EXPLORING THE PHONETIC VARIATION OF 'YEAH' AND 'LIKE'

BEN GIBB-REID, PAUL FOULKES, VINCENT HUGHES
University of York

Abstract

Polyfunctional words (often categorized as discourse-pragmatic features) vary in form depending on their discourse functions, prosodic contexts (such as utterance position), and usage across social groups. This study applies dynamic formant analysis to describe the detailed phonetic variation of two polyfunctional words, *like* and *yeah*, in a corpus of West Yorkshire English (WYE). Data was drawn from 16 male speakers, and analysis was conducted on F1 and F2 in the sequences /laɪ/ and /jɛ/, across different pause contexts and functions. Dynamic measures were taken to track the trajectory of formant movements across the phonetic sequences, whereas previous research has typically only utilised midpoint measurements of vowel nuclei. Results show that the realisations of these words vary based on pause position and function. It was found that *yeah* functioning as a feedback particle had less formant movement across /j/ and /ɛ/ and lower F2 at the start of the word, whereas *yeah* in direct response to a question had a wider range of formant movement.

This study expands previous findings of *like* (Drager, 2011; Schleef & Turton, 2016) showing that there is fine-grained phonetic variation between functional and prosodic variables and showing that a single lexical item can have different phonetic realisations based on functions.

There is also an indication that *like* and *yeah* have distinct formant trajectories across speakers, and low within-speaker variability. Though tentative, these results suggest that word-specific variation merits further analysis for application to the task of forensic voice comparison and that speakers utilise phonetic resources to indicate pragmatic meaning.

1. *Introduction*

Much of the research into polyfunctional words falls under the banner of 'discourse-pragmatic variation'. Discourse-pragmatic features are described by Pichler (2016, p.1) as "conventionalised, polyfunctional linguistic items and constructions". They are complex to define but include words (*like, yeah, just*), phrases (*you know, I mean, innit*), interjections *(ah, eh, oh*), and longer strings (*and stuff like that*). The broad category is united by the function of the features either to "express speakers' stance; [or] to guide utterance interpretation and to structure discourse" (Pichler, 2013 p.4). Other researchers have used the terms 'discourse marker' (Maschler & Schiffrin, 2015), 'pragmatic marker' (Fraser, 1996) or 'pragmatic particle' (Östman, 1982) to refer to the same phenomenon. Pichler (2016, p. 3) defends the label 'discourse-pragmatic feature' as "conceptually more neutral", avoiding the various applications of and meanings associated with other labels (see Schourup, 1999). Despite a growing body of work on discourse-pragmatic features, there are few studies which utilise acoustic methods to explore fine-grained phonetic detail on words; Drager (2011, 2016) and Schleef and Turton (2016) are notable exceptions in that they analyse the variation of *like* across different functions with reference to segment durations and vowel monophthongisation (see below).

One of the uniting factors of many discourse-pragmatic features is that they occur frequently in naturally-occurring speech, which only increases when they are homonyms with lexical items. For example, in the British National Corpus (Spoken BNC2014), *like* has a spoken frequency of 1.38 occurrences per 100 words across various functions (Love et al., 2017). Such frequent features could prove particularly useful in the task of forensic voice comparison (FVC) because highly-frequent features are more likely to occur across samples of speech in forensic contexts. Nolan (1997, p.746) describes the aim of FVC research as building a “speaker space” defined by linguistic and phonetic properties, or “features”, of the voice. For a feature to be useful in FVC, it should be frequent in speech (Nolan, 1983, p.11) and must have low within-speaker variability and high between-speaker variability (Rose, 2002, p.10)–ideally irrespective of accent differences. That said, it is also universally accepted that every utterance is unique in its fine detail; hence a perfect ‘match’ between samples is not possible. A specific difficulty often encountered in this respect is that recordings are unlikely to share many of the same words and so segmental units are often utilised. Recognising such difficulties, research in FVC is devoted to exploring those features that yield the best discriminatory power, as well as understanding the factors that affect a given sound in different word positions or prosodic contexts.

The present study uses the analysis of formant trajectories to examine the fine-grained phonetic properties of *like* and *yeah* in terms of their function and utterance position. We will also consider their potential as FVC features. The rest of this paper outlines previous research into relevant speech phenomena in the field of FVC (section 1.1), and work on *like* and *yeah* (section 1.2). The methodology is outlined in section 2, the results are presented in section 3, and a discussion of the implications of *like* and *yeah* variation across speakers, function and context is offered in section 4.

1.1. Forensic Voice Comparison (FVC)

Work on the value of individual words for forensic casework has been limited to a small handful of studies. This began with studies into filled pauses (FPs), which are hesitation markers, usually either *uh* or *um*, and in many languages are pronounced as a relatively central vowel close to [ɘː] (although see Lo 2021). Künzel (1997, p.51) claims that speakers are “quite consistent” in their use of FPs and research on them has supported this claim (Tschäpe et al., 2005; Hughes, Wood & Foulkes, 2016).

Previous FVC studies have generally not considered the prosodic and interactional context of tokens alongside their speaker-specificity. This is despite other studies finding that context generally affects phonetic form. Guaïtella (1993) found that hesitation markers were longer and had a larger f_0 fall when surrounded by pauses than when embedded in continuous speech. In a study of 12 monolingual speakers of Dutch, Swerts (1998) found that turn-initial FPs have a higher f_0 than non-initial ones, and that FPs were more likely to occur in places where participants rated breaks in the discourse as ‘strong’ (i.e. at a “paragraph transition”, p.488). Overall, FVC research on FPs extracts segments from broadly similar contexts, but rarely considers the effects that changing the context may have on the phonetic form of FPs.

1.2. Like and yeah

Yeah and *like* were chosen for analysis because they are very frequent in natural speech. In Spoken BNC2014 (Love et al., 2017), *like* is the 14th most frequent word at 1.38 occurrences per 100 words. *Yeah* is the 7th most frequent word, at 2.28 tokens per 100 words. Both *like* and *yeah* are also mostly composed of vowels and glides, making it easier to analyse them in respect of their key acoustic properties. These properties make *like* and *yeah* good candidates for being useful diagnostic features in forensic analysis.

1.1.1. Like

Like, in citation form, comprises three phonemes: /l/, /aɪ/ and /k/. /l/, a lateral approximant in all English accents, has a formant structure with antiformants when pronounced as 'clear', i.e. alveolar, [l] (Johnson, 2011, p.196–198). In West Yorkshire English (WYE, the accent of the speakers in this study), word-initial /l/ is often relatively velarized or 'dark', [ɫ], compared to standard British English, manifesting as a higher F1 and lower F2 (Sproat & Fujimura, 1993; Carter & Local, 2007; Kirkham & Wormald, 2015). The diphthong /aɪ/ is also called the PRICE vowel (in the keyword system devised by Wells 1982). In many British English accents, PRICE is realised as a diphthong; F1 is therefore expected to fall and F2 is expected to rise during the course of the articulation. WYE is, however, known to have a narrower diphthong, sometimes even using a monophthong vowel close to [a:] (Wilhelm, 2018). /k/ was not analysed in this study, for reasons discussed in the Methodology (section 2).

Looking at *like*, D'Arcy (2007) distinguished between what she calls 'grammatical' (verbs, adverbs, conjunctions, nouns and suffixes) and 'vernacular' (approximate adverbs, discourse markers, discourse particles and quotative *be like*) functions. Vernacular *like* tokens were found to be systematic in their syntactic position and used more by younger than older speakers in Toronto. Phonetic variation has been linked to these functions, with Drager (2011) analysing adolescent speakers of New Zealand English (NZE) and measuring the direction and amount of diphthongisation by taking two F1 and F2 measurements at the nucleus and offglide of the vowel in *like*. Schleef and Turton (2016) compared the acoustic properties of *like* for two groups of adolescents from London and Edinburgh using categories and methodologies comparable to those of Drager (2011), but also measuring boundary strength defined by levels of breaks between words, phrases and utterances. The findings of Drager (2011) and Schleef and Turton (2016) are consistent with each other: The vowel in quotative *like* tends to be more monophthongal than other categories, is followed by a weak boundary, and has a higher f_0. In contrast, grammatical *like* and discourse marker *like* display more vowel movement and a lower f_0. Drager (2011, p.694) explains this as reflecting a "relationship between phonetic reduction and an individual speaker's probability of using *like* when producing a quotative". Schleef and Turton (2016, p.54) conversely conclude that the functional categories are "very fine grained" pragmatically and that the properties of *like* are more closely tied to prosodic boundaries than grammatical function. By contrast, Drager (2011) interprets the link between the phonetic realisation of *like* and its pragmatic functions as revealing that the different functions of *like* are stored cognitively as distinct items.

1.1.2. Yeah

Yeah also has several functions in conversation, illustrated by examples (1)–(3):

(1) ***yeah*** so I did that

(2) A: so I said to them
B: ***yeah***
A: that is not acceptable

(3) A: you've played tennis before, right?
B: ***yeah*** once or twice

Yeah can indicate that someone is going to speak, as in (1). This is referred to as an incipient 'acknowledgment token' –described as a "marsheling [sic] of resources to assume speakership" (Drummond & Hopper, 1993a, p.206) and a shift "from recipiency to speakership" (Jefferson, 1984, p.200). Recipient *yeah* (2), however, does not indicate that someone is about to take the floor but that they are listening. Drummond and Hopper (1993b) highlight this as an example of backchanneling. *Yeah* can also enact agreement (3), the most transparent meaning, defined by Fuller (2003) as an "affirmative response to a question" (p.29). Although recipient *yeah* can indicate agreement, it doesn't always, necessitating the separate categories of recipient (2) and agreement *yeah* (3) (Fuller, 2003).

The citation form of *yeah* contains two phonemes: /j/ and /ɛː/ (OED Online, 2022). The palatal approximant [j] has similar formant structure to [i]: low F1 and high F2 (Raphael, Borden & Harris, 2007, p.116–118; Ladefoged, 2001, p.213–214). The /ɛː/ vowel, also referred to as SQUARE (Wells, 1982), has a lower F2 and higher F1 than /j/. SQUARE in WYE remains front [ɛː] (Wells, 1982, p.364) but Wilhelm (2018) found that younger speakers use diphthongs with a closer offglide. Therefore, across *yeah*, F1 should increase (and potentially decrease again), F2 should decrease.

Previous research on the functional and phonetic variation of *yeah* appears largely limited to turn-initial position. Some previous studies provide phonetic descriptions of *yeah*, but few refer to variations in vowel quality, focusing instead on prosody and duration. Based on a corpus of 1,034 *yeah* tokens across 14 recordings, Truong and Heylen (2010) found that when turn-initial *yeah* was used as an acknowledgement it was longer in duration than when it had a backchannel function. f_0, however, did not have a significant effect across contrasting turn-initial *yeah* functions. In a study of Colorado English with 93 tokens, Grivičić and Nilep (2004) also found that when a speaker produced *yeah* and did not follow this with speech, the token tended to have creaky voice.

This study will analyse all turn positions of *yeah* but code their functions according to these previous findings. As far as we are aware, no systematic work has considered the potential value of such affirmative 'acknowledgement tokens' for FVC. The only similar analysis in the FVC literature is found in Rose's (2013) case report involving *yes*. A likelihood ratio (LR) was calculated comparing the similarity and typicality of the formant dynamics of F1–F3 from /jɛ/ across the suspect and criminal samples in a voice comparison case. The result of the LR suggested that the formant pattern for /jɛ/ works well as a speaker discriminant.

1.3. The present study

The present study aims to explain the phonetic variation of *like* and *yeah* related to their pragmatic function and their prosodic context. In line with other work on discourse-pragmatic variation, there is an analysis of the overall frequency of *like* and *yeah* in the data, but also consideration of within- and between-speaker variation in pronunciation of the two words. The study builds on previous research by considering the functions and pause boundaries affecting *like* pronunciation, but goes further by taking multiple formant readings of the word tokens, allowing for a more detailed analysis of the quality of /l/ and the degree of vowel diphthongisation.

Two key questions drive this research:

(1) How do *like* and *yeah* vary functionally and phonetically?

(2) What is the individual variation of *like* and *yeah*?

2. Methodology

The data for the study is drawn from Task 3 of the West Yorkshire Regional Database (WYRED) (Gold, Ross & Earnshaw, 2018). The participants used for this study were sixteen males from Kirklees, aged 18–30. The task involved paired, spontaneous conversation with someone they had not met before. Praat (Boersma & Weenink, 2016, ver 6.0.17) was used to listen to the sound files, code tokens and take formant measurements. Previous research on *like* has relied on Euclidean distance measures of F1–F2, whereas the present study adopts formant trajectories as laid out in section 2. Also, though Drager (2011) and Schleef and Turton (2016) discuss /l/-to-vowel duration, there isn't a consideration of variation in /l/ quality, whereas in the present study, the consideration of dynamic formant readings across the whole token permits analysis of both /l/ and vowel quality and which, in turn, helps us understand more about the fine-grained phonetic differences between functions.

A Praat script (Lennes, 2002) was used to extract a sound file for every token from the segmented TextGrid file. A second script (Hughes, 2013) took nine time-normalised F1 and F2 measurements at +10% steps from every sound file. For each token, eighteen measurements were recorded. The formant settings were, by default, set to track 5 formants up to a maximum frequency of 5000Hz, although in some cases settings were manually changed to track 5.5 formants at a maximum of 5500Hz to ensure accuracy. R Studio (RStudio Team, 2016) was used to process, plot, and run statistical tests on the data. F1 and F2 formant dynamics were used to capture the phonetic quality of *like* and *yeah*. Dynamic properties of formants have previously been used in FVC analysis to measure FPs (Hughes et al., 2016) and /jɛ/ from *yes* (Rose, 2013). No /k/ segments were extracted from *like* tokens to ensure comparability between tokens (as a pilot study showed, the variation of /k/ realisation was across multiple places of articulation, and often /k/ was elided or glottalised). Vowels, nasals, and approximants can coarticulate with each other, making formant readings variable (Fabricius, 2002, p.219; Deterding, 1997, p.49), and therefore care was taken to avoid extracting tokens where formants would be influenced by coarticulation effects.

Figure 1 shows a token which was excluded. Here, *like* is preceded by /l/ from *school*, and there is no clear place to distinguish the preceding /l/ from the initial /l/ in *like*. Additionally, this token of *like* has an elided /k/ followed immediately by /m/, which affects the vowel quality.

Word counts were taken from the WYRED corpus and are described here in line with Pichler's recommendations (2010, p.594): filled pauses and minimal responses were included, as were false starts (listed as e.g. *lik-*). Clicks, uncertain transcriptions, laughter, breath, throat and sniff noises were not included in the counts.

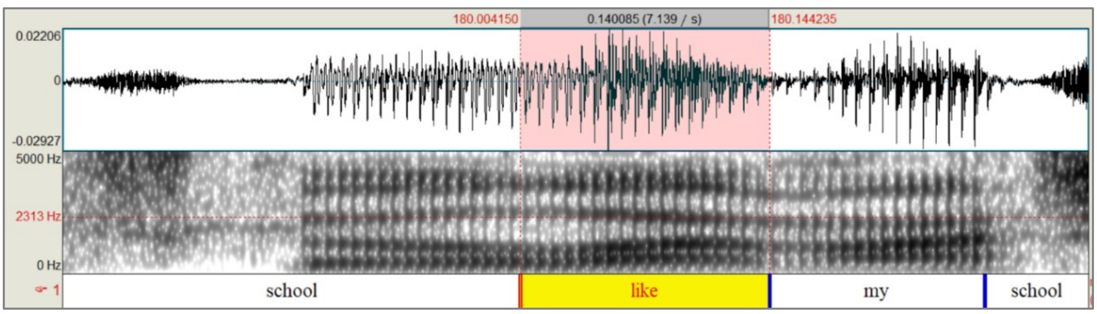

Figure 1: Spectrogram and waveform of a token of *like* omitted from analysis.

2.1. Coding

Pause positions were included in the analysis as previous research has found that they affect the phonetics of FPs. The acoustic properties of *like* are impacted by what Schleef and Turton (2016) call 'boundary strength': a higher strength boundary would include a pause, while a word followed immediately by speech would constitute a lower strength boundary. Any silent period lasting more than 100ms was defined as a pause, and therefore a higher-strength boundary (the 100ms threshold was adopted from Künzel, 1997, p.55). FPs such as *um* and *uh* were included in this pause measurement, as were breaths or other sounds such as clicks because they were seen to constitute a higher-strength boundary. Short pauses were defined as between 100ms and 1000ms and long pauses as >1000ms.

Four conditions of pause position are used in this study, scaled from low-to-high boundary strength:

1. No: when a token is not surrounded by any pauses i.e. it is surrounded by speech.
2. Pre: when a token comes before a pause.
3. Post: when a token follows a pause.
4. Both: when a token is surrounded by two pauses.

Functional categories of *like* are summarised below in Table 1. The broad groups of 'grammatical', 'quotative' and 'discourse' follow categories utilised by previous research (Drager, 2011; Schleef & Turton, 2016) to make the results more comparable. However, approximate adverb *like* contributes to the truth condition of the utterance, and therefore can align with 'grammatical' functions. Functions of *yeah* are listed in Table 2.

Category	***Grammatical function***	***Example from WYRED***
	Conjunction	*It's **like** you said, some people perceive you as a bit thicker*
Grammatical	Verb	*People **like** to move around a lot*
	Adverb	*Kopparberg and it's just **like** glorified fruit juice*
Quotative	Quotative	*You listen to them and you're **like**, 'what the hell?'*
	Approximate adverb	*He dropped me off at **like** seven o'clock*
Discourse	Discourse marker	***Like**, yours is actually a little sort of like*
	Discourse particle	*We also have, **like**, loads of different places*

Table 1: Summary of *like* functions used in this study. The grammatical functions are grouped into three broader categories, aligning with Drager (2011) and Schleef and Turton (2016).

Category	***Example from WYRED***
Feedback	*009: If I talk to someone outside of Yorkshire* *003:* ***Yeah*** *009: I'll enunciate better*
Response	*029: Have you always lived in xxx?* *026:****Yeah****, all my life*
Discourse	*It's so depressing but um* ***yeah****, so I did that*
Agreement	*015: God's own country, West Yorkshire* *019:* ***Yeah*** *true*

Table 2: Summary of *yeah* functions used in this study.

Spectrograms showing the coding are shown in Figure 2 for *like* and Figure 3 for *yeah*. Pauses are defined as short (S) or long (L).

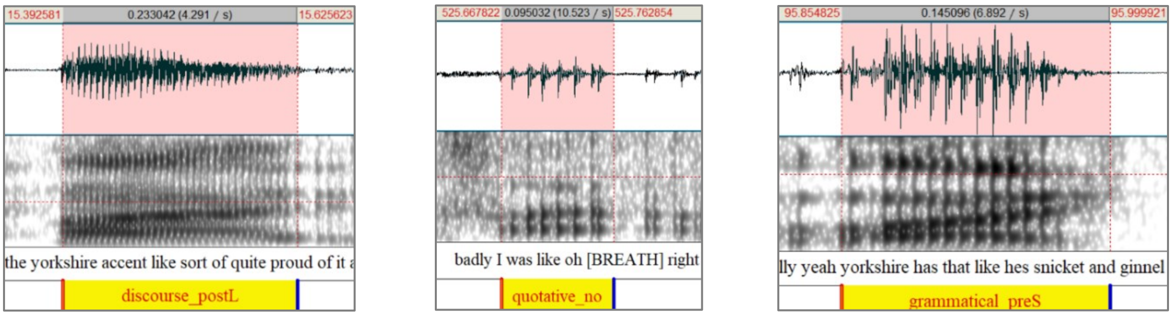

Figure 2: Spectrograms of *like* tokens with the following functions: discourse (left), quotative (mid) and grammatical (right).

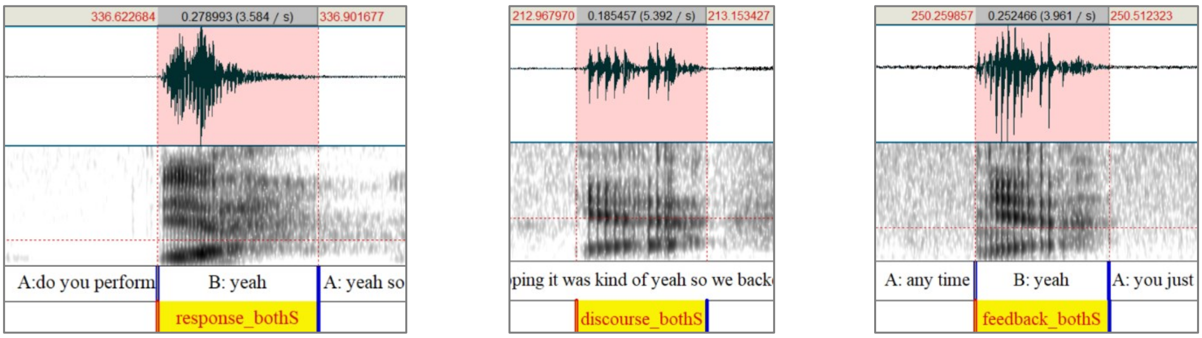

Figure 3: Spectrograms of *yeah* tokens with different functions: response (left), discourse (mid) and feedback (right).

2.2. Statistical Analysis

Generalised Additive Mixed Models (GAMMs) were fit to the data in RStudio (2016). GAMMs are similar to other mixed models (such as more common linear models) in that they allow for random effect structures to be included, which attempt to reduce the chance of significant effects which are driven by, for example, an individual speaker or word within a dataset. However, GAMMs are non-linear models and so are much better suited to the dynamic properties of speech as they allow for comparisons across both formant height (overall average) and trajectory shape. The analysis is informed by Sóskuthy (2017, 2021).

GAMMs were built using the built using the function `bam()` from the package `mgcv()` (Wood, 2006). Separate models for *like* and *yeah* F1 and F2 trajectories were fitted using *function* and *pause* as fixed effects. *Speaker* and *token.no* (used to align each measurement to its corresponding token) were added as a random smoothing term only. All other variables were added as fixed parametric and smoothing terms. Parametric coefficients relate to trajectory height, while smoothing coefficients are indicative of trajectory shape. A fixed smoothing coefficient is also present in every model called *measurement.no*. This coefficient contains the numbers 1–9, defining the individual formant measurement's place along the trajectory (i.e. 1 = the first measurement point, and so on). Without this term the model comparisons would be interpreting individual formant measurements, not trajectories.

Model comparisons were run to assess the impact of each variable on the formant trajectories of *yeah* and *like* as recommended by Sóskuthy (2017, 2021). This was done using the package `mgcv()` (van Rij et al., 2016) along with the script `gamm_hacks.r` (Sóskuthy, 2017). The most complex model with all possible variables is created first (M1), and then a model with one less term (M2) is made. For example, M1 would contain *like* F1 trajectories with *speaker* as a random effect, and *long pause* information, *short pause* information and *function* as both parametric and smoothing effects. M2 contains the same input as M1 but without the *function* smoothing coefficient. Therefore, a comparison of M1 and M2 shows the impact of *function* on *like* F1 trajectory shape. Next, a third model (M3) would have one less term than M2 and so on until all coefficients have been tested. In the tables presented in the Appendix (i) the first comparison is one which compares M1 with M2. As suggested by Sóskuthy (2021, p.8–9), model comparisons were conducted using the AIC difference (AICd).

3. *Results*

In this section the results are presented; firstly with reference to overall token frequency of use, then to dynamic formant measurements. Model comparisons used for statistical testing are listed in the Appendix Table 3–Table 10.

3.1. Overall Frequency

Figure 4 shows the distribution of *like* and *yeah* tokens per speaker. With a total of 666 tokens of *like*, and 869 tokens of *yeah* the 16 speakers selected for this study on average used them 2.17 and 2.83 times per 100 words respectively. There was also considerable variation in frequency across individuals from 0.67 to 7.72 for *yeah* and 0.69 to 3.81 for *like*.

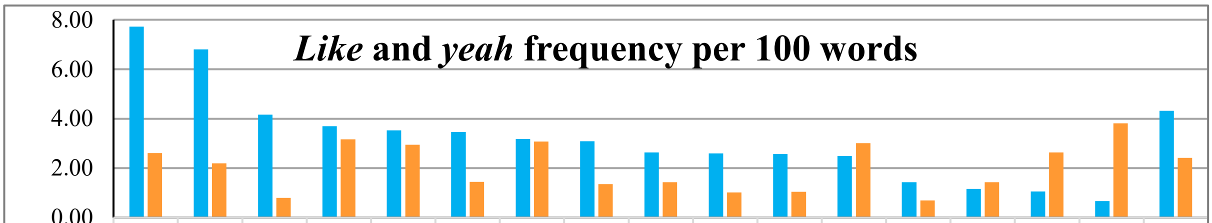

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
yeah	7.72	6.80	4.17	3.69	3.53	3.46	3.19	3.09	2.64	2.59	2.58	2.50	1.43	1.16	1.06	0.67	4.31
like	2.60	2.20	0.79	3.17	2.95	1.45	3.08	1.35	1.44	1.02	1.04	3.02	0.69	1.43	2.64	3.81	2.41

Figure 4: Average frequency of occurrence per 100 words across speakers. '#' refers to the average frequency across all 16 speakers.

On average, non-grammatical *like* (quotative and discourse, 64%) was more frequent than grammatical *like* (36%). *Like* is very unlikely to occur before or after a pause of more than 100ms (8%) or be surrounded by pauses of any length (4%). *Yeah* was most frequently used as a feedback token (54%), and compared with *like* was more often surrounded by long (48%) or short pauses (63%).

3.2. Like

3.2.1 Overall trajectories

423 *like* tokens were utilised for acoustic analysis, a lower number than the totals shown in Figure 4 as some tokens had coarticulation effects (discussed in section 2) and others did not have clear enough formants to allow readings to be taken.

Figure 5 shows the mean F1 and F2 across all tokens of *like*. *Like* generally has a rising F2 (as expected), and a rising-falling F1. A higher F1 and a lower F2 (here between 1000–1200 Hz) is found consistent with a relatively dark [ɫ] (Carter & Local, 2007, p.196).

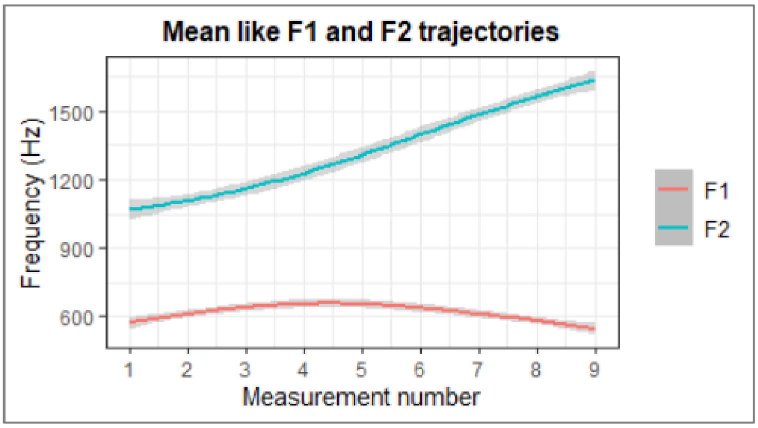

Figure 5: Mean F1 & F2 trajectories for *like*. SD ranges are indicated by line width.

3.2.2 Statistical modelling

Although both formants capture the same articulatory movement, it is necessary for them to be treated separately for statistical analyses. The discussion will draw the separate conclusions drawn about F1 and F2 together to interpret the effect of pause context and function on *like* and *yeah*. All variables for *like* were found to be significant for F1 and F2; therefore, the most complex model was utilised for the summary comparisons, and these are listed in Table 4 and Table 6 in the Appendix.

Short and *long pause contexts* and *functions* significantly improve the model fit of *like* F1 and F2 trajectories. Therefore, the most complex model is used to infer the phonetic differences between variants.

Firstly, for *pause contexts*, Figure 6 shows the F1 and F2 trajectories across short pause contexts (i.e. >100ms) for the four categories: Tokens surrounded by pauses on both sides (bo), tokens surrounded by speech (no), post-pausal tokens (po) and pre-pausal tokens (pr). Specifically, we see that *like* occurring before a pause (pr, green line) has the most movement in its trajectory. Pre-pausal *like* has a fairly flat F1 trajectory at first, followed by a sharp fall. This was significantly different from other pause positions.

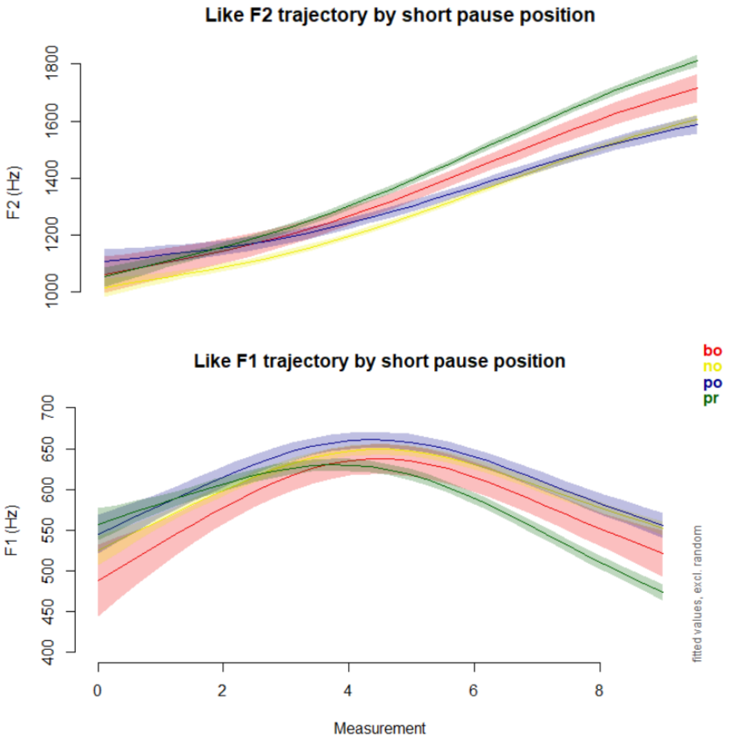

Figure 6: Mean F1 and F2 trajectories for *like* by short pause contexts. SD ranges are indicated by line width.

For F2, when tokens were surrounded by speech (no, yellow line), *like* had lower values, particularly over the initial part of the trajectory–corresponding to a backer articulation. Post-pausal *like* (po, blue line) also had a lower F2 overall but this was from the middle of the trajectory. These two effects were significant when testing overall trajectory averages. When *like* occurred between two pauses (bo, red line), it was more likely to show a linear increase; a movement from lower to higher F2 which was significant when testing trajectory shape. Pre-pausal *like,* however (pr, green line), had more F2 movement towards the latter part of the trajectory, potentially showing a fronter vowel. Again this was a significant effect in the model for trajectory shape.

There were only 39 tokens of *like* adjacent to a long pause in the dataset. These results showed a similar pattern of effects to short pause contexts as discussed above and exemplified in Figure 6. Although *long pause* information improves the model fit, *long pause* contexts did not significantly differ from one another. For these reasons they have not been plotted.

The F1 and F2 trajectories of *like* by function are shown in Figure 7, with the three categories: discourse *like* (d), grammatical *like* (g) and quotative *like* (q). Quotative *like* (blue line) has a higher F1 and the start of the trajectory moving to a lower value at the end compared with other functions. However, none of the F1 trajectory shapes were found to be significantly different across each function, and the same was true of trajectory averages.

F2 trajectories, however, yield a different result. Quotative *like* has a shallower rise in the latter half of the F2 trajectory than discourse or grammatical *like*. This was significant for trajectory shape, indicating less fronting in the diphthong and a clearer [l]. However, discourse and grammatical *like* are close to each other in the shape and height of F2, and no significant differences were found between them.

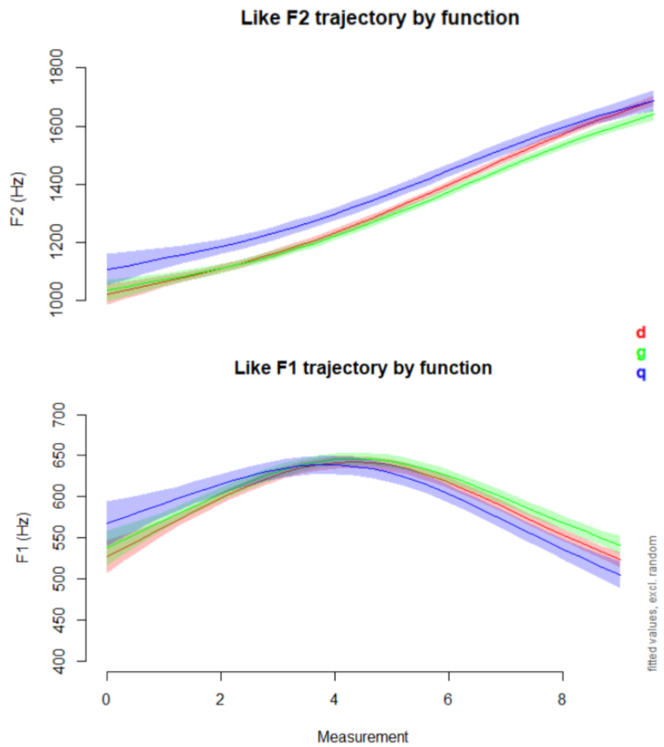

Figure 7: Mean F1 and F2 trajectories for *like* by function. SD ranges are indicated by line width.

3.2.3 Variability across speakers

Speaker significantly improves the models for both F1 ($X^2(3)=194.1$, $p<0.001$) and F2 ($X^2(3)=181.4$, $p<0.001$) trajectories of *like*, indicating a separation between speakers as shown in Figure 8. Visual inspection of the mean trajectories for individuals shows clear between-speaker variation in terms of both overall frequency and trajectory shape.

Standard deviation (SD) is linked to within-speaker variation. The SD of F1 is between around 48Hz and 65Hz (with one outlier, speaker 032 with 97Hz). F2 has a wider range of SD between around 100Hz and 160Hz. Generally, *like* F1 has a lower SD between measurement 3 and 6, while F2 has a lower SD between measurements 1 and 4. This means that *like* F2 has greater within-speaker variation.

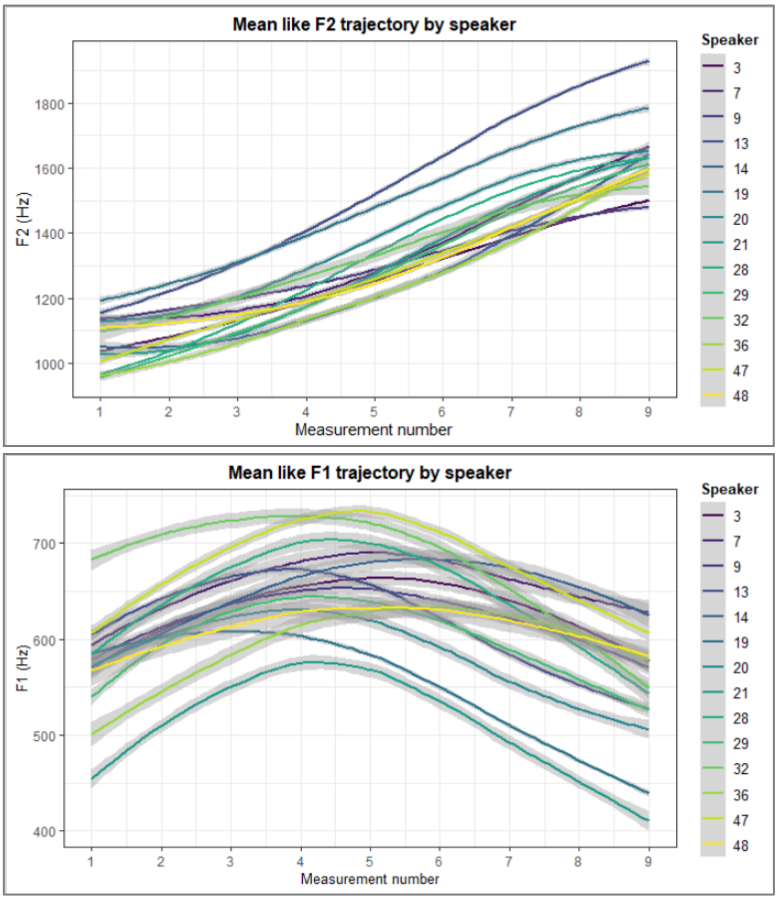

Figure 8: Mean F1 and F2 trajectories for *like* by individual speaker. Standard deviation represented by line thickness.

3.3 *Yeah*

3.3.1 Overall trajectories

The phonetic variation of *yeah* is discussed below. There were 453 *yeah* tokens utilised for acoustic analysis. Figure 9 shows the mean F1 and F2 trajectories with smoothing across all *yeah* tokens, showing a clear rise in F1 and fall in F2 across /j/ and /ɛː/.

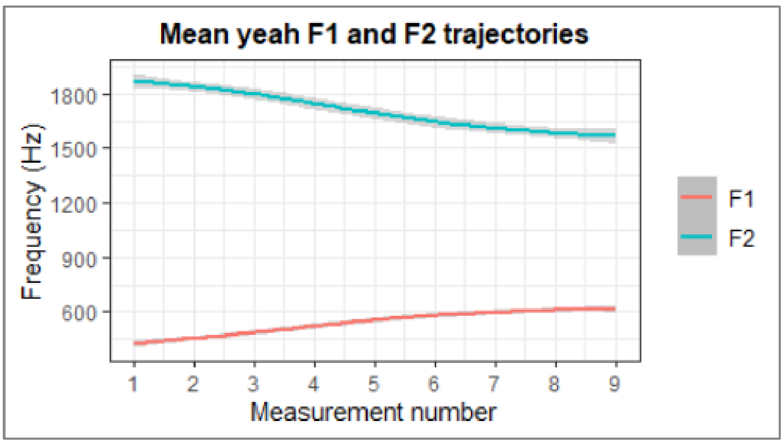

Figure 9: Mean F1 & F2 trajectories for *yeah*. SD ranges are indicated by line width.

3.3.2 Statistical modelling

As was the case for *like*, *short* and *long pause contexts* and *functions* again significantly improve the model fit of *yeah* F1 and F2 trajectories. Therefore, results will be discussed with reference to the most complex *yeah* F1 and F2 model summaries listed in Table 8 and Table 10 in the Appendix.

All pause contexts of *yeah* had a similar overall F1 average across the trajectory. This is reflected in the fact there were no significant differences found among comparisons for trajectory height. However, post-pausal *yeah* (po, blue line) has more F1 movement across the word, shown by its lower starting position compared with other contexts. This was significantly different from other contexts. For F2, pre-pausal *yeah* (pr, green line) has overall lower values across the trajectory, which was significant for trajectory height. In terms of shape, when *yeah* was surrounded by speech (no, red line) or post-pausal (po, blue line), it had a distinct high-to-low F2 shape with a sharp downward movement–again this was significantly different from other pause contexts.

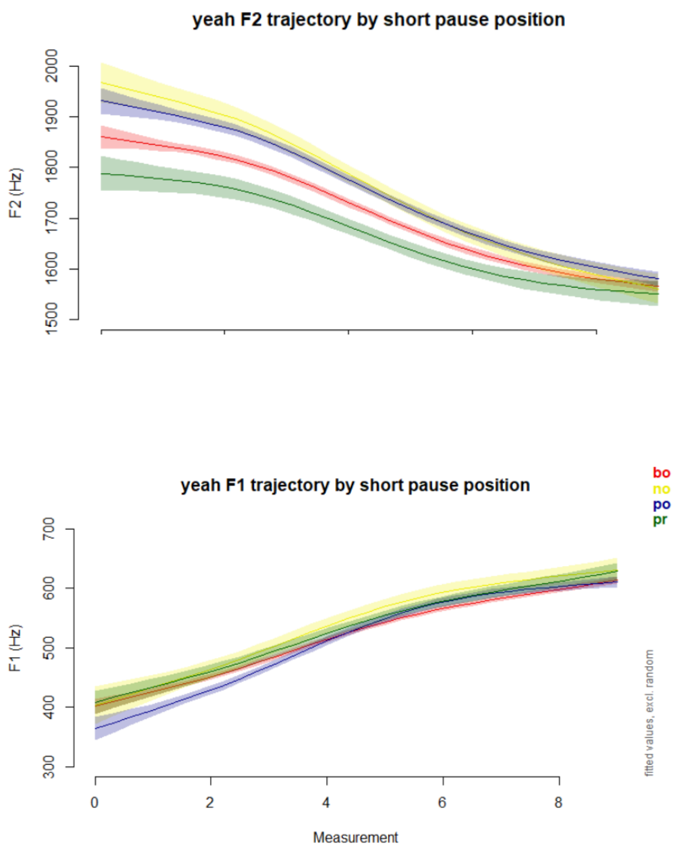

Figure 10: Mean F1 and F2 trajectories for *yeah* by short pause contexts. SD ranges are indicated by line width.

These effects correspond to post-pausal *yeah* having more F1 and F2 movement between /j/ and /ɛ/. *Yeah* in continuous speech has a similar amount of front-to-back movement to post-pausal *yeah*. Pre-pausal *yeah*, however, has a lower F2 indicating a backer articulation than other contexts. For *long pause contexts*, there were no significant differences across categories for F1 or F2 averages or shapes and therefore this is not plotted.

For *yeah functions,* Figure 11 shows the formant trajectories of the four categories: agreement (a), discourse (d), feedback (f), and response (r). All functions of *yeah* have a comparable average F1 height, shown by the narrow band of trajectories and reflected by the lack of significant differences between them. However, it seems that feedback *yeah* (f, blue line) has a more linear trajectory compared to the other functions, with a higher F1 at the start leading to a lower value at the end of the trajectory. This was significant for shape in the model comparisons.

Feedback *yeah* is also lower overall for F2 than other functions, particularly in the initial part of the trajectory. This correlates to a more back /j/ and therefore potentially more reduction. Again, this was significant. In terms of shape, response *yeah* is significantly different from other functions – potentially indicating a wider span of movement across the trajectory.

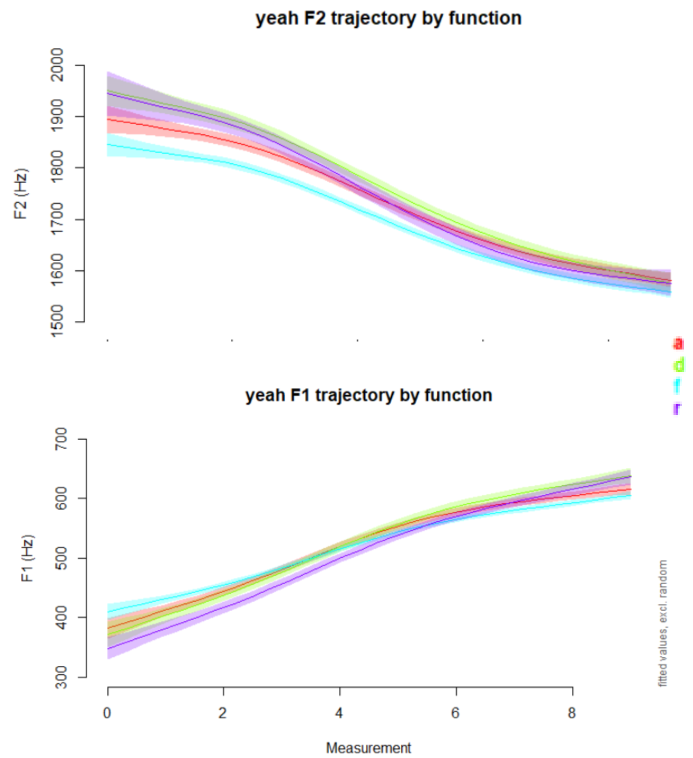

Figure 11: Mean F1 and F2 trajectories for *yeah* by function categories. SD ranges are indicated by line width.

Overall, this indicates that feedback *yeah* has less F1 and F2 movement across /j/ and /ɛ/ and a lower F2 over /j/. Response *yeah*, (r, blue line) however, has a wider range of movement between /j/ and /ɛ/–perhaps more sharply distinguishing one segment from the other.

3.3.3 Variability across speakers

Similarly to *like*, *speaker* is a significant variable in model comparisons for *yeah* F1 and F2 trajectory shape. Figure 12 reflects the mean F1 and F2 trajectories for each speaker. Again, there is considerable variation between-speakers. The standard deviation (SD represented by line thickness) shows that each speaker is relatively consistent. SDs for F1 and F2 are lower for *yeah* than for *like*, with F1 ranging between 30Hz and 69Hz, and F2 between 52Hz and 111Hz. Some speakers are showing a slight drop or plateau in F1 trajectories around measurement 8 or 9 which could indicate Wilhelm's (2018) finding of closer PRICE diphthong off-glides for younger speakers. There are also overall differences between speakers who also tend to have more formant movement and those who don't – suggesting between-speaker distinctions in vowel monophthongisation.

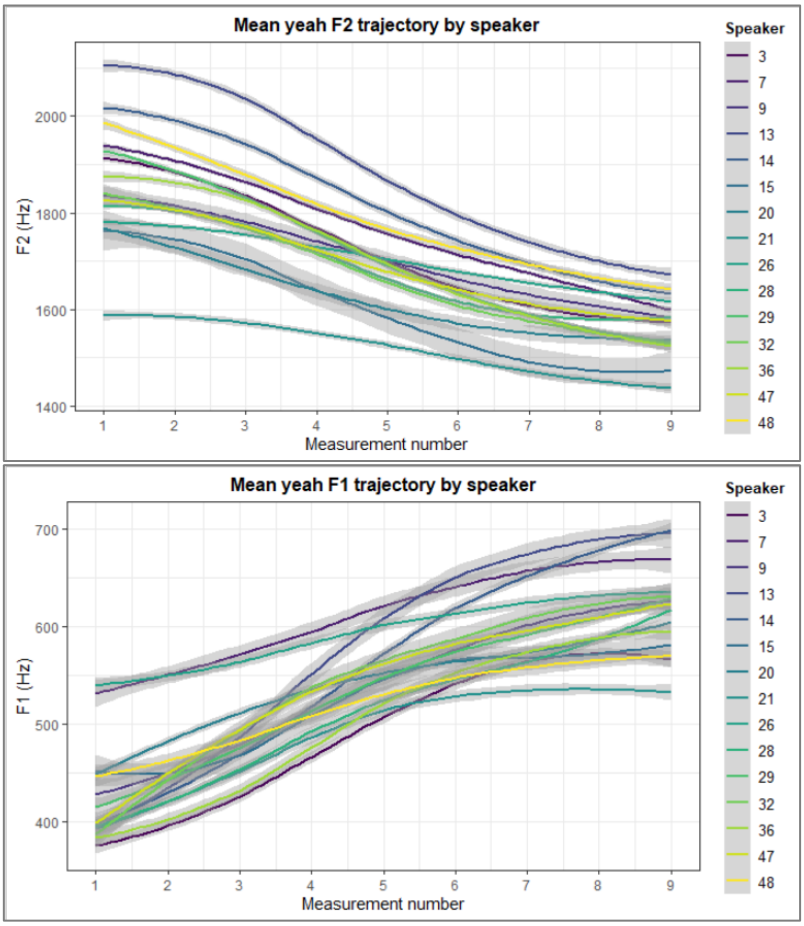

Figure 12; Mean F1 and F2 trajectories for *yeah* by individual speaker. SD is indicated by line width.

3.4 Summary

The various effects of function and prosodic context on *like* and *yeah* have been presented with reference to model comparisons and model summaries indicating differences across variants. The salient points found in these results are summarised below and are analysed further in the discussion.

- Pre-pausal *like* has a darker /l/ and a more diphthongal and closer vowel; *like* surrounded by pauses has a less diphthongal vowel and a clearer /l/; and post-pausal *like* and no-pause *like* have the least diphthongal vowel.
- Quotative *like* has less fronting of the diphthong than discourse or grammatical *like*.
- Post-pausal *yeah* has closer and front /j/ and both post-pausal and no-pause *yeah* have a lot of front-to-back movement across /j/ and /ɛ/. Pre-pausal *yeah*, however, has a lower F2 for /j/ and therefore less front-to-back movement across /j/ and /ɛ/.
- *Yeah* used primarily as a feedback function has less movement across /j/ and /ɛ/ and a lower F2 for /j/. Response *yeah* on the other hand has a wider range of formant movement across /j/ and /ɛ/. Discourse and agreement *yeah* sit somewhere in the middle.

4. *Discussion*

In this section, the results from section 3 are discussed with reference to previous findings and research on *like* and *yeah* in section 1. First, there is a discussion of phonetic variation according to function and prosodic context with reference to previous research and then an evaluation of the potential applications to FVC.

4.2 Like

The phonetic variation of *like* in the present data was found to be influenced less by its function and more by its prosodic context. *Quotative like* has a significantly different pronunciation than other functions of *like*–with less F2 movement over the latter part of the trajectory than *grammatical* or *discourse like*. This signals that there is less fronting over the course of the /aɪ/ vowel. As confirmed through the present findings, Drager (2011) and Schleef and Turton (2016) found that *quotative like* had a less diphthongal vowel than other functions of *like*. However, their methodology relied on the Euclidean distance between F2-F1 at the vowel nucleus and off-glide. The present study's analysis of formant trajectories allows for a more detailed analysis of the overall articulatory movement across the word. The differences in formant movement across /aɪ/ can also be influenced by the quality of the preceding sound, /l/. In the present data, higher F2 at the start of the *quotative like* trajectory can be seen to indicate a clearer /l/ - indicating that /l/ darkness is another phonetic correlate of *like* function.

Schleef and Turton (2016, p. 60) claim that *quotative like* is most frequently followed by a boundary with lower strength. This in turn explains its relatively monophthongal quality: it is reduced because it is generally located within connected speech. In the data used in this study, however, *quotative like* frequently appeared before a pause of at least 100ms (63% of tokens). Therefore, it seems that higher F2 across the *quotative like* vowel is not simply the result of its position in connected speech. It is possible that *quotative like* is undergoing phonetic reduction. Quotatives are generally short words (others include *said, be,* and *go*), and they have a clearly-defined function of reporting speech or attitudes. *Quotative like* may in some way be treated as a function word in speech, and therefore would be expected to show some phonetic reduction (Bybee, 2007).

Short pause contexts had a stronger effect on the trajectories of *like* than specific functions and long pause contexts, while not generating statistically significant effects, showed similar patterns. *Like* before a pause had a higher F1 at the start of the trajectory and more F1 and F2 movement from the middle to the end of the trajectory, ending with the lowest F1 and highest F2 values. This indicates a more diphthongal vowel and a darker /l/ for pre-pausal contexts. Tokens of *like* after a pause and surrounded by speech (i.e., not flanked by pauses) are similar, with a high F1 and low F2 trajectory and less movement across the latter part of the trajectories. This indicates a more open and back vowel with less movement across the nucleus and offglide. *Like* surrounded by pauses also has less formant movement across the vowel than pre-pausal *like* and has the lowest F1 values at the start of the trajectory. This shows that *like* surrounded by pauses has a clearer /l/ and a less diphthongal vowel than pre-pausal *like* (but not as monophthongal as post pausal *like* or *like* surrounded by speech).

Pre-pausal and *like* surrounded by pauses roughly correspond to 'high boundary strength' in Schleef and Turton's (2016) study, whereas *like* in post-pausal and no-pause contexts corresponds to a lower boundary strength. With this in mind, these results contrast with the variation seen in Edinburgh and London English. In our data, a 'strong boundary' after *like* generates not lower vowel movement, but higher vowel movement compared to other contexts.

While perhaps surprising, this finding is likely caused by differences in methodology. Only the 'following' context of *like* was considered in Schleef and Turton's (2016) analysis of *like*, so their 'high boundary strength' tokens may or may not come after a pause, whereas the present study considers both edges of a token. The methods behind defining boundary strength are also based on intonation phrase judgements, whereas the present study only considers pauses in speech. Additionally, our study combined 'discourse marker' and 'discourse particle' into one category, whereas Drager (2011) and Schleef and Turton (2016) treated them separately. Though there are distinctions between these categories, the present study had a smaller dataset, and fewer categories were more helpful for statistical modelling. It would be helpful for future research to collect more data in order to divide the categories in more detail.

One problem with the division of *like* functions into discourse versus grammatical groupings is that it makes it difficult to analyse the verb, approximant adverb, conjunction, discourse marker or discourse particle functions which sit within them. The functional categories were chosen to align with previous research (D'Arcy, 2007; Drager, 2011; Schleef & Turton, 2016). However, one presiding definition of 'discourse markers' and 'discourse-pragmatic variables' is that they are syntactically optional and can be removed from the sentence without affecting its truth propositional content (Schourup, 1999; Pichler, 2013). For instance, separating approximate adverbs (e.g. "I did that in *like* two days", taken from Drager, 2011 p.696) from the discourse function category may show a different effect in terms of formant trajectories.

Drager (2011) finds that phonetic realisations are linked to *like* functions and argues in favour of the view that there are distinct lemmas stored cognitively. By contrast, Schleef and Turton (2016) claim that neither their nor Drager's results supports the storage of multiple *like* lexical items. In the present study, *like* function is again closely related to fine-grained acoustic detail. However, along with Schleef and Turton (2016) we would not go so far as to claim that our data shows *like* functions as being distinct lemma categories: each functional category is not realised in a sufficiently unique way. For example, grammatical and discourse tokens of *like* are not significantly separable for F1 or F2, which is what one would expect if they were separate lemma categories. The same can be said of *yeah* functions which are discussed further below.

4.3 Yeah

The F1 and F2 trajectories of *yeah* were found to be affected by both function and short pause contexts. When *yeah* is used as a feedback token (as a non-turn-taking, often minimal backchannel), it has a lower F2 and a higher F1 over /j/, and less movement between the approximant and vowel. This 'recipient' *yeah* (Drummond & Hopper, 1993b) is not seen to indicate someone is going to take the turn, supported by the fact that these tokens tended to occur surrounded by short pauses (79% of cases). It is possible that the lack of formant movement across feedback *yeah* tokens emphasises a listening stance. Contrastingly, *response yeah* (tokens directly responding to a question) has higher F2 for /j/, which lowers across a wide range into /ɛ/. *Response yeah* may, in its wide F2 movement, align with 'speaker-incipient' *yeah*, which Truong and Heylen (2010) found to be longer in duration than *feedback yeah*.

If *yeah* is after a pause or surrounded by speech, then it has a higher F2 across /j/, which then lowers more sharply into /ɛ/. Post-pausal *yeah* also has a higher F2 and lower F1, indicting a more tense /j/. Pre-pausal *yeah,* on the other hand, has lower F2 for /j/ and less F2 movement across /j/ and /ɛ/. When *yeah* is post-pausal or not flanked by pauses it is followed by further speech (i.e., the turn continues). Thus, like *response yeah*, we can infer that the speaker is signalling an ongoing turn. The wider formant movement, in contrast to the limited movement of a turn-ceding *feedback yeah*, could provide a phonetic cue to signal turn holding. More work which analyses the turn position of tokens alongside duration may help to further differentiate functions of *yeah*.

4.4 Implications for forensic phonetics

The data in the present study reveal considerable between-speaker variation in the phonetic realisation of *yeah* and *like*, indicating potential as a speaker discriminant. *Yeah* in particular had lower within-speaker SD than *like* across F1 and F2 trajectories indicating that it may make a better discriminant.

Both *like* and *yeah* are frequent in naturally-occurring speech, which makes them, in principle, suitable for FVC (Nolan, 1983, p.11). As Figure 4 shows, *like* occurred 2.41 times per 100 words and *yeah* 4.31. This is much higher than in Spoken BNC2014 (Love et al., 2017) where *like* and *yeah* had a frequency of 1.39 and 2.28 respectively in the spoken data. The discrepancy may be because the data used in the present study (WYRED) was recorded in 2016-18 whereas Spoken BNC2014 was recorded earlier, and we know that *like* usage is increasing in time (D'Arcy, 2007). The higher use of *like* and *yeah* in WYRED could also be dialect-related, as Spoken BNC2014 involves various dialects across the UK whereas the present study focuses on West Yorkshire English. Comparing the two, it is likely that *yeah* occurs more than *like* because of its broader multifunctionality.

Extracting more phonetic features or variables from *yeah/like* may be useful to further test these results. f_0 of FPs has been shown to interact with turn-position (Swerts, 1998) and improve the speaker discriminatory power of filled pauses (Tschäpe et al., 2005; Braun & Rosin, 2015), while duration also makes some positive contribution to speaker-discriminant tests (Hughes et al., 2016). Voice quality (Grivičić & Nilep, 2004) has also been linked with *yeah* functions, as has duration (Truong & Heylen, 2010).

GAMMs are not very suitable for forensic phonetic research. It is complicated to compute them in numerical likelihood ratio (LR) methods, and they are not easy to measure, which violates one of the requirements for forensic voice comparison features (Nolan, 1983, p.11). Furthermore, the model comparisons and summaries from this study do not provide enough data to support their use in FVC casework. However, what we can conclude is that *like* and *yeah* do occur often, and the way which speakers utilise them may itself be speaker-specific. Some speakers stick out as using a higher-than-average number of one feature. Speakers 048 and 028 had a rate of 7.72 and 6.80 tokens of *yeah* per 100 words – far higher than the overall average of 4.31. For *like*, speakers 032 and 015 seem to be outliers in that they have a much lower frequency of use of only 0.69 and 0.79 respectively.

Yeah/like could be useful for FVC in this way as a grammatical variable. As D'Arcy (2007) claims, all functions of *like* are used systematically in speech. How a speaker uses *like* could be a discriminant. For example, speaker 19 uses *like* the most prolifically in the data (with 86 tokens) and uses grammatical, quotative and discourse *like* to a similar degree. Speaker 15, however, only produced 5 tokens, and all are grammatical *like*. The grammatical choices these speakers make as to where *yeah/like* can occur and the function they fulfil may serve as potential identifiers outside of the phonetic information.

Considering contextual information for phonetic comparison could additionally be applied to already established FVC features such as FPs. Applying function information to vowels may also help to better predict variation. Further testing of *yeah/like* within LR models may even show that contextual information such as pause context can improve numerical LRs and reduce error rates.

As expected, pronunciation of *like* varies by its function and pause context. *Yeah* follows a similar pattern, though this may have more to do with speaker incipiency (when speakers indicate they are about to begin a turn). There is more work to be done to assess *yeah* and *like* as potential speaker discriminants, but certainly their variation shows promise (Gibb-Reid, forthcoming).

5. *Conclusion*

In summary, *like* and *yeah* are frequently used words in WYE. Their acoustic variation indicates that they are pronounced differently based on pragmatic function. Quotative *like* has less F2 movement across its trajectory compared with other functions, indicating a more monophthongal vowel–as found by others–but also a clearer /l/. When *yeah* is used as a feedback token, it is more likely to have less F1 and F2 movement across /j/ and /ɛ/, whereas response token *yeah* is more likely to have a wider range of formant movement. These words are also affected by their prosodic context, though in a less-predictable way: when *like* occurs before a pause it is likely to have a more diphthongal vowel whereas when *yeah* occurs in the same position, it shows less formant movement across segments. The fact that these words vary in pronunciation across speakers, and their high frequency, suggests some merit to consider testing their usefulness for FVC analysis.

Acknowledgements

We are grateful to Erica Gold, Sula Ross and Kate Earnshaw for allowing us access to the WYRED corpus for this study. Thanks to Márton Sóskuthy for GAMMs advice wizadry and Claire Childs for insightful advice on discourse-pragmatic variation. Thanks also to the anonymous reviewers for their feedback and suggestions.

References

BOERSMA, P. & WEENINK, D. (2016). Praat: doing phonetics by computer [Computer program]. Version 6.0.17, retrieved from http://www.praat.org/

BYBEE, J. (2007). Frequency of use and the organization of language. Oxford University Press on Demand.

CARTER, P., & LOCAL, J. (2007). F2 variation in Newcastle and Leeds English liquid systems. Journal of the International Phonetic Association, 37 (2), 183–199.

D'ARCY, A. (2007). *Like* and language ideology: Disentangling fact from fiction. American Speech, 82(4), 386–419.

DETERDING, D. (1997). The formants of monophthong vowels in Standard Southern British English pronunciation. Journal of the International Phonetic Association, 27(1-2), 47–55.

DRAGER, K. (2011). Sociophonetic variation and the lemma. Journal of Phonetics, 39(4), 694–707.

DRAGER, K. (2016). Constructing style: Phonetic variation in discursive functions of like. In H. Pichler (Ed.), Discourse-Pragmatic Variation and Change in English: New Methods and Insights (pp. 232-251). Cambridge: Cambridge University Press.

DRUMMOND, K., & HOPPER, R. (1993a). Some Uses of *Yeah*. Research on Language and Social Interaction, 26(2), 203–212.

DRUMMOND, K., & HOPPER, R. (1993b). Back Channels Revisited: Acknowledgment Tokens and Speakership Incipiency. Research on Language and Social Interaction, 26(2), 157–177.

FABRICIUS, A. (2002). Weak vowels in modern RP: An acoustic study of happy-tensing and kit/schwa shift. Language variation and change 14(2). Cambridge University Press. p.211–237.

FRASER, B. (1996). Pragmatic markers. Pragmatics, 6(2), 167-190.

FULLER, J. M. (2003). The influence of speaker roles on discourse marker use. Journal of Pragmatics, 35(1), 23–45.

Gibb-Reid, B. (forthcoming) Comparing *like* with *like*: An analysis of the potential of the word *like* as a speaker discriminant.

GOLD, E, ROSS, S. & EARNSHAW, K. (2018). "The 'West Yorkshire Regional English Database': Investigations into the Generalizability of Reference Populations for Forensic Speaker Comparison Casework". Proceedings of Interspeech 2018, September 2-6 2018, Hyderabad, pp.2748-2752.

GRIVIČIĆ, T., & NILEP, C. (2004). When Phonation Matters: The Use and Function of *Yeah* and Creaky Voice. Colorado Research in Linguistics, 17(1), 1–11.

GUAÏTELLA, I. (1993). Functional, acoustical and perceptual analysis of vocal hesitations in spontaneous speech. In ESCA Workshop on Prosody. isca-speech.org. Retrieved from http://www.isca-speech.org/archive_open/prosody_93/pro3_128.html

HUGHES, V. (2013). save_dynamic_intervals.Praat

HUGHES, V., WOOD, S., & FOULKES, P. (2016). Strength of forensic voice comparison evidence from the acoustics of filled pauses. International Journal of Speech, Language and the Law, 99–132.

JEFFERSON, G. (1984). Notes on a Systematic Deployment of the Acknowledgement Tokens '*Yeah*' and 'Mmhm'. Papers in Linguistics 17 (2), 197–216

JOHNSON, K. (2011). Acoustic and Auditory Phonetics. John Wiley & Sons.

KIRKHAM, S. & WORMALD, J. (2015) Acoustic and Articulatory Variation in British Asian English Liquids. Proceedings of ICPhS XVIII. 1–5. Glasgow: University of Glasgow.

KÜNZEL, H. J. (1997). Some general phonetic and forensic aspects of speaking tempo. International Journal of Speech Language and the Law, 4(1), 48–83.

LADEFOGED, P. (2001). A Course in Phonetics (4th edn.). Harcourt College Publishers.

LENNES, M. (2002). save_intervals_to_wav_sound_files.PRAAT. Retrieved from https://lennes.github.io/spect/howto/slicing.html

LO, J. H. (2021) Between äh(m) and euh(m): The distribution and realization of filled pauses in the speech of German–French simultaneous bilinguals. Language and Speech 63(4): 746-768.

LOVE, R., DEMBRY, C., HARDIE, A., BREZINA, V. AND MCENERY, T. (2017). The Spoken BNC2014: designing and building a spoken corpus of everyday conversations. In International Journal of Corpus Linguistics, 22(3): 319-344. DOI: 10.1075/ijcl.22.3.02lov

MASCHLER, Y., & SCHIFFRIN, D. (2015). Discourse markers: Language, meaning, and context. The handbook of discourse analysis, 2, 189-221.

NOLAN, F. (1983). The phonetic bases of speaker recognition / Francis Nolan. Cambridge: Cambridge University Press.

NOLAN, F. (1997). Speaker recognition and forensic phonetics. The handbook of phonetic sciences. Cambridge: Cambridge University Press. 744–767.

OED ONLINE. (Ed.) (2022) "yeah ad.". Oxford University Press. https://www.oed.com/view/Entry/231470?redirectedFrom=yeah (accessed July 21, 2022).

ÖSTMAN, J.-O. (1982). The Symbiotic Relationship between Pragmatic Particles and Impromptu Speech. Paper presented at the Impromptu Speech: A symposium. Papers Contributed to a Symposium on Problems in the Linguistic Study of Impromptu Speech, Abo, Finland.

PICHLER, H. (2010). Methods in discourse variation analysis: Reflections on the way forward. Journal of Sociolinguistics, 14(5), 581-608.

PICHLER, H. (2013). The structure of discourse-pragmatic variation (Vol. 13). John Benjamins Publishing.

PICHLER, H. (2016). Discourse-pragmatic variation and change in English: New methods and insights. Cambridge University Press.

RSTUDIO TEAM (2016). RStudio: Integrated Development for R. Version 1.1.456. RStudio, Inc., Boston, MA. URL: http://www.rstudio.com/.

RAPHAEL, L. J., BORDEN, G. J., & HARRIS, K. S. (2007). Speech science primer: Physiology, acoustics, and perception of speech. Lippincott Williams & Wilkins.

ROSE, P. (2002). Forensic speaker identification. London: Taylor & Francis.

ROSE, P. (2013). Where the science ends and the law begins: likelihood ratio-based forensic voice comparison in a $150 million telephone fraud. The International Journal of Speech, Language and the Law, 20(2), 277–324.

SCHLEEF, E., & TURTON, D. (2016). Sociophonetic variation of *like* in British dialects: effects of function, context and predictability. English Language & Linguistics, 1–41.

SCHOURUP, L. (1999). Discourse markers. Lingua, 107(3-4), 227-265.

SÓSKUTHY, M. (2017). Generalised additive mixed models for dynamic analysis in linguistics: a practical introduction. arXiv [stat.AP]. Retrieved from http://arxiv.org/abs/1703.05339

SÓSKUTHY, M. (2021). Evaluating generalised additive mixed modelling strategies for dynamic speech analysis. Journal of Phonetics, 84, 101017. doi:https://doi.org/10.1016/j.wocn.2020.101017

SPROAT, R. & FUJIMURA, O. (1993). Allophonic variation in English /l/ and its implications for phonetic implementation. Journal of Phonetics 21(2), 291– 311.

SWERTS, M. (1998). Filled pauses as markers of discourse structure. Journal of Pragmatics, 30, 485-496.

TRUONG, K. P., & HEYLEN, D. (2010). Disambiguating the functions of conversational sounds with prosody: the case of "*yeah*." In Eleventh Annual Conference of the International Speech Communication Association. isca-speech.org. Retrieved from http://www.isca-speech.org/archive/interspeech_2010/i10_2554.html

TSCHÄPE, N., TROUVAIN, J., BAUER, D. & JESSEN, M. (2005). Idiosyncratic patterns of filled pauses. Paper presented at the 14th Annual Conference of the International for Forensic Phonetics and Acoustics Annual Conference, 3-6 August 2005. Marrakesh, Morocco.

VAN RIJ, J., WIELING, M., BAAYEN, R. H., & VAN RIJN, H. (2016). itsadug: Interpreting TimeSeries and Autocorrelated Data using GAMMs. R package version 2.2.

WELLS, J. C. (1982). Accents of English. Cambridge: Cambridge University Press.

WILHELM, S. (2018). Segmental and suprasegmental change in North West Yorkshire – a new case of supralocalisation? Corela, (HS-24). https://doi.org/10.4000/corela.5203

WOOD, S. (2006). Generalized additive models: an introduction with R. Boca Raton: CRC Press.

i. Appendix

Below are the model comparison summaries for each level of comparison across *like* and *yeah* F1 and F2 formant trajectories created using GAMMs. The process of making these is covered in section 2.2. The comparison column in each results table refers to the alternating coefficient being tested. The model summaries are listed in two tables for each combination of *like* and *yeah* F1 and F2 – the parametric coefficients correspond to the overall average trajectory (also called height) and the smooth terms correspond to the shape of the trajectory.

#	Comparison	Chisq	Df	p-value	AICd
1	Shape: long pause	9.8	6	0.003	-0.12
2	Height: long pause	7.2	2	<0.001	0.08
3	Shape: function	9.1	6	0.006	-3.03
4	Height: function	7.3	2	<0.001	-0.24
5	Shape: short pause	32.7	8	<0.001	9.71
6	Height: short pause	22.9	3	<0.001	-0.46
7	Shape: speaker	194.1	3	<0.001	11.91

Table 3: Model comparisons for *like* F1 trajectories.

Parametric coefficients (trajectory height)					
Coefficient	**Estimate**	**Std. Error**	***t***	***p*-value**	
(Intercept)	591.7	16.8	35.3	<0.001	***
Post-long pause	0.8	14.1	0.1	0.95	
Pre-long pause	-8.7	10.1	-0.9	0.39	
Grammatical function	2.8	5.3	0.5	0.59	
Quotative function	13.0	7.3	1.8	0.07	
No short pause	23.8	12.0	2.0	0.05	*
Post-short pause	30.1	13.3	2.3	<0.025	*
Pre-short pause	-0.8	12.5	-0.1	0.95	
Approximate significance of smooth terms (trajectory shape)					
	edf	**Ref.df**	***F***	***p*-value**	
No long pause	1.0	1.0	0.0	0.97	
Post-long pause	1.0	1.0	0.3	0.58	
Pre-long pause	<0.001	<0.001	0.0	1.00	
Discourse function	2.5	3.0	0.9	0.43	
Grammatical function	0.1	0.2	0.1	0.88	
Quotative function	1.9	2.2	0.6	0.60	
Both short pause	3.4	4.1	1.8	0.12	
No short pause	1.0	1.0	0.2	0.69	
Post-short pause	<0.001	<0.001	0.0	1.00	
Pre-short pause	4.3	5.3	4.2	<0.001	***
Speaker	46.9	68.0	8.9	<0.001	***

Table 4: Model summary for all variables of *like* F1.

#	Comparison	Chisq	Df	p-value	AIC difference
1	Shape: long pause	12.3	6	<0.001	-2.75
2	Height: long pause	9.8	2	<0.001	0.00
3	Shape: function	18.6	6	<0.001	-2.27
4	Height: function	8.3	2	<0.001	0.15
5	Shape: short pause	43.3	8	<0.001	0.56
6	Height: short pause	41.9	3	<0.001	1.14
7	Shape: speaker	181.4	3	<0.001	69.95

Table 5: Model comparisons for *like* F2 trajectories.

Parametric coefficients (trajectory height)					
Coefficient	**Estimate**	**Std. Error**	***t***	***p*-value**	
(Intercept)	1420.9	34.4	41.3	<0.001	***
Post-long pause	-50.2	32.5	-1.5	0.12	
Pre-long pause	-7.4	23.7	-0.3	0.75	
Grammatical function	-14.4	12.0	-1.2	0.23	
Quotative function	-22.0	16..677	-1.3	0.19	
No short pause	-113.8	28.0	-4.1	<0.001	***
Post-short pause	-68.8	30.2	-2.3	<0.025	*
Pre-short pause	-20.2	29.0	-0.7	0.49	
Approximate significance of smooth terms (trajectory shape)					
	edf	**Ref.df**	***F***	***p*-value**	
No long pause	1.4	1.6	0.9	0.34	
Post-long pause	1.0	1.0	0.0	0.95	
Pre-long pause	<0.001	<0.001	0.1	1.00	
Discourse function	1.2	1.3	0.8	0.31	
Grammatical function	<0.001	<0.001	0.1	1.00	
Quotative function	2.5	3.0	4.6	<0.004	**
Both short pause	1.0	1.0	7.4	<0.007	**
No short pause	2.6	3.2	2.0	0.11	
Post-short pause	<0.001	<0.001	0.3	1.00	
Pre-short pause	1.0	1.0	4.2	<0.001	***
Speaker	52.0	68.0	8.9	<0.001	***

Table 6: Model summary for all variables of *like* F2.

#	Comparison	Chisq	Df	p-value	AIC difference
1	Shape: long pause	13.9	8	<0.001	-1.88
2	Height: long pause	9.7	3	<0.001	-0.04
3	Shape: function	23.5	8	<0.001	-3.98
4	Height: function	9.5	3	<0.001	-0.07
5	Shape: short pause	28.8	8	<0.001	-2.42
6	Height: short pause	13.3	3	<0.001	0.79
7	Shape: speaker	176.3	3	<0.001	49.96

Table 7: Model comparisons for *yeah* F1 trajectories.

Parametric coefficients (trajectory height)					
Coefficient	**Estimate**	**Std. Error**	***t***	***p*-value**	
(Intercept)	533.4	10.6	50.2	<0.001	***
No long pause	-3.4	9.4	-0.4	0.72	
Post-long pause	-0.4	6.1	-0.1	0.95	
Pre-long pause	9.3	10.4	0.9	0.37	
Discourse function	8.0	8.0	1.0	0.32	
Feedback function	4.5	5.0	0.9	0.37	
Response function	11.5	7.2	1.6	0.12	
No short pause	9.2	11.3	0.8	0.41	
Post-short pause	-6.5	6.3	-1.0	0.30	
Pre-short pause	12.4	10.0	1.2	0.22	
Approximate significance of smooth terms (trajectory shape)					
	edf	**Ref.df**	***F***	***p*-value**	
Both long pause	1.0	1.0	0.4	0.54	
No long pause	<0.001	<0.001	0.1	1.00	
Post-long pause	2.4	2.8	1.0	0.40	
Pre-long pause	2.0	2.5	0.4	0.81	
Agreement function	1.9	2.2	2.9	0.61	
Discourse function	1.0	1.0	3.2	0.07	
Feedback function	2.1	2.5	8.4	<0.001	***
Response function	<0.001	<0.001	0.2	1.00	
Both short pause	1.0	1.0	0.0	0.90	
No short pause	3.6	4.4	1.9	0.09	
Post-short pause	4.4	5.3	3.8	<0.002	**
Pre-short pause	<0.001	<0.001	0.0	0.50	
Speaker	55.2	73.0	8.0	<0.001	***

Table 8: Model summary for all variables of *yeah* F1.

#	Comparison	Chisq	Df	p-value	AIC difference
1	Shape: long pause	14.5	8	<0.001	-0.97
2	Height: long pause	14.5	3	<0.001	0.22
3	Shape: function	20.6	8	<0.001	0.17
4	Height: function	19.9	3	<0.001	0.82
5	Shape: short pause	23.2	8	<0.001	-1.64
6	Height: short pause	29.4	3	<0.001	0.49
7	Shape: speaker	220	3	<0.001	46.53

Table 9: Model comparisons for *yeah* F2 trajectories.

Parametric coefficients (trajectory height)					
Coefficient	**Estimate**	**Std. Error**	***t***	***p*-value**	
(Intercept)	1715.2	23.6	72.6	<0.001	***
No long pause	-1.5	14.4	-0.1	0.92	
Post-long pause	17.7	9.2	1.8	0.08	
Pre-long pause	-28.9	16.5	-1.7	0.08	
Discourse function	-8.9	12.8	-0.7	0.49	
Feedback function	-24.2	8.2	-3.0	<0.004	**
Response function	11.9	11.7	1.0	0.31	
No short pause	-4.9	17.8	-0.3	0.78	
Post-short pause	15.9	10.3	1.5	0.12	
Pre-short pause	-39.9	16.1	-2.5	0.013	*
Approximate significance of smooth terms (trajectory shape)					
	edf	**Ref.df**	***F***	***p*-value**	
Both long pause	1.0	1.0	0.3	0.59	
No long pause	<0.001	<0.001	0.0	1.00	
Post-long pause	1.0	1.0	0.4	0.54	
Pre-long pause	1.0	1.0	0.9	0.36	
Agreement function	1.0	1.0	0.5	0.49	
Discourse function	<0.001	<0.001	0.1	1.00	
Feedback function	1.0	1.0	0.3	0.59	
Response function	3.6	4.4	2.8	<0.020	*
Both short pause	1.0	1.0	1.8	0.18	
No short pause	1.0	1.0	9.6	0.00	
Post-short pause	1.4	1.6	4.7	<0.025	
Pre-short pause	0.3	0.4	0.7	0.59	
Speaker	48.3	73.0	10.7	<0.001	

Table 10: Model summary for all variables of *yeah* F2.

Ben Gibb-Reid, Paul Foulkes, Vincent Hughes
Department of Language & Linguistic Science
University of York
Heslington
York
Email: ben.gibb-reid@york.ac.uk / benwillreid@gmail.com