

Testing claims of a usage-based phonology with Liverpool English *t-to-r*¹

LYNN CLARK
Lancaster University

and

KEVIN WATSON
University of Canterbury

(Received 15 November 2010; revised 30 March 2011)

The variable phenomenon in which /t/ can be realized as a tap or rhotic approximant in varieties of Northern British English (commonly referred to as *t-to-r*, Wells 1982: 370) has received some attention in English linguistics as debates have appeared over how best to model its phonology (e.g. Carr 1991; Docherty *et al.* 1997; Broadbent 2008). The occurrence of *t-to-r* seems to be constrained by the preceding and following phonological environment in a largely systematic way and so it is often accounted for within a rule-based model of grammar. Problematically, however, the rule does not apply blindly across the board to all words which fit the specified phonological pattern. Instead, *t-to-r* shows evidence of being lexically restricted, and this fact has recently encouraged a usage-based interpretation. Until now, there has been relatively little attempt to test the usage-based thesis directly with fully quantified data gleaned from naturally occurring conversation. This article investigates the extent to which certain usage-based predictions can account for variation attested in *t-to-r* in Liverpool English. Using oral history interviews with Liverpool English speakers born in the early 1900s, we examine the usage-based predictions first proposed by Broadbent (2008) that *t-to-r* is more likely in (a) high-frequency words and (b) high-frequency phrases. There is some support for the importance of lexical frequency as a motivating factor in the use of *t-to-r*, but our data do not fully support either of these claims wholesale. We suggest that *t-to-r* is not constrained simply by word frequency or phrase frequency alone, but by a combination of both. Finally, we explore the possibility of employing notions from Cognitive Grammar such as schema strength (e.g. Taylor 2002; Bybee 1995: 430) in our interpretation of these data.

1 Introduction

This article offers a contribution to the debate surrounding the nature of Northern English *t-to-r*, the phenomenon in which /t/ can be realised as a tap or rhotic approximant (e.g. ‘shut up’ [ʃʊɪ/r ʊp], ‘get off’ [gɛɪ/r ɒf]). *t-to-r*, which Wells (1982: 370) describes as ‘a widespread but stigmatised connected speech process

¹ We gratefully acknowledge the financial support of the Economic and Social Research Council, grant number RES-061-25-0458. We are also grateful to Paul Kerswill, Anastassia Loukina, Graeme Trousdale and Eivind Torgersen for their comments on an earlier draft of this article and we would like to thank the audience at the workshop for La Phonologie de l’Anglais Contemporain (PAC): usages, variétés et structure, where parts of this research were presented.

in the middle and far north (of England)', has received some attention in English linguistics as debates have appeared over how best to model the phonology of the phenomenon (cf. Carr 1991; Docherty *et al.* 1997; Broadbent 2008). The problem for phonological theory is this: on the one hand, *t-to-r* seems to be constrained by the preceding and following phonological environment in a largely systematic way, so it is often accounted for within a rule-based model of grammar. On the other hand, the rule does not apply blindly across the board to all words which fit the phonological pattern. Instead, *t-to-r* shows evidence of being lexically restricted. In order to account for this, recent analyses have invoked principles from usage-based models of grammar in an attempt to better understand the phenomenon (Asprey 2008; Broadbent 2008). However, there has been relatively little attempt to test the usage-based thesis directly with *t-to-r* data. This article investigates the extent to which the predictions of a usage-based model can account for the variation attested in *t-to-r*.

The vast majority of work on *t-to-r* to date has made reference to two varieties of English in Northern England: Tyneside (e.g. Carr 1991; Docherty *et al.* 1997) and West Yorkshire (e.g. Broadbent 2008; Richards 2008). In each of these varieties, *t-to-r* is complicated by the fact that /t/ can be realised as both a rhotic and a glottal stop, and so these variants are often in competition.² This article reports the results of the first systematic investigation of *t-to-r* in northwest England, and takes as its starting point a variety with very little occurrence of *t-glottalling*: Liverpool English (Watson 2007).³ By minimising the additional complication of *t-glottalling* and focusing exclusively on variation between the stop and rhotic variants of /t/, our objective is to reach a clearer understanding of the ways in which usage-based predictions could be operating on this variable phenomenon.

We begin in section 2 with a discussion of some of the attempts that have been made to account for *t-to-r* in different phonological theories. We show that while many of the claims surrounding *t-to-r* in generative phonology have now been empirically tested, many of the claims put forward from a usage-based perspective have not. In section 3, using a data set of almost 700 instances of *t-to-r* in Liverpool English, we examine the usage-based claims that *t-to-r* is more likely in (a) high-frequency words and (b) high-frequency phrases. As we will see, there is some support for the importance of lexical frequency as a motivating factor in the use of *t-to-r*, but our data do not fully support either of these claims wholesale. We suggest that *t-to-r* is not constrained by simple token lexical frequency or phrase frequency alone but by a combination of both. Finally, we explore the possibility of employing notions from Cognitive Grammar such as schema strength (e.g. Taylor 2002; Bybee 1995: 430) in our interpretation of these data.

² Broadbent (2008: 156) suggests that the rise of *t-glottalling* in West Yorkshire led to the restricted pattern of *t-to-r* in present-day use.

³ In the data analysed in this article, word-final intervocalic /t/ is realised as a glottal stop with a marginal frequency of just 4 per cent.

2 Modelling *t-to-r* in phonology⁴

2.1 Rule-based approaches to *t-to-r*

Wells (1982) offers the seminal description of the *t-to-r* process, noting that it can be found in many varieties of English in northern England (e.g. Greater Manchester, West Yorkshire, Lancashire, Merseyside, Tyneside). Using the terminology of rewrite rules, Wells (1982: 370) suggests that in these varieties, /t/ can be realised as a rhotic in the following phonological environment: $t \rightarrow r / [\text{short V}] _ \# \text{V}$. That is, when word-final /t/ is preceded by a short vowel and followed by a vowel-initial word, it can be realised as a rhotic (variably an approximant or a tap). However, since Wells notes that *t-to-r* also operates 'very occasionally' word internally, there is an indication even at this early stage in the story of *t-to-r* that the rule does not cover the full extent of variation. This is not a problem for rewrite rules per se (because the rule can be simply rewritten to include word-medial environments) but the rewrite rule format cannot capture the probability of variance for each of these environments; nor can it explain the variation.

Later rule-based accounts (e.g. Carr 1991) also admit exceptions but offer more detailed discussions of the derivation of the *t-to-r* rule, particularly in relation to the competing rule in which /t/ is realised as a glottal stop (or a glottally reinforced stop). Carr (1991) adopts a Lexical Phonology model of grammar and, using examples from Tyneside English, suggests that *t-to-r* affects phonological feet formed postlexically in derived environments, in contrast to *t-glottalling* and glottalisation which apply postlexically across the board. In order to make this claim, it is necessary for Carr to introduce a division in the postlexical component of the grammar. Carr's main theoretical problem is in trying to account for differences in apparent minimal pairs such as 'fitter' vs 'fit her' (with h-dropping) which share the same phonetic and metrical conditions but which occur with apparently different outputs (the claim is that 'fitter' occurs with glottalling in Tyneside, whereas 'fit her' occurs with *t-to-r*). Carr suggests that what differentiates these two examples is the level of grammar at which foot structure is formed: in 'fitter', foot structure is formed in the lexicon but with 'fit her', feet are formed postlexically under cliticisation. For Carr, *t-to-r* applies 'in a kind of postlexical, metrically derived environment', but does not apply to 'intervocalic /t/ with the appropriate metrical structure if that structure is formed in the lexicon' (Carr 1991: 48). Carr then suggests that $t \rightarrow r$ and $t \rightarrow ?$ are ordered rules. $t \rightarrow r$ is a more 'specific' rule (because it is more restricted) and so the application of the *t-to-r* rule takes place before the application of the 'across-the board' rule of *t-glottalling*, which has a much wider distribution. Moreover, it is argued that the application of the *t-to-r* rule effectively blocks the application of the *t-glottalling* rule and so the rhotic and glottal variants of /t/ should be in complementary distribution. Carr's (1991) explanation is an interesting and welcome attempt to account for features of non-standard phonology

⁴ This section is not intended to be an extensive review of all available *t-to-r* literature (see Broadbent 2008 for a more thorough overview, including sociolinguistic and historical accounts of *t-to-r*).

within a formal model of grammar. However, one of the drawbacks of the methods typically employed in (generative) theoretical linguistics, such as linguistic intuition and anecdotal data collection, is that they lack the means to empirically validate their claims. In the case of Tyneside English, Docherty *et al.* (1997) collected a sizeable corpus of conversation and word-list data from 32 native Tynesiders, stratified by age and socioeconomic class, and so were able to put to the test the issues raised in the theoretical phonology literature.⁵ Indeed, the proposals regarding the rule ordering of *t-to-r* and *t-glottalling* in Tyneside English are among several phonological claims which are tested – and ultimately refuted – by Docherty *et al.* (1997) using their corpus of production data. There is overwhelming evidence from natural conversation that speakers can and do realise word-final intervocalic /t/ as both a rhotic and a glottal stop; in other words, *t-to-r* and *t-glottalling* are not in complementary distribution in Tyneside English. The arguments set out in Carr (1991) with respect to the ordering of *t-to-r* and *t-glottalling* are therefore difficult to maintain.

Docherty *et al.* (1997) argue that many of the predictions of generative phonology (especially Carr 1991 and Harris & Kaye 1990) are problematic because they do not accurately predict variation. However, little is offered in the way of an alternative approach. They suggest that phonologists should work towards creating models of phonology that treat variation as central rather than peripheral and encourage future work to embed discussions of phonological theory in evidence of real language use. In what is now more than a decade since Docherty *et al.* (1997) first appeared, models of language which adopt these tenets as a central philosophy have been developing independently in various subdisciplines of linguistics (see e.g. network models of grammar (e.g. Bybee 2001, 2007; Hudson 2007), Exemplar Theory (e.g. Pierrehumbert 2001, 2002), various probabilistic models (e.g. Manning 2003; Bresnan & Nikitina 2003; Yang 2004) and also Cognitive Linguistic models of grammar (e.g. Langacker 1983, 1987, 1991)). These frameworks are now increasingly referred to as ‘usage-based’ models of language structure and have been offered as competing approaches to account for the *t-to-r* phenomenon (Broadbent 2008). In the following section, we briefly describe the main tenets of the usage-based approach before discussing its application to *t-to-r*.

2.2 Usage-based approaches to *t-to-r*

Frameworks which are described as usage-based (hereafter UB) typically share three defining characteristics: they are ‘maximalist’, in contrast to the ‘minimalist’ nature of generative grammars;⁶ they are ‘non-reductive’, because they are not guided by

⁵ See Docherty *et al.* (1997) on the methodological issues that surround using corpora of conversation to explore phonological knowledge.

⁶ Generative models of grammar are considered ‘minimalist’ because they traditionally try to minimise both the role of learning in language acquisition and the number of (language-specific) rules posited to account for language structure in the belief that the best grammar is the most economical one. Usage-based models are ‘maximalist’ because they assume a great deal of learning on the part of the speaker during language acquisition.

arguments of economy of storage, and they are ‘bottom up’, in that they only posit abstract structures in the grammar where there is good evidence for the existence of such structures from language use (Langacker 2000). This latter feature is perhaps the key characteristic of a usage-based model of language. Kemmer & Barlow (2000: ix) describe the relationship between language structure and language use as a ‘feedback loop’ since experience of language both results from and also continues to shape the speaker’s linguistic system. This is only possible, it is claimed, because humans possess the ability to register frequency effects in language (Kemmer & Israel 1994: 165). Frequency research has played a significant role in descriptions of usage-based models of language structure. The relationship between lexical frequency and ‘entrenchment’ (or lexical strength, cf. Bybee 2001) is particularly pertinent for the present discussion, as it has been directly invoked to account for patterns of *t-to-r* in northern English, most notably West Yorkshire English (hereafter WY, see Broadbent 2008).

It is clear from Broadbent’s (2008) initial description of *t-to-r* in WY that Wells’ rewrite rule ($t \rightarrow r / [\text{short } V] _ \# V$) provides only a partial picture of the phenomenon. While the rule largely predicts the phonological environment in which *t-to-r* can occur,⁷ it assumes that as long as the phonological environment is met, *t-to-r* is possible. However, as we will see below, this misses two key observations: (a) some words never allow *t-to-r*, even when the phonological criteria are met, and (b) *t-to-r* is more likely in some words than others, even when the words share the same phonological environment. Elaborating on this observation, Broadbent (2008: 45) proposes that there are four sets of *t-to-r* words in WY: (i) regular exhibitors of *t-to-r*, e.g. *but, put, get, go*; (ii) not so common exhibitors of *t-to-r*, e.g. *not, forget/forgot, shut(up)*; (iii) rare exhibitors of *t-to-r*, e.g. *lot, bet*; (iv) *t-to-r* is not possible, e.g. *it, sit, fit, cut, set, wet*.

In order to account for this, Broadbent (2008) draws on suggestions from UB models of phonology, in particular Bybee’s work on lexical strength. Bybee (2001, 2002, 2007) proposes a model in which frequently used instances (words, phrases, constructions) are stored in memory and reinforced (strengthened or ‘entrenched’) with repeated use. Speakers abstract similarity (‘schemas’) across instances of language stored in memory. Broadbent (2008: 161) proposes that the words in which *t-to-r* occurs in WY share certain similarities and so speakers abstract schemas relating to the phonological shape of these words. In particular, Broadbent claims that lexical items which allow *t-to-r* fall under one of three phonological shapes or schemas: [ʊɪ], [ɛɪ] and [ɒɪ].⁸ She suggests that one lexical item in each schema is particularly strong or entrenched ‘where strength is measured in terms of frequency’ (2008: 161) and that ‘the presence

Usage-based models also attempt, wherever possible, to reduce their reliance on language-specific cognitive structures and instead derive language structures from more general cognitive abilities.

⁷ As in Tyneside (Carr 1991; Docherty *et al.* 1997), there are occasional occurrences of *t-to-r* word medially in WY.

⁸ Broadbent (2008: 161) is not clear on how to handle the lexical item *that* which seems ‘only to be related to the other words by virtue of the fact that it can appear with *t-to-r*’.

of frequent exemplars may well keep this fossilised process productive for just such words' (2008: 161).

This is a welcome approach which begins to bring *t-to-r* under the microscope of the UB model, but a number of questions remain. For example, while an account based on lexical strength or frequency effects may explain why certain lexical items continue to occur with *t-to-r* in present-day WY, it remains unclear why certain words which have the required phonological shape do not allow *t-to-r*. Why, for instance, is 'get' a possible candidate for *t-to-r* but 'set' is not? Another question relates to schemas and language change. Schemas are proposed in UB models of phonology not simply as a means of describing the categorisation process; newness in a system may be created by partial sanction, where an innovation shares only part of the specifications of its sanctioning schema. In this way, the schema concept can be used to model language change because as schemas sanction new instances they allow category expansion. Indeed, Broadbent addresses this point explicitly, pointing out that 'a novel form may undergo lenition if that form closely resembles a form which regularly exhibits the lenition in question' (2008: 160). But, problematically for the WY data, the lexical item 'set' undoubtedly closely resembles the form 'get' which 'regularly exhibits the lenition in question'. This raises the question of why only the fossilising properties of lexical strength, and not also the sanctioning ability of such forms, seem to be at work in constraining WY *t-to-r*.

Another interesting line of enquiry is Broadbent's suggestion that the apparent lexical effect we see for *t-to-r* in WY could be related to the collocation in which the word occurs. Following Bybee's (2001) claim that very frequent collocations are often stored together in cognition as whole units, Broadbent proposes a similar analysis for high-frequency collocations involving *t-to-r*. By assuming that *t-to-r* is stored as part of the high-frequency collocation (e.g. *but I* or *get it*), 'we have a way of capturing the intuition that *t-to-r* is essentially becoming lexicalised' (Broadbent 2008: 163).

While Broadbent (2008) offers a useful alternative to the rule-based approach to *t-to-r* proposed by Carr (1991), the two approaches are similar in that they both make claims which invite further empirical testing. For instance, it is not explicitly clear what data collection methods were employed in the WY study, or how many tokens of *t-to-r* were considered; each point is exemplified by a handful of instances which 'come from a small number of conversations between speakers from Morley' (Broadbent 2008: 145). Also, because the data are presented as lists of examples, it is unclear whether all instances of the variable context were quantified (cf. Labov's (1972) Principle of Accountability). Broadbent argues that she does not employ such techniques because her aim is to 'track the phonology of the phenomenon' and leave 'any study of the interaction between the phonology and sociolinguistic factors to further work' (2008: 145). However, this separation between theory and data requires a cautious interpretation, as it can easily lead to overgeneralisation. Docherty *et al.*'s (1997: 275) critique of aspects of phonological theory showed explicitly that in order to track the phonology of a phenomenon it is possible – even necessary – to examine corpora of

conversation such as those typically used in sociolinguistic research. Just as Docherty *et al.* (1992) tested claims from generative phonology using stringent sociophonetic methods in their analysis of Tyneside English, so too it is necessary to employ a similar technique in order to test several of the claims surrounding the suitability of UB phonology to model *t-to-r*. This is especially important given the heavy emphasis often placed on empirically analysing corpus data within the growing UB tradition (e.g. Foulkes & Docherty 2006; Bybee 2001; Pierrehumbert 2003). While speaker intuitions and anecdotal accounts of variation are invaluable tools, any theoretical approach which proposes the existence of a relationship between language structure and language use must also employ empirically testable sources of language data as a source of evidence for understanding the structure of the linguistic system. The following discussion of *t-to-r* in Liverpool offers such an account.

3 *t-to-r* in Liverpool English

3.1 *Data*

The compilation of a large spoken corpus containing recordings from Merseyside is currently underway as part of a project to understand phonological diffusion and divergence operating in and around the city of Liverpool.⁹ The corpus will be comprised of interviews and elicitation-task data with 96 working-class speakers from three localities, stratified by age and sex. The corpus is also supplemented by archive recordings in the form of 45 oral history interviews conducted with speakers from these localities who were born in the early 1900s (and recorded in the 1970s and 1980s).¹⁰ Data collection is ongoing but a preliminary analysis of *t-to-r* among a subset of the archive recordings and adolescent recordings (with a time span of ninety years in between) suggests that *t-to-r* is changing in Liverpool English. In surrounding Lancashire (Watson & Clark 2010) and in West Yorkshire (Richards 2008), *t-to-r* has remained stable over the last fifty-or-so years. However, in Liverpool, adolescent speakers are using significantly more *t-to-r* forms than speakers who were born ninety years ago (see figure 1).

There is very little variation in *t-to-r* among the adolescent generation in Liverpool (most use *t-to-r* categorically in word-final, intervocalic position) so, for the purposes of this article, we focus on an analysis of *t-to-r* in the archive materials where *t-to-r* is still a variable process. Recordings of eight speakers from the Liverpool archive data (4 males and 4 females) were digitised using a sampling rate of 48000 Hz and a 16-bit resolution. The interviews were then orthographically transcribed using ELAN,¹¹ resulting in 70000 words of time-aligned conversation.

⁹ See www.lancs.ac.uk/fass/projects/phono_levelling/index.htm

¹⁰ These recordings were kindly donated by the North West Sound Archive.

¹¹ See www.lat-mpi.eu/tools/elan/




Figure 1. Percentage of variants of word-final, intervocalic (t) across ninety years of Liverpool English in apparent time (N = 1090). The difference in proportion of *t-to-r* across both groups of speakers is significant (Fisher's exact test: $P = 0.0001^{**}$)

3.2 *Circumscribing the variable context*

Previous accounts of *t-to-r* suggest that the variable can be realised either as a rhotic approximant or as a tap. In most Northern English varieties, the tap variant is marginal and so tap tokens are either conflated with approximant variants (see Broadbent 2008) or removed entirely (see Richards 2008). However, in Liverpool English, /t/ is commonly realised as [r] in intervocalic position (Watson 2007: 352). It became clear during the coding process that the tap realisation of intervocalic /t/ in Liverpool English was constrained in similar ways to the approximant variant in West Yorkshire. In other words, the tap variant of /t/ in Liverpool English is not indicative of a different lenition process; rather *t-to-tap* seems to be following the same pattern as *t-to-approximant* in other varieties of Northern English. For this reason, the small number of tokens of [ɾ] in Liverpool English (N = 29) were conflated with the majority [r] token (N = 301) and both were coded simply as tokens of *t-to-r*.

Some confusion continues to surround the nature of the variable context of *t-to-r*. Previous discussions largely follow Wells' description of *t-to-r* as being typical in the phonological environment $t \rightarrow r / [\text{short V}] _ \# \text{V}$, yet most include the proviso that there are a number of exceptions to this rule because *t-to-r* can also occasionally occur word-medially or following a long vowel. In Liverpool English intervocalic (t) contexts, *t-to-r* occurred word-finally following a long vowel in only 1.3 per cent of all instances and word-medially in only 1.5 per cent of all instances (always in the word *whatever*). This means that while there are exceptions to Wells' (1982) description of where *t-to-r* can occur, the great majority of rhotic tokens appear in the context covered by his rule:

[short V]_#V. The decision was therefore taken to delimit the envelope of variation to word-final, intervocalic (t) following a short vowel. This resulted in a data set of 669 tokens of potential *t-to-r* candidate words from eight Liverpool speakers born in the early 1900s.

3.3 *Testing Broadbent (2008)*

It is now possible to return to claims regarding the role of lexical frequency in motivating *t-to-r* previously set out in Broadbent (2008), and to further test these claims empirically. As we illustrated above, Broadbent makes two suggestions regarding the role of lexical frequency (and, hence, the suitability of the UB framework to model *t-to-r*), both of which are based on Bybee's extensive research on frequency effects in language variation and change. These are: (i) words with high token frequency are likely to show more evidence of *t-to-r* than words with low token frequency, and (ii) frequently occurring collocations with *t-to-r* candidate words are likely to favour *t-to-r* more than less frequently occurring collocations. In the following section, we deal with each of these claims in turn.

3.3.1 *Lexical frequency*

Following a number of cross-linguistic studies on the nature of language change, Bybee (2007) has suggested two main tendencies that emerge as general principles of sound change with respect to token frequency.¹² First, Bybee (2007) proposes the REDUCTION EFFECT: phonetic change proceeds more quickly in words which occur more frequently. Some examples of this include schwa deletion in American English (Hooper 1976), *t/d* deletion in American English (Bybee 2000) and [ð] deletion in Spanish (Bybee 2002). Bybee suggests that this is because phonetic reduction is directly linked to neuromotor processing; as neuromotor processes (e.g. the movement of the articulators during language production) become more efficient, there is an increase in overlap and reduction of the gestures involved which, in language, leads to reduction and assimilation processes. Higher-frequency words have more exposure to this reduction and so undergo phonetic reduction more rapidly. Second, Bybee (2007) proposes the CONSERVING EFFECT: high-frequency words tend to have higher levels of lexical strength/entrenchment and so are likely to resist certain types of change and preserve archaisms longer. For example, with respect to morphosyntactic constructions in English, Tottie (1991), cited in Bybee (2006), shows that there is variation between a 'negative incorporation' construction (e.g. *I know nothing about it*) and a 'not negation' construction (e.g. *I don't know anything about it*). The older negative incorporation construction is used in high-frequency contexts such as with existential constructions and constructions with possessive *have* and copular *be*. In other words, the older

¹² Bybee's work on lexical frequency effects in linguistics is vast and so in the interests of space, we only deal with those proposals which are discussed in detail by Broadbent (2008).

construction has remained in high token-frequency contexts but has been lost in lower token-frequency contexts.

t-to-r is a stable variable in present-day WY (empirical evidence for this comes from Richards 2008) and so Broadbent (2008) invokes Bybee's conserving effect to account for the apparent ability of the variety to retain *t-to-r* despite rapid changes in the realisation of /t/ in other environments (e.g. *t*-glottalling is less restricted than *t-to-r* and is fast replacing released variants of intervocalic /t/). In Liverpool, *t-to-r* is changing, but the predictions surrounding the effects of token frequency remain the same because of the nature of the change. In figure 1 above we presented evidence that the occurrence of *t-to-r* in the archive Liverpool English corpus represents the early stages of a lenition (a change from [t] to [r] in word-final, intervocalic position), which has almost reached completion in younger speakers today. We can expect the reduction effect to ensure that this change was more advanced in high-frequency items because *t-to-r* is a reductive sound change. In other words, the prediction in either case (whether *t-to-r* is stable or undergoing change) is that high-frequency words should favour *t-to-r* more than low-frequency words.

An additional prediction from the UB literature on lexical frequency concerns the gradient nature of frequency effects. In a UB model, language structure is emergent and so every instance of usage (production and perception) affects the representation of linguistic structure. This means that frequency effects should be manifest in gradient rather than categorical patterns because the difference is one of degree. Although in UB models frequency effects are assumed to be gradient, often in previous work lexical frequency has been treated as a categorical variable. Researchers typically group measures of lexical frequency into 'high' and 'low' frequency using a relatively arbitrary cut-off point that is rarely the same twice, making it virtually impossible to test the widely held assumption that frequency effects in language change are gradient (e.g. Bybee 2002). In order to avoid the inherent problems associated with creating discrete categories from continuous data and to use a method of analysis which more accurately models the predictions of the UB framework, we follow Hay (2001) in treating lexical frequency as a gradient phenomenon.

There is some debate surrounding the best way to measure lexical frequency when doing this type of analysis. For instance, it is possible to take the frequency value of a particular lexical item from a large corpus of English (e.g. Dinkin 2008; Abramowicz 2007) or from a list of frequency counts such as that provided by Baayen *et al.* (2005 [1995]) in the form of the CELEX lexical database (employed by Hay 2001), or from a locally constructed corpus of a single speech community (e.g. Clark & Trousdale 2009). Ideally, we would have preferred to treat lexical frequency as a local phenomenon and so gather frequency data from the locally based Liverpool corpus (see Clark & Trousdale 2009 for a discussion of the benefits of this approach). However, given the relatively small size of the Liverpool corpus in its present state, lexical frequency values were instead taken from the token frequency of each word as it appeared in the spoken portion of the British National Corpus (BNC). In order to try to assign a frequency value that was as close as possible to the likely frequency with which these words were




Figure 2. Correlation between (log) lexical frequency and percentage of *t-to-r* for each *t-to-r* candidate word in Liverpool English archive data

used by the Liverpool speakers, we used the demographically sampled subcorpus of the spoken BNC, and included only speakers from the north of England.¹³

First we examine whether *t-to-r* is more likely in high-frequency words. To do this, we calculated a correlation between (i) the percentage of /t/ as a rhotic for all words in the corpus which fit the phonological environment [short V]_#V, and (ii) the lexical frequency values for each word, which were first normalised using the Log transformation. The correlation is highly significant by both parametric (pearson's $r = 7.54$, $p < 0.01$) and non-parametric measures (spearman's $r = 0.797$, $p < 0.01$). In other words, *t-to-r* is statistically more likely to occur in high-frequency words. However, when we plot this correlation, a very unusual pattern emerges (see figure 2).

t-to-r occurs variably only in higher-frequency lexical items, but this pattern is categorical rather than gradient. There is a clear split in the data highlighted by the vertical dashed line in figure 2: words with a log lexical frequency below 2.76 never appear with *t-to-r*, and those with a higher log lexical frequency all have *t-to-r* to some

¹³ It is not possible to delimit the geographical location of speakers further than this in the demographically sampled subcorpus of the BNC.

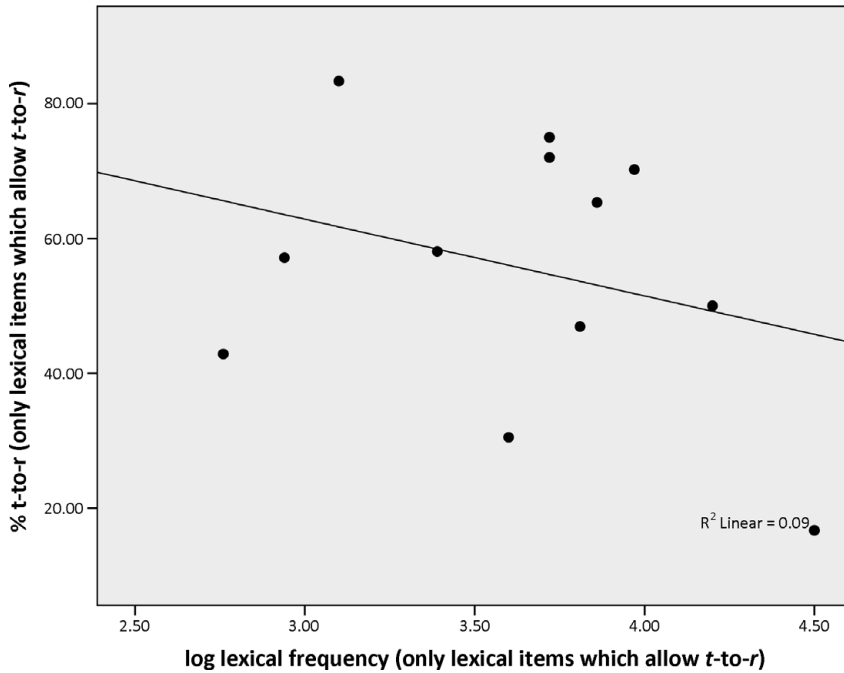


Figure 3. Correlation between (log) lexical frequency and percentage of *t-to-r* for only those words which allow *t-to-r* in Liverpool English archive data

extent.¹⁴ Interestingly, if the lexical items which never allow *t-to-r* are removed from the analysis, not only is the correlation no longer significant (pearson's $r = -3.00$, $p = 0.172$; spearman's $r = -1.44$ $p = 0.328$), but the tendency is towards a negative trend (see figure 3). In other words, if anything there is a tendency towards *t-to-r* being disfavoured by the highest-frequency items.

In testing Broadbent's (2008) suggestion that words with high token frequency are likely to show more evidence of *t-to-r* than words with low token frequency, it seems that we have stumbled upon something of a paradox. On the one hand, generative phonological theory struggles to account for the pattern of variation we find in *t-to-r* in Liverpool English. In standard generative accounts of phonology (see, for example, Kiparsky 1982; Kenstowicz 1994), the lexicon and the phonology are distinct and placed in separate modules of the grammar. The phonetic output of a lexical item is not stored in the lexicon but is arrived at once the lexical item has been retrieved from the lexicon and processed by the rules of the (phonological) grammar. This output is then fed to a phonetic implementation component which provides the

¹⁴ We explore in further detail precisely which words favour and disfavour *t-to-r* in section 3.3.3. As we will see, the words which favour *t-to-r* do not fall into natural linguistic categories (e.g. there are no nominal vs verbal or function vs lexical patterns to these data).

acoustic targets with which the word should be realized in real speech.¹⁵ However, this process applies in exactly the same way to all surface phonological representations. The *t-to-r* rule ($t \rightarrow r$ /[short V]_#V) should therefore apply blindly to all lexical items with this phonological structure. This is clearly not the case; *t-to-r* only applies to a small proportion of words with this phonological shape. Furthermore, the patterning of this rule is not random but is instead clearly governed by lexical frequency. A UB model of phonology can readily account for the fact that *t-to-r* is more likely to occur in high-frequency words than low-frequency words (regardless of whether this variable is a change in progress (reduction effect) or a stable variable (conserving effect)), but it is unclear how a UB framework would handle the categorical nature of these frequency effects. Emergentist models of language assume that frequency effects in production and perception will follow a continuous pattern because language structure is constantly updated as we experience and use language. Furthermore, a UB model of phonology would predict a positive correlation between *t-to-r* and lexical frequency in those lexical items which allow *t-to-r*, and yet in the Liverpool English data, the correlation, although not statistically significant, patterns in the opposite direction.

It seems there is partial support for the role of lexical frequency in the patterning of Liverpool English *t-to-r*, but that the picture is complex. In order to try to better understand the phenomenon in these data, we now move on to investigate the second of Broadbent's predictions.

3.3.2 Collocation frequency

Broadbent (2008) suggests that frequently occurring collocations with *t-to-r* candidate words are likely to show more evidence of *t-to-r* than less frequently occurring collocations. Similar effects have been found elsewhere. For instance, Bybee & Scheibman (1999) found that the vowel in *don't* is more reduced in the contexts in which this word occurs more often (e.g. *I don't know*), leading to the suggestion that high-frequency phrases 'change phonetically and semantically in a way that suggests storage as a single unit' (Bybee 2007: 283). It is not simply a matter of predicting phonetic change based on lexical frequency, however, because even within a particular lexical item, change may proceed more quickly or more slowly depending on the context in which the lexical item occurs. For Broadbent (2008), the assumption is that the realisation of /t/ as a rhotic may be more likely in high-frequency constructions because the rhotic variant may have become fossilised in particular high-frequency collocations (not simply in particular high-frequency words). To exemplify this claim, Broadbent provides lists of *t-to-r* candidate words in context and suggests that 'these examples illustrate that the words exhibiting *t-to-r* are commonly followed by *it*, *him*, *he*, *a* and *them*' (2008: 162). However, since the list of examples includes ten tokens of *t-to-r* preceding the word *it* but only one token of *t-to-r* preceding the word *him*,

¹⁵ This is an oversimplification and there are of course differences in the degree of abstractness that different phonologists allow for, but it is fair to say that most models of generative phonology continue to assume that morphemes and words which can be generated by automatic rules are not stored in the lexicon (see Booij 2009 for a detailed discussion of lexical storage in generative models of phonology).

the claim is ripe for further empirical testing. In order to test the claim that frequency of collocations (rather than simple lexical frequency) is an important predictor, we collected the raw frequency values for all bigrams (*t-to-r* word + following word) in which *t-to-r* was a potential variant. As before, these figures were taken from the demographically sampled northern component of the spoken BNC and then normalised using the log transformation. The frequency of occurrence we are dealing with is now rather small (for example, 37 per cent of the bigrams only occur once). Because a bigram which occurs only once will have a *t-to-r* value of either 0 or 100 per cent, it is useful to consider bigrams which occur several times in order to reach a reliable estimate of frequency of *t-to-r* occurrence. To allow this, only bigrams with four or more instances were considered in the following analysis ($N = 37$ different bigrams; a complete list of these bigrams can be found in appendix 1). Following the same procedure outlined above, the percentage of /t/ realised as a rhotic was correlated with (log) bigram frequency using both parametric and non-parametric tests. However, unlike before, these correlations did not produce significant results (pearson's $r = 0.22$, $p = 0.449$; spearman's $r = -1.24$ $p = 0.235$). *t-to-r* does not correlate statistically with bigram frequency in Liverpool English.

This is perhaps unsurprising. Bigram frequency is only an approximate indication of the extent to which two lexical items regularly co-occur. The problem with using only bigram frequency as an indication of lexical strength or unit status is that high-frequency words will appear together more often simply because of their high frequency and not necessarily because they are more lexicalised. For instance, in the Liverpool English data set, the bigram *it at* has a very high log frequency but this is merely a by-product of the high lexical frequency of the individual lexical items and has nothing to do with the extent to which these two forms are stored or accessed compositionally. Therefore, a more appropriate measure of unit status in this case would be to consider collocation strength rather than collocation frequency. A range of different methods are available to measure collocation strength statistically and these are all based on some consideration of observed vs expected frequencies of words and word pairs in a particular corpus. We consider collocation strength using five measures provided by BNCweb's collocation facility (see Hoffmann, Evert, Smith, Lee & Berglund Prytz 2008).¹⁶ The results of the correlations between *t-to-r* and the measures of collocation strength are presented below in table 1.

None of the measurements of collocation strength correlate significantly with *t-to-r* in Liverpool English. However, one final method of testing the extent to which multiple words may be stored as a single unit is to explore the constituent structure of each utterance and question whether *t-to-r* is more likely in bigrams which are, or have the potential to be, syntactic constituents. Each of the bigrams included in the previous

¹⁶ Each measure of collocational strength has its own advantages and drawbacks and so it is impossible to state which method is 'best' (e.g. some measures attach more importance to low-frequency collocations, others bias high-frequency collocates). For this reason, we measured collocational strength using five different tests: MI3, z-score, t-score, log likelihood, DICE (see Hoffmann *et al.* 2008).

Table 1. *Correlations between t-to-r and collocate strength using parametric and non-parametric measures*

Measure of collocate strength	Pearson's r	P value	Spearman's r	P value
MI3	0.026	0.441	0.098	0.286
z-score	0.146	0.198	0.115	0.252
t-score	0.162	0.173	0.101	0.280
Log-likelihood	0.185	0.140	0.110	0.262
Dice	0.115	0.252	0.112	0.257

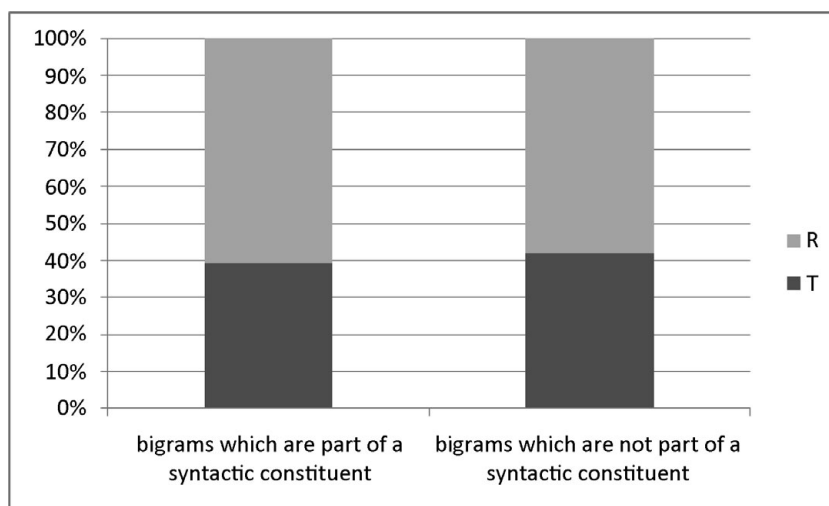


Figure 4. Frequencies of *t-to-r* in those bigrams which are part of a syntactic constituent and those which are not

analysis of collocate strength and frequency was also coded for possible constituency status using a set of standard constituency tests (replacement, movement, clefting, co-ordination and sentence fragment tests; see Carnie 2002: 51–3). For instance, the bi-gram *it is* in the utterance *just over here it is* (speaker Liv_ArchiveM04) passes the movement test (*'it is just over here'*). However, *it a* in *I'll give it a drink Bob* (Liv_ArchiveF01) clearly fails all of the standard constituency structure tests; *it a* is not a syntactic constituent. *t-to-r* was then correlated with syntactic constituency (see figure 4). A Fisher's Exact test reveals no significant difference in the frequency of *t-to-r* between these two conditions ($p = 0.5984$).

Thus, having found no relationship between bigram frequency, collocate strength or constituency structure and *t-to-r*, it is difficult to uphold, for Liverpool English, Broadbent's (2008) second claim that *t-to-r* occurs more often in frequently occurring (unit-like) collocations.

To summarise, the aim of this section was to empirically test two claims proposed by Broadbent (2008) on the likely relationship between *t-to-r* and lexical frequency. First, it

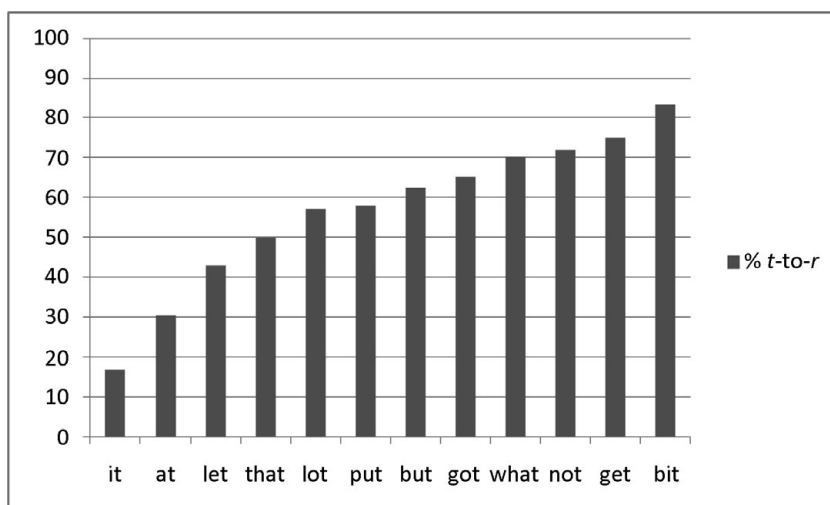


Figure 5. Frequency of occurrence of *t-to-r* for each lexical item in Liverpool archive data

was predicted that words with high token frequency are likely to show more evidence of *t-to-r* than words with low token frequency. There is statistical evidence to support the proposition that *t-to-r* occurs in high-frequency words more often than in low-frequency words when the variable environment is defined as [short V]_#V. However, on closer inspection, not all words in this environment allow *t-to-r*, suggesting a frequency effect that is categorical. A UB model of phonology can account for frequency effects in language production, but the data here are difficult to account for in UB terms because of the claim that frequency effects are gradient. Second, it was predicted that frequently occurring collocations with *t-to-r* candidate words are likely to show more evidence of *t-to-r* than less frequently occurring collocations. There is no evidence from Liverpool English that *t-to-r* occurs more often in high-frequency collocations or in collocations with high collocation strength.

Given the problems associated with confirming certain UB predictions in the data on *t-to-r* in Liverpool English, we must consider whether a UB model of grammar still provides a useful way to interpret variation in these data. In order to answer this question, and in the process reach a better understanding of *t-to-r*, it is necessary to examine in more detail the specific lexical items that favour *t-to-r* and the contexts in which they occur.

3.3.3 *Token/ type frequency and schema strength*

Figure 5 shows the percentage of *t-to-r* in words which allow this variant. These data are arranged along the x-axis by frequency of rhotic forms.

The most striking observation is that these words do not fall into patterns according to natural linguistic classes. Grammatical category seems unimportant (there are no nominal/verbal or function/lexical patterns to these data); *t-to-r* also seems to be

Table 2. *Frequency of occurrence of each t-to-r word and bigram in the Liverpool English corpus*

Word	Number of instances of each word in Liverpool corpus	Number of different bigrams in which each word occurs in Liverpool corpus	% <i>t-to-r</i>
<i>it</i>	83	27	16.7
<i>at</i>	36	9	30.5
<i>let</i>	7	5	42.86
<i>that</i>	43	16	50
<i>lot</i>	49	3	57.14
<i>put</i>	31	8	58.06
<i>but</i>	93	19	62.4
<i>got</i>	75	16	65.33
<i>what</i>	47	12	70.22
<i>not</i>	25	11	72
<i>get</i>	64	14	75
<i>bit</i>	18	2	83.33

unaffected by the quality of the vocalic element immediately preceding the /t/ (e.g. the KIT vowel appears in *it* with only 17 per cent *t-to-r* and in *bit* which occurs with 83 per cent *t-to-r*). An explanation for patterns of *t-to-r* which considers variation only in terms of the individual *t-to-r* candidate word is, therefore, unhelpful. However, if we consider the variation not just in terms of the individual word in which *t-to-r* may occur but also in terms of how many different bigrams the word appears in, an interesting picture emerges. Table 2 details the number of instances of each lexical item in the Liverpool English corpus along with the number of different bigrams in which each word appears (e.g. the word *bit* occurs 18 times, in two different bigrams – *bit of* [$n = 17$], *bit older* [$n = 1$]).

Because we are now dealing with corpus-internal frequency counts and not BNC frequency counts as before, the number of tokens are much smaller and so do not lend themselves to an analysis which employs statistical testing. Nevertheless, there does seem to be a special relationship between the number of instances of each of these lexical items in the Liverpool English data set and the number of different bigrams in which these instances occur. This is most apparent when we look at the two lexical items which occur with the most and least *t-to-r* in Liverpool English: *it* and *bit*. Both *bit* and *it* can be categorised as ‘high-frequency’ words, but notice that whereas all eighteen occurrences of the word *bit* appear in only 2 different bigrams, the word *it* occurs with a far greater number of words filling the bigram’s second slot ($n = 27$).

The relationship between lexical frequency and the constructions in which the lexical items appear has received attention in UB linguistics through discussions of morphological productivity, but, as we show below, the relationship is also of interest in the discussion of *t-to-r*. Productivity is defined as a pattern’s ability to apply to novel items in language. One generalisation that has emerged from the UB literature

is that patterns with higher type frequencies (where a ‘type’ is a more abstract pattern, abstracted over instances of language use) are more likely to be productive (Bybee 2000: 12–13). The classic example of this is the pattern of past-tense formation in English. The *-ed* past-tense form (or type) occurs with a very high number of instances in English and it is phonologically unrestricted. Irregular forms of the past tense, by contrast, are much more restricted, applying to individual verbs or small clusters of verbs with similar phonological structure (e.g. *tear/tore*; *bear/bore*; *wear/wore*). The *-ed* pattern of past-tense formation is much more productive and this, Bybee (2000: 12–13) argues, is a function of its high type frequency. Other examples include work on phonotactic productivity which has shown that nonce words are often judged to be more acceptable if they resemble high type-frequency patterns of real English words (e.g. see Pierrehumbert 1994). The proposed reason for this relationship between productivity and type frequency in UB linguistics is that both type and token frequency influence schema strength.

UB models of language, particularly those within the Cognitive Linguistic tradition (e.g. see Langacker 1987, 1991), have been concerned with understanding the ways in which speakers abstract and categorise linguistic structure from their experience of language in use. The relationship between more or less abstract patterns of linguistic knowledge has often been discussed in terms of the relationship between a *schema* and *instances* of the schema. The proposal that knowledge structures are organised around schemata was imported to linguistics from cognitive science and social psychology where schema theory is a well-established model of categorization (e.g. Mandler 1984; Rumelhart 1984; Rumelhart & Ortony 1977). In other disciplines, schemas are described as ‘abstract concepts that represent a general case, emphasising what is similar among a number of instances’ (Fiske & Morling 1996: 498). The same definition applies to linguistic knowledge where it is invoked by UB models to describe the relationship between instances of language use and the categories that speakers construct (at different levels of abstraction) over instances of language use (e.g. see Langacker 1987: 371).

Frequency effects in language are not only relevant to understanding lexical strength; frequency of occurrence is also important for schema strength because a schema gains strength (or becomes entrenched) in proportion to both the type and token frequency of the instances which elaborate it. Taylor (2002: 277) describes this relationship in the following way: ‘High token frequency of an instance entrenches the instance and weakens (or at least does not strengthen) the schema. High type frequency of an instance entrenches the schema and weakens (or at least does not strengthen) the instances.’¹⁷ The relationship between type/token frequency and schema strength is diagrammatised in figure 6.

The schema on the right in figure 6 has a number of instances and so the schema is strengthened each time one of the instances is used. The schema on the left has only a

¹⁷ This terminology can be confusing. ‘High type frequency of an instance’ simply means a high number of instantiations of a type in any given corpus (e.g. high number of past-tense words marked with *-ed*).




Figure 6. The relationship between frequency and schema strength (adapted from Taylor 2002: 276)

few instances, but these occur frequently and so each time the instances are used, the instances themselves are strengthened but the schema is not.

The notion of schema strength has become an important consideration in UB (especially Cognitive Linguistic) discussions of productivity because it is assumed that schemas which are strong/highly entrenched are more likely to be productive. In the case of highly entrenched instances, where the schema is weak, the pattern will not be productive. We can return to the past-tense formation in English as an illustration of this. The regular past-tense morpheme in English (+-ed) attaches to a great many lexical items and so the schema [PAST+ -ed] has a high type frequency. This means that the schema will be strongly entrenched and it is the reason why the pattern is productive (schema 2 in figure 6). Compare this with the pattern of irregular verbs which mark past tense by changing phonological form (e.g. *tear/tore*; *bear/bore*; *wear/wore*). The schema [PAST + e→o] is instantiated by only a few lexical items in English. Unless speakers have a large number of instances from which to abstract commonality, the schema will be weakly entrenched and not productive (schema 1 in figure 6).

How can the notion of schema strength help us to better understand the patterns of variation that we see in *t-to-r*? Although not explicitly discussed in linguistics, a great deal of work exists in the cognitive science and psychology literature which suggests that strong schemas are not only more productive; they are also more resistant to change. McNeil & Alibali (2002: 661) write:

When a schema is strong, it resists change (e.g. Allport, 1954, Bartlett, 1932, Bruner, 1957, Schutzwohl, 1998). Individuals who have a strong schema have been shown to resist learning new information when it is more specific than (Adelson, 1984, Thorndyke & Hayes-Roth, 1979), not applicable to (Voss, Vesonder, & Spilich, 1980), or discrepant with (Marchant, Robinson, Anderson & Schadewald, 1991, Markus, 1977) their current schema . . . Importantly, the stronger a schema is, the more resistant it is to changing (see Luchins, 1932).

Although these studies have all examined change with respect to behavioural schemas,¹⁸ the concept of a schema as a higher-level cognitive structure that both supersedes

¹⁸ McNeil & Alibali (2002) deal particularly with mathematical schemas and show that school children with a strongly entrenched maths additional schema ($3+4+5 = 12$) have more difficulty understanding maths equivalence problems ($3+4+5 = 3+_$) than those with a more weakly entrenched maths additional schema.

its constituent parts and endorses new environmental stimuli remains the same in linguistics. This means that the notion that strong schemas resist change should also be applicable to linguistic schemas. By combining ideas from schema theory in cognitive linguistics (i.e. high type frequency encourages strong schema) and cognitive science (i.e. strong schemas resist change), we should expect to find the following in the case of *t-to-r* in Liverpool English: (a) high-frequency lexical items which occur with high type/bigram frequency will be more resistant to *t-to-r* because their schemas will be strongly entrenched, (b) high-frequency lexical items which occur with low type/bigram frequency will show more evidence of the phonological reduction because their schemas will be more weakly entrenched.

This is what we find for *bit* and *it* in Liverpool English, illustrated in figure 7. *t-to-r* occurs with a frequency of 83 per cent in the lexical item *bit* in Liverpool English, whereas in *it* it occurs with a much lower frequency of 17 per cent. The frequency of *t-to-r* in these words is not simply a by-product of their high lexical frequency (see section 3.3.1), nor is it simply a result of the frequency or strength of collocations within which these words occur (see section 3.3.2). However, both type and token frequency can influence the likely strength of any generalisations or abstractions arising from the context in which these words regularly occur in language use. Generalisations (or schemas) which are weak are unlikely to resist change. Strong schemas, by contrast, will resist change. Thus, we can perhaps account for the fact that *it* resists the change from *t-to-r* despite the high lexical incidence of this word because both type and token frequency are important to the strength of more abstract schemas surrounding this lexical item. Because the [i(t)#V] schema is instantiated by so many bigrams, it is strong, and so is able to resist the change from *t-to-r* for longer.

4 Conclusions

This article set out to empirically test two particular hypotheses that had been proposed by Broadbent (2008) in the debate surrounding how best to model *t-to-r* in phonological theory. The first suggestion was that words with high token frequency are likely to show more evidence of *t-to-r* than words with low token frequency. This hypothesis was largely confirmed by statistical analysis and so it would appear that this particular prediction from UB phonology is borne out in Liverpool English. However, the *t-to-r* data from Liverpool do not show gradient frequency effects; the patterning of lexical frequency is categorical and this is difficult for a UB model to explain. The second hypothesis was that frequently occurring collocations would be more likely to favour *t-to-r* than less frequently occurring collocations. There was no evidence that bigram frequency, collocate strength or constituency status was a motivating factor in the realisation of *t-to-r* in Liverpool English.

Despite the categorical nature of the frequency effects found in these data, lexical frequency does seem to constrain variation in *t-to-r* in some way and so rather than abandoning our efforts to understand the relationship between phonological change

*Authors' addresses:**Linguistics and English Language**Lancaster University**Lancaster LA1 4YL**UK**l.clark6@lancaster.ac.uk**School of Languages, Cultures and Linguistics**University of Canterbury**Private Bag 4800**Christchurch 8140**New Zealand**kevin.watson@canterbury.ac.nz*

References

- Abramowicz, Łukasz. 2007. Sociolinguistics meets exemplar theory: Frequency and recency effects in (ing). *Penn Working Papers in Linguistics* 13(2), 26–37.
- Alba, Joseph W. & Lynn Hasher. 1983. Is memory schematic? *Psychological Review* 93, 203–31.
- Adelson, Beth. 1984. When novices surpass experts: The difficulty of a task may increase with expertise. *Journal of Experimental Psychology: Learning, Memory and Cognition* 10, 483–95.
- Allport, Gordon W. 1954. *The Nature of Prejudice*. Reading, MA: Addison-Wesley.
- Asprey, Esther. 2008. The sociolinguistic stratification of a connected speech process – the case of the T to R rule in the Black Country. *Leeds Working Papers in Linguistics and Phonetics* 13, 109–40.
- Baayen, R. Harald, Richard Piepenbrock & Leon Gulikers. 2005. *The CELEX lexical database (release 2)* [CD-ROM]. Linguistic Data Consortium, University of 333 Pennsylvania [Distributor]. Philadelphia, PA.
- Bartlett, Frederic Charles. 1932. *Remembering: A study in experimental and social psychology*. Cambridge: Cambridge University Press.
- Booij, Geert. 2009. Lexical storage and phonological change. In Kristin Hanson & Sharon Inkelas (eds.), *The nature of the word: Studies in honour of Paul Kiparsky*, 487–506. Cambridge, MA: MIT Press.
- Bresnan, Joan & Tatiana Nikitina. 2003. On the gradience of the dative alternation. MS available to download: www.stanford.edu/~bresnan/download.html
- Broadbent, Judith. 2008. *t to r* in West Yorkshire English. *English Language and Linguistics* 12(1), 141–68.
- Bruner, Jerome S. 1957. On perceptual readiness. *Psychological Review* 64, 123–52.
- Bybee, Joan. 1995. Regular morphology and the lexicon. *Language and Cognitive Processes* 10(5), 425–55.
- Bybee, Joan. 2001. *Phonology and language use*. Cambridge, Cambridge University Press.
- Bybee, Joan. 2002. Word frequency and context of use in the lexical diffusion of phonetically conditioned sound change. *Language Variation and Change* 14, 261–90.
- Bybee, Joan. 2006. From usage to grammar: The mind's response to repetition. *Language* 82(4), 711–33.

- Bybee, Joan. 2007. *Frequency of use and the organization of language*. Oxford: Oxford University Press.
- Bybee, Joan & Joanne Scheibman. 1999. The effect of usage on degrees of constituency: The reduction of *don't* in English. *Linguistics* 37(4), 575–96.
- Carnie, Andrew. 2002. *Syntax: A generative introduction*. Oxford: Blackwell
- Carr, Phillip. 1991. Lexical properties of post-lexical rules: Postlexical derived environment and the Elsewhere Condition. *Lingua* 85, 255–68.
- Clark, Lynn & Graeme Trousdale. 2009. The role of frequency in phonological change: Evidence from TH-fronting in east-central Scotland. *English Language and Linguistics* 13(1), 33–55.
- Dinkin, Aaron. 2008. The real effect of word frequency on phonetic variation. *Penn Working Papers in Linguistics* 14(1), 97–106.
- Docherty, Gerard J., Paul Foulkes, James Milroy, Lesley Milroy & David Walshaw. 1997. Descriptive adequacy in phonology: A variationist perspective. *Journal of Linguistics* 33, 275–310.
- Fiske, Susan T. & Beth A. Morling. 1996. Schemas/schemata. In Antony S. R. Manstead & Miles Hewstone (eds.), *The Blackwell encyclopedia of social psychology*, 489–94. Oxford: Blackwell.
- Foulkes, Paul & Gerard Docherty. 2006. The social life of phonetics and phonology. *Journal of Phonetics* 34(4), 409–38.
- Hammond, Michael. 1999. Lexical frequency and rhythm. In Mike Darnell, Edith A. Moravcsik, Frederick J. Newmeyer, Michael Noonan & Kathleen M. Wheatley (eds.), *Functionalism and formalism in linguistics: General papers*, 329–58. Amsterdam: Benjamins.
- Harris, John & Jonathan Kaye. 1990. A tale of two cities: London glottalling and New York City tapping. *The Linguistic Review* 7, 251–74.
- Hay, Jennifer. 2001. Lexical frequency in morphology: Is everything relative? *Linguistics* 39(6), 1041–70.
- Hoffmann, Sebastian, Stefan Evert, Nicholas Smith, David Lee & Ylva Berglund Prytz. 2008. *Corpus linguistics with BNCWeb – a practical guide*. Frankfurt am Main: Peter Lang.
- 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.
- Hudson, Richard A. 2007. *Language networks: The new word grammar*. Oxford: Oxford University Press.
- Johnson, Keith. 1997. Speech perception without speaker normalization: An exemplar model. In Keith Johnson & John W. Mullennix (eds.), *Talker variability in speech processing*, 145–65. San Diego: Academic Press
- Kemmer, Suzanne & Michael Barlow. 2000. Introduction: A usage-based conception of language. In Michael Barlow & Suzanne Kemmer (eds.), *Usage-based models of language*, vii–1. Stanford: CSLI.
- Kemmer, Suzanne & Michael Israel. 1994. Variation and the usage-based model. In Katherine Beals, Jeanette Denton, Robert Knippen, Lynette Melnar, Hisami Suzuki & Erica Zeinfeld (eds.), *Papers from the 30th Regional Meeting of the Chicago Linguistic Society: Parasession on Variation and Linguistic Theory*, 165–79.
- Kenstowicz, Michael. 1994. *Phonological theory*. Cambridge MA: Blackwell.
- Kiparsky, Paul. 1982. *Explanation in phonology*. Dordrecht: Foris.
- Labov, William. 1972. *Language in the inner city: Studies in the Black English Vernacular*. Philadelphia: University of Pennsylvania Press.
- Langacker, Ronald W. 1983. *Foundations of cognitive grammar: Theoretical prerequisites*, vol. 1. Stanford: Stanford University Press.

- Langacker, Ronald W. 1987. *Foundations of cognitive grammar: Descriptive application*, vol. 2. Stanford: Stanford University Press.
- Langacker, Ronald W. 1991. *Concept, image and symbol: The cognitive basis of grammar*. Berlin: Mouton de Gruyter.
- Langacker, Ronald W. 2000. A dynamic usage-based model. In Michael Barlow & Suzanne Kemmer (eds.), *Usage-based models of language*, 1–63. Stanford: CSLI.
- Luchins, Abraham S. 1942. Mechanization in problem solving. *Psychological Monographs* 54(6), whole no. 248.
- Mandler, Jean Matter. 1984. *Stories, scripts, and scenes: Aspects of schema theory*. Hillsdale, NJ: Erlbaum.
- Manning, Christopher D. 2003. Probabilistic syntax. In Rens Bod, Jennifer Hay & Stefanie Jannedy (eds.), *Probabilistic linguistics*, 289–341. Cambridge, MA: MIT Press.
- Marchant, Garry, John Robinson, Urton Anderson & Michael Schadewald. 1991. Analogical transfer and expertise in legal reasoning. *Organizational Behaviour and Human Decision Processes* 48, 272–90.
- Markus, Hazel. 1977. Self-schemata and processing information about the self. *Journal of Personality and Social Psychology* 2, 63–78.
- McNeil, Nicole M. & Martha W. Alibali (2002). A strong schema can interfere with learning: The case of children's typical addition schema. In Wayne D. Gray & Christian D. Schunn (eds.), *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society*, 661–6. Mahwah, NJ: Erlbaum.
- Pierrehumbert, Janet B. 1994. Knowledge of variation. In Katherine Beals, Jeanette Denton, Robert Knippen, Lynette Melnar, Hisami Suzuki & Erica Zeinfeld (eds.), *Papers from the 30th Regional Meeting of the Chicago Linguistic Society: Parasession on Variation and Linguistic Theory*, 232–56.
- Pierrehumbert, Janet B. 2001. Exemplar dynamics: Word frequency, lenition, and contrast. In J. Bybee & P. J. Hopper (eds.), *Frequency effects and the emergence of linguistic structure*, 137–57. Amsterdam: John Benjamins.
- Pierrehumbert, Janet B. 2002. Word-specific phonetics. *Laboratory Phonology VII*, 101–39. Berlin: Mouton de Gruyter.
- Richards, Hazel. 2008. Mechanisms, motivations and outcomes of change in Morley (Leeds) English. Unpublished PhD thesis, University of York.
- Rumelhart, David. E. 1984. Schemata and the cognitive system. In Robert S. J. Wyer and Thomas K. Srull (eds.), *Handbook of social cognition*, 161–88. Hillsdale, NJ: Erlbaum.
- Rumelhart, David. E. & Andrew Ortony. 1977. The representation of knowledge in memory. In Richard C. Anderson, Rand J. Spiro & William E. Montague (eds.), *Schooling and the acquisition of knowledge*, 99–135. Hillsdale, NJ: Erlbaum.
- Schützwohl, Achim. 1998. Surprise and schema strength. *Journal of Experimental Psychology: Learning, Memory and Cognition* 24, 1182–99.
- Taylor, John R. 2002. *Cognitive Grammar*. Oxford: Oxford University Press.
- Thorndyke, Perry W. & Barbara Hayes-Roth. 1979. The use of schemata in the acquisition and transfer of knowledge. *Cognitive Psychology* 11, 82–106.
- Tottie, Gunnel. 1991. Lexical diffusion in syntactic change: Frequency as a determinant of linguistic conservatism in the development of negation in English. In Dieter Kastovsky (ed.), *Historical English syntax*, 439–67. Berlin: Mouton de Gruyter.
- Voss, James F., Gregg T. Vesonder & George J. Spilich. 1980. Text generation and recall by high-knowledge and low-knowledge individuals. *Journal of Verbal Learning and Verbal Behaviour* 19, 651–67.
- Watson, Kevin. 2007. Liverpool English. *Journal of the International Phonetics Association*, 37(3), 351–60.

Watson, Kevin & Lynn Clark. 2010. 't-to-r' in British English: Word frequency, constructions, and sociolinguistics. Paper presented at *New Ways of Analysing Variation* (NWAY) 39, San Antonio.

Wells, John. C. 1982. *Accents of English*. Cambridge: Cambridge University Press.

Yang, Charles D. 2004. Universal Grammar, statistics, or both? *Trends in Cognitive Sciences* 8, 451–6.

Appendix 1

Word	Bigram frequency log	MI3	z-score	t-score	Log- likelihood	Dice	% r
at a	1.87	12.7	1.9	1.7	3.6	0.01	33.33
at all	2.37	19.7	56.4	14.4	874.0	0.05	0
at home	2.01	18.9	66.5	9.9	602.1	0.04	40
at it	2.11	14.4	2.6	2.4	6.7	0.01	83.33
bit of	2.5	21.8	104.4	17.4	1781.3	0.06	88.24
but as	1.36	9.8	2.5	2.0	5.9	0.01	25
but he	3.13	19.1	33.0	14.7	565.8	0.03	80
but I	3.05	22.7	62.9	27.2	2101.8	0.05	69.23
but in	1.43	8.4	-3.9	-5.8	-20.0	0.00	50
but it	2.78	20.4	34.5	17.9	737.8	0.03	50
but it's	2.42	20.4	68.8	15.4	1207.5	0.02	50
get a	2.66	20.2	42.6	17.6	883.9	0.04	87.5
get it	2.56	18.4	19.7	12.0	276.9	0.02	40
get out	2.12	17.2	30.1	10.2	347.6	0.03	80
get up	2.07	16.3	21.9	9.0	226.8	0.02	25
got a	3.03	23.5	94.3	29.6	3250.3	0.08	77.78
got in	1.88	12.7	1.2	1.1	1.4	0.01	50
got it	2.53	17.7	11.2	8.3	106.0	0.02	60
got older	0.3	1.2	0.1	0.6	0.5	0.00	100
got on	1.69	11.2	-0.2	-0.3	-0.1	0.01	42.86
got out	1.56	11.1	3.2	2.5	9.0	0.01	75
it a	2.1	12.1	-15.6	-29.9	-355.0	0.01	12.5
it all	2.42	17.2	12.9	8.8	130.2	0.01	0
it is	3.01	22.4	60.6	26.0	1904.5	0.05	22.22
it on	2.6	18.2	12.1	9.0	122.3	0.02	15.34
it out	2.44	17.8	18.9	11.0	244.6	0.02	33.33
it up	2.51	18.1	18.9	11.4	253.3	0.02	0
lot of	2.6	23.2	154.9	19.5	2722.8	0.08	59.57
not a	2.23	15.9	10.5	7.1	85.1	0.01	72.72
put in	1.93	14.8	13.5	6.9	107.7	0.01	83.33
put it	2.7	20.9	55.5	19.6	1306.7	0.03	61.54
put up	1.32	10.0	3.7	2.6	11.2	0.01	50
that had	1.3	8.4	-3.5	-5.1	-16.2	0.00	50
that I	1.48	15.4	-13.7	-20.2	-245.8	0.01	50
what a	2.04	13.2	-2.4	-2.8	-6.7	0.01	75
what I	2.83	20.1	21.9	14.6	375.3	0.03	75
what it	2.41	16.2	1.3	1.3	1.7	0.01	90.92