experience that is identical to an existing exemplar is mapped onto that exemplar, strengthening it. Tokens that are similar but not identical (differing in slight ways in meaning, phonetic shape, pragmatics) to existing exemplars are represented as exemplars themselves and are stored near similar exemplars to constitute clusters or categories.
>
> (Bybee (2006), p. 716)

> Thus, instead of a dictionary-like lexicon as in standard (generative) grammar, lexical items are stored in a network-like multi-dimensional organization: items that are similar are stored close to each other. This has psycholinguistic advantages, e.g. mispronunciations will often pick out a form which is close to the intended form. If a certain item is subject to variation, then both items will be stored, roughly in the proportion of the frequency with which the items are encountered. Variation is thus a natural part of the lexicon in an ET grammar.

Gahl & Yu (2006) write that:

> The central idea behind exemplar-based models is that mental representations consist of memory traces of specific tokens. This idea runs counter to the goal of developing maximally simple, redundancy-free representations, a goal that has been central to many proposals within linguistic theory.

On this kind of UBP approach, each lexical item has its own exemplar cloud

- when speakers come to speak, there is no derivation from UR to SR – rather, one of the exemplars from the exemplar cloud for the relevant lexical items is chosen and implemented in speech
  - there is no UR → SR / input → output computation
  - phonological categories do not have any obvious existence, but can only be thought to exist as ad hoc generalisation over forms in the lexicon
- the forms that are generalised over (exemplar clouds) are gradiently different
- speakers are changing their phonology all the time, as they hear new exemplars
  - this contrasts fundamentally with generative models, where the basic assumption is that most of phonology is fixed at the end of the critical period
  - while there may be modification around the edges, fundamental, 'deep' aspects of phonology are fixed

Why does UBP predict frequency effects?

There is constant update of phonology as new exemplars are added to the exemplar clouds that instantiate each word

- more exemplars (as with more frequent words) = different to less exemplars
- this means that UBP has a means of storing the different frequency of occurrence of each word – with different sized exemplar clouds
  - but in and of itself this doesn't provide a potentially explanatory link

As we saw above,

For high-frequency effects, proponents have proposed two types of explanation
(i) articulation-based – speakers store lots of 'practised' exemplars for some words
(ii) listener-based – speakers know that listeners have lots of exemplars for some words

(ii) is easiest to describe

- Kiparsky sums it up thus:

  "LISTENER-BASED explanations say that frequent words are more predictable, so speakers can put less effort into their articulation without risk of being misunderstood (Jurafsky et al. 2001)."

  - this should mean that more frequent words can be more reduced by a speaker because listeners have more chance of predicting what a more frequent word will be
  - it should be less of a problem if a higy frequency word like *felt* loses its [t] than if a low frequency word like *built* loses its [t]

(i) Bybee and others (Phillips, Pierrehumbert etc) tend towards an articulatory approach

Bybee (2001)

> If sound changes are the result of phonetic processes that apply in real time as words are used, then those words that are used more often have more opportunity to be affected by phonetic processes. If representations are changed gradually, with each token of use having a potential effect on representation, then words of high frequency will change at a faster rate than will words of low frequency.[3] The streamlining of high-frequency words and phrases has the effect of automatizing production. Any motor activity that is repeated often becomes more efficient. The first effect of frequency, then, is to automate production (Boyland 1996).

Bybee (2006) writes further that:

- "the articulatory representation of words and sequences of words are made up of neuromotor routines. When sequences of neuromotor routines are repeated, their execution becomes more fluent. This increased fluency is the result of representing the repeated sequence at a higher level as a single unit"
  - the 'storage as a single unit' is possible as this is an exemplar model, where full pronunciations of words are stored – whole reduced phrases are stored as exemplars

Tamminga (2014) also explains the link that Bybee sees between high frequency and 'high frequency effects' in phonological phenomena:

> Given a tendency for reduction during production, the phonetic representation of a word will gradually accrue more exemplars that are reduced, and these exemplars will become more likely to be chosen for production, where they may undergo further reduction, gradually moving the words of the language in a consistent direction. The more frequent words will have more chances to undergo online reduction and thus will change more rapidly. (Bybee, 2002:271)

Bermudez-Otero (2016) clearly sets out how any change or process that it due to any kind of 'phonetic bias' (assimilations, palatalisations, reductions, everything?) is predicted to affect high frequency words more:

Postulated mechanism

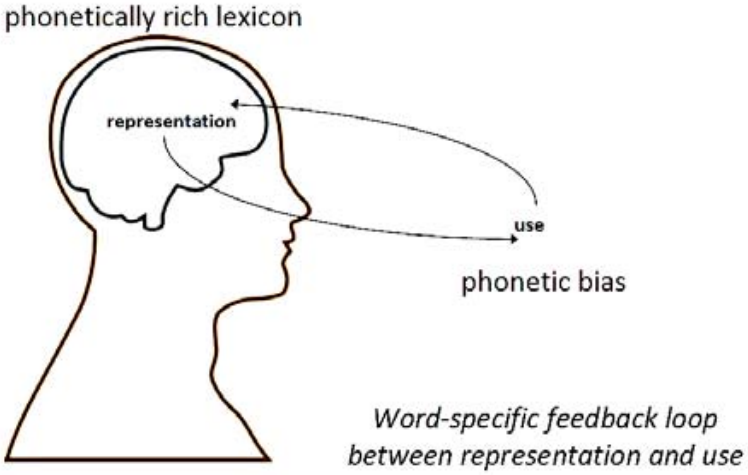

high-frequency words
⇩
many iterations of the loop

low-frequency words
⇩
few iterations of the loop

- High-frequency words undergo greater exposure to reductive phonetic biases during use.
- The gradient effect of these biases is registered separately for each word in its own cloud.

For low-frequency effects, there is more agreement in terms of explanation
(i) and (ii) – less frequent forms store less exemplars

Bybee argues that there is an explanatory link between 'low frequency effects' and exemplar theory:

- "frequency strengthens the memory representations of words or phrases making them easier to access whole and thus less likely to be subject to analogical reformation"
  - infrequent words lack the 'conserving effect' of this result of exemplar entrenchment and may more easily undergo change
  - high frequency can inhibit changes or processes in this way, so low frequency words are most likely to undergo them

Bybee (2001)

> The frequency effects mentioned briefly in Chapter 1 find a natural expression in this model. Since tokens of use map onto existing representations, high-frequency items grow strong and therefore are easier to access. The other side of the coin is that little-used items will tend to fade in representational strength and grow more difficult to access. The conserving effect of high token frequency, which protects high-frequency items from regularization on the basis of productive schemas, is represented as lexical strength. Since regularization only occurs when existing forms are difficult to access, high-frequency forms are not prone to regularization.

If these predictions for different types of frequency effects hold true, it would be evidence in favour of the UBP approach to phonology.

Why does FP predict that they shouldn't exist?

There is no obvious way to express the fact that the rules seem to affect some words more than others in generative frameworks

- as Pierrehumbert (2001) writes, this could have substantial implications:

> These results challenge standard models of phonology and phonetics at two levels. First, in all standard models, the lexicon is distinguished from the phonological grammar. The exact phonetic details of a word's pronunciation arise because the word is retrieved from the lexicon, and processed by the rules or constraints of the grammar whose result (the surface phonological form of the word) is fed to a phonetic implementation component. The phonetic implementation component computes the articulatory and/or acoustic goals which actualize the word as speech. The phonetic implementation component applies in exactly the same way to all surface phonological representations, and the outcome depends solely on the categories and prosodic structures displayed in those representations. As a result, there is no way in which the phonetic implementation can apply differently to some words than to others. If a phonetic implementation rule is variable and gradient, then the same probability distribution of outcomes would arise for all words which meet the structural description of the rule. This generic feature of modular generative models with phonetic implementation rules is developed at more length in Pierrehumbert (1994).
>
> A second challenge arises from the fact that the differential phonetic outcomes relate specifically to word frequency. Standard generative models do not encode word frequency. They treat the word frequency effects which are so pervasive in experiments involving priming or lexical decision tasks as matters of linguistic performance rather than linguistic competence. Thus the intrusion of word frequency into a traditional area of linguistics, namely the conditioning of allophony, is not readily accommodated in the classical generative viewpoint.

A classical model of the grammar the generative phonology is:

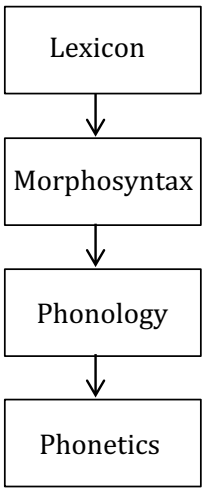

The components shown above are modular and cannot 'refer back' up the grammar

- this is a 'modular feed-forward' model
- a module can only influence another if they have an interface
  - syntax and phonology have an interface, but lexical information does not have an interface with phonology

Different models of things affecting pronunciation:

UBP

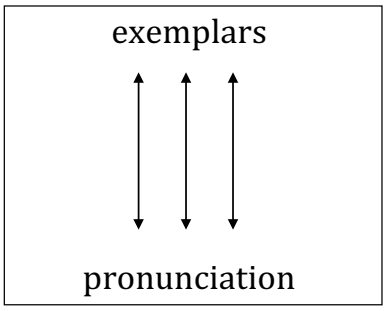

'I = E'

FP

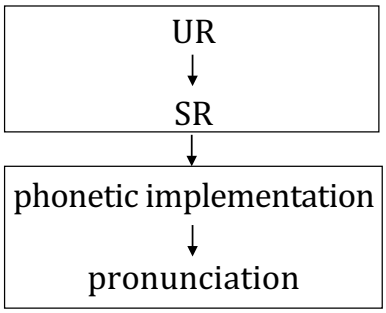

'I knows nothing of E'

The frequency with which a word is used is not an I-linguistic fact

- usage is canonically an E-linguistic thing
- FP is a feed-forward model: the output of phonology is fed to speech production
- usage-related phenomena should not be able to impact on phonological behaviour
- URs (and SRs) are categorical phonological forms

Should FP just roll over and give up?

Reasons to be cautious...

Verner (1875) didn't give up when trying to figure out an explanation for what seems like exceptions to 'Grimm's Law':

> "the best proposal is Scherer's explanation in the fine section concerning the sound-shift [...]: 'I now assume that all irregularly shifted tenues were first shifted regularly to voiceless spirants, that these, particularly in frequently used words (like fadar, môdar), were under the influence of the surrounding voiced elements also produced with voice ...' [...] it is enough for us to have determined that the irregular shifts also followed at one time the sound stage of the regular shifts; from there, however, they progressed further. [...] And we can now phrase the question of the etymological explanation thus: Why did the sound current of the shift in some cases stop with the voiceless fricative and in other cases progress further through the voiced fricative to the voiced stops?

> "The only person who has sought an answer to this question, as far as I know, is Scherer in the passage just cited. He assumes that the shift to voiced stops occurs 'in frequently used words (like fadar, môdar)' consequently the regular shift occurs in less frequently used words. I believe that the venerable author did not wish to attach great weight to this attempt at explanation and that he permitted himself to mention it only as a conceivable possibility. A careful scrutiny of the Germanic vocabulary is not favorable to his thesis. Is it probable that fadar and môdar were used more frequently than brôþar? [...]
>
> Could fehu-, the Germanic epitome for material well-being, cattle, money, wealth, possessions and the like, have been a more infrequently occurring word than, for example, lagu- 'lake' (ON lög-r, OE lagu = Lat. lacu-s)? May one assume that our Germanic ancestors used the numbers 4 and 100 (fedvór, hund) more frequently than the number 10 (tehan)? More such examples could be cited, I will, however, find occasion in what follows to demonstrate the improbability of that thesis."

It has been claimed that "Verner's may be the single most influential publication in linguistics." Lehman (1967)

- perhaps we should listen to him...

The very fact that many changes are exceptionless might give us pause

- as might the fact that there are not spontaneous segmental splits in phonology

Is variation an issue? Should FP give up because there's variation?

All of the cases of frequency effects that we have considered involve variation

- [hɛl] *vs* [hɛld]
- [mɛməri] *vs* [mɛmri]
- *time* [tʰaɪm] x msec *vs* [tʰaɪm] x-5 msec

Variation is not, however, a problem for FP in and of itself...

For example, Smith & Holmes-Elliott (2018) present the following graph from their fieldwork in Buckie, Scotland, showing how often speakers realise /t/ as [ʔ]

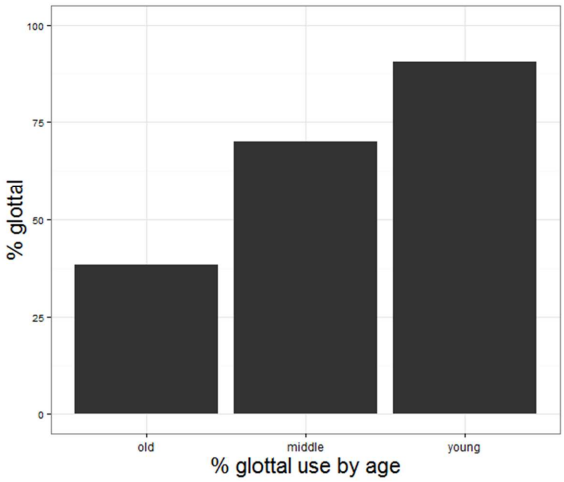

FP can recognise the existence of a process of glottalling, and describe the environment in which stops are glottalled, but can *leave the details of the intra-speaker variation* in (how often the process occurs, in different social groups) to sociolinguistics:

- glottalling = t → ʔ / ́tσ [of the kind considered here]
- this is a *variable* or *optional* rule; many rules are *obligatory*, and do not show variation

The proportions of variation (how often a rule applies) are interesting (and can be dependent on age, class, sex, geography), but are of a different nature and status to the phonological factors that condition formal processes (eg, phonological features, syllable structure, stress placement)

- this simply needs to give a role to the sociolinguistically-influenced implementation of phonological phenomena

The simple, early way to incorporate and model variation in phonological theory was to use optional or variable rules; there is a difference

- optional rules simply state that a variant *can* be a realisation of an underlying segment
- variable rules attempt to incorporate factors which govern variation into the rule

Coetzee & Pater (2011, 3)

> In rule-based phonology, variation is standardly handled by simply marking a rule as [+optional] (see recently Vaux 2007). Labov (1969) suggests that this could be formalized by writing parentheses around the structural change of a rule, as in (3).
>
> (3) Labov (1969:737)
>
> X → (B) / Y _ Z

Variability can be modelled in OT in a range of ways, including 'partially ordered constraints'

- this assumes that while most constraints are ranked normally, certain constraints are only partially ordered – not ranked with reference to each other
- each time that the grammar is used, the partially-ordered constraints are fully ranked consistent with the partial ordering – the ranking is randomly chosen

Coetzee & Pater (2011) show this in an abstract form:

Grammar: $C_1 >> C_2$, $C_1 >> C_3$

a. First possible ranking: $C_1 >> C_2 >> C_3$

/input$_1$/	C_1	C_2	C_3
☞ *cand*$_1$			*
cand$_2$		*!	
cand$_3$	*!		

b. Second possible ranking: $C_1 >> C_3 >> C_2$

/input$_1$/	C_1	C_3	C_2
cand$_1$		*!	
☞ *cand*$_2$			*
cand$_3$	*!		

Other types of variation can be accounted for using differences of how forms are stored

- this can account for lexical variability in a phenomenon where it looks like some words have a particular type of phonological behaviour, and some don't

For example, is there *n*-deletion in English?

- the behaviour of the indefinite article might lead us to think so
  [ə] *give me a cat* [ən] *give me an egg*
- the behaviour of all other words should give us second thoughts, however:
  [tɛn] *give me ten cats* [tɛn] *give me ten eggs*

It seems more likely that that some words have one UR, and some words have two

- ‹indefinite ARTICLE› = /ə, ən/ – with phonological-controlled allomorphy
- ‹NUMERAL 10› = /tɛn/

This accounts for the variation straightforwardly, but variation in such cases in not always phonological controlled – this is 'optional' variation as we can use either UR

- ‹THIRD PERSON oblique PRONOUN› = /ðəm, əm/
  give them a cat *give 'em a cat*

Phonology can allow social factors to determine whether a rule applies or which UR is used, but variation shouldn't be governed by frequency, should it?

What is the *full* balance of predictions?

The two basic approaches to phonology clearly make different predictions in terms of what should occur in the relationship between frequency and phonology

- the full details of those predictions need to be thought about carefully, however, and are more subtle than is sometimes claimed:

1. High frequency effects should exist

- UBP – yes
- FP – no

What is the *full* balance of predictions?

The two basic approaches to phonology clearly make different predictions in terms of what should occur in the relationship between frequency and phonology

- the full details of those predictions need to be thought about carefully, however, and are more subtle than is sometimes claimed:

1. High frequency effects should exist

- UBP – yes
- FP – no

2. High frequency effects should *always* exist in 'natural' changes/rules

- UBP – yes
- FP – no

Pierrehumbert (2002) writes that

- "Any systematic bias on the allophonic outcome would incrementally impact high frequency words. In short the model is applicable to any Neogrammarian sound change..."

'High frequency effects' should therefore be omnipresent in N-changes, and the synchronic processes connected to them

- there is no reason why any such phenomenon should *not* be affected by frequency, according to Pierrehumbert

Tamminga (2014) explains how Pierrehumbert (2002) sets this out:

> Pierrehumbert explicitly extends the claim that frequent words lead sound change to any kind of gradient phonetic change, stating that "any systematic bias on the allophonic outcome would incrementally impact high frequency words at a greater rate than low frequency words" (2002:118). Just as frequent words that undergo reduction in speech should end up being more reduced in the phonetics inherent to their representation, frequent words that are undergoing non-reductive sound change (for example, the raising of /ey/ along the front diagonal in Philadelphia (Labov et al., 2013)) should accumulate advanced tokens more quickly than their less-frequent counterparts.

What is the *full* balance of predictions?

1. High frequency effects should exist

- UBP – yes
- FP – no

2. High frequency effects should *always* exist in 'natural' changes/rules

- UBP – yes
- FP – no

3. Low frequency ('frequency conserving') effects should exist

- UBP – yes
- FP – ...?

What are the predictions for low frequency effects?

Although this doesn't seem to be recognised frequently, low frequency effects are not a problem for formal phonology

- formal models are, in fact, fully compatible with low frequency effects, if the role of acquisition and change are recognised

English preterite regularization

If learners are more likely to hear *ic drāf* 'I drove' than *ic bād* 'I bided', because *drāf* is more frequent than *bād*...

- then they are more likely to fix irregular *drāf/drove* in their grammar as the past tense of *drīfan/drive* than they are to fix irregular *bād/bode* in their grammar as the past tense of *bīdan/bide*
- if speakers do not acquire an irregular past tense for *bīdan/bide* because they don't get enough input to do so (because the word is less frequent), then its past tense will be derived regularly by the rule for past tense formation (add *-ed*)
- formal phonology is compatible with the conserving effect of frequency through diachrony driving phonological acquisition
- this kind of frequency effects can be seen in FP as a diachronic frequency effect

How about Diatonic Stress Shift?

This, too, is not a problem for FP.

DSS involves the creation of diatonic pairs from monotonic pairs

coonvíct$_N$	~	*convíct*$_V$	>	*cónvict*$_N$	~	*convíct*$_V$
reecórd$_N$	~	*recórd*$_V$	>	*récord*$_N$	~	*recórd*$_V$

- the "words which have undergone the Diatonic Stress Shift have lower frequency than those which have not" (Sonderegger 2010)
- this is exactly what we would expect

The change involved in the creation of diatonic pairs:

- σσ́$_N$ > σ́σ$_N$

Diatones exemplify some of the basic patterns of English stress:

- a final syllable is typically unstressed in nouns
- = extrametricality, high-ranked NONFINALITY *etc*
- a heavy final syllable is unproblematically stressed in verbs

The pre-change situation was:

discount

V /dɪskaʊnt/ → σσ́ (a verb with heavy final syllable gets accent)

N /dɪskáʊnt/ → σσ́ (nouns normally do not have final stress, so it must be lexically stored)

The post-change situation was:

discount

V /dɪskaʊnt/ → σσ́ (a verb with heavy final syllable gets accent)

N /dɪskaʊnt/ → σ(σ) → σ́(σ) (stress assigned to only non-extrametrical syllable)
(noun final extrametricality)

Diatonic Stress Shift involves individual words submitting to the general pattern of English phonology, removing the exception that the Noun forms allow final stress

- children acquiring words need strong evidence for the exception marker
  - low frequency words are less likely to provide that evidence, so are more likely to be (re)analysed in acquisition as non-exceptional, in which case the Noun will receive non-initial stress

What is the *full* balance of predictions?

1. High frequency effects should exist
- UBP – yes
- FP – no

2. High frequency effects should *always* exist in 'natural' changes/rules
- UBP – yes
- FP – no

3. Low frequency ('frequency conserving') effects should exist
- UBP – yes
- FP – yes

What is the real full balance of predictions?

The two basic approaches to phonology clearly make different predictions in terms of what should occur in the relationship between frequency and phonology

- the full details of those predictions need to be thought about carefully, however, and are more subtle than is sometimes claimed:

1. High frequency effects should exist
- UBP – yes
- FP – no

2. High frequency effects should *always* exist in 'natural' changes/rules
- UBP – yes
- FP – no

3. Low frequency ('frequency conserving') effects should exist
- UBP – yes
- FP – yes

4. Frequency effects, like all phonological generalisations, should always be gradient
- UBP – yes
- FP – no

On a UBP approach, each lexical item has its own exemplar cloud

- when speakers come to speak, there is no derivation from UR to SR – rather, one of the exemplars from the exemplar cloud for the relevant lexical items is chosen and implemented in speech

On a UBP approach, phonological categories do not have any obvious existence, but can only be thought to exist as ad hoc generalisations over forms in the lexicon

- the forms that are generalised over (exemplar clouds) are gradiently different
- so there should be *no sharp edges to categories* – the difference between 'parts of words' should be gradient

Therefore:

UBP expects gradience
exemplar clouds do not have sharp edges

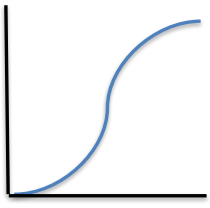

FP expects categoricality
formal categories *do* have sharp edges

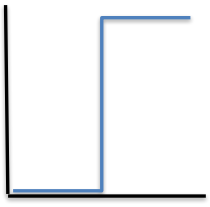

Fundamentally, frequency is a gradient phenomenon...

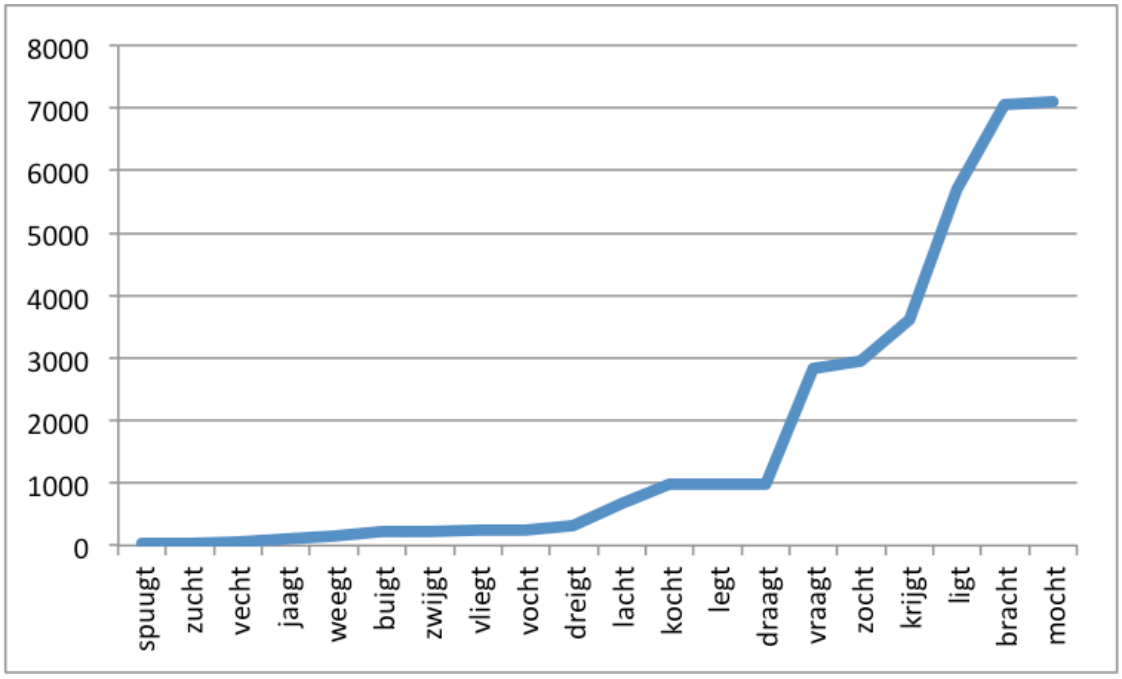

...so exemplar effects should be gradual

How do things stand in terms of the full balance of predictions...?

1. High frequency effects should exist
- UBP – yes
- FP – no

2. High frequency effects should *always* exist in 'natural' changes/rules
- UBP – yes
- FP – no

3. Low frequency ('frequency conserving') effects should exist
- UBP – yes
- FP – yes

4. Frequency effects, like all phonological generalisations, should always be gradient
- UBP – yes
- FP – no

Let's see...