CAMBRIDGE
UNIVERSITY PRESS

The intersection of ethnicity and social class in language variation and change

Catherine E. Travis[1] (iD) and Qiao Gan[2]

[1]School of Literature, Languages and Linguistics, The Australian National University, Canberra, ACT, Australia and [2]New Zealand Institute of Language, Brain and Behaviour, University of Canterbury, Christchurch, Canterbury, New Zealand
Corresponding author: Catherine E. Travis; Email: Catherine.Travis@anu.edu.au

Abstract

In the wealth of literature on ethnic variation, ethnicity is often considered independently of other social characteristics. However, prioritizing ethnicity in this way risks overlooking the potential impact of other social factors. In this study, we demonstrate an intersection between ethnicity and social class based on a sociolinguistic corpus of Australian English, representing some of the country's largest ethnic groups (Australians of Anglo-Celtic, Italian, Greek, and Chinese backgrounds), stratified according to age, gender, and social class. Rather than beginning with the social groupings, we first identify linguistic groupings to then consider how these groupings align with social dimensions. Cluster analyses of speaker random intercepts derived from independent regression analyses of 10 linguistic variables in recordings from 159 speakers reveal primary divisions for age, reflecting change over time, and secondary divisions for ethnicity in conjunction with social class, highlighting the interconnected nature of these social dimensions in linguistic variation.

Introduction

There has been longstanding attention paid to the role of ethnicity in shaping language variation and change, from as early as Labov's founding studies in Martha's Vineyard (Labov, 1963) and New York (Labov, 1966), such that ethnicity is recognized as one of the fundamental social variables alongside gender and social class. However, while gender and social class are regularly considered in relation to each other (Labov, 1990), ethnicity is often given privileged status and considered independently of other social factors, reflecting an assumption that ethnically marked ways of speaking serve to index ethnicity only (as observed by Britain, 2022:332; Eckert, 2008:26).

In this paper, we challenge this assumption by demonstrating an intersection between ethnicity and social class. The analyses presented here build on a body of work on ethnic variation in Australian English that has demonstrated not only that

social class is a strong explanatory factor for language variation among Australians of Anglo-Celtic background and of minority ethnic backgrounds, but also that the patterning of different ethnic groups is impacted by corresponding social class affiliations, such that apparent ethnic differences are often diminished once social class is factored in (as summarized in Travis, 2024:177-178). The evidence for this has come from the linguistic behavior observed across different ethnic and social class groups, following the standard approach in sociolinguistics. However, an approach that relies on predetermined social categories is not ideal if the purpose is to explore the nature of those categories (cf. Horvath & Sankoff, 1987:182-184). First, it assumes a priori the reality of these social groupings, which is problematic because social class is notoriously difficult to define and has not been widely studied in Australian English, and because ethnicity is not a monolithic concept. Second, it can be difficult to obtain sufficient representation across ethnic and social class groups to fully test interactions, given that ethnic groups often show skewed social distribution. In order to address this, here we reverse the lens, to first identify groupings defined according to linguistic behavior and then explore how these groupings align with social dimensions, following the pioneering work of Horvath and Sankoff (1987).

We consider 10 linguistic variables spanning phonetic, morphophonological, morphosyntactic, and discourse levels, occurring in approximately 750,000 words of speech recorded with 159 Australians across two time periods (the 1970s and 2010s). These variables are the FLEECE and FACE vowels, word-final -er, prevocalic the, word-final unstressed -ing, existential there's, modals of obligation have got to and need to, and quotatives go and be like. Speakers come from the majority Anglo-Celtic population and three of Australia's largest migrant groups (Italian, Greek, and Chinese Australians), and are further stratified by age, gender, and social class. To identify linguistic groupings, we conduct cluster analyses using speaker random intercepts extracted from independent regression analyses that included the linguistic conditioning only for each of the 10 variables. Cluster analyses for the 1970s and the 2010s groups point to an intersection between ethnicity and social class, in that ethnic groups cluster with their age and social class cohorts such that no clusters are defined by ethnicity alone. This methodology provides an empirical, quantitative foundation for exploration of social groupings based on linguistic behavior, and the result evidences the interconnectedness of these social dimensions in shaping linguistic variation.

Ethnic variation in its social context

A key question that has been explored in relation to ethnic variation relates to the "integration of ethnic groups into the social system" (Labov, 1966:vi). One particularly marked case of ethnic variation is the development of "ethnolects" (e.g., Hoffman & Walker, 2010:42) or "multiethnolects" (e.g., Cheshire et al., 2011:153), that is, new language varieties that arise in contexts where people from different ethnic and linguistic backgrounds come together. Such varieties have been interpreted as being used as "a means of expressing linguistic identity" (Clyne et al., 2001:226).

Another line of research has been around the participation of ethnic minorities in wider patterns of variation and change, and here also, ethnic identity has been drawn on as an explanation for any differences observed. For example, young people

of migrant background in Gothenburg, Sweden, were reported to be differentiated from the Swedish majority in that they were leading in a change, being the first in the city to adopt a shift taking place in Stockholm towards a more open realization of the /ɛ/ vowel (Gross et al., 2016:238), whereas Turkish and Moroccan youths in the Netherlands were reported to be differentiated by lagging in a change towards a diphthongized variant for Dutch /ɛi/, preferring the older monophthongal realizations "characteristic of the 'traditional' urban dialects" (van Meel et al., 2014:68). In both cases, this patterning was interpreted as serving to index ethnicity (Gross et al., 2016:243; van Meel et al., 2014:69). Hoffman and Walker identified similar linguistic conditioning among British, Italian, and Chinese Canadians who grew up in Toronto for both the Canadian Vowel Shift, a change in progress, and (t/d)-deletion, a stable variable (2010:52, 55), but found some differences in rates of participation and deletion, leading them to propose that participants "may be using overall rates of use to construct and express ethnic identities" (2010:58). In Australia, differential patterning for Lebanese Australians' realization of /l/, Voice Onset Time, and prosody has also been tied to ethnic identity (Clothier, 2019:1891; Clothier & Loakes, 2018:15-16; Cox & Palethorpe, 2011:531).

It is important to acknowledge that not all studies find ethnic differences. In the Swedish study reported above, ethnic differences were observed in Gothenburg, but not in Stockholm (Gross et al., 2016:238). In Toronto, Hoffman and Walker's finding regarding shared linguistic conditioning across ethnic groups in Toronto corresponds with other work by Hoffman that has reported no significant differences in participation in the Canadian Vowel Shift between British, Italian, and Chinese Canadians, as all groups were "acquiring and accommodating to the local norm" (2010:136). Another study comparing Chinese Americans in San Francisco and New York found that they were participating in change in the ʙᴏᴜɢʜᴛ vowel as it was taking place in the city in which they lived, and differently from each other, such that their patterning aligned not with ethnicity, but with regional variation (Wong & Hall-Lew, 2014:33).

Furthermore, where differences do emerge, assuming a priori that these index ethnicity may fail to account for differences within the one ethnic group and may risk overlooking other factors that may be at play. For example, one study identified differences in participation in the Northern Cities Vowel Shift by Jewish Americans in Chicago across more and less affluent neighborhoods, leading the authors to conclude that "productions were influenced by local place- and class-based meanings" (Benheim & D'Onofrio, 2024:152). Class-based meanings were also drawn on to account for participation in the Northern Cities Vowel Shift by some African Americans in Rochester, New York, where it was argued that a group of "Mobile Black Professionals" were adopting supralocal sound changes as part of their external orientation (King, 2021:174).

While these studies have suggested that ethnic minorities draw on the potential indexicality associated with class-based differences, it has also been shown that ethnic patterning cannot be separated from social class characteristics of the relevant ethnic groups. This was demonstrated in early variationist work for Jewish communities in North America. For example, in comparing the linguistic behavior of Italian and Jewish Americans in New York City, though Labov found that ethnicity was a "more powerful factor" than social class for the ᴛʀᴀᴘ vowel (Labov, 1966:306), the distribution across

the ethnic groups was skewed, with the Jewish participants having an overall higher class distribution than the Italian Americans (Labov, 1966:293). In Boston, to interpret her finding that Jewish Americans were ahead of Italian and Irish Americans in a shift away from regionally specific patterning for the NORTH vowel, Laferriere noted that Jewish Americans were more likely than both Italian and Irish Americans to be university educated and work in white collar occupations, rendering them "the first to adopt the standard variant, associated with the values of the larger society, through educational exposure" (Laferriere, 1979:612). Many years later, Boberg attributed his finding of greater participation in Canadian norms by Jewish than Italian Canadians to Jewish Canadians' "access to higher education, and therefore exposure to mainstream, middle-class, Canadian English" (Boberg, 2004:563). A parallel in the United Kingdom may be what has been named "British Asian English," spoken by people of Punjabi heritage in a middle class area in West London, which, unlike Multicultural London English (spoken in underprivileged areas), is described as aligning with the norms of "high prestige Standard Southern British English," while also incorporating some features from "elite Indian English" (Sharma, 2020:56). In Australia, work that compares speech in more versus less ethnically diverse regions in Sydney is also faced with a skewed distribution, in that the more linguistically diverse areas (specifically, Western Sydney) tend to be more socio-economically disadvantaged (Cox & Penney, 2024:204). Most recently, Walker found that level of education helped explain differences in patterning of *-ing* among Chinese, Italian, and Portuguese Canadians in Toronto: the Chinese Canadians (all of whom were university educated) had the lowest rate of vernacular *-in'* and no gender difference, and the gender differences among the Italian and Portuguese Canadians corresponded to the different educational backgrounds of the men and women (2024:302).

Some of the first studies to directly test the intersection between ethnicity and social class were conducted in Australia, using a corpus collected in the 1970s and 1980s. In this early work, Horvath found that Australian-born teenagers of Italian and Greek backgrounds were leading a change in Australian English vowel realizations, away from both "broad" and "cultivated" realizations (associated with working class speech and British Received Pronunciation, respectively), towards more "general" realizations (Horvath, 1985:91-94). She interpreted this as an effort by the teenagers to differentiate themselves from their Italian- and Greek-born parents, whose accented vowels were described as "ethnic broad" (Horvath, 1985:69), in that they were qualitatively distinct from realizations of the Anglo-Celtic majority (being influenced by the speakers' first language) and were closest to "broad" vowels, consistent with their work in lower-status occupations, for example factories, the construction industry, and family businesses such as milk bars or vegetable shops (Horvath, 1985:46; cf. Jupp, 2001:83). The vowel realizations of the Greek and Italian teenagers, in contrast, were not qualitatively distinct from those of their Anglo peers, though these groups were socially distinct—neither the Greeks nor the Italians exhibited the same gender differences as the Anglos, and the Italians did not exhibit the same class differences (Horvath, 1985:81-83). Horvath hypothesized that, at this time, "the safest ground for 'sounding' Australian, whether Greek, Italian or Anglo, [was] the middle of the Broad-General-Cultivated continuum" (1985:176), and in this way, the Italian and Greek teenagers

drove forward a change that was underway in Australian English. Using the same corpus, Guy, Horvath, and colleagues found no linguistic differentiation in High Rising Terminal pitch by the Anglo, Italian, and Greek teenagers, with, again, the only distinction being in relation to the social conditioning, specifically in how the ethnic groups respond to class differentiation (1986:40). Our own recent research, in which we have reanalyzed these same data together with a comparable contemporary corpus (outlined in the following section), corroborates this general pattern.

What is clear from this body of work is that ethnic variation cannot be assumed to be directly indexing ethnicity and must be considered in relation to the broader social context. Britain has observed that what are assumed to be "ethnic" variants may not "merely (or possibly even mainly) index ethnicity per se, but potentially a localized, gendered, classed, age-constrained identity (too)," leading him to call for a shift away from an "asocial perspective" to pay greater attention to potential intersectionality (Britain, 2022:332). This paper seeks to address this call, and to do this at scale, looking across multiple variables, different ethnic groups, and over time.

Sydney Speaks

The data for this study come from the Sydney Speaks corpus (Travis et al., 2023), a sociolinguistic corpus collected for the purpose of examining language variation and change in one of Australia's largest and most diverse cities. The corpus totals over 1.2 million words of transcribed speech from 130 hours of recordings with 265 Australians taken from three sub-corpora—recorded in the 1970s, 1980s, and 2010s—with birth years spanning from 1889 to 2001 (see Travis, 2024, for a corpus overview).

Participants

For this paper, we draw on two sub-corpora: the Sydney Social Dialect Survey recorded in the 1970s with Australians of Anglo-Celtic, Greek, and Italian backgrounds (Horvath, 1985), and the ANU Corpus of Sydney Speech, recorded from 2014 to the present with Australians of Anglo-Celtic, Greek, Italian, and Chinese backgrounds. Participants are stratified for age, gender, socio-economic class, and ethnicity, allowing for ethnic variation and change to be examined in both real and apparent time. A total of 159 participants are included in this study (restricted to those for whom data for all the variables is available; see the section below on the linguistic variables). As summarized in Table 1, these participants represent four time points, based on time of recording and year of birth: Adults and Teens recorded in the 1970s (born around the 1930s and 1960s, respectively), and Adults and Young Adults in the 2010s (born around the 1960s and 1990s, respectively).[1] As there is evidence of language change over this time, for the purposes of the analyses presented here, we analyze each group with their cohort recorded in the same time period.[2]

All participants grew up in Australia (either being born in Australia or having arrived before the age of six) and live in Sydney. The different ethnic groups included represent the largest groups in Australia at the time of recording. Anglo-Celtic Australians are defined as those whose parents and grandparents were born in Australia, and who grew up in an English-speaking household; this represents the

majority community in Australia, and they are included at each time point. Greek and Italian Australians are included in the Sydney Speaks corpus from the 1970s Teens on. In the analyses conducted here, we include Italian Australians for three time points, but Greek Australians for the 1970s Teens only, due to low numbers of participants for the other age groups. The 1970s Teens and 2010s Adults are the children of the first large wave of migration from Greece and Italy following the Second World War (Jupp, 2001:83). The 2010s Italian Young Adults are primarily second-generation Australians; three of the 12 are third-generation Australians. The Chinese Australians are second-generation Australians whose parents migrated from Hong Kong and Cantonese-speaking areas in mainland China, primarily under business migration schemes that ran in the 1980s, following the opening up of Australian migration policy (Jupp, 2001:218). The Greek, Italian, and Chinese participants all grew up in homes in which the respective language was spoken; they thus all have exposure to the language, but not all of them speak it.

Table 1. Participants by age, ethnicity, and gender

Year of recording	1970s				2010s				
Age group	Adults		Teens		Adults		Young adults		
Approx. birth year	1930s		1960s		1960s		1990s		
Approx. age range	32-64 yrs		12-18 yrs		42-61 yrs		18-31 yrs		
	Men	Women	Men	Women	Men	Women	Men	Women	**Total**
Anglo-Celtic	7	9	12	11	6	9	8	10	72
Italian			11	12	8	8	5	7	51
Greek			11	8					19
Chinese							9	8	17
Total	7	9	34	31	14	17	22	25	159

Participants are stratified for social class, which was determined from a composite score made up of three independent factors, obtained from a demographic questionnaire that was carried out immediately following the sociolinguistic interview: occupation, for the participants themselves or for their parents for those who were still in high school (five points, scored according to the Australian Socio-Economic Index, AUSEI; McMillan et al., 2009); highest level of education (five points, distinguishing between non-completion of high school, completion of high school, technical college [known as TAFE], university, and post-graduate degree); and type of high school, based on historical and current government funding as well as fees and resources typically contributed by families for attendance (four points, distinguishing between State, Catholic, Selective [prestigious, government-funded high schools], and Independent schools). These scores were summed together to provide a continuous measure, on a scale from three to 14. This operationalization allows us to capture, in a single measure, the complexity of social class, by taking into account different elements that make it up.

The different ethnic groups are not evenly spread across social class, as can be seen in the density plot in Figure 1, which shows the distribution of the participants by class and ethnicity for the different age groups. Class is presented on the x-axis, moving up in social class score as we move farther to the right, and the four age groups are presented

on the y-axis, from oldest to youngest as we move down. The different ethnic groups are captured in color, and height represents the number of speakers in each group. First, of note here is the overall higher social class scores for the 2010s participants relative to the 1970s participants, seen in that the peaks fall farther to the right as we move down the chart, especially for the 2010s Young Adults. This is largely due to changes in higher education tuition funding from 1974 onwards, which made university more accessible, resulting in increased participation rates (Norton, 2023:32-33).

Of more interest to us here is the distribution of the ethnic groups according to social class. For the 1970s Adults and Teens, the Anglo Australians cover the full range; the Italian and Greek Teens, on the other hand, fall slightly lower on the range (the peaks are to the left of the Anglos), a reflection of the working-class occupations of
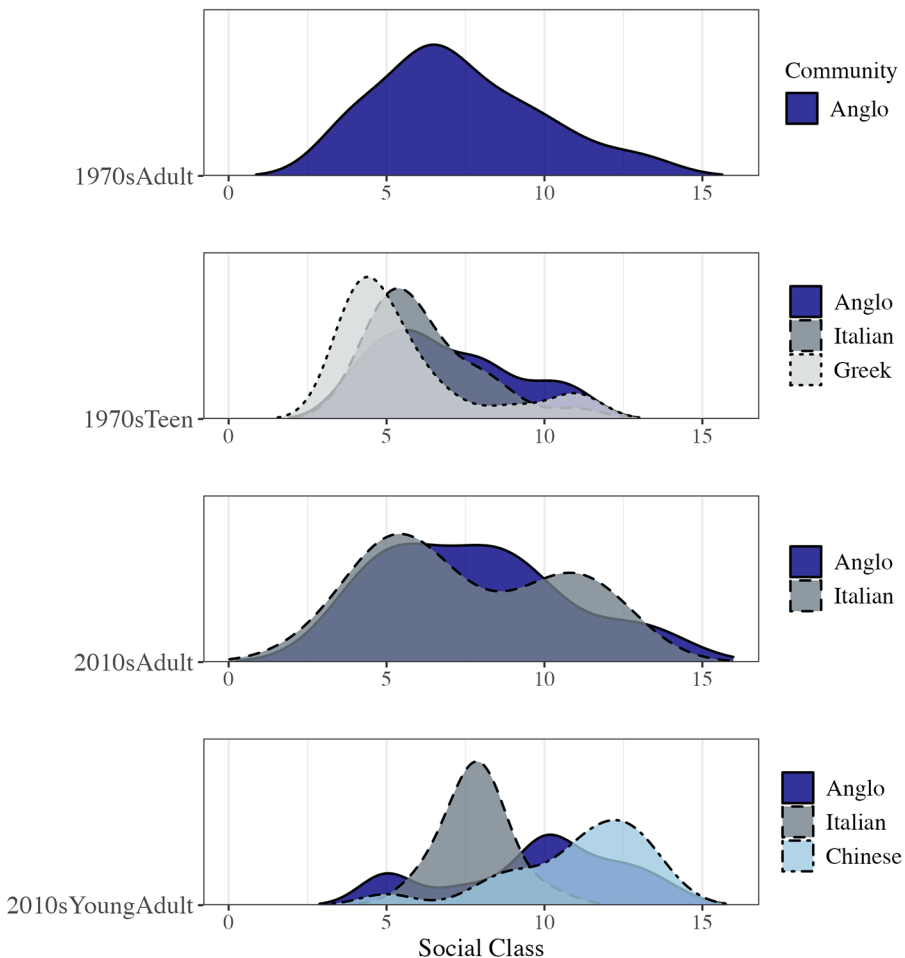

Figure 1. Social class scores for participants by ethnicity over time (*n* = 159).

their parents. The 2010s Adult Italians, however, have a wider distribution, reflecting the upward mobility this group experienced over this 40-year time period (Ricatti, 2018:41). For the 2010s Young Adults, the Italian-background participants have limited distribution, most falling between six and nine (a distribution that we are working to broaden through ongoing data collection). The Chinese Australians, in contrast, are overall higher on the scale, something which is not a sampling issue but a reflection of the nature of the community of second-generation Australians of Cantonese background today, a group that tends to attend selective or private high schools, achieve high levels of education, and work in professional occupations (Jupp, 2001:221).

Sociolinguistic interview data

The recorded speech comes from sociolinguistic interviews (Labov, 1972), recorded in the 1970s by Horvath and associates, and in the 2010s by research assistants who were members of the same ethnic communities as the participants. From interviews lasting between 45 and 90 minutes, approximately 30 minutes were selected for transcription, which amounted to around 4,500 words of speech from each participant. The data were transcribed orthographically in *ELAN* (Lausberg & Sloetjes, 2009) and force-aligned in *LaBB-CAT* (Fromont & Hay, 2012). Vowel alignments were manually corrected and processed in *Praat* (Boersma & Weenink, 2024) to obtain estimates of F1 and F2 values for analysis (more details on corpus preparation available in Travis, 2024:172-175).

Though the topics covered vary across the two time periods and across individual interviews, common themes that came up include participants' life experiences (e.g., growing up in Sydney, travel, work, and hobbies), language use, network, and identity, all of which can be helpful for understanding the social background of the participants and interpreting the linguistic behavior observed. To illustrate, Example 1 presents an excerpt from an interview with a Young Adult Chinese Australian, Lathan, who is talking about accent differentiation across Sydney and outlining his belief that, more than ethnic differences, there are class-driven regional differences, in what turns out to be a very astute observation.

(1)

Lathan: *it- it's not necessarily an .. Asian thing.*
it --
.. it might be a factor,
but it could be,
.. you know,
there's an Eastern Suburbs accent.
((33 intervening Intonation Units))
.. It's --
it could be .. drawn --
...(1.0) i- it's --

it's kind of like a --
.. what do you call those.
like a plane graph.
So I think,
... you know,
.. ethnicity could be the X-axis of whatever.
and then,
you know,
.. socio-economic factors or,
like,
geography or whatever is the Y,
.. axis.
And then you've got a Z-axis.
and all that kind of stuff.

[SydS_CYM_072: 00:58:53-59:12][3]

Linguistic variables

As this work is concerned with the social patterning for multiple variables, we do not have space for detailed explanation of the conditioning of each individual variable. We therefore utilize a set of variables for which analyses are already available, allowing us to focus our attention on the patterning across these variables as a set. The variables included are listed in Table 2, along with the nature of the variation observed in prior Sydney Speaks work and the number of tokens included in the regression analyses

Table 2. Linguistic variables included

Variable	Variation	*n* 1970s	*n* 2010s
Phonetic			
FLEECE	raising and fronting	2,701	4,236
FACE	raising and fronting	2,904	5,091
Morphophonological			
Word-final -*er*	lengthening	1,456	1,967
the + Vowel	FLEECE vs. schwa	—	1,440
Word-final unstressed -*ing*	velar vs. alveolar	2,824	7,366
Morphosyntactic			
there's	*there's* vs. *there are*	615	934
Modals of obligation	*have got to* vs. *have to*	548	560
	need to vs. *have to*	—	471
Discourse pragmatic			
Quotatives	*go* vs. *say/think/tell*, etc.	761	—
	be like vs. *say/go/think/tell*	—	3,335

(the model summaries of which are produced in the Supplementary Materials). It is necessary to consider the 1970s and 2010s data separately because the variation is not identical at the two time points, and because we include different ethnic groups at each time point. We now briefly describe each of the variables in turn.

The variables come from different linguistic levels. Two phonetic variables are the vowels FLEECE and FACE, which have been demonstrated to be undergoing raising and fronting over time, moving away from realizations traditionally associated with men and the working class, though with different timing, with FLEECE shifting from the 1970s and FACE from the 2010s (Grama et al., 2021:300, 306; Purser et al., 2020:284). As noted above, Italian and Greek teenagers were at the forefront of this change in the 1970s (Horvath, 1985:94), whereas today, it is the Chinese Australians who are the most ahead (Grama et al., 2021:307). For these variables, as an overall measure, we use the position along the front diagonal, calculated by subtracting F1 from F2, using the Lobanov normalized values converted back to Hz (as was done in Grama et al., 2021:302).

We include three variables at the morphophonological level: word-final *-er*; prevocalic *the*; and word-final unstressed *-ing*. Word-final *-er* refers to unstressed schwa in word-final position followed by an orthographic "r" that is typically not pronounced in Australian English, which is a non-rhotic variety (e.g., *teacher, culture*). Grama et al. (2020) found lengthening (and accompanied backing and lowering) of this form beginning with the 1970s Teens, and in particular the Greek Teens, for whom this appears to have been an ethnolectal marker (cf. Clyne et al., 2001:228; Kiesling, 2005:20). Today, *-er* is no longer differentiated by ethnicity, with the lengthening having been taken up across the community (Grama et al., 2020:358-359; Sheard, 2022). Prevocalic *the* (e.g., in *the other, the Italian*) is currently undergoing a change from FLEECE to schwa in Australian English, as in other varieties (e.g., Cox et al., 2023:15; Fox, 2015:167; Hay et al., 2012). In Australia, this is a recent change, seen in apparent time comparisons between the 2010s Adults and Young Adults, and is led by middle-class women and by Chinese Australians (Gan & Travis, 2022:58-59). The variation between velar and alveolar *-ing* (e.g., *going/goin'*) follows a similar pattern in Australia as elsewhere: it is generally stable, with the alveolar variant favored by men and the working-class. The highest alveolar rates are found here with the 1970s Teens, largely consistent with an adolescent peak, and the 2010s Adult Italians, attributed to the working-class associations of this group; the lowest rates are found with the 2010s Young Adult Chinese (Travis et al., 2023:446).

At the morphosyntactic level, we consider existential *there's* and modals of obligation. Existential *there's* refers to the alternation of the singular and plural verbs with plural arguments (*there's* versus *there are*; e.g., *there's lots of local breweries around,* [SydS_AOF_005]). There was a marked increase in *there's* from the 1970s Adults to Teens, led by the working class, and today, this change has largely stabilized, with 2010s Adults and Young Adults exhibiting a similar rate of *there's* as the 1970s Teens (Gan, 2024:151, 159). For modals of obligation, we consider the long-standing variant *have to* in variation with *have got to,* for which there is largely stable variation, with a favoring of *have got to* (also in the reduced form *gotta*) by the working class and men. In the 2010s a new variant has arisen, namely *need to,* which is not socially distinguished from

have to but which is favored by the middle class and by Chinese Australians relative to *should* (Travis & Torres Cacoullos, 2023:364, 369).

Finally, we include quotatives as a discourse-pragmatic variable. Variation across quotatives in Australia (as elsewhere) includes *say, be like, think, go,* and zero quotative, as well as other miscellaneous verbs (*tell, wonder,* etc.). In the 1970s, quotative *go* increased as a competing variant (e.g., *he goes mate, come on let's go for a drive.* [SydS_IOM_058]), and in the 2010s, *be like* has arisen and taken over as the majority variant among Young Adults (e.g., *I'm like, I don't even have this much food,* [SydS_IYF_037]) (Lee, 2020:34). The 1970s rise of *go* was led by the Greek Teens, and by the Anglo and Italian middle and lower middle class men, while 2010s *be like* is favored by women, the middle class, and Anglos and Italians over Chinese Australians (Lee, 2020:62, 76).

This set of variables includes some that exhibit change that is ongoing over the period examined (FLEECE, *-er*) and some that are stable (*-ing, have to* versus *have got to*), as well as others that undergo change in one period and are stable in the other (e.g., FACE stable in the 1970s, *there's* stable in the 2010s), and others that are only relevant in one time period (quotatives, *go* in the 1970s and *be like* in the 2010s; *the* + Vowel and *need to* in the 2010s). Some have clear class and gender differentiation (e.g., *-ing, have got to*), while others do not (e.g., *-er* in the 2010s), and some exhibit related ethnic differentiation, with the Italian and Greek Teens and Italian Adults tending to favor forms that are more associated with men and the working class, and Chinese Young Adults forms that are more associated with the middle class (e.g., *-ing*). The individual analyses carried out already paint a picture of class-ethnicity intersections for many variables, but we might wonder whether this result has been promoted by assuming defined ethnic and social class groups. Thus, we now test this further by examining how speakers are clustered when we consider only their linguistic behavior. If this reveals clear ethnic groupings, this would support the primacy of ethnicity in driving variation; if not, then the emerging groups should help us better understand the relationship between ethnicity and other social factors.

Cluster analyses

The method we employ to group speakers according to their linguistic behavior is the clustering technique of Divisive Analysis (DIANA), following previous work assessing social groupings in New York (Haddican et al., 2021, 2022).[4] Prior to presenting the results, we discuss the preparation of the data for entering into the clustering algorithm, the nature of the clustering performed, and the methods used to interpret the clusters.

Speaker random intercepts as a measure of speaker behavior

In order to be able to group speakers according to their linguistic behavior, the first step is to obtain comparable data for each of the linguistic variables included. The raw values are on different scales (Hz measurements for the vowels, duration in milliseconds for *-er*, and rates of a certain variant for the categorical variables). Though we could scale these to make them comparable, relying on raw values is still potentially problematic because they are impacted by the linguistic context in which the instances occur, and

the distribution of those contexts is not even across speakers in these spontaneous speech data (cf. Torres Cacoullos & Travis, 2018:148). For example, a speaker who produces a high proportion of quotatives in the historical present is likely to have a higher rate of the *go* quotative than others who use less historical present because this is an environment that strongly favors *go* (Lee, 2020:57). In such a case, a higher rate of *go* may reflect the use of historical present more than a general favoring of *go*.

In order to control for the impact of predicting factors (such as tense for quotatives), we use by-speaker estimates from independent regression models run for each variable as a measure of "how much that individual's trend diverges from the predicted trends set forth in the statistical model" (Drager & Hay, 2012:62). These models are based on the patterning already identified in our prior studies in the literature referred to in the previous section on the linguistic variables but exclude the social predictors to ensure that the clusters are formed on the basis of the linguistic conditioning only. For continuous variables (FLEECE, FACE, -*er*) we used the lmer function, and for categorical variables (*the* + Vowel, -*ing, there's*, modals of obligation, quotatives) we used the glmer function (in the *lme4* package; Bates et al., 2019), including speaker as a random intercept in all models. (See Supplementary Material for model summaries for all variables.)

We then extracted the speaker random intercepts from each model and used these as the dependent variable in the divisive clustering analysis. These speaker random intercept scores are on different scales, coming from both continuous and categorical variables, and thus we scaled them in order to make them comparable (by centring each score and then dividing it by the standard deviation for each set to produce a *z*-score, using the scale function in *R*) (cf. Haddican et al., 2022:516).

The reliability of the random intercepts in capturing speaker behavior is supported by a strong correlation with raw values. Figure 2 visualizes this correlation for one continuous (FLEECE) and one categorical variable (-*ing*), in the 1970s (top panes) and 2010s (bottom panes). The x-axis represents the scaled speaker random intercepts, and the y-axis the raw values (normalized F2 minus F1 for FLEECE and rate of velar for -*ing*). Farther to the right on the x-axis represents a positive intercept (and thus a favoring of a higher and fronter FLEECE, or of velar over alveolar -*ing*), and higher on the y-axis represents a higher and fronter FLEECE as measured in Hz, or a higher rate of velar -*ing*. As can be seen, the values are strongly correlated, indicating that the overall picture would be similar regardless of the measure used. In taking account of the linguistic conditioning, however, the random intercepts distinguish the speakers better. For example, for -*ing* there are several categorically velar speakers who differ according to the random intercept, such as Jade (a 1970s Anglo Adult) and Lisa (a 1970s Italian Teen), who are both categorically velar, but whose scaled random intercepts differ (1.947 for Jade and 0.838 for Lisa). Examination of their data reveals that Jade produces proportionally more tokens preceding a coronal (e.g., *going to school, something like that*), an environment that is highly propitious to the alveolar realization (Travis et al., 2023:448). The high intercept reflects her strong favoring of the velar variant overall, even in contexts where the alveolar variant is favored. In this way, we consider the random intercept to be a more reliable measure than the raw values.

We obtain speaker random intercepts for each of the variables included, and thus for the 1970 analysis, we obtain seven measures for each speaker, and for the 2010s nine measures (see Table 2); this is the data on which we perform the clustering.
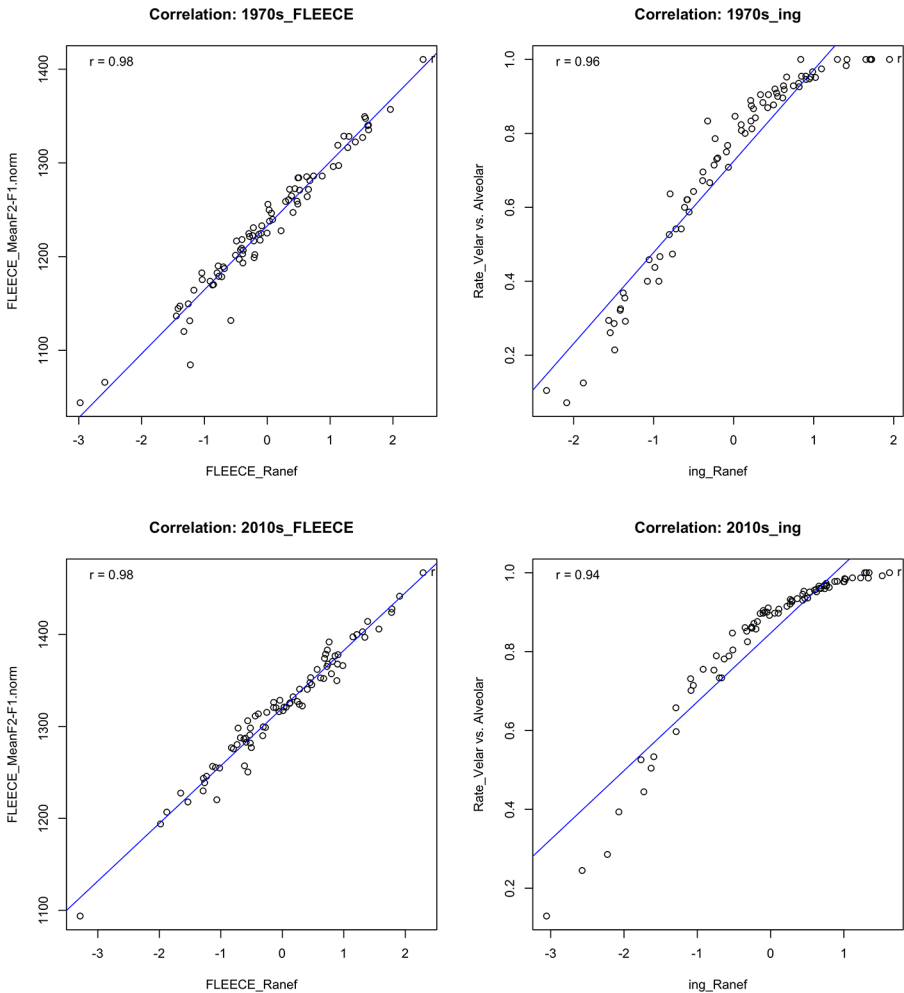
Figure 2. Correlations between the by-speaker random effects and raw values.

Identifying the social makeup of the linguistically defined clusters

The divisive clustering technique we use (DIANA) is particularly well suited for identifying potentially intersecting groups for a number of reasons, as outlined by Haddican and colleagues (2021:149, 2022:516). This is a top-down clustering approach, which starts with the entire dataset (here, all speakers, with the corresponding random intercept scores for each variable) in a single cluster and recursively splits the data (the speakers) into smaller clusters based on their similarity across these variables until reaching individual speakers. Clusters are determined based on the overall similarity between sets of speakers across all variables, measured in Euclidean distance (the square root of the sum of differences). Unlike other clustering algorithms, the number

of clusters is not pre-defined but is determined on the basis of the data, which is important here, as how many clusters there are in the data is one of the questions that we are asking. The hierarchical structure of this clustering method identifies splits in the dataset at different levels, thus bringing to the fore intersections across social categories, and the results are visualized in dendrograms, which facilitates the identification of groupings at different levels of the hierarchy.

Clustering algorithms will produce clusters because that is their function, but this does not mean that the clusters are meaningful. The validity of the clusters formed here can be demonstrated both socially and linguistically. To interpret their social makeup, we examine the distribution according to the age, gender, social class, and ethnicity of cluster members through visualizations in the form of dendrograms and scatterplots, and to interpret their linguistic makeup, we assess the relative contribution of the different linguistic variables to the clusters via random forest models.

The social makeup of the clusters

1970s

Figure 3 displays the dendrogram that is output from the cluster analysis of the 1970s data. The points along the bottom of the dendrogram represent individual speakers, clustered according to their linguistic behavior across these variables. The length of each branch indicates the relative similarity between speakers and clusters, thus the higher up the tree, the greater the difference between the groups. These clusters are entirely agnostic as to their social makeup, but we have visualized age (in shape) and ethnicity (in color) to give an initial sense of the patterns. It is by considering the social groupings that come out of the linguistic patterning that we can track both change over time and how that is taken up according to gender, class, and ethnicity.

This model initially identifies three main clusters. Within Group 1, the largest cluster (57 members), we observe two subgroups that are distinct socially (labelled here Group 1.1 and 1.2, with 41 and 16 members, respectively), which will be examined separately. Groups 2 and 3 each have smaller numbers of members (18 and six, respectively), and we will not break them down further.

The order in which the clusters at different levels of the hierarchy are placed is arbitrary (that is, Group 1.2 is no more similar linguistically to Group 2 than Group 1.1 is), but the social patterns that emerge suggest that they are distinguished first and foremost by age, in that most of the Adults (empty squares) fall into Groups 1.1 and 1.2, while Group 2 has only two adults, and Group 3 has none. To better visualize the social composition of these groups, we reproduce the clusters from the dendrogram in scatter plots (Figure 4), re-ordering them in a socially relevant way, with the cluster containing the highest proportion of the oldest age group (1970s Adults) on the left, and the youngest age group (1970s Teens) on the right, thus ordering the panes in a way that captures the apparent time change. We further distinguish the members of each cluster according to age, gender, ethnicity, and social class.
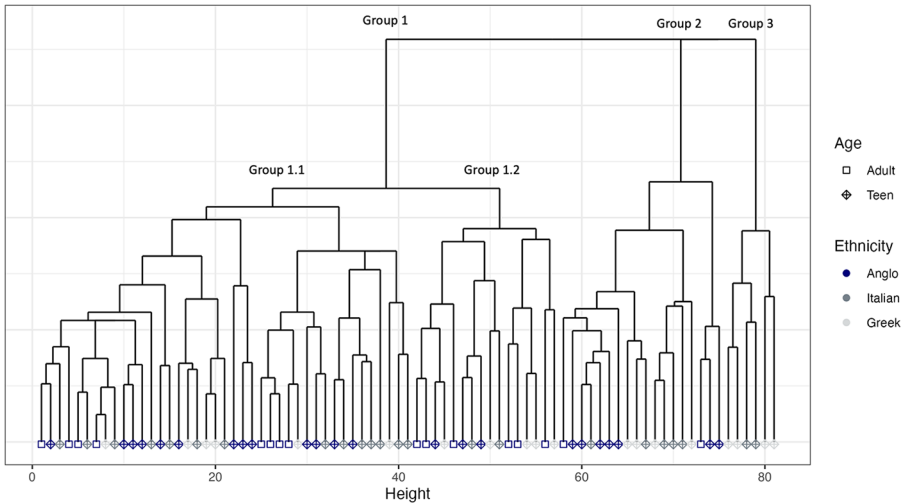
Figure 3. Dendrogram of a DIANA cluster analysis: 1970s (81 speakers, 7 variables: FLEECE; FACE; *-er; -ing; there's; have got to;* quotative *go*).
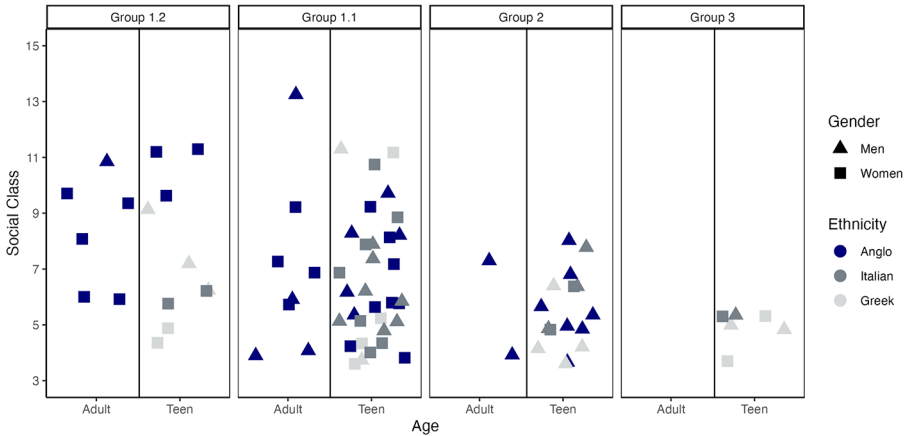

Figure 4. Scatter plots for social groupings: 1970s.

In Figure 4, each of the four panes represents one of the clusters identified. We have split each pane according to age group (Adults on the left, Teens on the right), and the y-axis shows social class, with a higher score corresponding to higher social class. Men are marked by triangles and women by squares, and ethnicity is coded by color.

Here, it can more easily be seen that almost all the Adults occur in the first two panes (Groups 1.2 and 1.1, the two groups making up the first cluster in Figure 3). The first pane (Group 1.2) includes both Adults and Teens and spans a wide range for social class, but it is distinguished by gender, comprising more women than men. For

ethnicity (relevant for the Teens, as the Adults are all Anglos), we find that the Anglos, Greeks, and Italians are all represented. The second pane (Group 1.1) is distinguished by having proportionally more Teens and proportionally more men than the first pane; again, it covers the range of social class, and for the Teens the three ethnic groups are represented.

The third pane (Group 2) comprises only two Adults, and only two women, and overall corresponds to lower social class than the previous two groups. We still see no ethnically driven clustering, with Anglo, Greek, and Italian Teens occurring together.

It is in the fourth pane (Group 3) that we do observe some ethnic differentiation. This small cluster, consisting of only six participants, is made up entirely of Italian and Greek Teens with low social class scores. This marks the first time we see ethnicity playing a defining role in the clustering of linguistic behavior, and it does so in conjunction with low socio-economic status.

For the 1970s, then, the clusters formed on the basis of linguistic behavior are primarily distinguished by age, reflecting change in apparent time. The distribution of the social groups across the clusters suggests that the changes are led by the working class and by men, seen in the gradual shift down in social class as we move from left to right across the panes, and in the shift from predominantly women to gradually more men for all but the final pane, which has both men and women. Crucially, ethnicity plays a minor role here, emerging as a distinguishing factor only in a small cluster and with participants of lower social class.

2010s

Figure 5 presents the dendrogram output from the 2010s analysis. Here, the initial split delineates two primary groups (Group 1 and Group 2), quite markedly distinguished by
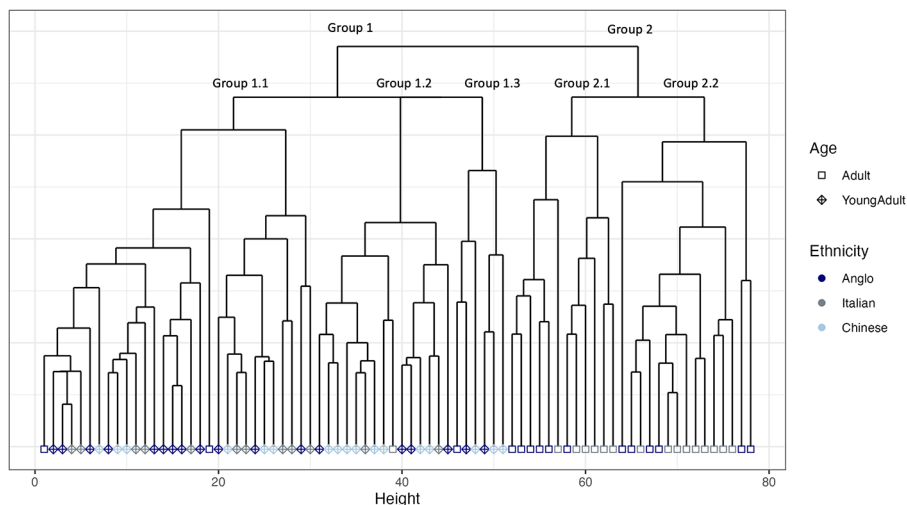

Figure 5. Dendrogram of a DIANA cluster analysis: 2010s (78 speakers, 9 variables: FLEECE; FACE; *-er; the* + Vowel; *-ing; there's; have got to; need to;* quotative *be like*).
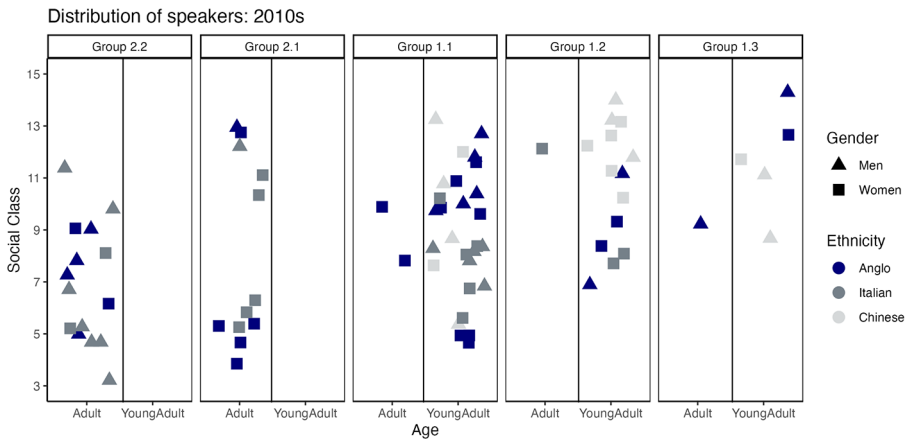
Figure 6. Scatter plots for social groupings: 2010s.

age, with almost all Adults (empty squares) and no Young Adults occurring in Group 2. Group 1 is divided into three subgroups (1.1, 30 members; 1.2, 15 members; 1.3, six members), and Group 2 into two subgroups (2.1, 12 members, and 2.2, 15 members). We consider the social makeup of these five groups presented in the scatterplot in Figure 6, in which (as for Figure 4) we have re-ordered the clusters to capture the progression of change, presenting the two sub-groups (made up of Adults) in Group 2 on the left, and then the three sub-groups (made up of Young Adults) in Group 1 on the right.

The two groups where almost all of the older speakers cluster, the first two panes in Figure 6, are distinguished primarily by gender, with the first pane (Group 2.2) having proportionally more men and the second (Group 2.1) having proportionally more women. Both groups cover a broad social class distribution and, importantly, Anglos and Italians are similarly distributed in both. Thus, the separate grouping of some Italian and Greek Teens with lower social class scores that we saw for the 1970s has dissipated for this age group.

The next three panes (Groups 1.1, 1.2, and 1.3) represent the younger speakers, who are not differentiated by gender but are differentiated by social class. Group 1.1 has a similarly broad class distribution as the two Adult groups, while in Groups 1.2 and 1.3, there are progressively fewer speakers with low class scores. There is a corresponding shift in the ethnic makeup of these clusters; Anglo Young Adults are distributed across the three clusters, though the majority are situated in 1.1 and 1.2, with the Italian Young Adults, whereas the majority of the Chinese Young Adults are situated in Groups 1.2 and 1.3 (11 out of the 17 Chinese Australians occur in these groups, as opposed to six out of the 18 Anglo Australian Young Adults and two out of the 12 Italian Australian Young Adults, both in 1.2).

For the 2010s, we thus see clustering according to age, as was also seen for the 1970s, with gender playing a role in further differentiating the age groups. For ethnicity, the Anglo and Italian Australians pattern similarly, and though the Chinese Australians

are distinct, this is again in conjunction with a class distinction, this time patterning with higher social class.

Summary: social class, ethnicity, and ethnic orientation

In sum, in neither time period do we get clustering according to ethnicity alone, but rather the ethnic patterning that we observe is tied to social class, specifically to lower social class for the 1970s Greek and Italian Teens, and higher social class for the 2010s Chinese Young Adults. It is also of note that the nature of the set of changes in each time period is distinct: in the 1970s, the change is led by men and the working class, whereas in the 2010s, it is led by women and the middle class. This may be due to the specific set of variables and ethnic groups considered here, rather than an overall shift in the leaders of language change over time, a question which we leave for further analysis. What is of more interest to us here is that the intersection between class and ethnicity remains, independent of the leaders of change: the ethnic groups pattern with their corresponding class affiliations.

Though here we have highlighted class distinctions, other work has highlighted distinctions in the degree to which members of an ethnic community orient to their ethnic heritage, and such "ethnic orientation" has been reported to impact linguistic behavior (e.g., Hoffman & Walker, 2010:59). We have also considered ethnic orientation in Sydney Speaks, compiling information from the sociolinguistic interview and from the demographic questionnaire about the makeup of participants' social network, language ability and use, visits to and connections with the countries of their parents, engagement with cultural traditions, and so on (see Travis, 2024:171). However, ethnic orientation has typically been found to have no effect (e.g., Gan, 2024:189; Lee, 2020:94-95; Sheard, 2023:177-178, 212, 285), and in the few cases where an effect has been identified, this occurs with variables that are highly constrained by social class (lengthening of word-final -er for the Greek 1970s Teens; Sheard, 2023:249; and use of schwa versus FLEECE in prevocalic *the* for the Chinese 2010s Young Adults; Gan, 2024:189). Perhaps even more telling is the fact that, for these groups, ethnic orientation correlates with social class: for the 1970s Greek Teens, higher ethnic orientation correlates with lower social class (and longer -er), and for the 2010s Chinese Young Adults, higher ethnic orientation correlates with higher social class (and higher use of schwa) (Gan, 2024:192; Travis, 2021). Thus, ethnic orientation, just like ethnicity, is tied to social class.

The relative contribution of the linguistic variables to the clustering

As a final consideration, we turn to the linguistic makeup of the clusters. The linguistic variables that form the basis of the clustering are for the most part structurally independent (an exception may be FLEECE and FACE, and the two sets of modals of obligation), ensuring that the results are not restricted to a specific aspect of language. However, we might ask whether the different linguistic features equally impact the clustering or whether the variation is driven more by any particular subset of variables. To determine the relative contribution of each variable to the clusters, we applied random forest analysis, which ranks the variables according to their relative importance in predicting an outcome. Random forests are a tree-based modeling method that use

recursive partitioning, assessing the likelihood of each variant within a specified set of predictors and making binary splits in progressively smaller data subsets until no further significant splits are found. Random forests average the results across multiple conditional inference trees produced in this way to determine the overall importance of each predictor, based on randomly generated subsets of the data (Tagliamonte & Baayen, 2012:159-160). For the random forests produced here, the dependent variable was the cluster in which the speaker was placed (four for the 1970s and five for the 2010s), and the independent variables were the speaker random intercepts for each of the linguistic variables in the analysis (seven for the 1970s and nine for the 2010s). We used the *ranger* package in *R* (Wright & Ziegler, 2017) and set it to produce 500 decision trees and to consider three linguistic variables at each split.

Figure 7 presents the results of the random forest analyses, displaying the importance ranking of the variables in determining the clusters for the 1970s data on the left and the 2010s data on the right.[5] The x-axis shows the weights of the variables, which reflect their importance, and the y-axis lists each variable in order of its ranking. All variables have scores greater than zero, indicating that each plays a role in shaping speaker groupings, but the wide range shows that they do this to varying degrees. The range is larger for the 1970s (2.7 to 10.6) than for the 2010s (2.9 to 7.8), indicating that the contribution of the different variables is more similar in the 2010s than it is in the 1970s data.

The ranking of the variables in each model is slightly different: for the 1970s, FLEECE contributes most to the speaker groupings, followed by *there's* and *-ing*, and then (contributing progressively less) *-er*, quotative *go*, FACE, and modal *have got to*. For the 2010s, quotatives have the largest impact, followed by FLEECE, *(have) got to,* and FACE, then (with progressively smaller contributions) *there's, -ing, the* + Vowel, and finally, *-er* and modal *need to*.

It has been proposed that phonetic variables are more widely available for social indexing than morphosyntactic and discourse variables (e.g., Cheshire, 2005:479). If that applied here, the phonetic variables should be ranked higher in importance than the grammatical or discourse variables, but this is not the case. For the 1970s, FLEECE is the most important but FACE is among the least important, and for the 2010s, FLEECE again has a higher ranking than FACE, but the strongest effect is with the discourse-pragmatic variable, *be like*. The primary factor that would appear to be driving the relative importance is the timing of the change. For example, though FLEECE and FACE have both been undergoing change for some time, FLEECE began from the 1970s Teens, whereas FACE began later; thus, FLEECE distinguishes the 1970s groups more than FACE does. In the 2010s, we see the meteoric rise of *be like* among the Young Adults (with a rate of use of just 10% for the Adults and 70% for the Young Adults), rendering this the most important variable. Likewise, *there's, -ing, -er,* and quotative *go* are variables that show a marked shift from the 1970s Adults to the Teens, whereas the variation between *(have) got to/have to* is more similar across the two groups. For the 2010s, there is a drop in use of *(have) got to* versus *have to* from the Adults to the Young Adults, in contrast to *there's, -ing,* and *need to,* for which there is less apparent time change. The low importance of *-er* and *the* + Vowel for the 2010s is surprising, given that they do undergo change in this time period. It may be that the overall change is smaller for these forms, thus lessening their importance relative to the other variables.
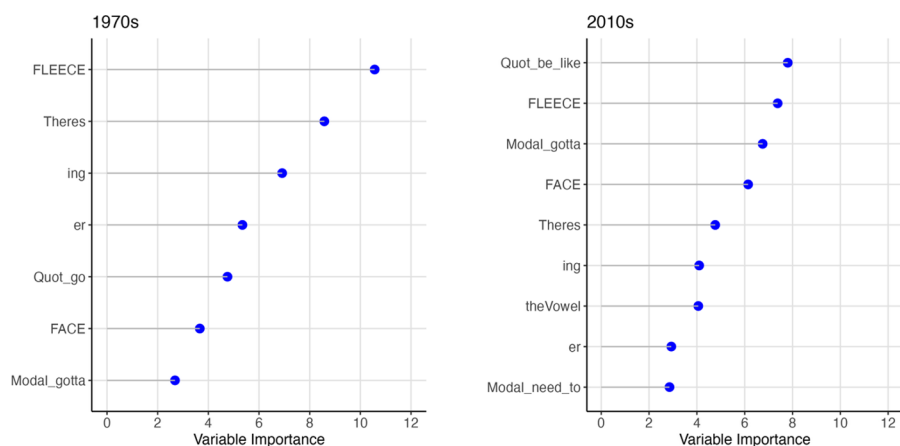
Figure 7. Two independent analyses testing the importance ranking for linguistic variables in the 1970s (L) and 2010s (R).

(Regarding the timing of these changes, see section on the linguistic variables above and references therein.)

These analyses show, then, that first, in each time period, all variables contribute to the clustering observed, and second, the main factor that determines their relative importance is the difference between the two age groups within each time period, reflecting the changes that have occurred. This is consistent with the social distribution of the clusters, for which age was the primary factor in the first division made, thus confirming that these clusters are not random, but meaningfully capture social and linguistic patterning in the data.

Conclusion

The results presented here align with observations in prior Sydney Speaks work that interpreting ethnic variation in the Sydney speech community requires consideration of social class. This intersection is likely a reflex of multiple things, including, socially, the network in which people move, and linguistically, the prestige attached to different variants (cf. Travis et al., 2023:461-463). The current study presents a further level of evidence for an ethnicity-social class intersection in treating the linguistic behavior as the organizing factor for the social groups, rather than comparing the linguistic behavior of pre-defined social groups. Doing this demonstrates that this ethnicity-social class intersection holds across different time points and for different ethnic groups. This is hardly surprising, given what we know about intersections between other social variables, but it remains the case that many studies overlook this, considering ethnicity in isolation.

There are several risks with prioritizing ethnicity in this way. From a linguistic perspective, it may give an incorrect impression of the patterns of variation observed and impede our understanding of the nature of language change and variation across society. There are also social implications, as considering ethnicity independently

from other social factors may be taken to suggest that the linguistic behavior of migrant groups (or ethnic minorities more generally) can be explained by their ethnic background alone, and thereby that they are solely defined by their ethnicity, thus potentially "idealising or glamorising a rather undifferentiated monolithic perspective on the international migrant" (Britain, 2022:332).

In contrast to such an essentialist approach, the intersectionality we put forward here recognizes that "no one category […] is sufficient to account for individual experience or behavior," just as has been recognized for gender (Levon, 2015:295), and thus better captures the multidimensionality of social groups. Such an approach brings to the fore the fact that someone like Lathan (from Example 1 above) is not just a Chinese Australian who speaks some Cantonese and Mandarin, has a very Asian social network, and has visited Hong Kong. He is also a 24-year-old man who attended a selective high school in Sydney and is now completing an Arts/Law degree and working as a paralegal. Categorizing him solely as a Chinese Australian ignores these other factors relating to social class that exist alongside his Chinese identity. Here, we have taken them into account and find that the social class measures are closely tied in with ethnicity (and with ethnic orientation). By considering all these factors together, we can gain a better understanding of what it means to be a "Chinese Australian" or a "Greek Australian," which in turn allows for a more informed interpretation of the linguistic variation observed in ethnically diverse settings, and in particular in urban environments like Sydney.

Notes

1. We use capitalization for age groups when used in a technical sense, as defined for this project.

2. For example, *be like* as a quotative and *need to* as a modal of obligation are largely absent for the 1970s Teens but present for the 2010s Adults (Lee, 2020:34; Travis & Torres Cacoullos, 2023:367).

3. In parentheses are the corpus name, the speaker social code (ethnicity, age group, gender) and speaker number, and time stamps of the beginning and end of the excerpt. All names given are pseudonyms. Backchannels from the interlocutor have been deleted in this example for readability purposes. Transcription conventions: each line is an intonation unit, and the following conventions are used: . final intonation contour; , continuing intonation contour; -- truncated intonation contour; = lengthened syllable; .. short pause (0.2 secs); …(1.0) 1 second pause (cf. Du Bois et al., 1993).

4. We are grateful to Bill Haddican for making the *R* code for the DIANA clustering available to us.

5. The 1970s model achieves a classification accuracy of 73.3%, which, though not significantly higher than the baseline accuracy of 53.3% (obtained by always selecting the most frequent group), does approach significance ($p = 0.096$), indicating reasonable performance in distinguishing between the four levels of the dependent variable. The 2010s model achieves a classification accuracy of 93.3%, significantly higher than the baseline of 40.0% ($p < 0.001$), indicating excellent discrimination between the five groups.

References

Bates, Douglas, Mächler, Martin, Bolker, Ben, Walker, Steve, Christensen, Rune Haubo Bojesen, Singmann, Henrik, Dai, Bin, Scheipl, Fabian, Grothendieck, Gabor, Green, Peter, & Fox, John. (2019). lme4: Linear mixed-effects models using 'Eigen' and S4. R package, version 1.1-121.

Benheim, Jaime & D'Onofrio, Annette. (2024). Local features, local meanings: Language ideologies and place-linked vocalic variation among Jewish Chicagoans. *Language in Society* 53(1):129-155. https://doi.org/10.1017/S0047404522000690.

Boberg, Charles. (2004). Ethnic patterns in the phonetics of Montreal English. *Journal of Sociolinguistics* 8(4):538-568. https://doi.org/10.1111/j.1467-9841.2004.00273.x.

Boersma, Paul & Weenink, David. (2024). *Praat: Doing phonetics by computer [Computer program] Version 6.1.16*: retrieved 1 December 2024 from http://www.praat.org/.

Britain, David. (2022). Ethnolects, multiethnolects and urban contact dialects: Looking forward, looking back, looking around. In P. Kerswill & H. Wiese (eds.), *Urban contact dialects and language change: Insights from the Global North and South*. London: Routledge. 325-336.

Cheshire, Jenny. (2005). Syntactic variation and beyond: Gender and social class variation in the use of discourse-new markers. *Journal of Sociolinguistics* 9(4):479-508. https://doi.org/10.1111/j.1360-6441.2005.00303.x.

Cheshire, Jenny, Kerswill, Paul, Fox, Sue, & Torgersen, Eivind. (2011). Contact, the feature pool and the speech community: The emergence of Multicultural London English. *Journal of Sociolinguistics* 15(2):151-196. https://doi.org/10.1111/j.1467-9841.2011.00478.x.

Clothier, Josh. (2019). A sociophonetic analysis of /l/ darkness and Lebanese Australian ethnic identity in Australian English. In S. Calhoun, P. Escudero, M. Tabain, & P. Warren (eds.), *Proceedings of the 19th International Congress of Phonetic Sciences (ICPhS)*. Melbourne, Australia: ASSTA. 1888-1892.

Clothier, Josh & Loakes, Deborah. (2018). Coronal stop VOT in Australian English: Lebanese Australians and mainstream Australian English speakers. In J. Epps, J. Wolfe, J. Smith, & C. Jones (eds.), *Proceedings of the 17th Australasian International Conference on Speech Science and Technology*. Sydney: ASSTA. 13-16.

Clyne, Michael, Eisikovits, Edina, & Tollfree, Laura. (2001). Ethnic varieties of Australian English. In D. Blair & P. Collins (eds.), *English in Australia*. Amsterdam: John Benjamins. 223-238.

Cox, Felicity & Palethorpe, Sallyanne. (2011). Timing differences in the VC rhythm of standard Australian English and Lebanese Australian English. In W.S. Lee & E. Zee (eds.), *Proceedings of the 17th International Congress of Phonetic Sciences (ICPhS)*. Hong Kong: City University of Hong Kong. 528-531.

Cox, Felicity & Penney, Joshua. (2024). Multicultural Australian English: The new voice of Sydney. *Australian Journal of Linguistics* 44(2-3):200-219. https://doi.org/10.1080/07268602.2024.2380680.

Cox, Felicity, Penney, Joshua, & Palethorpe, Sallyanne. (2023). Fifty years of change to prevocalic definite article allomorphy in Australian English. *Journal of the International Phonetic Association* 53(3):804-834. https://doi.org/10.1017/S002510032200007X.

Drager, Katie & Hay, Jennifer. (2012). Exploiting random intercepts: Two case studies in sociophonetics. *Language Variation and Change* 24(1):59-78. https://doi.org/10.1017/S0954394512000014.

Du Bois, John W., Schuetze-Coburn, Stephan, Cumming, Susanna, & Paolino, Danae. (1993). Outline of discourse transcription. In J. Edwards & M. Lampert (eds.), *Talking data: Transcription and coding in discourse*. Hillsdale: Lawrence Erlbaum Associates. 45-89.

Eckert, Penelope. (2008). Where do ethnolects stop? *International Journal of Bilingualism* 12(1-2):25-42. https://doi.org/10.1177/13670069080120010301.

Fox, Susan. (2015). *The New Cockney: New ethnicities and adolescent speech in the traditional East End of London*. New York: Palgrave Macmillan.

Fromont, Robert & Hay, Jennifer. (2012). LaBB-CAT: An annotation store. *Proceedings of the Australasian Language Technology Workshop*:113-117. https://aclanthology.org/U12-1015

Gan, Qiao. (2024). *Ethnic variation in its social context: Evidence from the English of Chinese Australians*. PhD thesis, School of Literature, Languages and Linguistics, Australian National University. http://hdl.handle.net/1885/311094

Gan, Qiao & Travis, Catherine E. (2022). Ethnicity and social class in pre-vocalic *the* in Australian English. In R. Billington (ed.), *Proceedings of the 18th Australasian International Conference on Speech Science and Technology*: ASSTA. 56-60. https://sst2022.files.wordpress.com/2022/12/qiao-travis-2022-ethnicity-and-social-class-in-pre-vocalic-the-in-australian-english.pdf

Grama, James, Travis, Catherine E., & Gonzalez, Simon. (2020). Ethnolectal and community change ov(er) time: Word-final (er) in Australian English. *Australian Journal of Linguistics* 40(3):346-368. https://doi.org/10.1080/07268602.2020.1823818.

Grama, James, Travis, Catherine E., & Gonzalez, Simon. (2021). Ethnic variation in real time: Change in Australian English diphthongs. In H. Van de Velde, N. H. Hilton, & R. Knooihuizen (eds.), *Studies in Language Variation*. Amsterdam: John Benjamins. 292-314. https://www.jbe-platform.com/content/books/9789027259820-silv.25.13gra

Gross, Johan, Boyd, Sally, Leinonen, Therese, & Walker, James A. (2016). A tale of two cities (and one vowel): Sociolinguistic variation in Swedish. *Language Variation and Change* 28(2):225-247. https://doi.org/10.1017/S0954394516000065.

Guy, Gregory R., Horvath, Barbara M., Vonwiller, Julia, Daisley, Elaine, & Rogers, Inge. (1986). An intonational change in progress in Australian English. *Language in Society* 15(1):23-52. https://doi.org/10.1017/S0047404500011635.

Haddican, Bill, Cutler, Cecelia, Newman, Michael, & Tortora, Christina. (2022). Cross-speaker covariation across six vocalic changes in New York City English. *American Speech* 97(4):512-542. https://doi.org/10.1215/00031283-9616153.

Haddican, Bill, Newman, Michael, Cutler, Cecelia, & Tortora, Christina. (2021). Aspects of change in New York City English short-a. *Language Variation and Change* 33(2):135-163. https://doi.org/10.1017/S0954394521000120.

Hay, Jen, Walker, Abby, McKenzie, Janye, & Nielsen, Daniel. (2012). The changing realisation of 'the' before vowels in New Zealand English. *New Zealand English Journal* 26:23-32. https://doi.org/10.26021/4

Hoffman, Michol F. (2010). The role of social factors in the Canadian Vowel Shift: Evidence from Toronto. *American Speech* 85(2):121-140. https://doi.org/10.1215/00031283-2010-007.

Hoffman, Michol F. & Walker, James A. (2010). Ethnolects and the city: Ethnic orientation and linguistic variation in Toronto English. *Language Variation and Change* 22(1):37-67. https://doi.org/10.1017/S0954394509990238.

Horvath, Barbara. (1985). *Variation in Australian English: The sociolects of Sydney*. Cambridge: Cambridge University Press.

Horvath, Barbara & Sankoff, David. (1987). Delimiting the Sydney speech community. *Language in Society* 16(2):179-204. https://doi.org/10.1017/S0047404500012252.

Jupp, James. (2001). *The Australian People: An encyclopedia of the nation, its people and their origins* (2nd ed.). Cambridge: Cambridge University Press.

Kiesling, Scott Fabius. (2005). Variation, stance and style: Word-final -er, high rising tone, and ethnicity in Australian English. *English World-Wide* 26(1):1-42. https://doi.org/10.1075/eww.26.1.02kie.

King, Sharese. (2021). Rethinking race and place: The role of persona in sound change reversal. *Journal of Sociolinguistics* 25(2):159-178. https://doi.org/10.1111/josl.12454.

Labov, William. (1963). The social motivation of a sound change. *Word* 19:273-309. https://doi.org/10.1080/00437956.1963.11659799

Labov, William. (1966). *The social stratification of English in New York City*. Washington, DC: Center for Applied Linguistics.

Labov, William. (1972). Some principles of linguistic methodology. *Language in Society* 1(1):97-120. https://doi.org/10.1017/S0047404500006576.

Labov, William. (1990). The intersection of sex and social class in the course of linguistic change. *Language Variation and Change* 2(2):205-254. https://doi.org/10.1017/S0954394500000338.

Laferriere, Martha. (1979). Ethnicity in phonological variation and change. *Language* 55(3):603-617. https://doi.org/10.2307/413319.

Lausberg, Hedda & Sloetjes, Han. (2009). Coding gestural behavior with the NEUROGES-ELAN system. Behavior Research Methods, Instruments, & Computers (Max Planck Institute for Psycholinguistics, The Language Archive, Nijmegen, The Netherlands). 41(3):841-849. http://tla.mpi.nl/tools/tla-tools/elan/, https://doi.org/10.3758/BRM.41.3.841.

Lee, Esther. (2020). *Quotatives over time: A study in ethnic variation*. Honours thesis, School of Literature, Languages and Linguistics, Australian National University. http://hdl.handle.net/1885/298816

Levon, Erez. (2015). Integrating intersectionality in language, gender, and sexuality research. *Language and Linguistics Compass* 9(7):295-308. https://doi.org/10.1111/lnc3.12147.

McMillan, Julie, Beavis, Adrian, & Jones, Frank L. (2009). The AUSEI06: A new socioeconomic index for Australia. *Journal of Sociology* 45(2):123-149. https://doi.org/10.1177/1440783309103.

Norton, Andrew. (2023). Mapping Australian higher education. 6th, Methods: Centre for Social Research and. ANU. https://csrm.cass.anu.edu.au/research/publications/mapping-australian-higher-education-2023. (accessed January 2025)

Purser, Benjamin, Grama, James, & Travis, Catherine E. (2020). Australian English over time: Using sociolinguistic analysis to inform dialect coaching. *Voice and Speech Review* 14(3):269-291. https://doi.org/10.1080/23268263.2020.1750791.

Ricatti, Francesco. (2018). *Italians in Australia: History, memory, identity*. London: Palgrave Macmillan.

Sharma, Devyani. (2020). Prestige factors in contact-induced grammatical change. In K. V. Beaman, I. Buchstaller, S. Fox, & J. A. Walker (eds.), *Advancing Socio-grammatical Variation and Change: In Honour of Jenny Cheshire*. New York: Routledge. 55-72.

Sheard, Elena. (2022). Longevity of an ethnolectal marker in Australian English: Word-final (er) and the Greek-Australian community. In: R. Billington (ed.), *Proceedings of the 18th Australasian International Conference on Speech Science and Technology*: ASSTA. 51-55. https://sst2022.files.wordpress.com/2022/12/sheard-2022-longevity-of-an-ethnolectal-marker-in-australian-english-word-final-er-and-the-greek-australian-community.pdf

Sheard, Elena. (2023). *Explaining language change over the lifespan: A panel and trend analysis of Australian English*. PhD thesis, School of Literature, Languages and Linguistics, Australian National University. http://hdl.handle.net/1885/292110

Tagliamonte, Sali & Baayen, R. Harald. (2012). Models, forests, and trees of York English: *Was/were* variation as a case study for statistical practice. *Language Variation and Change* 24:135-178. https://doi.org/10.1017/S0954394512000129

Torres Cacoullos, Rena & Travis, Catherine E. (2018). *Bilingualism in the community: Code-switching and grammars in contact*. Cambridge: Cambridge University Press. https://doi.org/10.1017/9781108235259

Travis, Catherine E. (2021). Ethnolectal variation in real time: Ethnicity, gender and class in Sydney, Australia. Paper presented at the *Abralin Ao Vivo Series*. July 14. https://www.youtube.com/watch?v=DLEI67i9E-w.

Travis, Catherine E. (2024). Sydney Speaks Corpus: An overview. *Australian Journal of Linguistics* 44(2-3):163-181. https://doi.org/10.1080/07268602.2024.2386387.

Travis, Catherine E., Grama, James, Gonzalez, Simon, Purser, Benjamin, & Johnstone, Cale. (2023). *Sydney Speaks Corpus*. Canberra: ARC Centre of Excellence for the Dynamics of Language, Australian National University. https://doi.org/10.25911/m03c-yz22.

Travis, Catherine E., Grama, James, & Purser, Benjamin. (2023). Stability and change in (ing): Ethnic and grammatical variation over time in Australian English. *English World-Wide* 44(3):435-469. https://doi.org/10.1075/eww.22043.tra.

Travis, Catherine E. & Torres Cacoullos, Rena. (2023). Form and function covariation: Obligation modals in Australian English. *Language Variation and Change* 35(3):351-377. https://doi.org/10.1017/S0954394523000200.

van Meel, Linda, Hinskens, Frans, & van Hout, Roeland. (2014). Variation in the realization of /ɛi/ by Dutch youngsters: From local urban dialects to emerging ethnolects? *Dialectologia Et Geolinguistica* 22(1):46-74. https://doi.org/10.1515/dialect-2014-0004.

Walker, James A. (2024). Soundin(g(k)) ethnic in Toronto. *English World-Wide* 45(3):283-310. https://doi.org/10.1075/eww.24014.wal.

Wong, Amy Wing-Mei & Hall-Lew, Lauren. (2014). Regional variability and ethnic identity: Chinese Americans in New York City and San Francisco. *Language and Communication* 35:27-42. https://doi.org/10.1016/j.langcom.2013.11.003

Wright, Marvin N. & Ziegler, Andreas. (2017). ranger: A fast implementation of Random Forests for high dimensional data in C++ and R. *Journal of Statistical Software* 77(1):1-17. https://doi.org/10.18637/jss.v077.i01.