Selectional Effects in Allomorph Competition

Samantha Gordon Danner
University of Southern California

1   Introduction

Allomorphy is a complex and multifaceted phenomenon in linguistics that has been studied within a variety of theoretical frameworks and subfields. Part of the problem in describing allomorphy, even within a single language, is that there appears to be numerous underlying mechanisms driving the process(es) of allomorphy, and various linguistic contexts that condition alternations in the realization and/or pronunciation of a morpheme. In the allomorphy literature, there are two main theoretical motivations for allomorphy: one is purely morphologically conditioned, e.g. verbal stem alternation in Latin: *armare* (inf. 'to arm') can take three verbal stems, *arm–a–, arma–v–*, or *arma–t–* (example from Aronoff 1994, cited in Booij, 1997). Booij argues that the choice of the final segment in the stem is purely morphological and not conditioned by phonology, because there is no phonological trigger for the use of stem-final *–v* as opposed to stem-final *–t* in the perfect stems. The other motivation for allomorphy is purely phonologically conditioned, e.g. example 1:

(1) Korean subjunctive suffixation
   a.   aɪ-**wa**   'child and X'
         sɛ-**wa**   'bird and X'
   b.   pap-**kwa**   'rice and X'
         san-**kwa**   'mountain and X'

This example shows clear phonological conditioning in the allomorphy of the subjunctive: following a vowel, –**wa** is used; following a consonant, –**kwa** is used (Lapointe, 1999). There are also examples which are driven by phonology but which do *not* improve phonological structure with respect to structural markedness or faithfulness to the base. An example from Turkish, cited by Underhill (1976), is the passive marker. The marker is realized as *-n* following a V-final or l-final stem, and *–l* elsewhere.

(2) Turkish passive marking
         a.   ara-n-   'be sought'
               oku-n-   'be read'
         b.   çal-ın-   'be struck'
               bil-in-   'be known'
         c.   kullan-ıl-   'be used'
               kaybed-il-   'be lost'

The vowels preceding the marker in b and c are epenthetic, and do not contribute to the passive marker allomorphy of the underlying forms *–n* and *–l*. The examples show that, other than the dissimilatory use of *–n* in the examples in (2b), there is no clear phonological benefit from this distribution of allomorphs.

Occasionally, the morphological and phonological impulses for allomorphy are considered together (Wolf; 2008, 2013), but the examination of multiple conditioning factors in allomorphy remains sparse.

Proceedings of *AMP 2014*
Completed February 27, 2015

More work that considers the concurrent effects of phonological and morphological (and other) conditioning will serve to develop a stronger theory of allomorphy that accounts for all available evidence.

In this paper, I present the case of the *demonym* in English as an argument in favor of allomorphy conditioned by phonological and non-phonological factors. This will expand the present work on phonology-morphology interactions by showing that multiple factors interact in deriving a surface form, particularly when the stem, the affix, or their combination is unfamiliar to the speaker. Section 2 lays out the case for demonyms as a type of allomorphy explained by multiple conditioning factors. Section 3 describes a corpus investigation of demonyms in real usage contexts. Section 4 reviews a behavioral experiment that extends predictions of demonym usage to novel and unfamiliar forms. Section 5 sketches an analysis of demonym allomorphy. Finally, section 6 closes the paper with a consideration of the theoretical implications of this work and other concluding remarks.

2   Demonym Allomorphy in English

A demonym is "a name given to a people or inhabitants of a place;" (U.S. Government Printing Office, 2008) i.e., a resident of Iceland is an Iceland*er*, or a resident of Chicago is a Chicago*an*. Demonyms present a unique case in the derivational morphology of English: there are many allomorphs available to represent the same demonym intent (the set considered here is shown in Table 1), and these forms have a variety of etymological backgrounds and phonological profiles. The bases to which these demonym allomorphs affix are also etymologically and phonologically diverse, and novel place name bases are constantly added to the lexicon. Although formed from phonetically similar *and* dissimilar morpheme strings (*–an*, *–ard*, *–ene*, *–er*, *–ese*, *–i*, *–ian*, *–iot*, *–ish(wo)man*, *–ite*, *–o*, etc.), English demonym allomorphs generally surface in suffix position. With so many allomorphs, the phonological distribution of morphemes overlaps in many situations. In appending a demonym suffix to a place name base, phonological repair strategies such as base truncation, epenthesis, stress shift and vowel alternation may be employed. The result is that, depending on which strategies a speaker uses when forming a demonym, there may be many phonologically acceptable demonyms for a given place name. There is no reason that *Icelander* ~ [ˈaɪs.lən.ˌdɹ], e.g., may be phonologically preferred to *Icelandite* ~ [ˈaɪs.lən.ˌdaɪt] (both demonym suffixes add a syllable to the base, do not require stress shift in the base, and cause resyllabification of the base-final consonant), yet *Icelander* is the accepted form. As discussed below, certain phonological properties of place name bases trend towards a preference for a given demonym allomorph, but these trends are not absolute, and they are not always adopted by speakers in the formation of new demonyms.

Allomorph	Phonetic Realization	Demonym Example
–an	[ən]	American [ə.ˈmɛ.ɹɪ.ˌkən]
–er	[ɹ]	New Zealander [nu.ˈzi.lən.ˌdɹ]
–ese	[iz]	Japanese [ˌd͡ʒæ.pə.ˈniz]
–ian	[ni.ˌjən/i.ˌjən]	Canadian [kə.ˈneɪ.di.ˌjən]
–ite	[aɪt]	Brooklynite [ˈbɹʊk.lɪ.ˌnaɪt]

Table 1: Selected Demonym Allomorphs of English

When language users are given so many allomorph possibilities, and when more than one of the options give a phonologically acceptable output, how does a language user decide the 'winning' surface form? The goal of this work is to describe both the phonological and the non-phonological (i.e. morphological, semantic, and usage-based) factors at play when a speaker forms a novel demonym, and then to resolve how a speaker selects an output form from a set of phonologically acceptable competitors.

This paper will investigate a restricted set of productive demonym suffix allomorphs, most of which maintain substantial base identity for ease of analysis. Here I will also note that this paper investigates strictly the nominal uses of demonyms, as the claims I will make later may not hold for different 'intents' in the case of homophony (e.g, -er can be both a demonym allomorph as in 'Berlin*er*,' or an agentive suffix as in 'writ*er*'). This means that, in general, demonyms should be a grammatically acceptable answer to the question "What do you call a person from…?," e.g., "What do you call a person from Bulgaria?" can be answered with "A Bulgar*ian*."

2.1   *Phonological and Non-Phonological Conditioning of Allomorphs*   One of the difficulties of understanding demonym allomorphy within existing theory is that there are so many demonym allomorphs, and it is not clear that use of the different allomorphs is conditioned by a defined set of factors. For example, the demonym suffix –*ese* [iz] is used with *Vietnam* [ˌvi.jɛt.ˈnɑm], creating a suffixed form [vi.ˌjɛt.nə.ˈmiz], which shifts stress and vowel identity from the base form to the affixed form. *Suriname* [ˌsʊ.ɹɨ.ˈnɑm], though phonologically quite similar to *Vietnam*, takes the suffix –*er*/[ɹ]. In the affixed form ([ˌsʊ.ɹɨ.ˈna.mɹ]), both stress identity and vowel quality are maintained from the base to the affixed form and the only difference is a resyllabification to accommodate the suffix. This example demonstrates that demonym allomorphy is not clearly conditioned by just phonology, because different allomorph suffixes are used in bases that are phonologically similar. If phonological conditions in the base were solely driving selection of the suffix, we might expect that *Vietnam* would also get the allomorph –*er*, to form Vietnam*er*, which results in well-formed syllables *and* maintains stress and segment identity with the base. Comparing this case to the Korean subjunctive examples in (1), wherein allomorphy is conditioned solely by the identity of the base-final segment, it is clear that demonym allomorphy must not rely on some single conditioning factor. While the phonological well-formedness of a demonym is surely a factor in allomorph selection, many other conditioning factors are likely to influence allomorph choice as well.

Usage effects such as frequency and familiarity are often proposed as conditioners of allomorph alternation (Bybee, 2001). Bybee notes that lexical variation is intertwined with frequency, because language change – in her example, schwa reduction/deletion – affects lexical items differently depending on their frequency. A very common word like *every* [ɛ.vɹi] is therefore very consistently pronounced without a schwa, a mid-frequency word like *memory* is variably pronounced with *or* without a schwa ([mɛ.mɹi] or [mɛ.mə.ɹi]), and a low-frequency word like *mammary* [mæ.mə.ɹi] never undergoes schwa-reduction. As research shows, it is highly likely that frequency – of a base, of an allomorph, and of their combination – is a potential conditioning factor in the selection of a demonymic form. This research also motivated the corpus study, reported below, as a way of better understanding where these frequency effects might be felt in the items under consideration in the present work.

Prior familiarity with a word may also be an influence in allomorph selection. Familiarity with a word and any of that word's morphologically related forms is often referred to as a *lexical entry* or as being *listed* in a speaker's mental lexicon. Work on lexical variation in Tagalog (Zuraw; 2000, 2010) proposes that in an OT framework, familiarity can be represented within a derivation as a factor conditioning lexical variation. Other proposals (Becker, 2009; Pater, 2007) can access individual words in the lexicon through lexical indexation of constraints wherein variable rankings are possible depending on the indexation. Novel words are distinguished from those words already in the lexicon by not being indexed initially. The several different proposals are all ways of noting that individual lexical items can have different specifications, even for the same general process. This implies that a speaker's knowledge of the different specifications is linked with their knowledge of lexical items, and lexical variation arises from being unfamiliar with 'special' processes that apply to certain lexical items. A person unfamiliar with Western Canada, for example, may guess that a Vancouver resident is called a *Vancouverian* (and surely the intended meaning is still conferred), but unbeknownst to the speaker, frequent usage of the term *Vancouverite* has made this latter variant the dominant form. If enough speakers adopt the *Vancouverian* form, this may lead to a lexical variant that could supplant the originally dominant *Vancouverite* form.

The Vancouver example suggests another consideration to be accounted for; that is the case of completely novel or unfamiliar words that may not have an extant representation in the mental lexicon. Becker (2009) presents evidence that speakers of Hebrew, Turkish and English will apply knowledge of the phonological processes in their language to novel lexical items in a way that maximizes observed trends in the base language. However, Becker's examples considered phonological processes that only required a binomial choice. In the context of choice among multiple options, speakers' reliance on lexical trends may still apply, but more specific knowledge of the conditioning factors will be necessary. If the novel words are quite dissimilar compared with existing words in the lexicon, it may be difficult for a speaker to apply existing phonological knowledge and to maximize the observed trends of the base language. The behavioral experiment described below was developed to test such questions in the application of English demonym allomorphs to novel/unfamiliar forms. The experiment expands this prior work on trends in allomorph selection to a situation that requires choice between more than two available allomorphs.

The final non-phonological factor that will be examined here is a potentially *dominant* allomorph that

might emerge as a default or tie-breaker when multiple options are phonologically acceptable. The Priority constraint (Mascaró, 2007) was formulated for this issue, and represents an alternative to lexically indexed constraints (Becker, 2009; Pater & Coetzee, 2005; Pater, 2007). Lexical indexation can make problematic predictions for the wider grammar, as indexed constraints are highly specific to individual morphemes. Priority instead considers a set of allomorphs to be lexically listed and ranked according to markedness. This ranking generates a *dominant* allomorph (or a dominant set of allomorphs) that applies in cases of underdetermination, and subordinate allomorphs are (partially) ordered beneath the dominant allomorph. The Priority constraint is ideal for multiple exponence in allomorphy because it can provide a derivation when phonology is not the deciding factor in allomorph selection. Additionally, Priority may be a way of resolving conflicts when a novel base is far afield from other bases in the lexicon. This is especially relevant for demonyms, because when a speaker encounters a foreign-sounding word (such as a place name), it may not fit easily with phonological expectations formed on words from one's native language. The summary of factors that are relevant to the conditioning of demonym allomorph is thus 1) phonological conditioning (details discussed below); 2) frequency; 3) familiarity; 4) analogy to known forms; and 5) the application of a default allomorph. The following section, which presents a corpus analysis of demonyms and place name bases, will speak to the first 3 of these factors.

3   Corpus Study

This corpus analysis was undertaken to better understand the effects of frequency and familiarity described in the previous section. The goals of the corpus study were: 1) to corroborate prior research about which demonymic suffixes are used with certain place names; 2) to estimate the different levels of familiarity that speakers have with different demonymic forms based on type and token frequency of the attested demonyms; and 3) to identify shared phonological properties of bases that select the same demonym suffix.

The Corpus of Contemporary American English, or COCA (Davies, 2008) the corpus used in this research. It is ideal for present purposes because it is recently (2012) updated, and because it is specific to American English. All results reported below are based on frequency per million words across all the sections of COCA except 'Fiction,' and in all cases the searches were limited to proper nouns (to avoid inclusion of adjectival results, e.g. to select results like "every Estonian" while excluding results like "Estonian universities"). Results are summed over neutral, singular and plural uses, so the results for 'Mexican,' e.g., actually include the frequency for 'Mexican' *and* 'Mexicans.' Alternate spellings are collapsed to the most frequent variant, e.g. the frequency per million for 'Argentinean' is 39 and the frequency per million for 'Argentinian' is 17, so these spelling variants will be combined to 'Argentinean' with a frequency per million of 56. One further decision used to develop this analysis was the classification of place names ending in –*ia* (Alban*ia*, Ind*ia*, Californ*ia*), as taking the –*an* affix rather than the –*ian* affix. This decision arose from the observation that forms which take –*ian* generally involve base stress shift in an affixed form, such as Italy [ˈɪ.tə.li] → Ital*ian* [ɪˈ.tʰæl.jən], Hungary [ˈhʌn.gə.ɹi] → Hungar*ian* [hən.ˈgæ.ɹi.ˌjən] and Canada [ˈkʰæ.nə.də] → Canad*ian* [kʰə.ˈneɪ.dɪ.ˌjən] where base stress identity and often segment identity is clearly *not* maintained, unlike the Alban*ian* [æl.ˈbeɪ.nɪ.ˌjən], Ind*ian* [ˈɪn.di.ˌjən], Californ*ian* [ˌkʰæ.lɪ.ˈfɔɹ.ˌnʲən] cases. A few forms involving lenition at morpheme boundaries also have special classification. Paris → Paris*ian* and Haiti → Haiti*an* are analyzed as taking the –*an* suffix because of the greater surface similarity to the Albania/India/California examples.

3.1   *Corpus Methods*   A list of 50 demonyms was compiled to be included in the behavioral experiment described below. The place names to be included in this study had several requirements, as follows: 1) A variety of place name bases ending with vowels and consonants; having different stress patterns, numbers of syllables, and etymological backgrounds was necessary to obtain a mix of phonological conditions. 2) Place names from both North America (where this study was conducted) and abroad, and names of large regions like countries and continents, and smaller regions like states and cities, were needed to produce varying degrees of frequency and expected familiarity with place name bases. 3) Demonym forms for these 50 place names should be distributed as evenly as possible among the five demonym allomorphs under investigation here, –*an*, –*er*, –*ese*, –*ian*, and –*ite* (see Table 1 for details). Although this observation is not researched further in this study, –*ite* did not occur with place names of

countries or larger regions, so of the 25 country place names investigated, seven took the *–ian* allomorph, six took *–an*, six took *–er*, and six took *–ese*. Locality place names were divided evenly among selection for each of the five allomorphs.

The place names developed for this study, and their resultant demonym forms, were checked in the corpus for frequency (per million words), as well as for alternate spellings or variation in their demonym forms. A subsequent corpus analysis looked at the type and token frequencies of *any* stems occurring with each of the five demonym allomorphs under consideration. Finally, the stems resulting from that analysis were evaluated on a variety of phonological parameters to better understand if any phonological characteristics of the place name bases might be related to a pattern of demonym allomorph selection.

3.2   *Corpus Results*

Place Name (Country)	Official Demonym	Token Freq. (per Mil.)	Place Name (Locality)	Official Demonym	Token Freq. (per Mil.)
Senegal	Senegalese	93	Saigon	Saigonese	1
Mexico	Mexican	2848	Brooklyn	Brooklynite	46
Panama	Panamanian	232	Sparta	Spartan	641
Sudan	Sudanese	4	Vancouver	Vancouverite	12
Italy	Italian	5098	Bern	Bernese	2
Sri Lanka	Sri Lankan	65	Bangkok	Bangkokian	1
Iceland	Icelander	180	Nashville	Nashvillian	9
Virgin Islands	Virgin Islander	2	Berlin	Berliner	415
America	American	83342	Portland	Portlander	20
New Zealand	New Zealander	181	Boston	Bostonian	136
Estonia	Estonian	147	Detroit	Detroiter	38
Luxembourg	Luxembourger	9	Vienna	Viennese	1
Ukraine	Ukrainian	376	Chicago	Chicagoan	195
Jordan	Jordanian	436	Dallas	Dallasite	7
India	Indian	17591	Budapest	Budapester	0
Cuba	Cuban	2109	Houston	Houstonian	386
Morocco	Moroccan	311	Rome	Roman	2626
Canada	Canadian	3514	New York	New Yorker	3351
Japan	Japanese	8017	Milan	Milanese	68
Greenland	Greenlander	48	Fresno	Fresnan	2
Vietnam	Vietnamese	312	Atlanta	Atlantan	562
Surinam	Surinamer	10	Toronto	Torontonian	8
Hungary	Hungarian	532	Delhi	Delhiite	1
Taiwan	Taiwanese	2	Shanghai	Shanghainese	33
China	Chinese	2561	Austin	Austinite	11
Total Tokens		**128020**		**8572**	

Table 2: Place Names and Demonym Frequencies in COCA

Table 2 lists all of the place names developed for this project, divided by whether they are countries or localities. The attested demonym form is also listed for each place name, as well as that demonym's frequency per million in the corpus. A very small number of attested alternate forms were found (for seven of the 50 place names searched, e.g. Brooklyn*er* instead of Brooklyn*ite*), suggesting a fair amount of stability in English demonym usage. Only one demonym form, *Budapester*, was unattested in the corpus. The rest of the sample showed a Zipfian distribution, which indicates that the set of place names and demonyms used in the study described below are a good representative sample of the larger set of demonyms in a typical speaker's mental lexicon.
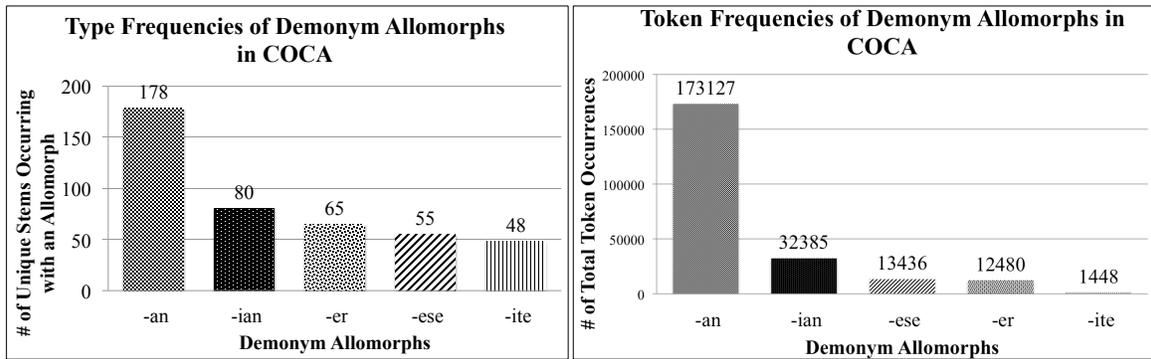
Figure 1: Type and Token Frequency of Demonym Allomorphs in COCA

Type and token frequencies of all of the demoynm suffix allomorphs were also collected. With respect to this set of demonym data, the term 'type frequency' corresponds to the number of unique stems in the corpus with which the demonym occurs. The term 'token frequency' refers to the sum of all occurrences of the demonym suffix in the corpus, summed across all unique stems with which that demonym suffix occurs. Figure 1 represents this data graphically. In this data, the *–an* allomorph has, by a large margin, the greatest number of type (178) and token (173,127) occurrences in the corpus. Following *–an*, *–ian* has the next greatest number of type (80) and token (32,385) occurrences. Type and token frequencies diverge slightly with *–er* and *ese*; *–er* has a higher type frequency count (65) than *–ese* (55), but *–ese* has a higher token frequency (13,436) than *–er* (12,480). Finally, *-ite* had the lowest type and token frequency of 48 and 1,448 respectively. The token frequencies suggest that *–an* would make a good candidate for a *dominant* allomorph by sheer frequency of occurrence in the corpus: its' type frequency is more than double the next closest type frequency, and its' token frequency is over five times more frequent than the next most frequent allomorph, *–ian*. Perhaps because COCA is a corpus of American English, demonym frequency counts for the term *American* were especially numerous. There were 83,342 tokens of *American* found in the corpus, accounting for 48.14% of the token frequency count of *–an*; however, each of the other demonym allomorphs also had one overwhelming frequency contributor, suggesting that *American* is not unique in this respect. It is also possible that the majority of place names in English happen to favor a certain demonym; for example, place names derived from Latin and other Romance languages may favor demonym allomorphs like *–an, –ian,* while place names derived from Germanic languages might favor *–er*.
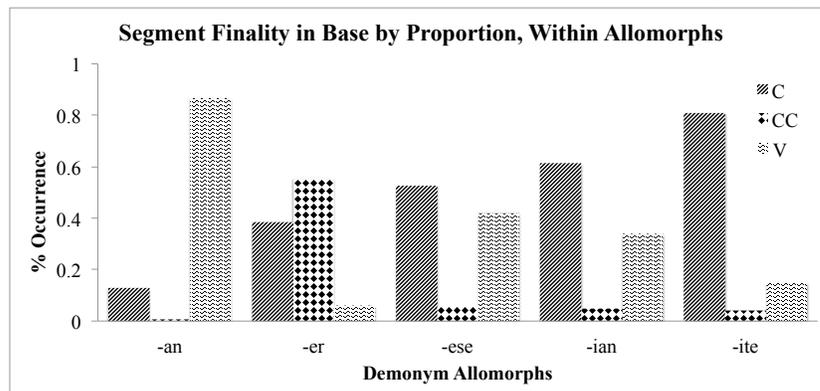
A final intention of the corpus study is to determine if the different place name bases that were found for each of the demonym suffix allomorphs share any important phonological properties. Individual place names were examined to understand what phonological factors of the base (e.g. final syllable shape, base stress, and number of syllables in the base) related to demonym allomorph selection.

One phonological factor that was hypothesized to condition distribution of allomorphs in previous work is the quality of final stress in the base (Smith, 2013). The bases that had a demonym form in the corpus were therefore analyzed and labeled as having final primary stress, final secondary stress, or no final stress in the base. Examination of these results shows that bases with no final stress often selected the *–an* allomorph, and selected the *–ese* allomorph far less frequently. Primary final stress was more evenly distributed across each of the different demonym allomorphs, but it occurred most often with *–er* and least often with *–ite*. Secondary final stress was also distributed throughout the corpus, but was found most often with *–ian* and least often with *–ite*. Within allomorphs, *–an*, *–er*, and *–ite* largely prefer bases that do not have any stress in their final syllable. *–ese* prefers bases with secondary stress in the final syllable, and *–ian* has a near-equal distribution of bases with secondary and no stress in the final syllable.

In addition to the stress pattern of the base syllable being of phonological importance, foot structure may also be implicated in allomorph selection. In his chapter on phonologically-conditioned allomorphy, Nevins (2011) provides several examples of languages that exhibit allomorphy conditioned on maintenance of ideal foot structure, and argues that some languages exhibit allomorphy based on syllable-counting to maintain appropriate foot structure. Analysis of the data shows that foot structure and/or syllable count in the base does not appear to be a relevant conditioning factor for English demonym allomorphs. Perhaps because the place name bases considered in this study have various linguistic origins, and because many of the demonym allomorphs considered here have similar syllable structure, it is unlikely that demonym formation is linked to foot structure or syllable counting. However, one important observation is that the

suffix –*ian* is sometimes analyzed as disyllabic whereas all of the other allomorphs considered here are monosyllabic. This fact was apparently less important than the base stress in determining allomorph distribution, though –*ian* appears to have a general preference for two- and three-syllable bases. No allomorph appeared to have a strong preference for odd- vs. even-numbered syllable counts in the base. It is also possible that this reflects some conflation with the quality of the final segment(s) in the base (consonant, vowel or cluster) or with the base truncation that selected demonyms exhibit. Both of these possibilities are investigated subsequently.

The quality of the final base segment (either C, CC, or V) is another key phonological conditioner in many cross-linguistic examples of allomorphy. In English demonym allomorphy, final segment quality proves to be one of the more obvious sources of phonological conditioning in allomorph selection. Across allomorphs, C-final bases could be said to preferentially attach to –*ian* and –*ite*. CC-final (cluster-final) bases have a clear preference for –*er*, while V-final bases display a preference for –*an*. Within each allomorph (Figure 2), this tendency of V-final bases to take –*an* is even more evident. –*an* is also very infrequently attached to cluster-final bases. –*er*, conversely, shows a strong dispreference for V-final bases but will take either C-final or CC-final bases. Both –*ese* and –*ian* do not prefer cluster-final bases, but will take either C- or V-final bases. Recalling the

Figure 2: Final Segment Type as Phonological Conditioning Factor

previous analysis of syllable count distribution, these findings for –*ian* are interesting because the final segment (and optional resyllabification) determines whether –*ian* will contribute one or two additional syllables to the base. For example, the base 'Kingston' which is disyllabic, ends in C, and takes –*ian*. Here, –*ian* adds two additional syllables with resyllabification of the coda consonant to the onset of the demonym suffix. But in a trisyllabic base which ends in a V, such as 'Florida,' it appears that –*ian* causes truncation of the final vowel and resyllabification of the new 'last' C to the onset of the following syllable, for 'Floridian.' In this case the base had three syllables and the suffixed form has four syllables, meaning –*ian* contributed only one additional syllable. Other bases that end in a vowel and take –*ian* are similar, which helps to explain why, in the previous syllable count analysis, -*ian* did not have a strong preference for even or odd syllable counts. Lastly, -*ite* prefers to attach to C-final bases, tolerating V-final bases occasionally and only rarely admitting cluster-final bases.

Finally, the corpus analysis queried whether any of the demonym suffixes was likely to cause/select for base truncation in stems. Plag (1998) argues that truncation should be considered a *phonological* rather than a *morphological* effect, because in suffixation, bases with similar foot structure are truncated predictably. Truncated bases were rarely found with –*er* and –*ite*, which may be expected due to the earlier observations that these two demonyms were far less likely to take V-final bases (V-final bases seem to be stronger candidates for truncation here than C- or CC-final bases). *Within* allomorphs, each example has a strong tendency toward selecting untruncated bases, which implies a fairly powerful preference for base-affixed identity (Benua, 1997), but –*ese* and –*ian* seemed to prefer truncated bases more than any of the other affixes (perhaps due to syllable well-formedness constraints, though it is unclear why –*ese* and –*ian* would pattern differently from the other vowel-initial suffixes). Again the –*er* and –*ite* bases exhibit a dispreference for truncated bases, but at this level of analysis –*an* also disfavors truncation. This is surprising in light of the findings for segment finality, where –*an* attaches to V-final bases extremely frequently. In the corpus results for the –*an* group, many of the base place names end in –*ia*, which suggests that preservation of that particular sequence is especially important, perhaps as a cue to the relationship of the suffixed form with the place name it is derived from.

–an	*–er*	*–ese*	*–ian*	*–ite*
–an prefers 3-syllable bases that end in a V and do not end in final stress. Many place name bases ending in *–ia* select *–an*.	*–er* most frequently found with 2-syllable bases that end in a consonant cluster. *–er* has a strong preference for untruncated bases.	*–ese* prefers bases that are two syllables long and have final secondary stress. This allomorph takes final stress and causes stress shift when affixed.	*–ian* is often found with 2- or 3-syllable bases that do *not* have final primary stress, and strongly prefers consonant-final bases (but not clusters).	*-ite* is the least frequent allomorph. It is often found with 2-syllable bases that do not end in final stress. *–ite* strongly prefers C-final, untruncated bases.

Table 3: Summary of Possible Phonological Conditioning Factors by Allomorph

3.3 *Corpus Discussion*   The previous paragraphs discussed findings for individual phonological conditions, some of which are more relevant to demonym allomorph selection (segment finality, stress finality) than others (syllable count, truncation). A summary of findings about phonological conditioning by allomorph is presented in Table 3. No two allomorphs discussed here have the exact same preferences on all of the phonological conditioning categories that were considered, which makes a strong case for the argument that these affixes are partially conditioned by phonological factors. Conversely, there are no completely exceptionless phonological distributions. This suggests both that phonological factors are not acting alone in the conditioning of demonym allomorphy, and that lexicalization of various demonym forms may be the cause of some exceptions to the distributions noted in the table.

Previous work has suggested that phonological processes can also be 'lexicalized' over time: "At an early stage, phonological similarity among allomorphs, which is often found in external allomorphy, is a reflex of the origin of the phenomenon; at later stages, phonological similarity tends to diminish gradually" (Mascaró, 2007). A demonym form for a new or fictional place could initially exhibit free variation but stabilize over time, *and/or* the set of bases to which a certain demonym allomorph can attach could gradually expand or contract over time, changing the basic set of phonological processes used in forming new demonyms. Precisely because demonym formations are highly productive with novel and fictional place names, it is necessary to ask if the findings from the corpus study are salient enough to speakers that these conditions are generalizable to novel or unfamiliar forms just entering the language. This question will be taken up in the next section, which describes a behavioral experiment on demonym allomorph selection.

4   Survey on Real & Fictional Demonym Suffix Preferences

This experiment investigates English speakers' demonym allomorph selection preferences in a variety of base contexts differing in expected *familiarity* (as described above). The goal of the experiment was to determine which aspects of the unaffixed base, and which aspects of the potential base-affixed form relationship, are important to a speaker when they are selecting from a set of possible demonym suffixes. This is relevant in cases where the participant is *unfamiliar* with the demonym form, because they are more likely to rely on phonological factors (such as consonant vs. vowel finality, assumed word stress, syllable count/identity of the base, and assumed base truncation properties) to decide among different possible demonym allomorphs. The following section will describe the experiment and resulting findings specific to the *unfamiliar* forms participants encountered in the experiment. Analyzing these unfamiliar forms will allow an interpretation of the results that shows speakers' preferences for conditioned allomorph selection independent of lexicalized or very familiar forms.

4.1 *Experiment Design*   100 total items were included in the experiment, 25 in each of four familiarity conditions: Real Country ~ RC, Real Locality ~ RL, Fictional Literary ~ FL and Fictional Screen/non-literary ~ FS. The real place names are the 50 items discussed in the corpus study above. The fictional place names are target items, meant to determine the phonological distribution of the suffix allomorphs given properties of the base such as final segment, base syllable count, and assumed stress pattern in the base. Most fictional place names were unattested in the corpus, so they did not have expected suffix choices. The fictional bases were selected from existing books, movies, TV shows, and video games to have a variety of

stress patterns, final segments, and syllable counts comparable to the variety found in the real place names. Some examples of fictional place names used in the experiment are *Zamunda*, *Wormulon*, *Naboo*, and *Duckburg* in the FS condition, and *Terabithia*, *Hogsmeade*, *Asgard*, and *Narnia* in the FL condition.
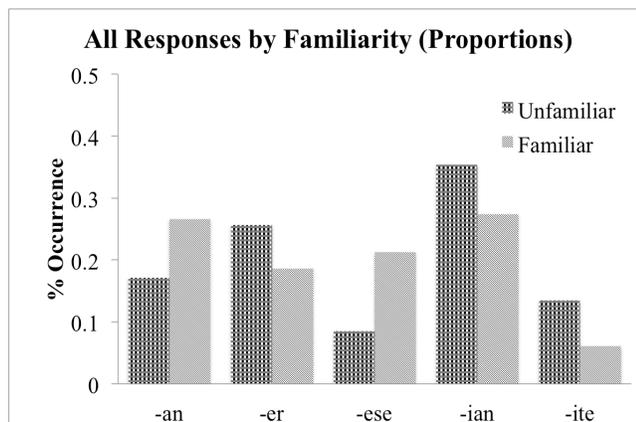
4.2 *Methods*   The experiment was conducted using Qualtrics online survey software (Qualtrics, 2014). Data for 44 native English speakers is reported here. The mean age of participants is 30.721, ranging from 19 to 63. Participants were shown experimental items presented orthographically in a multiple-choice format. Each item screen contained three questions. Responses to 'A resident of X is called:' are the main dependent variable in the following analyses. Responses to 'have you heard the term you selected used before' are called 'familiarity responses' and were used to group items as familiar/unfamiliar. Responses to 'please rate your confidence in the term you selected from one to five stars' are 'confidence responses.'

In the experiment, participants were randomly assigned to one of three lists, each with 16 items from each of the four conditions. Due to the orthographic presentation of items, decisions about the composition of affix forms were made to constrain the multiple choice questions to just five options with as little deviation from base identity as possible. At the end of the survey, participants were asked additional demographic questions, such as their native English speaker status, their home country (or the country they most identify with), the demonym they use as a resident of their home country, their home city/locality (or the city they most identify with), and the demonym they use as a resident of their home city/locality.

4.3 *Predictions*   For a demonym that a participant has seen or heard before, the participant is expected to make a decision based on stored lexical knowledge of the place name and the expected demonym. These *familiar* items are expected to be from real conditions (RC and RL; see section 4.1). RC should be more familiar than RL, based on corpus results showing nearly 15 times more country demonyms (128,020 tokens) than locality demonyms (8,572 tokens). Participants were expected to be less familiar with fictional bases *and* were expected to not have encountered the demonym form before. Among fictional items, FS items are expected to have lower familiarity than FL items, because fictional place names from literary sources may be more familiar from prior orthographic presentation than non-literary fictional place names. Overall familiarity is predicted to be ordered, from most familiar to least familiar, RC >> RL >> FL >> FS.

In unfamiliar items, participants are expected use knowledge of properties of the base to select appropriate demonym allomorphs. Segment finality is the least ambiguous property of the bases, so it is expected to show a strong association with certain allomorphs. Stress placement in the base is also likely to condition allomorph selection, based on corpus results. For unfamiliar words, the participant may not *know* the intended stress of the base. This last point is key; it motivates the necessity of gathering stress judgments for the unfamiliar items in the experiment (a follow-up experiment is planned). Analysis of stress finality, based on expected stress judgments, is still included for illustrative purposes in the analysis.

4.4 *Experiment Results*   The results presented in this section are based on observed/expected results taken from chi-square goodness of fit tests, and descriptive statistics from the raw data. Both results are included because these data violate assumptions of standard statistical tests: the chi square test assumes independence of observations, but this data has repeated measures of subjects and items; however, if t-tests were performed on proportion and count statistics, these would be inaccurate because these data violate assumptions of continuity and normal distribution. Though neither statistic is fully satisfactory, their comparison

Figure 3: All Experiment Responses by Familiarity

gives an idea of whether the observations are generally going in the same direction.

Descriptive statistics show a pattern of differences between experimental conditions, wherein average familiarity for the fictional conditions is considerably lower than for the real conditions. As expected, the

breakdown of familiarity between conditions is RC >> RL >> FL >> FS. Figure 3 displays the differences between those items noted as familiar vs. unfamiliar by participants.

There are generally different distributions of allomorph selection for unfamiliar and familiar items. This reflects the way the conditions were designed: bases in the two real conditions were chosen for their even distribution among the allomorph choices. The unfamiliar items, mostly comprised of items from the fictional conditions, are not hypothesized to have expected demonym allomorphs because the participant may not have seen them before. This could explain why selection of different allomorphs was not evenly distributed for unfamiliar items, but it does not explain the preference for *–ian* in the unfamiliar conditions.

Stress-finality in the base was also clearly shown to be different in familiar vs. unfamiliar conditions. Primary stress-final bases in familiar items patterned very similarly to the corpus results. Primary-stress final bases in unfamiliar items were more evenly distributed across all choices, but showed a preference for *–ian*, contrary to corpus results and the familiar items. In unfamiliar responses, secondary final stress was observed for bases that selected *–ian* and *–er*, and to a lesser degree with *–an*. Final *unstressed* bases in both familiar and unfamiliar items showed a similar distribution.

Results in the experiment, as in the corpus results, did not indicate that syllable count of the base was likely to condition selection of demonym allomorphs. However, experiment results differed from the corpus results, even when comparing *only* familiar items to the corpus. The distributions of syllable base counts are different in unfamiliar vs. familiar responses, and both of these are different from the corpus results. This may be due to the greater number of bases in the corpus.
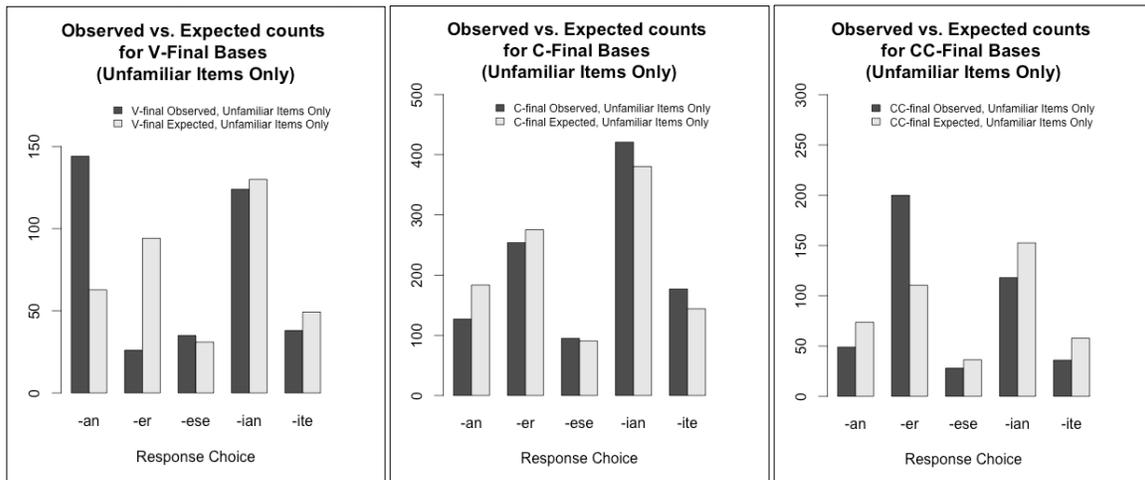
Figure 4: Observed vs. Expected Segment Finality in Unfamiliar Forms

In accord with the corpus results, base segment finality is shown to select for different demonym allomorphs in unfamiliar experiment items. The graphs in Figure 4 show that V-final bases clearly favored the use of *–an*, and to a lesser extent, *–ese* beyond expectation, while C-final bases preferred *–ian* and *–ite* at a higher than expected rate. Cluster-final bases clearly preferred *–er*. In general, these results provide evidence for base segment finality as being an important conditioning factor for demonym allomorphy.

–an	*–er*	*–ese*	*–ian*	*–ite*
–an was the third most frequent allomorph for unfamiliar items, but the most common allomorph in the corpus. It has more limited distribution than *–ian*, suggesting lexicalized place names optimized for *–an* over time.	*–er* was second most frequent after *–ian*, but was observed less frequently in the corpus results. It is associated with CC-final place names.	*–ese* was used least frequently with unfamiliar items. This may be due to its distribution among real forms (it is often used with African and Asian place names).	*–ian* was used most frequently with unfamiliar items. It will work with a variety of forms distributed across the different phonological conditioning factors identified here.	*–ite* was second least frequent in unfamiliar items, but was least frequent in the corpus. The distribution of its phonological properties here is less restricted than those for *–ese*.

Table 4: Summary of Experimental Findings for Unfamiliar Items

4.5   *Experiment Discussion*   A summary of findings from the unfamiliar items in the experiment is presented above in table form, comparable to the table summarizing findings for each demonym allomorph in the corpus results. The conditioning factors have a considerably different distribution than many of those same features in the corpus study. This could imply that 1) unfamiliar bases rely more on phonological conditioning factors, because other factors like familiarity and frequency are unavailable; and/or 2) that familiar bases are optimized for use with given allomorphs by the time they are lexicalized, and thus have a different distribution than novel place names that are not yet lexicalized; and/or 3) the difference observed in the corpus results and the unfamiliar results is a task-based effect. However, results show that certain sets of phonological conditions in the base favored the use of less frequent allomorphs, and even with unfamiliar items. This may counter the suggestion in (3) above that these findings are task-based, but further experimentation is needed to show this conclusively. If speakers do rely on different sets of information when using familiar vs. unfamiliar demonym forms, the distinctions between the distribution of allomorphs in the corpus findings and the unfamiliar items in the experiment are an encouraging result.

5   General Discussion

Corpus-based and experimental results supported the observation that demonym allomorphy is not conditioned completely by phonology, a result already acknowledged by multiple theories of lexical variation (Wolf, 2008; Zuraw, 2000). This work adds to prior theories of lexical variation with a case of multiple exponence in allomorphy, and describes how conditioning factors affect lexical items differently.

In light of findings in the present work, consider the following case. There is no clearly phonological reason why the demonym form for Houston is *Houstonian*, but the demonym form for Austin is *Austinite*. The bases are clearly phonologically similar, yet each takes a different demonym allomorph. Under the analysis presented here, we might first examine frequency information. *Houstonian* appears 386 times in COCA, while *Austinite* appears only 11 times. Referring back to the previous section, we can see that the non-phonological factor *frequency* might be expected to play a role here. Consider the possibility that at an earlier point in the history of the language, *Houston* selected the demonym allomorph *-ite*. This would be consistent with the distribution of phonological factors in the unfamiliar forms from the experiment, wherein –*ite* is strongly predicted to be found with 2-syllable, non-stress-final, C-final bases, of which *Houston* is an example. However, perhaps due to lexicalization of different demonym forms, the –*ite* allomorph eventually became less frequent and –*ian* came to be the dominant allomorph (applied at the expense of stress and vowel identity from the base to the suffixed form: ˈhju.stɨn + i.ˌjən → hju.ˈstoʊ.ni.ˌjən). The same erroneous lexicalization could eventually happen to *Austin*, if at a later point its frequency increases past some threshold. This story, though speculative, makes contact with all of the findings of this paper: phonological factors are indispensible in determining allomorph selection, but they do not act alone; additional factors such as frequency, familiarity, allomorph dominance, and analogy all play important roles that help explain the sometimes unexpected distribution of allomorphs in corpora.

6   Implications & Conclusion

The case of demonym allomorphy presented here supports a new perspective on the topic of conditional allomorphy; that is, selection of an allomorph is necessarily conditioned by both phonological and non-phonological factors. The sources of conditioning must also be able to interact with each other to reproduce observations derived from behavioral data for unfamiliar items. Furthermore, this study of demonym allomorphy expands theories of phonologically-conditioned lexical variation, which have been considered primarily with reference to binomial choices and solely phonological conditioning (Becker, 2009; Nevins, 2011; Raffelsiefen, 1999; Zuraw, 2000; etc.), to multiple exponence contexts with many possible sources of conditioning, including some phonological and some extra-phonological factors.

Through the comparison of familiar and unfamiliar demonym forms, two distinct sets of distributions were discovered, implying a notable distinction between familiar (lexicalized) forms, which are relatively stable members of the lexicon, and novel forms, which do not yet have a stable lexical representation and which rely on conditioning from multiple sources to achieve allomorph selection. An aim of this work is to model these conditioning factors in an optimality theoretic framework (Prince & Smolensky, 1993/2003).

Cases of multiple exponence in allomorphy have not previously been analyzed as predictable

distributions systematically conditioned by both phonological and non-phonological factors, so it has been the goal of this work to lay out some possibilities in the development of such a theory as a principled and expected process in the evolution of a grammar. The use of English demonym allomorphy as a case of multiple exponence is certainly extreme (with more than five allomorphs for a single intent), but it is likely that this degree of extremity drives diachronic language change. The consideration of a variety of cross-linguistic cases will be necessary to further this line of theoretical work, but the process can already be started by reconsidering cases of allomorphy that are not well understand, and which had previously been labeled with vague terms such as "morphological allomorphy" (Booij, 1997) or "non-optimizing allomorphy" (Paster, 2009). English demonyms and other cases of multinomial allomorphy have previously been minimized as diachronic artifacts that will be lost in the history of language; yet they also represent synchronic processes that provide opportunities for the study of what it means to 'know' the sounds and processes of one's language and to apply them to an ever-changing lexicon.

References

Becker, Michael. (2009). Phonological trends in the lexicon: The role of constraints (Doctoral Dissertation). University of Massachusetts, Amherst.

Benua, Laura. (1997). *Transderivational Identity: Phonological Relations between Words* (Doctoral Dissertation). University of Massachusetts, Amherst.

Booij, Geert. (1997). Allomorphy and the Autonomy of Morphology. *Folia Linguistica*, *31*(1-2), 25–56.

Bybee, Joan. (2001). *Phonology and Language Use*. New York: Cambridge University Press.

Davies, Mark. (2008). The Corpus of Contemporary American English: 450 million words, 1990-present. Retrieved from http://corpus.byu.edu/coca/.

Lapointe, Steven G. (1999). Stem Selection and OT. *Yearbook of Morphology 1999*, 263-297. Springer Netherlands.

Mascaró, Joan. (2007). External allomorphy and lexical representation. *Linguistic Inquiry*, *38*(4), 715–736.

Nevins, Andrew. (2011). Phonologically-Conditioned Allomorph Selection. In C. Ewen, E. Hume, M. Van Oostendorp, & K. Rice (Eds.), *The Companion to Phonology* (pp. 2357–2382).

Paster, Mary. (2009). Explaining phonological conditions on affixation: Evidence from suppletive allomorphy and affix ordering. *Word Structure 2*(1), 18-37.

Pater, Joe. (2007). *The locus of exceptionality: Morpheme-specific phonology as constraint indexation.* (No. 32). *University of Massachusetts Occasional Papers in Linguistics* (Vol. 32). Amherst, MA.

Pater, Joe, & Coetzee, Andries. (2005). Lexically Specific Constraints: Gradience, Learnability, and Perception. In *Proceedings of the 3rd Seoul international Conference on Phonology)*. 85–119.

Plag, Ingo. (1998). Morphological haplology in a constraint-based morpho-phonology. *Phonology and morphology of the Germanic languages,* 199-215.

Prince, Alan., & Smolensky, Paul. (1993). *Optimality Theory: Constraint interaction in generative grammar*. John Wiley and Sons.

Qualtrics. (2014). Qualtrics Research Suite. Provo, UT: Qualtrics.

Raffelsiefen, Renate. (1999). Phonological constraints on English word formation. *Yearbook of Morphology 1998*, p. 225-287. Springer Netherlands.

Smith, Brian W. (2013). *-(a)licious: a new morpheme obeys existing constraints (Ms)*. University of Massachusetts, Amherst.

U.S. Government Printing Office. (2008). *Style Manual: An official guide to the form and style of Federal Government printing*. Washington, DC: U.S. Government Printing OFfice.

Underhill, Robert. (1976). *Turkish grammar*. Cambridge, MA: MIT Press.

Wolf, Matthew. (2008). *Optimal Interleaving: Serial Phonology-Morphology Interaction in a constraint-based model* (Doctoral Dissertation). University of Massachusetts, Amherst.

Wolf, Matthew. (2013). Candidate chains, unfaithful spell-out, and outwards-looking phonologically-conditioned allomorphy. *Morphology*, 23(2), 145–178.

Zuraw, Kie. (2000). *Patterned Exceptions in Phonology* (Doctoral Dissertation). University of California, Los Angeles.

Zuraw, Kie. (2010). A model of lexical variation and the grammar with application to Tagalog nasal substitution. *Natural Language & Linguistic Theory*, *28*(2), 417–472.