The changing realisation of 'the' before vowels in New Zealand English[1]

Jen Hay[+], Abby Walker[*], Jayne McKenzie[+] and Daniel Nielsen[+]
+NZILBB, University of Canterbury
*The Ohio State University

Abstract

This paper reports on a study of the realisation of the word 'the' before vowels in New Zealand English. New Zealand English speakers use both 'thuh' (/ðə/) and 'thee' (/ði:/) in this position. We show that younger speakers and 'non-professional' speakers are more likely to use a 'thuh' realisation than older, 'professional' speakers. 'thuh' realisations and 'thee' realisations are both attested with and without glottalisation on the following vowel. When glottalisation is involved, a 'thuh' realisation is much more likely in infrequent phrases. When glottalisation is not involved, the frequency effect is absent. We speculate that these differences in frequency effects and in glottalisation may indicate that two separate processes are at work: one involving reduction and the other analogy.

1. Introduction

This paper investigates the changing pronunciation of the determiner 'the' in New Zealand English. In most varieties of English, it is standard to pronounce the definite determiner as /ði:/ before a vowel and /ðə/ before a consonant (see, e.g. textbooks such as Swan (1995)). While many speakers of New Zealand English (NZE) also show this distribution, we noticed that some speakers appear to produce /ðə/ prevocalically. This has also been informally noted in New Zealand media. For example, in a Christchurch newspaper column on language, prescriptivist Frank Haden writes:

Michael Burch, who teaches speech to a group of 30 in a profession where correct speech is essential, says "thuh" has taken over from "thee" almost entirely among the under-30s…Indeed a few of them even appeared to be unaware "thee" even existed. One of them even asked "what's a vowel" ("thuh over thee", 2003).

This paper documents a study designed to investigate this phenomenon and the degree to which it is a recent innovation. The data come from two sources: a corpus of spontaneous speech from 39 New Zealand English speakers and a corpus of 56 speakers describing a cartoon designed to elicit tokens of 'the' before vowel-initial words. The results indicate that the pronunciation of 'the' is undergoing change in New Zealand English, which is conditioned by social and usage-based factors.

2. Methodology

The data come from two corpora. The first is a corpus of spontaneous speech which is held at the University of Canterbury and known as the Canterbury Corpus. The second is a corpus of speakers describing a cartoon, which was specifically designed to elicit tokens of 'the' in appropriate environments. We will refer to this latter corpus as the Cartoon Corpus.

The Canterbury Corpus consists of recordings held at the University of Canterbury. These recordings have been made by students of a third-year New Zealand English course and, at the time of analysis, consisted of interviews with speakers born between 1933 and 1981. Speakers in the corpus are categorised by their age, gender and social class. A binary gender code is assigned according to self-report. Participants are categorised as 'older' or 'younger' depending on the age they are at the time of recording: participants aged 20-30 are 'younger' and participants aged 45-60 are 'older'. Note that birthdate information is also collected. Social class categorisation is also binary; speakers are informally classified as 'non-professional' speakers if they have had no tertiary education, have a manual or unskilled occupation, and their parents also have manual/unskilled occupations, with the opposite characteristics applying to 'professional' speakers. A more formal post-hoc check on these classifications has been conducted using a social class index, incorporating both occupation and education levels. This reveals a good deal of separation between the two groups, indicating that they do, indeed, represent two very different social groups (Gordon et al., 2007).

Our dataset taken from the Canterbury Corpus consists of a sample of 39 speakers – across 8 different cells – female/male, professional/non-professional, older/younger speakers. Younger speakers were born between 1965 and 1981. Older speakers were born between 1933 and 1965. We analysed five speakers in each cell, with the

exception of younger male professionals, for which cell we analysed four. Between them, these speakers produced 493 tokens of 'the' before a vowel-initial word. These tokens were reasonably well spread across the 8 cells, as shown in Table 1.

Table 1: Distribution of tokens from the Canterbury Corpus

		Older	Younger
Female	Professional	71	63
	Non-professional	53	59
Male	Professional	71	44
	Non-professional	86	46

We also designed a short elicitation task for friends, family and colleagues, in which the speaker is asked to describe the series of events in a cartoon strip designed to elicit tokens of 'the' before vowels (see Figure 1). The cartoon featured objects such as an alien, an apple, and an ice-cream, and because each object was shown more than once, we hoped that although the speaker might use an indefinite article 'a' or 'an' on the first mention of a particular object, they would use a definite article 'the' on subsequent mentions.

Figure 1: Cartoon used in Elicitation Task

Speakers were categorised according to age, sex and professional status using the same criteria as the Canterbury Corpus. From 56 speakers, the cartoon elicited a total of 327 tokens of 'the' before vowel-initial words. The distribution of these tokens across the eight cells are displayed in Table 2.

Table 2: Distribution of tokens from the Cartoon Corpus

		Older	Younger
Female	Professional	55	33
	Non-professional	19	35
Male	Professional	35	53
	Non-professional	42	55

Both datasets were analysed in the same way. Coders binarily assessed each prevocalic token of "the" as being either /ðə/ or /ði/, using both acoustic and auditory cues. For each token, two analysts coded independently, and any discrepancies in judgements were checked by a third analyst. We also recorded whether the boundary involved any glottalisation or a glottal stop. In a small number of cases, the determiner did not contain a vowel which could be isolated from the onset of the following word (e.g. "thother"). Such tokens were excluded from the final dataset. We also excluded cases in which a substantial pause intervened between the determiner and the vowel initial word, as these may be more prone to an /i/ realisation, as discussed by Fox Tree and Clark (1997).

As well as considering the social factors of speaker age, social class and gender, we coded for the following linguistic factors: the status of the following word (head noun, filler, other), whether the following vowel was high/low, back/front, monophthong/diphthong, whether the 'the' was stressed, whether the following syllable was stressed, the frequency of the following word, whether the head noun has been previously mentioned (i.e., discourse given), and whether there was glottalisation on / following the determiner.

We also considered the log frequency of the following word lemma, as given in the CELEX Lexical Database (Baayen et al., 1995). Additionally, we were interested in the degree to which the word may be overrepresented in its co-occurence with 'the' - that is, when the word occurs, how frequently does it occur following 'the'? In order to assess this, we considered the frequency of the word in our database of words following 'the' in our sample of Canterbury corpus interviews. We centred this frequency distribution so that it was comparable to the CELEX counts and then took the difference between the two log frequencies (henceforth, "relative 'the' frequency"). Positive values of this frequency difference relate to words which are over-represented with 'the' (compared to CELEX), and negative values represent words which are under-represented with 'the'. Examples of words which frequently occur with 'the' are *Olympics*, *eldest* and *airwaves*. Examples of words which infrequently occur with 'the' are *expression, education and arm.*

3. Results

Rates of /ðə/ usage prevocalically were similar across both corpora (18% in the Canterbury Corpus and 23% in the Cartoon Corpus). In the following analysis we therefore collapse the two datasets into a single corpus of 820 tokens.

Tokens of both variants were variably realised with and without glottalisation and with glottal stops. The glottal and glottalised variants were more frequent with /ðə/ (henceforth, THUH) than /ði/ (henceforth, THEE). This distribution is shown in Table 3.

Table 3: Proportion of THEE and THUH forms which are realised with different degrees of glottalisation

	THEE	THUH
GLOTTAL STOP	0.11	0.34
GLOTTALISATION	0.05	0.33
OTHER	0.84	0.34

44 tokens were followed by fillers (such as *um* and *ah*). 39% of these were realised as THUH (as opposed to 19% of non-filler items). We excluded these tokens from the following statistical analysis.

To examine how social and linguistic factors influence 'the' realisation, we hand-fit a binomial mixed effects model to the data (Baayen et al., 2008), with speaker and word as random effects. The mixed effects model ensures that no particular speakers or words have undue influence on the results. We tested a range of social and linguistic factors (see previous section). The resulting model is shown in Table 4.

Table 4: Results of a logistic mixed effects model, modelling the likelihood of a THUH realisation. Speaker and word were random effects

	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-3.28664	0.51739	-6.352	<.0001
relative 'the' frequency	0.0684	0.07814	0.875	.38
glot=yes	2.16019	0.34849	6.199	<.0001
age=young	1.90811	0.52526	3.633	<.001
class=professional	-1.41736	0.51582	-2.748	<.01
relative 'the' frequency x glot=yes	-0.22869	0.09699	-2.358	<.02

This model shows that there is a significant effect of class, as pictured in Figure 2 (left panel). Non-professional speakers are significantly more likely to use THUH

than professional speakers. The right panel of Figure 2 shows the model prediction for age. Younger speakers are significantly more likely to produce THUH.

A difference between younger and older speakers could indicate a change-in-progress or a stable age-graded phenomenon. We find it much more likely that what we are dealing with is a change-in-progress towards the THUH realisation. Older reports of New Zealand English do not feature this phenomenon, and the increase of THUH before vowels has been reported for other dialects (Cheshire, 2011). And the realisation of 'the' before vowels is not something that appears to be above the level of consciousness in NZ. We find it unlikely that it is manipulated as a stylistic variable (which would leave open the interpretation that older speakers revert to a more 'formal' realisation). Some evidence in favour of the phenomenon being non-stylistic is the lack of observed difference in the cartoon and corpus materials in this dataset.
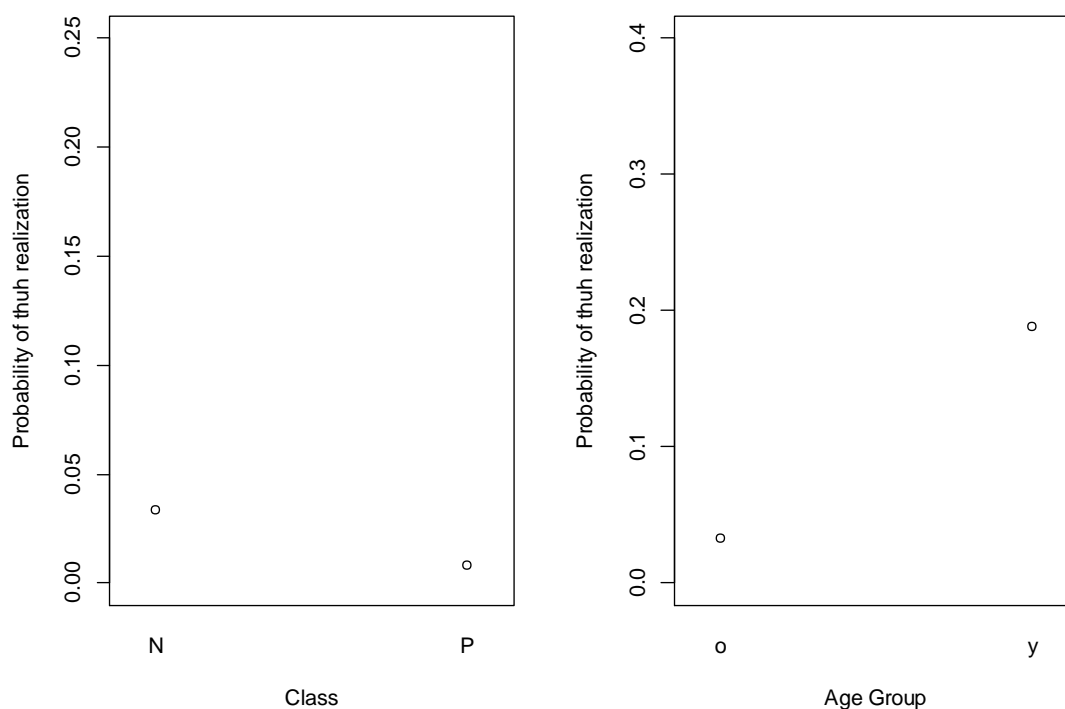
Figure 2: Model predictions for the effect of Class (left panel – N (non-professional) and P (professional)) and Age (right panel – O (older) and Y (younger)) on the likelihood of a THUH realisation

29

The model also includes a significant interaction between the degree to which the word is a relatively frequent collocation with 'the' and whether it is produced with glottalisation (collapsing together glottalisation and glottal stops). The effect is plotted in Figure 3.
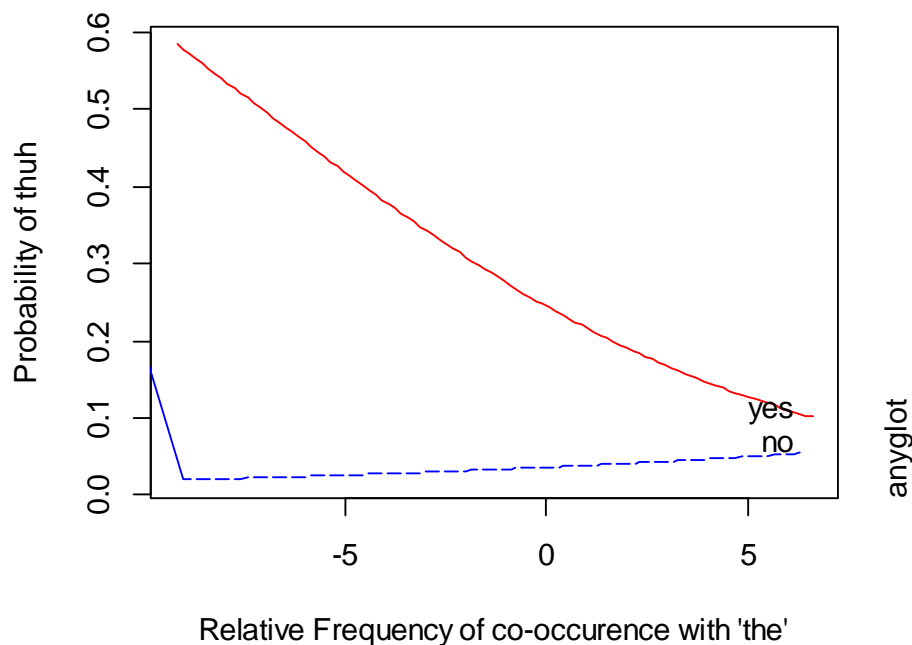
Figure 3: Model prediction of the effect of the relative frequency of co-occurence with 'the' on THUH realisations. High values represent words which occur disproportionately with 'the'. Low values represent words which occur disproportionately in other contexts. The dashed line represents those tokens containing no glottalisation or glottal stop. The solid line represents tokens containing glottalisation/glottal stop

What this shows is that the relative frequency of collocation with 'the' plays an important role, but that this is only true in cases where we have some glottalisation. If there is no glottalisation, the probability of a THUH realisation is not only much lower, but it is also unaffected by the relative frequency. That is, the interaction is driven by the glottalised forms. In a model of these forms alone, the relative frequency remains significant, but this is not true of models of non-glottalised forms alone. In cases with glottalisation we get more THUH when the word frequently occurs together with 'the'. As a reminder, some examples of words which frequently occur with 'the' are *Olympics*, *eldest* and *airwaves*. Examples of words infrequently occurring with 'the' are *expression, education and arm.*

Our interpretation of the interaction in this figure is that it teases apart two underlying mechanisms driving the change to THUH. There are two ways in which

THEE could become THUH: one of these is through reduction and one is through analogy. We argue that the latter is more likely to involve a glottal stop than the former. In the case of reduction, we would be dealing with reduced articulatory effort and possibly with the erosion of the boundary between the words, circumstances which do not particularly facilitate the insertion of a glottal stop.

However, the regularisation of THUH presumably involves some analogy to the realisation of 'the' before consonants. 'The' occurs more frequently before consonants, and children acquiring definite articles have been shown to use THUH both preconsonantally and prevocalically until reaching adult patterns by the age of 7 (Newton and Wells, 1999). In their discussion of the increase in prevocalic THUH in London English, Cheshire et al. (2011: 190) call prevocalic THUH a "latently present feature" of English. In the case of analogy, then, the insertion of a glottal stop maintains the parallelism with the more common preconsonantal tokens, while preserving the boundary between 'the' and the following word.

In cases of analogical change, it is well attested this targets low frequency forms first (Bybee 1985). In the regularisation of the English past-tense, for example, we get regularisation of low frequency forms (*weeped* > *wept*), but not of high frequency forms (*kept* > **keeped*). High frequency forms have robust stored representations and so resist analogical forces longer. Frequent "the+X" forms are likely to be stored and so resist the analogical change to THUH. For example, a form which frequently occurs with 'the', like '*the Olympics*', might resist analogical change more than a combination which is low frequency (and thus unlikely to be stored), such as '*the education*'.

A further reinforcing factor at play may be that we are dealing with a boundary phenomenon here – an effect which is driven by the relationship between the form of 'the', and the initial segment of the following form. Any change in behaviour at the boundary is going to be slower to affect frequent collocations which are stored (and for which the boundary is thus less boundary-like). This is the argument put forth by Hay and Sudbury (2005) in their discussion of linking /r/ in Early New Zealand English. They document a scenario in which the realisation of linking /r/ before vowel-initial words (as in, e.g. *neve_r_ again*) was waning; however, frequent collocations were not nearly so affected by the change as infrequent collocations. They interpret this as evidence for the storage of some frequent word pairs and for stored collocations being relatively robust against changes involving boundaries.

In cases without glottalisation, we may be dealing with an entirely different process – or possibly a different stage in the process. This would appear to be a reductive process - a breakdown of the boundary between the article and the following word, in which the 'the' form becomes more reduced and schwa-like, something closer to THUH than THEE. When viewed as a reductive process, we might expect more

erosion of the boundary in cases involving high frequency collocations such as, e.g., *The Olympics* (cf Phillips 1984). Indeed, for these forms we do see an increase in probability of reduction with increased likelihood of co-occurrence with THUH. But this increase is not, by itself, statistically reliable.

What does seem clear is that there are different things going on with the glottalised and non-glottalised forms. The glottalised forms are in the majority, and these seem to be the forms that are embodying the analogical levelling of the definite article.

4. Conclusion

This paper provides a simple analysis of a change in progress involving the production of 'the' before vowels in New Zealand English. While historically 'the' has been realised as /ði:/ before vowels and /ðə/ before consonants, this distribution appears to be changing in NZE, with young non-professional speakers paving the way. Forms with and without glottalisation are observed, and these behave quite differently with respect to frequency. This different sensitivity to lexical frequency factors suggests that two separate forces may be at play – one involving analogy and one involving reduction.

References

Baayen, R. Harald, Richard Piepenbrock and Leon Gulikers 1995. The CELEX Lexical Database (release 2), CD-ROM, Linguistic Data Consortium, University of Pennsylvania (Distributor), Philadelphia, PA.

Baayen, R. Harald, Douglas J. Davidson and Douglas M. Bates 2008. Mixed-effects modeling with crossed random effects for subjects and items. Journal of Memory and Language 59, 390–412.

Bybee, Joan 1985. Morphology: a study of the relation between meaning and form. Amsterdam, Philadelphia: John Benjamins.

Cheshire, Jenny, Paul Kerswill, Sue Fox, and Eivind Torgersen 2011. Contact, the feature pool and the speech community: The emergence of Multicultural London English. Journal of Sociolinguistics, 15(2), 151–196.

Fox Tree, Jean E. and Herbert H. Clark 1997. Pronouncing "the" as "thee" to signal problems in speaking. Cognition, 62, 151-167.

Gordon, Elizabeth, Margaret Maclagan and Jennifer Hay 2007. The ONZE Corpus. In Joan C. Beal, Karen P. Corrigan and Hermann L. Moisl (Ed.), Creating and Digitizing Language Corpora. Volume 2: Diachronic Databases. Basingstoke: Palgrave Macmillan, 82-104.

Hay, Jennifer and Andrea Sudbury 2005. How rhoticity became /r/-sandhi. Language 81(4): 799-823.

Newton, Caroline and Bill Wells 1999. The development of between-word processes in the connected speech of children aged between 3 and 7 years. In B. Maasen and E. Groenen (eds.), Pathologies of Speech and Language Advances in Clinical Phonetics and Linguistics. London: Whurr Publishers, 67-75.

Phillips, Betty 1984. Word frequency and the actuation of sound change. Language 60, 320–342.

Swan, Michael 1995. Practical English Usage, Oxford: Oxford University Press.