

RESEARCH

Category-specific effects in Welsh mutation

Michael Hammond¹, Elise Bell², Skye Anderson¹, Peredur Webb-Davies³, Diane Ohala¹, Andrew Carnie¹ and Heddwen Brooks¹¹ U. of Arizona, Tucson, AZ, US² U. of Texas–El Paso, El Paso, TX, US³ Bangor University, Bangor, Gwynedd, UKCorresponding author: Michael Hammond (hammond@u.arizona.edu)

In this paper we investigate category-specific effects through the lens of Welsh mutation. Smith (2011) and Moreton et al. (2017) show that English distinguishes nouns and proper nouns in an experimental blending task. Here we show that Welsh distinguishes nouns, verbs, personal names, and place names in the mutation system. We demonstrate these effects experimentally in a translation task designed to elicit mutation intuitions and in several corpus studies. In addition, we show that these effects correlate with lexical frequency. Deeper statistical analysis and a review of the English data suggests that frequency is a more explanatory factor than part of speech in both languages. We therefore argue that these category-specific effects can be reduced to lexical frequency effects.

Keywords: mutation; Welsh; Celtic; frequency; part-of-speech; morphology

1 Introduction

In this paper, we use data from Welsh and English to demonstrate category-specific phonological effects and derive them from frequency effects.

Smith (2011) reviews a number of category-specific phonological effects, showing how different parts of speech exhibit differing degrees of faithfulness to the input. In cruder terms, the phonology of a language can affect some parts of speech more than others. Among other effects, she shows that nouns generally exhibit greater faithfulness to the input than other parts of speech. Being more faithful means that nouns resist operations that would make them less like their input. It also means that they are more varied phonologically than other categories.¹

Moreton et al. (2017) expand on this result demonstrating emergent category effects in English that also distinguish proper names; specifically, proper names are more faithful to the input than other nouns. They do this experimentally, using a word-blending task. For example, subjects were asked about the acceptability of nonce blends involving items like *soprano* and *preening* as either *sopreening* or *sopraning*. Moreton et al. found that subjects were more inclined to accept *sopraning* over *sopreening* when *soprano* was interpreted as referring to the TV program *The Sopranos* than if it referred to a type of singer. Loosely, more of the word is preserved in blending if it is a proper noun than if it is a common noun.

¹ There have been a number of other approaches to the formalization of category-specific effects and to how such systems might be learned, e.g. Itô & Mester (1999), Alderete (2001), Inkelas & Zoll (2007), Albright (2008), Itô & Mester (2009), Shih & Inkelas (2015), Becker & Gouskova (2016), etc.

In this paper we report on a behavioral study using a translation task, designed to elicit Welsh mutation (a process where the initial consonant of a word changes in different morphosyntactic contexts; more on this below). First, we replicate the effect for English nouns showing that Welsh nouns are more faithful than verbs in initial consonant mutation. We go on to show that place names exhibit an intermediate status between proper nouns and common nouns in terms of mutation. We follow this up with several corpus studies that demonstrate the same effects.

We demonstrate that all of these distinctions are unexpectedly correlated with lexical frequency. Specifically, more frequent items undergo mutation more readily. We then go back to Moreton et al.'s data and show that they are correlated with lexical frequency in the same way. Specifically, more frequent items are more likely to simplify in the blending task. Statistically, once lexical frequency is in the model, there is no need for lexical category.

We attribute our results to a well-known frequency effect whereby reduction or lenition processes apply more readily to more frequent items. This observation goes way back to Hooper (1976) who cites the case of syncope in English, i.e. that syncope applies more readily in high-frequency items like *memory* [mɛm(ə)ri] vs. low-frequency items like *mammary* [mæm(ə)ri]. A similar observation is made by Fidelholtz (1975) with respect to vowel reduction in English. For example, the corresponding syllable of a relatively high-frequency form like *astronomy* [əstr'anəmi] is more likely to reduce than the initial syllable of a relatively low-frequency form like *gastronomy* [gæstr'anəmi]. This has been studied more recently by, e.g. Hammond (1999), Hammond (2004), Coetzee (2009), Coetzee & Kawahara (2013), etc.

We thus establish three principal effects:

- i. There are superficial category effects for Welsh mutation, similar to those of English.
- ii. The Welsh effects also correlate with lexical frequency.
- iii. In fact, the English effects correlate with lexical frequency as well.

This suggests that frequency is the guiding force here rather than lexical category per se. (Note that we are *not* arguing that all of grammar follows from frequency effects, just that target category effects do.)

The organization of this paper is as follows. First, we review the facts of Welsh mutation, with particular attention to how it interacts with lexical category. Next, we go on to report the results of our behavioral study, showing how it replicates the noun and proper noun effects noted by Moreton et al. As just described above, these behavioral effects *also* show an effect of lexical frequency and we next probe this more closely with a series of corpus investigations. In our corpus investigations, we show how mutation is less likely with less frequent forms and we show how the different parts of speech correlate with lexical frequency as we would expect. Specifically, lexical categories with higher lexical frequency undergo mutation more readily. We confirm this frequency effect by looking back at Moreton et al.'s experimental results with respect to English. We then provide a formal analysis showing how the frequency effects we've demonstrated can be incorporated into the grammar. Finally, we conclude with a discussion of how lexical frequency and lexical category can become intertwined as shown.

2 Mutation in Welsh

Initial consonant mutation in Welsh is a typologically rare process where the first sound of a word changes in specific morphological and syntactic contexts (Morris-Jones 1913; Ball & Müller 1992; King 2003). Mutation has been analyzed extensively in the lin-

guistic literature, e.g. Awbery (1973), Lieber (1983), Tallerman (1990), Kibre (1997), Pyatt (1997), Green (2006), Mittendorf & Sadler (2006), Wolf (2007), Stammers (2009), Tallerman (2009), Iosad (2010), Hammond (2011), Hannahs (2011), Hannahs (2013), Prys (2015), etc. The facts of this section are consistent with standard descriptive and pedagogical sources on Welsh, e.g. King (2003), except where noted. The Welsh mutation system is quite complex and we cannot hope to treat all of it here; our description treats those aspects of the system relevant to the behavioral and corpus studies in this paper.

Welsh has at least three distinct mutations, but we focus on the *soft mutation* here. Basically, the process can be triggered in two ways. First, various preceding elements induce it. In the following examples, mutation is triggered by the definite article *y* [ə] when the following noun is feminine singular, the possessive marker *dy* [də] ‘your’, the preposition *am* [am] ‘about’, the disjunction *neu* [neu] ‘or’, and the prenominal adjective *hen* [he:n] ‘old’.² Examples appear in Table 1.

The examples in Table 1 are all nouns, but the soft mutation applies to other lexical categories as well. Table 2 gives examples of verbs and Table 3 gives examples of adjectives. Many other triggers of soft mutation appear in the language.

Second, the soft mutation is triggered by certain syntactic contexts. For example, the object of an overtly inflected verb undergoes soft mutation. In the example below the direct object *cathod* [k^haθɔd] ‘cats’ does not undergo mutation because the verb *gweld* [gwɛld] ‘see’ is not directly inflected. In the present tense, the auxiliary verb *bod* ‘be’ marks person and number.

- (1) dw I 'n gweld cathod [du i n gwɛld k^haθɔd]
 am I PRT see cats
 ‘I see cats’

Table 1: Particles.

a.	cath	vs.	y gath
	[k ^h a:θ]		[ə ga:θ]
	‘cat’		‘the cat’
b.	tad	vs.	dy dad
	[t ^h a:d]		[də da:d]
	‘father’		‘your father’
c.	mêl	vs.	am fêl
	[me:l]		[am ve:l]
	‘honey’		‘about honey’
d.	brawd	vs.	neu frawd
	[brawd]		[neu vrawd]
	‘brother’		‘or a brother’
e.	gwin	vs.	hen win
	[gwi:n]		[he:n wi:n]
	‘wine’		‘old wine’

² Here and following, we transcribe our examples in the northern dialect of Welsh. Note that we transcribe diphthongs ending in the high back unrounded glide with [uɨ] rather than the more usual [i]. We do this as this captures the fact that these are falling diphthongs and the element on the right is properly a glide, rather than a full vowel.

Table 2: Verbs.

a.	mynd	vs.	newydd fynd
	[mind]		[nɛwɪð vɪnd]
	'go'		'just went'
b.	canu	vs.	am ganu
	[kʰanɨ]		[am ganɨ]
	'sing'		'about singing'
c.	dringo	vs.	i ddringo
	[drɪŋɔ]		[i ðrɪŋɔ]
	'climb'		'to climb'

Table 3: Adjectives.

a.	pert	vs.	rhy bert
	[pʰɛrt]		[rɨ bert]
	'pretty'		'too pretty'
b.	diddorol	vs.	yn diddorol
	[dɪðɔɾɔl]		[ən ðiðɔɾɔl]
	'interesting'		'interestingly'
c.	diflas	vs.	hynod o ddiflas
	[divlas]		[hənɔd o ðivlas]
	'miserable'		'extremely miserable'

Compare this with the following past tense form. Here, the verb is directly inflected for person and number (in the past tense) and the direct object appears in the soft mutation.³

- (2) gwelais i gathod [gwɛlajs i gaθɔd]
 saw-1SG I cats-SOFT
 'I saw cats'

Verbs in certain embedded clauses with an overt subject will also display the soft mutation. In the following example, the verb *mynd* [mind] 'go' does not undergo soft mutation since there is no overt subject in the embedded clause.

- (3) dw i eisiau mynd [du i iʃɔ mind]
 am I want go
 'I want to go'

Compare this with the following example where the embedded subject is overt and *mynd* mutates to *fynd* [vind]:

- (4) dw i eisiau i ti fynd [du i iʃɔ i tʰi vind]
 am I want to you go-SOFT
 'I want you to go'

³ A more comprehensive and theoretically-aware characterization would be to say that the syntactic soft mutation happens after an XP (Tallerman 2009).

Finally, we've already seen that mutation can be dependent on grammatical gender; we saw in Table 1 that the definite article triggers soft mutation on feminine singular nouns. Adjectives with feminine singular nouns also mutate. Compare:

- (5) a. ci du [k^hi: di:]
 dog-MASC black
 'black dog'
- b. cath ddu [k^haθ ði:]
 cat-FEM black
 'black cat'

Table 4 gives the orthographic and phonetic effects of soft mutation. Other consonants do not change in mutation contexts, e.g. [s, n, v, l, r, ʃ, ç, χ, θ, f].

Interestingly, personal (or family) names do not generally undergo mutation. Compare the names in Table 5 with Table 1 above. In fact, some Welsh personal names also exist as common nouns with distinct meanings. This results in minimal pairs in mutation environments depending on whether the word is used with its literal meaning or as a name as in Table 6.

In very rare circumstances, personal names can undergo soft mutation. As a measure of how rare this is, there is not a single example in the CEG corpus, a corpus of written Welsh of over a million words (Ellis et al. 2001). When this does occur, in some cases it seems to correlate with treating the name as if it were a common noun. For example, in

Table 4: Soft mutation.

Unmutated		Mutated	
Spelled	Pronounced	Spelled	Pronounced
p	p ^h	b	b
t	t ^h	d	d
c	k ^h	g	g
b	b	f	v
d	d	dd	ð
g	g	∅	∅
ll	ł	l	l
rh	ɾ	r	r
m	m	f	v

Table 5: Names.

a.	Mair	[majr]	'your Mair'
	dy Mair	[də majr]	
	*dy Fair	[də vajr]	
b.	Llinos	[łiɲɔs]	'about Llinos'
	am Llinos	[am łiɲɔs]	
	*am Linos	[am liɲɔs]	
c.	Bronwen	[brɔɲwɛn]	'or Bronwen'
	neu Bronwen	[neɥ brɔɲwɛn]	
	*neu Fronwen	[neɥ vrɔɲwɛn]	

the following examples from Twitter, the names *Dafydd* [davið], *Caradog* [kʰaradɔg], and *Geraint* [gerajnt] are used as if they were common nouns and undergo mutation. In the first case *Dafydd* takes a postnominal adjective; in the second case *Caradog* takes a definite article; and in the third case, *Geraint* takes a definite article and a number.⁴

- (6) gan Ddafydd arall [gan ðavið araɫ]
by David-SOFT other
'by another David'
- (7) yr hen Garadog [ər hən garadɔg]
the old Caradog-SOFT
'the old Caradog'
- (8) gwahaniaethu 'r ddau Eraint [gwahanjeuθiɾ ðau ɛrajnt]
distinguishing the two Geraint
'distinguishing the two Geraints'

Place names exhibit a more complex pattern. The prescriptive rule is that Welsh place names and certain non-Welsh place names mutate. Other non-Welsh place names do not mutate. All three cases are given in Table 7. *Bangor* and *Conwy* are the names of towns in Wales that do mutate. *Paris* and *Califfornia* are foreign place names that do mutate.⁵ *Taiwan* and *Berlin* are place names that do not generally mutate.

Table 6: Names or nouns.

a.	Llinos	female name meaning 'finch'		
		i linos	[i linɔs]	'to a finch'
		i Llinos	[i ɫinɔs]	'to Llinos'
b.	Glyn	male name meaning 'valley'		
		i lyn	[i li:n]	'to a valley'
		i Glyn	[i gli:n]	'to Glyn'

Table 7: Place names.

Welsh	mutating	Bangor	i Fangor
		[banɔr]	[i vanɔr]
		Conwy	i Gonwy
		[kʰɔnwi]	[i ɔnwi]
Non-Welsh	mutating	Paris	i Baris
		[pʰarɪs]	[i barɪs]
		Califfornia	i Galiffornia
		[kʰalɪfɔrnja]	[i galɪfɔrnja]
	non-mutating	Taiwan	i Taiwan
		[tʰajwan]	[i tʰajwan]
		Berlin	i Berlin
		[berɪn]	[i berɪn]

⁴ See Morgan (1952) for a general discussion of mutation of proper names in the literary language.

⁵ Note that *Califfornia* is the Welsh spelling for the name of the state.

Ball & Müller (1992) maintain that non-Welsh place names mutate when they are “considered to be common enough to be brought into the system” (Ball & Müller 1992: 205). Prys (2015) establishes a more general result, demonstrating with corpus data that more frequent place names generally mutate more readily.

Place names are rather sporadic in their mutation and can often go unmutated in mutation contexts in more casual styles. For example, we can also find *i Bangor*, *i Conwy*, *i Paris*, and *i California* in Twitter data.⁶

There are related frequency effects with verbs as well. Stammers (2009) establishes that more frequent verbs occur more frequently in mutation contexts. Stammers & Deuchar (2012) establish that more frequent verbs also mutate more often.⁷ We return to this below.

The prescriptive rules thus support the idea that lexical category can affect morphological processes. Specifically, personal names exhibit greater faithfulness by resisting soft mutation. On the other hand, we’ve seen that place names exhibit a more complex pattern, one that we examine more closely in the following section.

3 Behavioral experiment

In this section, we describe a behavioral experiment that examines more closely the role of lexical category in the Welsh mutation system. In addition, we examine lexical frequency and hypothesize, following Prys (2015), that it is what is responsible for the distinction above between mutating and non-mutating place names.

In the experiment, subjects were asked to translate very simple English sentences into conversational Welsh. We chose this task because it’s been used before in the documentation of Scottish Gaelic (Dorian 1973; Dorian 1978; Dorian 1981; Hammond et al. 2014; Hammond et al. 2017). The logic for this choice is that we wanted a simple method for eliciting intuitions about the contexts for mutation. Translation items were chosen such that subjects would not be able to deduce that we were interested in mutation, as mutation, typically shows a high degree of style-shifting (Prys 2015). Moreover the expected statistical distribution of mutation in our items was essentially equivalent to what’s seen in normal Welsh conversation.

For example, one of our prompts was “Dewi went to a new brewery”. This was designed to elicit a sentence that would test whether the noun for brewery mutates as expected after the mutating preposition *i* [i] ‘to’. We would expect a response like:

- (9) Aeth Dewi i fragdy newydd [auθ dəwi i vragdi newið]
 went Dewi to brewery new
 ‘Dewi went to a new brewery’

Subjects were allowed a lot of latitude in their responses, except for the key parts we were interested in, in the case above, the preposition *i* and the noun *fragdy*. If they used different words for those elements, they would be prompted for whether they could say the sentence in another way, using the relevant items. For example, if the subject said *o fragdy* ‘from a brewery’ instead, we would ask if they could say ‘to a brewery’ (in English). Similarly, subjects might code-switch or say they didn’t know the word for brewery. We would then offer the item *bragdy* and ask if they knew it and could use it in the sentence.

⁶ Interestingly, in our Twitter corpus, when California is spelled as a Welsh word as *Califfornia*, it always mutates in mutation context; when it is spelled *California* as in English, it may or may not mutate. In other words, the decision to treat it as a Welsh word orthographically seems to implicate treating it as a Welsh word with respect to mutation.

⁷ This latter result is summarized and amplified in Deuchar et al. (2018).

We then noted whether the target item, in this case the word for brewery, was mutated *fragdy* [vragdi] or not *bragdy* [bragdi]. We would also note if a prompt was necessary and, if so, whether they then used the desired construction.

The experiment was conducted at Bangor University in Bangor, Wales. There were 84 items and 36 subjects. Items were presented in a single pseudo-random order first to last or last to first; half the subjects received the items in one order and the other half saw them in the reversed order. All items are given in the appendix. For subject responses, mutated items are coded as 2; unmutated items are coded as 1.

The experiment was designed to test various factors all designed to tap into the role of lexical category in mutation: i) lexical category of the triggering element, i.e. prepositions vs. adjectives; ii) lexical category of the element undergoing mutation, i.e. common nouns, verbs, and place names; and iii) frequency of place name as targets. In addition, though not relevant to our hypothesis here, triggers were selected so as to vary in terms of whether they ended with a vowel or consonant and mutation targets varied in terms of whether they began with a single consonant or a consonant cluster.

Our omnibus design is not suitable for a single analysis as not all factors interact. We therefore report several separate analyses. We have two random variables, subjects and items, so mixed effects modeling is appropriate. Since the dependent variable, mutation status, is a binary one, the data were analyzed using mixed effects logistic regression (Jaeger 2008).⁸ In all of our analyses, we follow the recommendations of Barr et al. (2013) using maximal design-based models with random slopes as appropriate.⁹

3.1 Lexical category of the trigger

Our first analysis examines the lexical category of the triggering item, specifically whether it is an adjective or a preposition. The means are given in Table 8 and plotted in Figure 1 (where again mutated items are coded as 2 and unmutated items are coded as 1 in both). Mutation is slightly more likely with a preceding adjective than with a preceding preposition.

Table 8: Effect of trigger type.

adjective	preposition
1.87	1.81

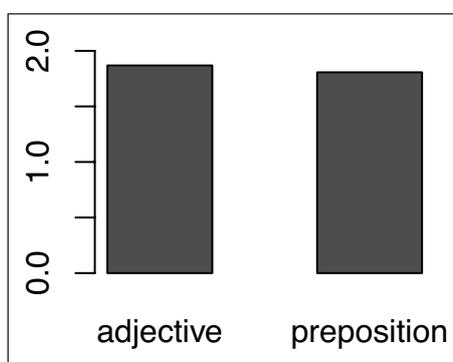


Figure 1: Effect of trigger type.

⁸ These were performed using the R (R Core Team 2014, version 3.4.3) `lme4` package (version 1.1-15).

⁹ We thus include all random slopes possible given our fixed effects. This also entails that we do not adjust models incrementally in the face of preliminary statistical analyses.

The effect of trigger part of speech is not significant as seen in the second row of Table 9.¹⁰ Based on the facts reviewed in Section 2, we did not anticipate an effect here.

3.2 Lexical category of the target

The next analysis is to determine if there is an effect of lexical category in terms of the target of mutation contrasting nouns, verbs, and place names. We used only Welsh place names, ones that the prescriptive rules say should mutate. We see in Table 10 that place names exhibit the least mutation, followed by nouns, and then verbs. This is plotted in Figure 2. With nouns as the reference level, the comparisons with place names and verbs are both significant as seen in rows two and three of Table 11.¹¹ This factor has three levels, but the relatively low rate of mutation with place names stands out.

Table 9: Effect of trigger part of speech.

	coef. est.	st. error	Pr (> z)
Intercept	4.412	0.818	0.00000
Preposition	-0.935	0.838	0.26495

Table 10: Effect of target type.

noun	place name	verb
1.90	1.59	1.95

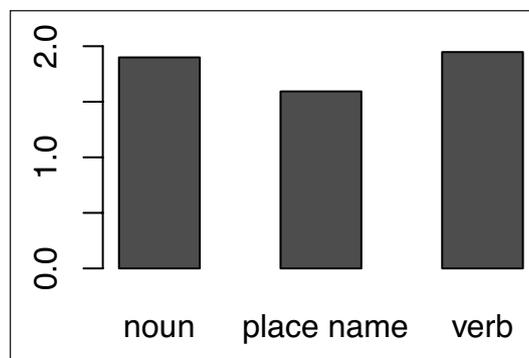


Figure 2: Effect of target type.

Table 11: Effect of target part of speech.

	coef. est.	st. error	Pr (> z)	
Intercept	4.75	0.733	0.00000	
Place name	-3.81	0.822	0.00000	*
Verb	3.78	1.331	0.00449	*

¹⁰ Here the reference level for trigger part of speech is *verb*. We provide the R equation for all mixed effects analyses here. The R equation used for this specific analysis is:

```
mut ~ trigger-pos + (1|items) + (1 + trigger-pos|subjects)
```

¹¹ The R equation used is:

```
mut ~ target-pos + (1|items) + (1 + target-pos|subjects)
```

3.3 Frequency of place names

We saw above that place names exhibit sharply reduced rates of mutation. We sought to probe this further by considering the potential role of lexical frequency. Our items included two classes of Welsh place names: relatively high-frequency items and relatively low-frequency items. See Table 12.

Note that frequency was assessed in terms of northern Welsh speakers. Thus, for example, Tremadog is a fairly small town, but quite well-known in the north.¹²

The infrequent places are all small towns in central and southern Wales. All the names are well-formed morphologically and are phonotactically unobjectionable. To make sure that subjects treated them as Welsh, all subjects were informed in advance that the experiment included the names of small towns in south Wales that they might not know.¹³

Table 13 shows the rate of mutation for high- and low-frequency place names and how that variable interacts with the lexical category of the mutation trigger (which we already treated on its own in Subsection 3.1 above). This is plotted in Figure 3. High-frequency place names exhibit a higher rate of mutation than low-frequency place names. Prepositions appear to trigger less mutation than adjectives.

As we see in Table 14, the overall effect of frequency is significant (row 2) but not the effect of the lexical category of the trigger (row 3) or the interaction (row 4).¹⁴ The

Table 12: Frequent and infrequent place names.

Frequent	Cymru	[kʰəmri:]	‘Wales’
	Prydain	[pʰrədajɪn]	‘Britain’
	Bangor	[baŋgɔr]	
	Caerdydd	[kʰauɾdi:ð]	‘Cardiff’
	Cricieth	[kʰɾikjɛθ]	
	Tremadog	[tʰrɛmadɔg]	
Infrequent	Penbryn	[pʰɛnbɾɪn]	
	Cribyn	[kʰɾɪbɪn]	
	Talsarn	[tʰalsarn]	
	Trefenter	[tʰrɛvɛntɛr]	
	Glanrhyd	[glanɾi:d]	
	Cilcennin	[kʰɾɪlkʰɛnɪn]	

Table 13: Effect of lexical frequency on mutation with place names.

	adjective	preposition
high	1.88	1.86
low	1.48	1.27

¹² Frequency in the north was assessed informally by the Welsh-speaking authors who live or have spent time in north Wales.

¹³ Note that this leaves open the possibility that some subjects may have treated these as nonce forms. This would mean that subjects assumed we were not telling the truth about these being actual places in south Wales. If subjects did do this, then the risk is that subjects might treat nonce forms differently from low-frequency items, that they are not just the extreme end of low frequency. This, of course, could be tested with a follow-up experiment that presented subjects with more degrees of frequency. Thanks to an anonymous reviewer for drawing this possibility to our attention.

¹⁴ Here, the reference level for frequency is *high* and the reference level for trigger part of speech is *adjective*. The R equation used is:

$$\text{mut} \sim \text{freq} * \text{trigger-pos} + (1|\text{items}) + (1 + \text{freq} * \text{trigger-pos}|\text{subjects})$$

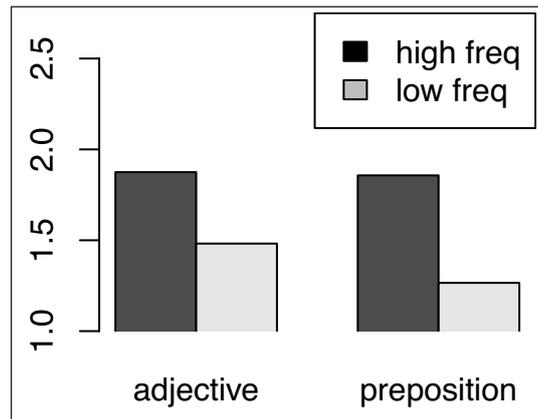


Figure 3: Effect of lexical frequency on mutation with place names.

Table 14: Effect of frequency and trigger part of speech on mutation of place names.

	coef. est.	st. error	Pr (> z)	
Intercept	3.419	1.35	0.01140	
Low	-3.607	1.76	0.03997	*
Preposition	-0.199	1.55	0.89789	
Interaction	-1.911	1.98	0.33532	

latter are perhaps unsurprising as we saw no main effect of the lexical category of the trigger either.

Interestingly, the frequency effect for place names is affected by how many times subjects hear an unfamiliar place name over the course of the experiment. Several of the place names we used were repeated over the course of the experiment with different triggers: *Cymru*, *Bangor*, *Caerdydd*, *Cilcennin*, *Penbryn*, and *Talsarn*. Since we presented the whole experiment in a single pseudo-random order and then in that order reversed, we can examine whether repeating an item increases the likelihood of mutation. In Figure 4, we plot the mean mutation values for place names separated by frequency and whether it was the first, second, or third repetition. Solid lines show high-frequency items and dashed lines show low-frequency items. The two different orders are indicated with color: black lines give one order and red lines give the reversed order. The shape of the lines in general is not itself meaningful here as different triggers were involved. However, the change in the shape of that line under the two presentations is meaningful. We see that the lines shift position as a function of presentation order which tells us that repetition does seem to affect mutation. We also see a difference in how extreme that shift is as a function of frequency so we may see a different effect in the two cases.

Turning now to significance testing, in Table 15 we see significant main effects of order (row 2) and frequency (row 4).¹⁵ We also see a significant interaction between order and presentation (row 5) confirming that the number of times subjects hear a place name has an effect on the likelihood of mutation.

Summarizing, we see three main effects in our behavioral study. First, there is an effect of lexical category with place names mutating the least, followed by nouns, followed by verbs. Second, frequent place names mutate more readily than infrequent place names.

¹⁵ Here, the reference level for frequency is *low* and the reference level for presentation is *forward*. The R equation used is:

```
mut ~ order + pres + freq + order:pres + order:freq + pres:freq + (1 +
pres|items) + (1 + order + freq + order:freq|subjects)
```

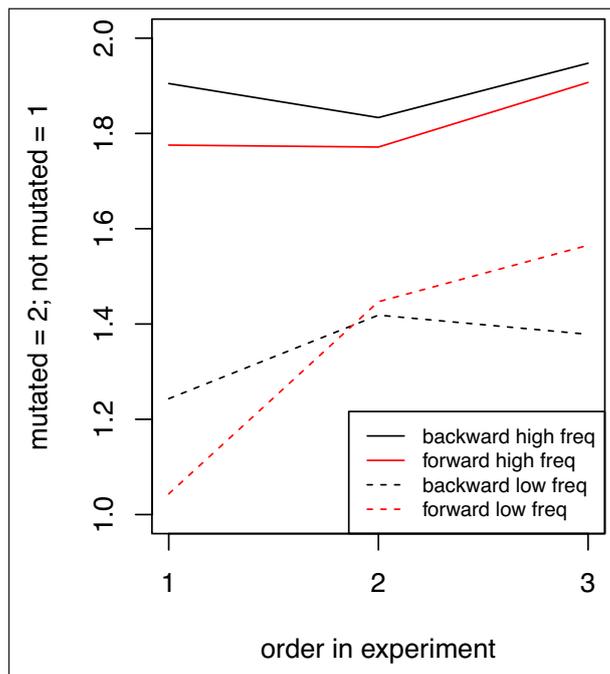


Figure 4: Effect of repetition on mutation for place names.

Table 15: Effect of order, presentation and frequency on mutation of place names.

	coef. est.	st. error	Pr (> z)	
Intercept	-6.341	2.351	0.00698	
Order	2.308	0.999	0.02083	*
Presentation	2.907	1.812	0.10878	
Frequency	6.456	2.933	0.02770	*
Order and Presentation	-1.419	0.598	0.01759	*
Order and Frequency	-0.797	1.382	0.56394	
Presentation and Frequency	-0.067	1.395	0.96149	

Third, number of repetitions in the experiment also affects mutation such that when a place name is repeated more, it is more likely to mutate.

4 Corpus analysis

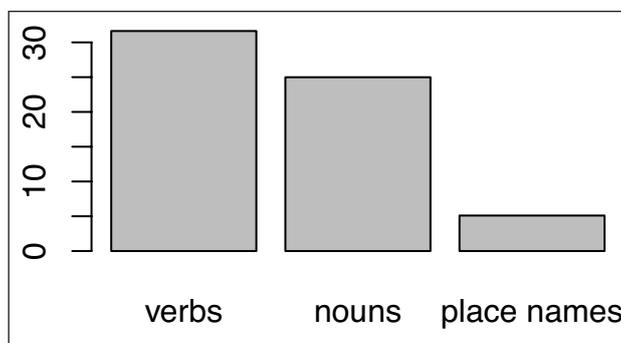
We now turn to corpus data to see if we can make sense of the patterns we saw in our experimental data. Do we see the same effect of part of speech on targets in corpus data that we saw in the experimental data? Do we also see an effect of frequency? And the key question: are these two effects distinct? We will see that, in fact, frequency effects drive the apparent category effects we’ve seen in our behavioral data.

Can we disassociate frequency and part of speech with corpus data? In fact, using data from the CEG corpus (Ellis et al. 2001), these variables are strongly associated. In Table 16 we see the mean counts for verbs, nouns, and place names in the CEG corpus, calculated as the average frequency for all words in each of those categories. This is plotted in Figure 5.

The difference between nouns and verbs is not significant, $t(3247.898) = 1.620$, $p = 0.105$, but the difference between verbs and place names is: $t(2963.295) = 6.618$, $p < .001$. The difference between nouns and place names is also significant: $t(10771.748) = 16.925$, $p < .001$. The upshot of this is that frequency correlates with target part of

Table 16: Mean counts for verbs, nouns, and place names in the CEG corpus.

Part of speech	Mean count
Verbs	31.633
Nouns	24.985
Place names	5.102

**Figure 5:** Mean counts for verbs, nouns, and place names in the CEG corpus.

speech, such that more frequent items are more likely to undergo mutation, consistent with the general claim that frequency drives the effect, not part-of-speech per se.

But is frequency an *independent* effect from part of speech?

We go to data from Twitter now to test this. Twitter is a much more unedited and unfiltered corpus and we can expect to see more variation in the distribution of mutation than in the CEG corpus. The corpus we use contains over 7 million Welsh-language tweets collected over several years (Jones et al. 2015).

To get a sense of how the language of Twitter differs from other sources, here are a few tweets from the beginning of the corpus. Even without translations, you can see that there are a few obvious differences. First, there are bits of text typical for the internet and twitter: URLs, hashtags, responses to another twitter user (indicated with @). Another difference is that there is a fair amount of code switching, which is also fairly typical of the spoken language. Finally, there's a fair amount of slang, misspellings, and non-standard dialect forms.

- i. Sut mae unigolion a gwladwriaethau yn ymateb i risg? Dewch i ddysgu mwy am hinsawdd, dŵr, bwyd a haen osôn y byd: <http://t.co/w3MK8gZRkz>
- ii. pob lwc i bawb sy'n derbyn eu canlyniadau heddiw! #LefelA
- iii. RT @lliwiol: dyna ni @yrawrgymraeg drosodd am wythnos arall, noswaith dda #yagym
- iv. @DelythMairEvans awks sai yn e....
- v. RT @Elenmair28: Ar @S4Carlein ishe dangos mwy o dalent pobol ifanc CFFI, pantos,eisteddfod, hanner awr adloniant a RHAID arddangos 'Nyth
- vi. @Elenmair28 @gwawrwilliams @sianowilliams co waint mor ciwt i fi guys, xxxxxxxxx <http://t.co/518s9bAz>
- vii. @rhodriorgan @Osian_Davies nos da rhodersss :)
- viii. @rhodriorgan @Osian_Davies joio sing along ni nithwr guys, #Ifeltit yee ar ol yr anthem, pam nin plano hwn? #gays
- ix. @rhodriorgan @Osian_Davies ok, cwpwl o jog's wthnos ma nare I dwymo lan ;) na naaa, ar y diwedd, rhag ofan gewn ni'n twli mas ;) hahaa
- x. @Osian_Davies @rhodriorgan bydd rhaid stretcho cyn trial clyrio'r ffences na, ddimishetynnumuscle move over elin fflur....

- xi. @rhodriorgan @Osian_Davies llanddarog, I'm there! Ok rhoders? #nojibbing ti sy rhi fast i ni, #toofastforusbro
- xii. @Osian_Davies cawn weld, dybynnu ar yr hwyl ;) @rhodriorgan jwmpo ar ben stage job? Hahaha #creiziii

We searched through the corpus for any of the targets we used in our behavioral study and counted the number of occurrences in the same mutation contexts we used in our experiment. We also collected the total counts for each target item in its unmutated or mutated form. The total counts are not normally distributed, so we log-scale them as in Figure 6.

Figure 7 shows the mean mutation frequency for our items in the Twitter corpus by log total and part of speech. Here we split log count into two categories: high and low. One can see that rate of mutation increases for both part of speech and for log count. This would seem to suggest that our basic behavioral effects show up in corpus data too.

If we run this as a regression, then log total is significant, part of speech is not, and the interaction is not. This is given in Table 17: $R^2 = 0.37$, $F(5, 29) = 3.39$, $p = 0.016$. This is consistent with the effect being driven by log count rather than by part of speech.

We can also test this with a likelihood ratio test. If we put both part of speech and log total into a model and then drop part of speech, there is no significant effect: $X^2(3) = 4.18$, $p = 0.12$. On the other hand, if we drop log total, then the effect is significant: $X^2(4) = 6.37$, $p = 0.01$. This is consistent with our conclusion that frequency drives the effect.

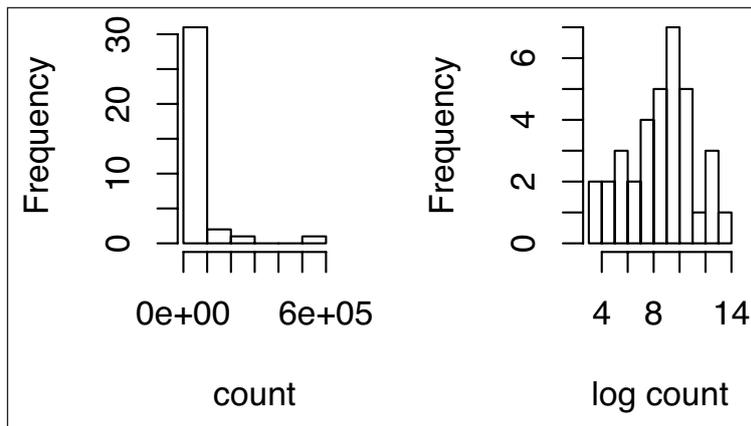


Figure 6: Log transforming total counts from Twitter corpus.

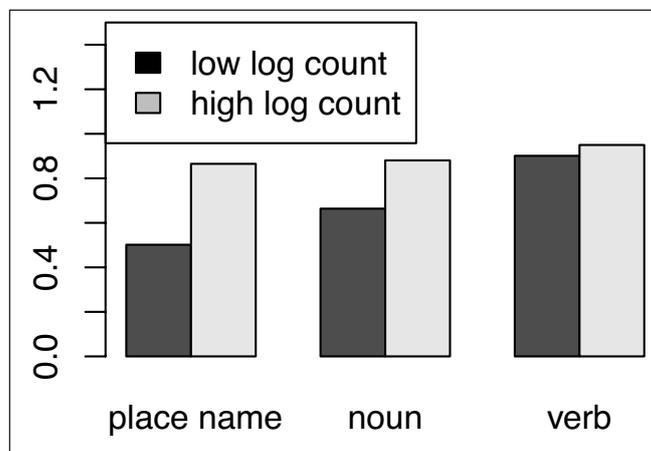
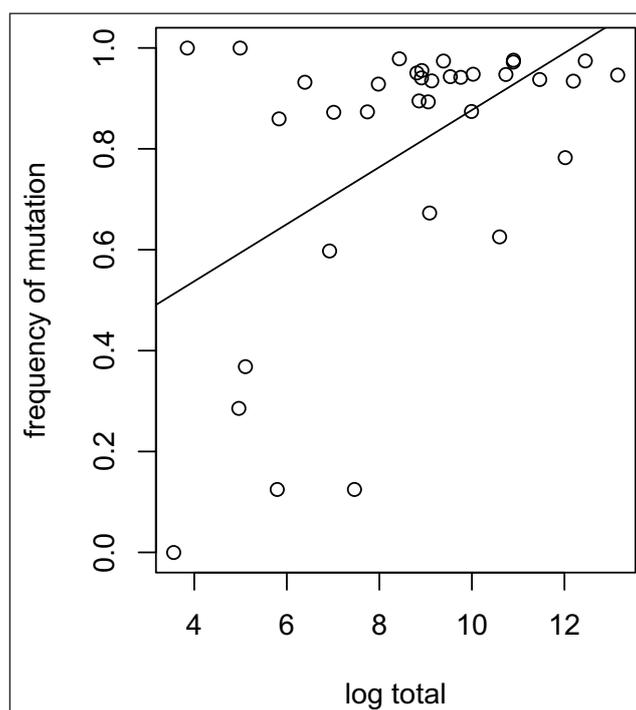


Figure 7: The effect of log count and part of speech on frequency of mutation.

Table 17: Regression for log count and part of speech on twitter data.

	b	t	p	
Intercept	0.19	1.03	0.31	
Log count	0.06	2.62	0.014	*
POS, noun	0.42	1.19	0.242	
POS, verb	0.65	1.14	0.264	
Log count & POS, noun	-0.04	-0.97	0.341	
Log count & POS, verb	-0.05	-0.82	0.417	

**Figure 8:** Effect of log total on mutation.

To fully appreciate the relationship, we now drop part of speech from the regression model in Table 17 and plot the regression line for log total against rate of mutation in Figure 8. Here each point represents an individual item showing the effect of log total on the relative frequency of mutation.

In summary, our corpus data also show category effects and frequency effects. Closer analysis shows that frequency is the driving factor and that part of speech does not contribute significantly to the model.

5 English

Given that we've seen that frequency seems to be a stronger predictor of mutation than part of speech, it's worth looking back at Moreton et al.'s effects and see if there is potentially a frequency effect there as well. In other words, are their effects actually due to lexical category or to frequency?

Recall that Moreton et al. constructed blends that varied in terms of how much of each word appears in the blend. They demonstrate a number of effects with this task including the category effects we explore here. For example, as already noted above, they found that subjects were more inclined to accept *sopraning* over *sopreening* when *soprano* was

interpreted as referring to the TV program *The Sopranos* than if it referred to a type of singer. Loosely, more of the word is preserved in blending if it is a proper noun than if it is a common noun.

This is a different sort of process than consonant mutation, most obviously because the input is comprised of two words. What do we expect if frequency drives the category effects? The most reasonable interpretation would be that more frequent words should play a bigger role in the blend. In other words, more of a frequent form should appear in the ultimate blend form. In the example above, we would expect the TV program interpretation to be more frequent than the singer interpretation.

Moreton et al. report several studies. We set aside their studies with respect to constituency, branching, and position of stress. They also report two studies that compare nouns and verbs and two that compare common nouns and proper nouns.

Let's look first at the studies comparing nouns and verbs, specifically their experiments 3a and 3b. Experiment 3b involved blends of either a verb or a noun with another noun. The dependent variable is how much of the first or second word is preserved in the blend. The relevant independent variable is whether the first word is interpreted as a verb or a noun. For example, subjects were asked to judge the acceptability of *floatex* vs. *flatex* as a blend of *float* and *latex*. Subjects were told that the blend meant either 'latex that is used to waterproof a parade float' (N + N) or 'latex that is light enough to float' (V + N). What they find is that the verbal interpretation biases subjects toward the blend form that preserves less of the verb, i.e. *flatex* in this case.

To check for a frequency effect, we examined all their experimental items in the first 100 million words of the Wacky corpus (Marco Baroni & Zanchetta 2009). This corpus is useful here because it is extremely large and all words are tagged for part of speech. We can therefore get fairly accurate relative counts for all the items Moreton et al. use. These are given in Table 18.

Table 18: Items from experiments 3a and 3b.

word	noun	verb
drain	444	250
drag	447	446
brood	87	1
creep	1	211
plot	2144	145
club	9761	2
spot	2891	625
break	4108	2674
storm	1253	4
fling	56	23
float	418	275
slip	699	363
spell	800	391
clog	23	42
crop	1590	96
block	2968	738
grouse	110	8

Mean values are given in Table 19 and plotted in Figure 9. Nouns are more frequent than verbs. Since these are count data, they are not normally distributed and we therefore log-transform them. The difference is significant in a paired t -test: $t(16) = 2.266, p = 0.038$.

It's fair to conclude that Moreton et al.'s results with respect to the distinction between nouns and verbs is consistent with the frequency-based story we've developed here. Their experimental items are more frequent when they are tagged as nouns than when they are tagged as verbs. Hence we expect them to be preserved in blending more when they are nouns.

We can also look at their experiments that involve proper nouns vs. nouns. Items and counts are given in Table 20. Note that their proper noun category includes what we would term place names.

Table 19: Means for nouns and verbs in experiments 3a and 3b.

Noun	Verb
1635.29	370.24

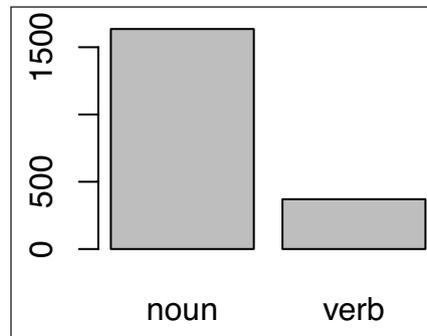


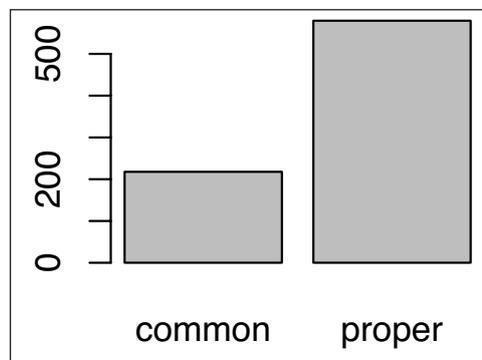
Figure 9: Means for nouns and verbs in experiments 3a and 3b.

Table 20: Items from experiments 5a and 5b.

word	common	proper
bohemian	7	2
soprano	192	48
cologne	15	196
canary	44	421
chihuahua	12	37
superior	2	180
independence	1813	315
crusade	223	10
narcissus	28	44
turkey	298	1581
jersey	239	788
sparrow	78	253
buffalo	81	153
china	288	4925
hamlet	314	218
potter	60	962
boulder	215	94
homer	10	202

Table 21: Means for nouns and proper nouns in experiments 5a and 5b.

Noun	Proper noun
217.72	370.24

**Figure 10:** Means for nouns and verbs in experiments 5a and 5b.

Mean values are given in Table 21 and plotted in Figure 10. Proper nouns are more frequent than nouns. Again, the counts are not normally distributed and we log-transform them. The difference here only trends in a paired t -test: $t(17) = -1.830$, $p = 0.085$.¹⁶

To conclude, the blending facts from Moreton et al. (2017) with respect to nouns vs. proper nouns and verbs vs. nouns are consistent with the frequency story developed here.

6 Formalizing the role of frequency

In this section, we propose an account of these facts within a version of Optimality Theory (McCarthy & Prince 1993; Prince & Smolensky 1993) that makes use of lexically-conditioned constraints (Hammond 1999; Pater 2000) and weighted constraints as in Harmonic Grammar (Smolensky 2006; Pater 2009; Potts et al. 2010).¹⁷ We'll need lexically-conditioned constraints to capture the fact that lexical items behave differently. We'll need weighted constraints to capture trade-offs and the gradient nature of the system.

We have seen that frequency plays a significant role in the distribution of mutation in Welsh and in the pattern of blending in English. In the case of Welsh, mutation is more likely with more frequent forms. In the case of English, retention of material in blends is greater with more frequent forms.

These effects seem to go in opposite directions. In the case of Welsh, more frequent forms are *less* likely to be preserved, because they are more likely to be mutated. Thus *Cymru* [k^həmri] 'Wales' is more likely to mutate than *Cribyn* [k^hribin] 'Cribyn'. In the case of English, more frequent forms are *more* likely to be preserved in blends. Thus *sopranning* is more like to be preferred over *sopreening* when *soprano* refers to the more frequent name of the TV program, than to a type of singing voice.

This argues against a treatment in terms of lexical faithfulness in Optimality Theory. The basic idea is that there would be separate faithfulness constraints for individual lexical items. These constraints are ranked in terms of the frequency of the lexical item. Thus relatively infrequent items would have high-ranked faithfulness constraints, while

¹⁶ We use the term "trend" to refer to a p -value less than .1 and greater than .05.

¹⁷ The account could also be tweaked for Maxent modeling (Hayes & Wilson 2008) or noisy Harmonic Grammar (Coetzee & Pater 2011; Coetzee & Kawahara 2013).

more frequent items would have low-ranked faithfulness constraints. We can schematize this as in Figure 11.

In the case of Welsh, *Cribyn* resists mutation because its lexical faithfulness constraint outranks the pressure to mutate; *Cymru* mutates more readily because its faithfulness constraint is outranked by the pressure to mutate. In Welsh, the faithfulness constraint corresponding to the less frequent form is the higher-ranked. In the case of English, *sopraning* is preferred to *sopreening* for *Sopranos* because the corresponding faithfulness constraint outranks the pressure to blend. Here the faithfulness constraint corresponding to the *more* frequent form is higher-ranked.

If the ranking is to be a consistent consequence of lexical frequency, we must rule out an account in terms of lexical faithfulness. Instead, we develop an account in terms of surface correspondence (McCarthy & Prince 1995). The basic idea is that we have correspondence constraints with respect to *surface* forms and these are weighted with respect to the grammatical constraints of the two systems.

In the case of English, the system is unchanged: high-ranked correspondence to more frequent forms cause them to be more preserved in blends. In the case of Welsh though, we have correspondence constraints for both mutated and non-mutated forms where the more frequent the form is, the higher-ranked the corresponding correspondence constraint is. In the case of all Welsh words, the correspondence constraint for un-mutated forms must have a greater weight than that for mutated forms to capture the fact that in the absence of the pressure to mutate, the form surfaces as unmutated. The difference between forms like *Cribyn* and forms like *Cymru* is that in the latter case, the correspondence constraint for *Gymru* is ranked high enough to sometimes tip the balance.

For *Cribyn* in non-mutation context, we would have something like Table 22. Here we've provided weights that capture the intuitions expressed above. The winning pronunciation is the one with the lowest weighted sum of violations. The correspondence constraint for *Cribyn* is stronger than that of *Gribyn*. What we see then is that in a non-mutation context, it is better to not mutate.

For a low-frequency item like *Cribyn*, it's also better not to mutate in a mutation context as in Table 23.

Welsh: FAITH(*Cribyn*) ≫ MUTATE ≫ FAITH(*Cymru*)
 English: FAITH(*Sopranos*) ≫ BLEND ≫ FAITH(*soprano*)

Figure 11: How lexical faithfulness would work.

Table 22: *Cribyn* in non-mutation context.

	Cribyn	C(Cribyn) <i>w</i> = 6	MUTATE <i>w</i> = 4	C(Gribyn) <i>w</i> = 1	Total
☞	Cribyn			1	1
	Gribyn	6			6

Table 23: *Cribyn* in mutation context.

	i Cribyn	C(Cribyn) <i>w</i> = 6	MUTATE <i>w</i> = 4	C(Gribyn) <i>w</i> = 1	Total
☞	i Cribyn		4	1	5
	i Gribyn	6			6

For a high-frequency item like *Cymru*, the key difference is that the weight for the correspondence constraint for *Gymru* is higher. This has no effect in non-mutation context as in Table 24.

Finally, we see the need for finite constraint weights as in Harmonic OT when we consider *Cymru* in mutation context as in Table 25. In this case, the weight of the correspondence constraint for *Gymru* when added to the weight of the general mutation constraint is sufficient to force mutation.

We’ve used specific weights above to get the effects desired, but other weights are possible. For this story to go through, there are two constraints on the weights. First, the weight for the unmutated form must exceed the weight for the mutated form. This corresponds to the fact that unmutated forms are generally more frequent than mutated forms and guarantees that the unmutated form will show up in the absence of mutation. The second property that must hold is that, for mutating forms like *Cymru*, the weight of the correspondence constraint for *Gymru* must be greater than the difference between the weights for the unmutated form and the general mutation constraint.

The account makes several interesting predictions. First, what happens if the weight for the mutated form should exceed the weight for the unmutated form? In such a case, the unmutated form will never show up. This is effectively reanalysis. This, in fact, seems to be happening for some speakers with respect to the item *tref* [t^hre (v)] ‘village’. In our Twitter data, there are a number of examples of the mutated form of this showing up in non-mutation contexts. Following are examples of tweets with *dre* (*f*) showing up after the preposition *mewn* [mewn] ‘in’ which does not trigger mutation.¹⁸

- i. RT @BethanWalkling: #hoffdafarn @ClwbYBont clwb y bont- lle fach llawn pobol hyfryd, dihangfa fach cymreig *mewn dre* eitha seisnigedd! #j ...
- ii. @idrischarles theatr dda iawn *mewn dref* llawn siopau bach gwydch.
- iii. @gaitoms @dylmei fysa, wast fel just ty tafarn arall *mewn dref* hefo gormod o rhai gwag yn barod!
- iv. Caernarfon Hanesyddol: y Rhufeiniaid: Yn amlwg mi ges i fy magu *mewn dre* llawn o hanes gyda raenau o wahanol g... <https://t.co/3jQLCXtcJj>
- v. Parti/gig chweched wedi ei gyhoeddi am Hydref 30ain yn The Scene, *mewn dre*. Mynediad £3 ar y drws, croeso cynnes i bawb xxxxx

Table 24: *Cymru* in non-mutation context.

	Cymru	C(Cymru) w = 6	MUTATE w = 4	C(Gymru) w = 3	Total
☞	Cymru			3	3
	Gymru	6			6

Table 25: *Cymru* in mutation context.

	i Cymru	C(Cymru) w = 6	MUTATE w = 4	C(Gymru) w = 3	Total
	i Cymru		4	3	7
☞	i Gymru	6			6

¹⁸ This general phenomenon has been noted before. See, for example, Thomas (1984).

- vi. @gaitoms @dylmei fysa, wast fel just ty tafarn arall *mewn dref* hefo gormod o rhai gwag yn barod!
- vii. Dwi'n byw *mewn dre* bach del. <http://t.co/pCjbKufjgX>

Notice that a side effect of this analysis is that the actual input form stays the same, so that while the apparently mutated form occurs in non-mutation contexts, the same form occurs in mutation contexts. In other words, this reweighting of constraints makes the correct prediction that we do *not* see superficially doubly mutated forms like *ddre(f)* [ðre(v)]. Indeed, there are no such forms in our Twitter data.¹⁹

Another prediction made by this account is that the frequency effect is driven not by the overall frequency of the word, but by the frequency of the mutated form. This, in fact, is testable with our corpus data. If we go back to our Twitter data and do a simple regression from log total to mutation rate, we get a significant effect as in Table 26: $R^2 = 0.26$, $F(1,33) = 11.35$, $p = 0.002$.²⁰

However, if we do a regression from the log count of mutated forms only, as in Table 27, the result is still significant, but we get a much higher R^2 : $R^2 = 0.47$, $F(1, 33) = 29.66$, $p < .001$. The greater R^2 for the second analysis supports the formal analysis we've proposed above.

We've seen that we can implement the role of frequency using well-established tools in the grammatical sphere: constraint weighting and lexical constraints. This formal analysis is not inextricably tied to the particular formal system we've used however. It would be possible to express the same ideas using other constraint-based formalisms, e.g. Maxent modeling, Stochastic OT, or Noisy Harmonic Grammar.

Setting aside the formal system, however, the analysis is quite intuitive: the likelihood of mutating a form depends on how often we've heard the mutated form itself.

7 Conclusion

Summarizing, we saw in our behavioral study that Welsh mutation is indeed subject to lexical category effects. We saw that lexical category effects extend beyond major categories like nouns and verbs, and beyond proper names, to also include place names. We saw that there are also frequency effects.

In our corpus study, we replicated these category and frequency effects. We also saw that the category effects do not contribute significantly beyond their role in frequency effects.

Table 26: Regression for log count.

	b	t	p	
Intercept	0.31	2.06	0.047	
Log count	0.06	3.37	0.002	*

Table 27: Regression for log mutation count.

	b	t	p	
Intercept	0.41	5.24	<.001	
Log count	0.07	5.45	<.001	*

¹⁹ Thanks to an anonymous reviewer for very helpful discussion here.

²⁰ In this and the following analysis we do add-one smoothing on the independent measure to avoid taking the log of 0.

We then turned to the noun vs. verb and proper noun vs. noun distinctions treated in Moreton et al. (2017) and saw that, given the specific items used in the relevant experiments, those results are also consistent with a frequency effect.

We can now hypothesize that similar lexical category effects others have seen also correlate with frequency effects.

Smith (2011) observed that category effects are not always the same. In some languages nouns are more faithful than verbs and in other languages the reverse is true. We hypothesize that this occurs when lexical frequency relationships reverse as well. Our conclusion that lexical frequency drives the category effects of Welsh thus provides a potential solution to this previously unexplained aspect of the category-based treatment.

Note that we are not claiming that all grammatical effects follow from frequency: our analysis deals only with target part-of-speech effects. We are also not claiming that *all* category effects necessarily derive from frequency effects. The mutation and blending effects treated here both involve degree of application or whether some process applies and it is straightforward to see this in terms of frequency of the relevant targets. Other category effects, like nominal vs. verbal stress in English (Chomsky & Halle 1968; Hayes 1981; Hayes 1995), are difficult to see in these terms. We hypothesize that category effects like these are not due to frequency.

Finally, a treatment in terms of lexical frequency makes good sense theoretically. This accords with the general principle that high-frequency items participate in the grammar of a language more fully than low-frequency items. We've shown how it is possible to implement the central intuition of the analysis using constraint weighting and lexical constraints. We've also shown how the particular formal analysis we developed makes additional correct predictions.

Abbreviations

1SG = first singular, SOFT = soft mutation, MASC = masculine, FEM = feminine, PRT = particle

Acknowledgements

Earlier versions of this research were presented at the annual Welsh Linguistics Seminar at Gregynog and benefited from much useful feedback there. The following individuals helped us run our experiments in Wales: Sara Huws, Sam Johnston, Nick Kloehn, Daniel Olson, Nia Parry, and Anna Weesner. We very gratefully acknowledge all the support we received from faculty, students, and staff at Bangor University. Special thanks also to Kerry McCullough. Thanks to Maggie Tallerman, several anonymous reviewers, and the editor for very useful feedback. All errors are our own.

Funding Information

This research was supported by NSF grant BCS-1453724 (Michael Hammond, Diana Archangeli, Heddwen Brooks, Andrew Carnie, Diane Ohala, Adam Ussishkin, Peredur Webb-Davies, and Andy Wedel).

Competing Interests

The authors have no competing interests to declare.

References

Albright, Adam. 2008. How many grammars am I holding up? Discovering phonological differences between word classes. In Charles B. Chang & Hannah J. Haynie (eds.), *Proceedings of the 26th west coast conference on formal linguistics*, 1–20. Somerville, MA: Cascadilla Proceedings Project.

- Alderete, John D. 2001. Dominance effects as trans-derivational antifaithfulness. *Phonology* 18. 201–253. DOI: <https://doi.org/10.1017/S0952675701004067>
- Awbery, Gwenllian M. 1973. Initial mutation in a generative grammar of Welsh. *Phonetics Department Report* 4. 28–46. University of Leeds.
- Ball, Martin J. & Nicole Müller. 1992. *Mutation in Welsh*. London and New York: Routledge.
- Barr, Dale J., Roger Levy, Christoph Scheepers & Harry J. Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language* 68. 255–278. DOI: <https://doi.org/10.1016/j.jml.2012.11.001>
- Becker, Michael & Maria Gouskova. 2016. Source-oriented generalizations as grammar inference in Russian vowel deletion. *Linguistic Inquiry* 47. 391–425. DOI: https://doi.org/10.1162/LING_a_00217
- Chomsky, Noam & Morris Halle. 1968. *The sound pattern of English*. New York: Harper & Row.
- Coetzee, Andries W. 2009. Learning lexical indexation. *Phonology* 26. 109–145. DOI: <https://doi.org/10.1017/S0952675709001730>
- Coetzee, Andries W. & Joe Pater. 2011. The place of variation in phonological theory. In John Goldsmith, Jason Riggle & Alan Yu (eds.), *Handbook of phonological theory*, 401–434. Cambridge, MA: Blackwell. DOI: <https://doi.org/10.1002/9781444343069.ch13>
- Coetzee, Andries W. & Shigeto Kawahara. 2013. Frequency biases in phonological variation. *Natural Language and Linguistic Theory* 31. 47–89. DOI: <https://doi.org/10.1007/s11049-012-9179-z>
- Deuchar, Margaret, Peredur Webb-Davies & Kevin Donnelly. 2018. *Building and using the siarad corpus*. Amsterdam/Philadelphia: John Benjamins. DOI: <https://doi.org/10.1075/scl.81>
- Dorian, Nancy C. 1973. Grammatical change in a dying dialect. *Language* 49. 413–438. DOI: <https://doi.org/10.2307/412461>
- Dorian, Nancy C. 1978. The fate of morphological complexity in language death: Evidence from East Sutherland Gaelic. *Language* 54. 590–609. DOI: <https://doi.org/10.1353/lan.1978.0024>
- Dorian, Nancy C. 1981. *Language death: The life cycle of a Scottish Gaelic dialect*. Philadelphia, PA: University of Pennsylvania Press. DOI: <https://doi.org/10.9783/9781512815580>
- Ellis, N. C., C. O’Dochartaigh, W. Hicks, M. Morgan & N. Laporte. 2001. Cronfa electroneg o Gymraeg (CEG): A 1 million word lexical database and frequency count for Welsh. <http://www.bangor.ac.uk/canolfanbedwyr/ceg.php.en>.
- Fidelholtz, James. 1975. Word frequency and vowel reduction in English. In *Chicago Linguistic Society* 11. 200–213.
- Green, Antony D. 2006. The independence of phonology and morphology: The Celtic mutations. *Lingua* 116. 1946–1985. DOI: <https://doi.org/10.1016/j.lingua.2004.09.002>
- Hammond, Michael. 1999. Lexical frequency and rhythm. In M. Darnell, E. Moravcsik, F. Newmeyer, M. Noonan & K. Wheatley (eds.), *Functionalism and formalism in linguistics*, 329–358. Amsterdam: John Benjamins.
- Hammond, Michael. 2004. Frequency, cyclicity, and optimality. *Studies in Phonetics, Phonology, and Morphology* 10. 349–364.
- Hammond, Michael. 2011. Welsh mutations and statistical phonotactics. In Andrew Carnie (ed.), *Formal approaches to Celtic linguistics*, 337–358. Newcastle upon Tyne: Cambridge Scholars Publishing.
- Hammond, Michael, Natasha Warner, Andréa Davis, Andrew Carnie, Diana Archangeli & Muriel Fisher. 2014. Vowel insertion in Scottish Gaelic. *Phonology* 31. 123–153. DOI: <https://doi.org/10.1017/S0952675714000050>

- Hammond, Michael, Yan Chen, Elise Bell, Andrew Carnie, Diana Archangeli, Adam Ussishkin & Muriel Fisher. 2017. Phonological restrictions on lenition in Scottish Gaelic. *Language* 93. 446–472. DOI: <https://doi.org/10.1353/lan.2017.0020>
- Hannahs, S. J. 2011. Celtic mutations. In Marc van Oostendorp, Colin Ewen, Keren Rice & Elizabeth Hume (eds.), *The Blackwell companion to phonology* 5. 2807–2830. Malden, MA: Wiley-Blackwell. DOI: <https://doi.org/10.1002/9781444335262.wbctp0117>
- Hannahs, S. J. 2013. *The phonology of Welsh*. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/acprof:oso/9780199601233.001.0001>
- Hayes, Bruce. 1981. *A metrical theory of stress rules*. New York: Garland. 1980 MIT doctoral dissertation.
- Hayes, Bruce. 1995. *Metrical stress theory*. Chicago, IL: University of Chicago Press.
- Hayes, Bruce & Colin Wilson. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39. 379–440. DOI: <https://doi.org/10.1162/ling.2008.39.3.379>
- Hooper, Joan. 1976. Word frequency in lexical diffusion and the source of morphophonological change. In William Christie (ed.), *Current progress in historical linguistics*, 96–105. Amsterdam: North-Holland.
- Inkelas, Sharon & Cheryl Zoll. 2007. Is grammar dependence real? A comparison between cophonological and indexed constraint approaches to morphologically conditioned phonology. *Linguistics* 45. 133–171. DOI: <https://doi.org/10.1515/LING.2007.004>
- Iosad, Pavel. 2010. Right at the left edge: Initial consonant mutations in the languages of the world. In Michael Cysouw & Jan Wohlgemuth (eds.), *Rethinking universals: How rarities affect linguistic theory*, 105–138. Berlin: Mouton de Gruyter. DOI: <https://doi.org/10.1515/9783110220933.105>
- Itô, J. & A. Mester. 1999. The phonological lexicon. In Natsuko Tsujimura (ed.), *The handbook of Japanese linguistics*, 62–100. Malden, MA: Blackwell.
- Itô, Junko & Armin Mester. 2009. Lexical classes in phonology. In Shigeru Miyagawa & Mamoru Saito (eds.), *Handbook of Japanese linguistics*, 84–106. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/oxfordhb/9780195307344.013.0004>
- Jaeger, T. Florian. 2008. Categorical data analysis: away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language* 59. 434–446. DOI: <https://doi.org/10.1016/j.jml.2007.11.007>
- Jones, D. B., P. Robertson & A. Taborda. 2015. Corpws Trydariadau Cymraeg. <http://techiaith.cymru/corpora/twitter>.
- Kibre, Nicholas J. 1997. *A model of mutation in Welsh*. Bloomington, IN: Indiana University Linguistics Club Publications.
- King, Gareth. 2003. *Modern Welsh: A comprehensive grammar*. London: Routledge. DOI: <https://doi.org/10.4324/9780203987063>
- Lieber, Rochelle. 1983. New developments in autosegmental phonology: Consonant mutation. In *WCCFL* 2. 165–175.
- Marco Baroni, Adriano Ferraresi, Silvia Bernardini & Eros Zanchetta. 2009. The wacky wide web: A collection of very large linguistically processed web-crawled corpora. *Language Resources and Evaluation* 43. 209–226. DOI: <https://doi.org/10.1007/s10579-009-9081-4>
- McCarthy, John & Alan Prince. 1993. Prosodic morphology: Constraint interaction and satisfaction. Amherst: University of Massachusetts.
- McCarthy, John & Alan Prince. 1995. Faithfulness and reduplicative identity. In Jill Beckman, Laura Walsh Dickey & Suzanne Urbanczyk (eds.), *Papers in optimality theory*, vol. 18 (U. Mass. Occasional Papers in Linguistics). 249–384. [ROA].

- Mittendorf, Ingo & Louisa Sadler. 2006. A treatment of Welsh initial mutation. In *Proceedings of the LFG06 conference*, 343–364. Stanford, CA: CSLI Publications.
- Moreton, Elliott, Jennifer L. Smith, Katya Pertsova, Rachel Broad & Brandon Prickett. 2017. Emergent positional privilege in novel English blends. *Language* 93. 347–380. DOI: <https://doi.org/10.1353/lan.2017.0017>
- Morgan, Thomas John. 1952. *Y treigladau a'u cystrawen*. Caerdydd (Cardiff): Gwasg Prifysgol Cymru.
- Morris-Jones, John. 1913. *A Welsh grammar: Historical and comparative*. Oxford: Clarendon.
- Pater, Joe. 2000. Non-uniformity in English secondary stress: The role of ranked and lexically specific constraints. *Phonology* 17. 237–274. DOI: <https://doi.org/10.1017/S0952675700003900>
- Pater, Joe. 2009. Weighted constraints in generative linguistics. *Cognitive Science* 33. 999–1035. DOI: <https://doi.org/10.1111/j.1551-6709.2009.01047.x>
- Potts, Christopher, Joe Pater, Karen Jesney & Rajesh Bhatt. 2010. Harmonic Grammar with linear programming: From linear systems to linguistic typology. *Phonology* 27. 77–117. DOI: <https://doi.org/10.1017/S0952675710000047>
- Prince, Alan & Paul Smolensky. 1993. *Optimality Theory*. University of Massachusetts, Amherst and University of Colorado at Boulder.
- Prys, Myfyr. 2015. *Style in the vernacular and on the radio: Code-switching and mutation as stylistic and social markers in Welsh*. Bangor, Wales: Bangor University dissertation.
- Pyatt, Elizabeth J. 1997. *An integrated model of the syntax and phonology of Celtic mutation*. Cambridge, MA: Harvard dissertation.
- R Core Team. 2014. *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. <http://www.Rproject.org/>.
- Shih, Stephanie S. & Sharon Inkelas. 2015. Morphologically-conditioned tonotactics in multilevel maximum entropy grammar. In Gunnar Ólafur Hansson, Ashley Farris-Trimble, Kevin McMullin & Douglas Pulleyblank (eds.), *Proceedings of the 2015 annual meeting on phonology*. *Linguistic Society of America*. DOI: <https://doi.org/10.3765/amp.v3i0.3659>
- Smith, Jennifer L. 2011. Category-specific effects. In Marc van Oostendorp, Colin Ewen, Beth Hume & Keren Rice (eds.), *The Blackwell companion to phonology*, 2439–2463. Malden, MA: Wiley-Blackwell.
- Smolensky, Paul. 2006. Harmony in linguistic cognition. *Cognitive Science* 30. 779–801. DOI: https://doi.org/10.1207/s15516709cog0000_78
- Stammers, Jonathan Roy. 2009. *The integration of English-origin verbs in Welsh*. Bangor, Wales: Bangor University dissertation.
- Stammers, Jonathan R. & Margaret Deuchar. 2012. Testing the nonce borrowing hypothesis: Counter-evidence from English-origin verbs in Welsh. *Bilingualism: Language and Cognition* 15. 630–643. DOI: <https://doi.org/10.1017/S1366728911000381>
- Tallerman, Maggie. 1990. VSO word order and consonantal mutation in Welsh. *Linguistics* 28. 389–416. DOI: <https://doi.org/10.1515/ling.1990.28.3.389>
- Tallerman, Maggie. 2009. Phrase structure vs. dependency: The analysis of Welsh syntactic soft mutation. *Journal of Linguistics* 45. 167–201. DOI: <https://doi.org/10.1017/S0022226708005550>
- Thomas, Peter Wynn. 1984. Variation in South Glamorgan consonant mutation. In Martin J. Ball & Glyn E. Jones (eds.), *Welsh phonology*, 208–236. Cardiff: University of Wales Press.
- Wolf, Matthew. 2007. For an autosegmental theory of mutation. In Leah Bateman, Michael O'Keefe, Ehren Reilly & Adam Werle (eds.), *University of Massachusetts occasional papers in linguistics 32: Papers in optimality theory iii*, 315–404. Amherst, MA: GLSA.

How to cite this article: Hammond, Michael, Elise Bell, Skye Anderson, Peredur Webb-Davies, Diane Ohala, Andrew Carnie and Heddwen Brooks. 2020. Category-specific effects in Welsh mutation. *Glossa: a journal of general linguistics* 5(1): 1.1–26. DOI: <https://doi.org/10.5334/gjgl.1007>

Submitted: 22 May 2019 **Accepted:** 24 October 2019 **Published:** 03 January 2020

Copyright: © 2020 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

 *Glossa: a journal of general linguistics* is a peer-reviewed open access journal published by Ubiquity Press.

OPEN ACCESS 