



Open Research Online

Citation

Gupta, Ayan-Yue and Gupta, Suman (2024). Catchwords: An Essay on Research Methods. Key Words: A Journal of Cultural Materialism (In Press).

URL

<https://oro.open.ac.uk/100561/>

License

(CC-BY-NC-ND 4.0) Creative Commons: Attribution-Noncommercial-No Derivative Works 4.0

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Policy

This document has been downloaded from Open Research Online, The Open University's repository of research publications. This version is being made available in accordance with Open Research Online policies available from [Open Research Online \(ORO\) Policies](#)

Versions

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding

Catchwords: An Essay on Research Methods

Ayan-Yue Gupta (University of Milan, Italy) and Suman Gupta (The Open University, UK)

Keywords and catchwords

The study of collective life¹ has often taken recourse to ‘keywords’ as formulated and examined influentially by Raymond Williams (1976).² The journal in which this essay appears is itself testimony to the continuing influence of Williams’s formulation. In his view, keywords have two overlapping aspects. Since they are in general usage and are widely regarded as ‘significant’, they enable active engagement with or participation in collective life. At the same time, such words are also ‘significant’ because they concentrate attempts to understand and explain collective life, and therefore serve to conceptualise collectivity. In Williams’s words:

they are significant, binding words in certain activities and their interpretation; they are significant, indicative words in certain forms of thought. Certain uses bound together certain ways of seeing culture and society, not least in these two most general words. Certain other uses seemed to me to open up issues and problems, in the same general area, of which we all needed to be very much more conscious.³

Focusing on a selection of these significant terms, with historical shifts and current cross-connections noted, could be taken as articulating collective life – culture and society – from its midst.

‘Culture and society’ are keywords that are used to refer to collective life itself. By focusing on a selection of keywords, Williams sought to articulate society and culture without subscribing to narrow definitions of either. As a keyword, ‘culture’ interested him particularly. He pegged it as ‘used for important concepts in several distinct intellectual disciplines and in several distinct and incompatible systems of thought.’⁴ Approaching culture via keywords underscores a salutary disinvestment from normative ascriptions in the sense of ‘high’ or ‘low’ culture or of being ‘cultured.’ The apprehension of culture in this politically and sociologically expedient, unprejudiced, non-normative manner was the thrust of Williams’s larger project⁵; focusing on keywords seems to methodologically ground that thrust. Unsurprisingly, apart from being frequently employed and cited in scholarly writings, updated lists of keywords following or building upon Williams’s method have been produced regularly since.⁶

However, Williams’s method for foregrounding and examining keywords, or for discerning and attributing significance to them, came with certain implicitly undermining limitations. Alluring though the method and its rationale have proven, these limitations have invariably rendered studies along those lines a bit shaky. This is the case for Williams’s own treatment of keywords and also for subsequent treatments which have tried to finetune his method. In this essay, we suggest that the purposes that Williams’s keywords served are better addressed now by focusing on *catchwords*, with an analogous lexical emphasis. Here we describe what catchwords are in contradistinction from keywords, outline the distinct conceptual principles involved, and consider the available and potential methods for analysing them. To that end, we pause first on the limitations in studying keywords mentioned above, which Williams worked with but was well aware of.

To put Williams’s and Williams-inspired concepts of keywords into perspective, let us turn to two senses of the term ‘keyword’ which do not gel with those. The first and older usage

is picked up in the Oxford English Dictionary's (OED) sense 1a, which dated from 1762: 'A word that serves as the key to a cipher or code.'⁷ In this sense, the *keyness*, so to speak, of a keyword is not in its connotations or denotations but in the randomness and arbitrariness of its appointment as such for, usually, a transposition or substitution cipher. For cryptography, a keyword which can be used to decode encrypted texts serves their concealing purpose only if the keyword cannot be guessed. The other and more recent usage at odds with Williams's approach is noted in OED sense 2 (dating from 1827), particularly in the later usage therein (from 1967 onwards): 'A word (usually one of several) chosen to indicate or represent the content of a larger document, text, record, etc., in an index, catalogue, or database. Later also: any word entered as a search term in a database or search engine.'⁸ Search engine 'keywords' probably have the upper hand over the other senses in everyday exchanges now and that is the one which sharpens the distinction from Williams's approach. Such keywords indicate the searcher's individual intention in searching through content accessible through the Internet. The Internet – especially insofar as grounded in the World Wide Web – facilitates the fulfilment of that intention by enabling the individual to make associative links and connections through online context (using hypertextual devices).⁹ Of late, ways of anticipating those individual intentions are being designed into the content of the Internet, by pre-structuring the data field of content to anticipate searches – that is, by honing metadata layers or building in knowledge ontologies. The function of the 'keyword' in searches, however, remains ultimately grounded in the individual intention.

Williams's and Williams-inspired approaches are contrary to regarding keywords as either being random/arbitrary or being expressive of individual intent. For those examining them, these words are taken as being *key* in an active and explanative sense. These are fleshed out to substantiate the usage defined in OED sense 1b, especially in the final clause: 'A word or idea that serves as a solution or explanation for something; a word, expression, or concept of particular importance or significance.'¹⁰ The significance in question is emphatically collective, understood as revealing of social and cultural formations and mores. As such, a couple of methodological limitations have dogged this approach.

Firstly, some early respondents to Williams's *Keywords* suspected that he had overendowed the suggestiveness of his chosen words, and possibly of words in general.¹¹ Williams had in fact been very circumspect in this respect, as his use of the OED shows.¹² On the one hand, he eschewed any idea of an essential philological or etymological core in OED definitions and unpacked his keywords as constantly renegotiated according to historical circumstances. On the other hand, he found that the OED usefully offered material for such historicising, in its characterisation and especially exemplifications of different senses. Those were usually the springboard for his own accounts of keywords. Nevertheless, his treatment of keywords does continue to suggest that they have a kind of accretive significance which makes them weighty in themselves. In relation to their place in semiotic systems – i.e., within a network of other words and registers and their shifting significances – one feels that Williams leaves much out, and without a clear rationale for what is left out. Williams was keenly aware of this, and put it down to unavoidable limitations of presentation:

In writing about a field of meanings I have often wished that some form of presentation could be devised in which it would be clear that the analyses of particular words are intrinsically connected, sometimes in complex ways. The alphabetical listing on which I have finally decided may often seem to obscure this, although the use of cross-references should serve as a reminder of many necessary connections. The difficulty is

that any other kind of arrangement, for example by areas or themes, would establish one set of connections while often suppressing another.¹³

Whether because of Williams's indexical ordering or his penchant for focalising lexical emphases, a certain overdetermination of keywords as words-in-themselves seemed difficult to avoid – and has remained so for those building upon his project.

Secondly, it has always been unclear why he chose these particular words and not others.¹⁴ Arguably, a great many words (especially, almost any abstract noun) could have been recruited to similar treatment and found to be similarly significant and revealing. Williams's selection appeared to be free of words used widely in colloquial exchanges and popular registers; Stephen Heath noted in passing that 'most of the *Keywords* words are abstract nouns of Latin rather than Anglo-Saxon provenance, representative of what Williams would call "the vocabulary of learning and power".'¹⁵ Many words in the selection seemed to have an academic resonance, which Williams tried to explain in the Introduction.¹⁶ An unavoidable degree of subjectivity ultimately underpinned the selection, which worked awkwardly against the collective or general significance attributed to them. Williams could only choose from what he himself encountered, in terms of his own perceptions and circles, and structured by his ideological investments (significantly by his research for *Culture and Society: 1780-1950*). Williams was variously upfront about this, in observations like (I have added the emphases): 'The vocabulary *I have* selected is that which *seems to me* to contain the key words in which both continuity and discontinuity, and also deep conflicts of value and belief, are in this area engaged.'¹⁷ He didn't eschew the subjective dimension, and hoped that nonetheless the *idea* of collective significance would be conveyed and could be developed further.

Subsequent studies in a similar vein have not gone much beyond Williams. The limitations in Williams's approach have been refracted through attempts at honing the methodology for contemplating keywords. Let us pause on two recent instances. Lesley Jeffries and Brian Walker's *Keywords in the Press* (2018) employs a distinctive statistical method for identifying socio-politically significant keywords from a corpus of news articles covering the period of Tony Blair's Labour government in the UK, 1997-2007.¹⁸ This involved identifying relevant words (categorised as sociopolitical) which were used more frequently or less frequently in this corpus, compared to a reference corpus (in this instance one of news articles from the preceding period, 1991-1996).¹⁹ The usual practice in this sort of approach has been to use a larger general-language reference corpus, such as the British National Corpus. It is assumed that the higher frequency of a word's occurrence in the given corpus relative to the reference corpus indicates that it is a keyword and has particular sociopolitical significance, which could then be interpreted. By doing this Jeffries and Walker hoped to circumvent the possibility of investigators' intuitions and perceptions interfering with the choice. The attribution of sociocultural significance to what seems statistically significant, however, does not quite do that. At best, sociocultural significance is inserted (interpretively conferred) after the fact on terms which are simply statistically thrown up; there is no necessary rationale between the statistical method of selection and the interpretive attribution of significance. With reference to the book, the various limitations of this method have been carefully demonstrated in a paper by Martin Montgomery and Carol Ting (2023).²⁰ With a somewhat different turn, a study in Chinese produced by Shao Yanjun (邵燕君) and Wang Yusu (王玉琿), *Breaking the Wall: Keywords in Chinese Internet Subcultures* (破壁书: 网络文化关键词),²¹ identifies keywords (关键词) as a subset of popular words or catchwords (流行语) in a corpus of distinctive Internet usage (mostly

neologisms). As Wang Yusu says in a helpful blog posting,²² the catchwords are understood as such in loosely statistical terms as being high frequency, and the keywords are a selection therefrom because of their ‘explanatory power, representativeness, and durability (解释力, 代表性, 持久性), along with their ability to work across communities.’ While explaining what those qualities mean, Wang maintains that ‘a certain degree of trust in researchers in the field and their academic judgement’ is needed, and that ‘We – the *Breaking the Wall* author team – use "keywords" to name catchwords that are more important according to our standards.’ This approach has the merit of straightforwardly embracing intuition and insight for recognising the cultural significance of the chosen keywords. However, that does not dispel possible doubts about the method for attributing key-like significance.

The way in which we conceptualise *catchwords* here is not meant to replace attempts to get to social and cultural analysis via *keywords* in Williams’s mould. Ours is an alternative approach which sidesteps those methodological limitations and enables social and cultural analysis in a different way. This approach also has a lexical emphasis, and also seeks to articulate culture and society via that emphasis. We feel that the limitations in Williams’s method were intrinsic to his way of understanding and attributing significance – therefore *keyness* – to his selection of keywords. A different conception is needed of the words that can be analysed in more methodologically robust ways to articulate culture and society in our time. Accordingly, a concept of *catchwords* is foregrounded here, and the remainder of this paper elaborates on its scope and methodological possibilities for this purpose.

Our use of the term ‘catchword’ here includes phrases insofar as the integrity of their sequence is maintained in usage, i.e., insofar as such phrases are used in the way a lexical unit is (e.g., ‘new normal’, ‘artificial intelligence’, ‘cost of living’). ‘Catchword’ in the sense relevant here has been around for somewhat longer than ‘keyword’ in the relevant sense. In the OED, this is defined as sense 2, ‘A frequently used word or phrase, esp. one associated with a particular group or fashionable at a particular time; a topical slogan; a buzzword,’ and dates its usage from 1715 onwards.²³ Notably, the recent (from 1946 onwards) term of American provenance, ‘buzzword,’ is pegged as synonymous. To us the implications of the metaphor ‘buzz’ seem different from ‘catch,’ but let us not digress here. As such, ‘catchwords’ have for long seemed meaningful and recognisable. Dictionaries of (or including) catchwords have been published periodically, usually in a popular vein.²⁴ However, we are yet to find any rigorous methodological account for studying catchwords. In the available dictionaries of and texts on catchwords, they are widely regarded as somewhat blithe and readymade, interesting to contemplate but in little need of careful elaboration. Methods for identifying and commenting on them are consequently uneven. This essay proposes that they should be taken seriously – as seriously as keywords, and possibly more seriously.

By way of a brief statement of our approach, we begin with the obvious: catchphrases *catch*. The defining feature of these words (and phrases, insofar as phrases hold their integrity and are used as lexical units) is that they are taken up and dispersed quickly. They are identified by the frequency of their usage, which can be estimated and, better yet, tracked. The only reason for attending to them is that there is some evidence that they have caught on. That does not mean that they are therefore intrinsically culturally or socio-politically significant or contain a particular explanative power; in fact, they might seem connotatively and denotatively quite trivial. However, the fact that they have caught on within some scope calls for explanation. Attempts to explain why they catch and with what scope they catch offer a mapping of and connections within prevailing cultural and socio-political formations. It is not the catchwords

that are significant or key for understanding culture and society, it is the circumstance of their catching and consideration of the reasons for their doing so that may help develop our understanding of culture and society. Catchwords do not have an innate *keyness* (cultural significance, explanatory power, representativeness, etc); they are conceptually empty signifiers which may be filled in according to contexts of usage. Catchwords are not identified in terms of perceived or intuited significance by some authority or savant; they are thrown forth by the evidence of their growing frequency of usage.²⁵ Charting and analysing the process of their proliferating usage and dispersal across contexts is a way of mapping, describing, and explaining social and cultural phenomena. Doing so may also open possibilities for intervening effectively in prevailing cultural and social formations.

Understood thus, three methodological steps are called for in the study of catchwords: first, how should these be defined to enable investigation; second, how should they be identified for undertaking investigation; and third, how should the reasons for their catching be investigated. In other words: what are catchwords?; how are catchwords identified?; why do they catch?

What are catchwords?: definition

With *catchiness* in mind, we have the following three points as definitive features of catchwords/phrases.

First, the connotations with which catchwords start catching on can usually be traced to a specific social juncture (an ‘origin’). The initiation of catching on is likely to involve the following:

- An existing word or phrase acquires a new sense at a social juncture within a context (a distinct discourse area, linguistic domain, social group), which then leads to it catching on in a significant manner over a subsequent time-period. In standard dictionaries, word-senses are enumerated separately, and new word-senses added as they settle.
- A neologism or new phrase-form appears at a social juncture within a context and then catches on in a significant manner over a subsequent time-period. Such a neologism is more than a ‘nonce’ word/phrase or a context-specific coinage for one-off or occasional purposes.²⁶

Second, ‘catching on in a significant manner’ means that, starting from that ‘original’ social juncture, catchwords have steadily growing usage-frequencies over a time-period, *extensively* (spread across contexts) and/or *intensively* (within a context). They are to be considered catchwords for the *time-period* over which there is steadily growing usage-frequency; they are to be considered catchwords insofar as such growing usage-frequency is *extensive* or *intensive* or both. The relativistic terms are italicised, and may be sharpened as follows:

- *Time-period*: For a catchword to be recognised as such, the time-period over which its usage-frequency grows should be sustained enough and yet concentrated enough to become discernible. If the word or phrase were used, for instance, enthusiastically for a day or two in some context and then stopped being used, it cannot be considered discernible as a catchword (it is in the province of ‘nonce’ words). Equally, if usage of the word or phrase grew in frequency very slowly, with rises and dips, over several decades, then too it cannot be considered discernible as a catchword.

- *Extensive* usage: The catchword has caught on across a range of domains of usage, e.g., across different regions, registers, cultural contexts.
- *Intensive* usage: The catchword has caught on in that its usage-frequency is high within particular domains, e.g., within a specific region, register, cultural context.
- Combinations of *extensive* and *intensive* usage: If a catchword is used both extensively and intensively, it is likely to be a *popular* catchword, which may well enter general everyday usage. If the catchword is used extensively but not particularly intensively, it is likely to be an *organizational* catchword, as may be used by academics, businesspeople, bureaucrats, activists, and the like within their circles in different locations. If the catchword is used intensively but not particularly extensively, it is likely to be an *in-group* catchword, as may be used in a school, workplace, or neighbourhood at a specific location.
- In principle, a *degree of catchiness* of catchwords should be calculable as a function of the time-period of steady rise in usage-frequency and of the combination of extensive and intensive usage.

Third, when the period of ‘catching on in a significant manner’ ends, the word or phrase in question ceases to be a catchword. That could mean either diminishing usage-frequency or a levelling of usage-frequency. In descriptive terms, that could indicate that the word or phrase has:

- entered a stable level of everyday or professional or in-group usage (has become ordinary usage in those circles); or
- come to be considered idiomatic (its sense is stabilised, but it is understood as linguistically distinctive, perhaps with non-logical implicature); or
- come to be regarded as a cliché (used not so much to state a proposition as to distance from the statement of such a proposition, usually ironically or disparagingly); or
- become outmoded.

How are catchwords identified?: Estimation and measurement

Perceptual frequency estimation

Simple recognition appears to be an obvious basis for discerning what is or is not a catchword. If we randomly select a number of persons from our immediate circle, give them a few examples of what we consider to be catchwords and ask them to list five current ones, they will usually come up with a list. Generally, they will do this confidently, unhesitatingly, and without asking searching questions about what ‘catchword’ means. ‘Catchword’ is, after all, a familiar everyday term.

Such immediate recognition might be based on rule-of-thumb frequency estimations. If we encounter a word or phrase more often than before in news media, official documents, popular publications, everyday conversations, etc. we peg it as a current catchword. That is not just because we feel called upon to start using it ourselves, but because we know that each of those occurrences in news media, official documents, etc. are also within the remits of many others like us. Simple recognition is a reasonably robust indication of what is dubbed *perceptual frequency estimation* here.

Reasonably robust as it is, perceptual frequency estimation is loose enough to make for significant divergences. When requested, the lists of current catchwords that our chosen persons present may have some overlaps and significant differences. Moreover, if we ask them

to put their five choices in order (such as, more to less popular), the orderings are often different too. We may say that perceptual frequency estimation is robust to the extent that there are overlaps and some similarity in orderings. It may be possible to enhance the robustness of perceptual frequency estimation by asking our chosen persons to have a look at each other's lists and agree upon one list – to come up with a consensual listing.

Much scholarship in this area has ploughed perceptual frequency estimations, usually by beginning with a declaration that such-and-such terms have obviously caught on and then analysing them accordingly.²⁷

Empirical frequency measurement

The obvious way to obtain a measured sense of the usage-frequencies of words and phrases over specific time-periods is to:

- take a word/discourse corpus of meaningful size for a given domain;
- use a standard scale (percentage is common) to work out how often a word or phrase was used relative to all the words in the corpus (i.e., that word's or phrase's usage-frequency) in consecutive time-periods (let us say, at daily/monthly/annual/decadal intervals);
- draw a graph of the usage-frequency of that word or phrase [x-axis] in relation to consecutive time-periods [y-axis].

If the graph shows a notable and steady upward curve over a significant time-period for a word or phrase, we may say that we have a measurement for that word or phrase 'catching on in a significant manner' – so, that is a catchword for that time-period.

A *popular* catchword, as described above, might be trackable in a large language corpus for a given language (e.g., a corpus for general English usage).

The most accessible facility based on a large corpus for obtaining such graphs at present is the Google Ngram Viewer.²⁸ It draws upon the Google Book corpus of around 200 billion American and British English words. The various grey areas of the facility notwithstanding,²⁹ this offers immediately processed graphic representations of usage-frequencies over time for words and phrases. For ease of discussion, let us confine the rest of this section to words that catch by using this tool (usefully, the Google Ngram Viewer can be used for word as well as phrase graphs).

If a Google Ngram tracing of usage-frequency for a word over a given time-period shows a continuous increase, we may consider it as catching on. For example, the word 'austerity' in the decade 2008-2018 (when we suspect it became a catchword following the 2007-2008 financial crisis and consequent 'austerity' measures) shows a steady growth in usage-frequency as follows:




Figure 1: Google Ngram of the word ‘austerity’ in the period 2008-2018 (smoothing 3, English [2019], case-insensitive), showing continuous growth in the Google Book Corpus

This confirms that the word did indeed catch on in this period. Typically, we may expect to see a continuous logistic curve and at times an exponential curve for catchwords. However, to gauge whether this is indeed a catchword, we need rather more than this tracing. We need, first, some equipment to understand how this word’s usage-frequency is placed in relation to those of other words (especially other relevant words). And second, we also need some way of pegging when such growth in usage-frequency of a word could be considered of a scale that can be understood as being a *catchword*. It is possible that a continuous increase in usage-frequency is nevertheless not significant enough to be categorizable as a *catchword*.

To understand how a word’s usage-frequency is placed in relation to those of other words (especially other relevant words), it is useful to have some kind of orientation device. In other words, it is useful to have a general sense of the usage-frequencies of different kinds of words to understand what the usage-frequency of a specific word in a given time-period suggests. One loose orientation device is provided by the online Oxford English Dictionary’s (OED) Frequency Bands.³⁰ The OED draws upon the Google Books 1500-2010 corpus and other corpora to come up with a rounded average figure of usage-frequencies for the words it covers (as occurrences per million). Usefully, the OED divides these rounded average figures into eight Bands from low to high usage-frequency (Band 1 to Band 8), thus:

Band	Frequency per million words	% of entries in OED
8	> 1,000	0.02%
7	100 – 1,000	0.18%
6	10 – 100	1%
5	1 – 10	4%
4	0.1 – 1	11%
3	0.01 – 0.1	20%
2	< 0.01	45%
1	-	18%

Table 1: OED frequency range for bands. Source: OED ‘Frequency’ webpage

The table shows that the higher usage-frequencies (i.e., higher Bands) are found for a relatively small number of words. Only 5.20% of all words in the OED have a usage-frequency of 1 and above (>1) occurrences per million, in Bands 5-8. The catchwords that are likely to come up

in our time are very likely to be in Bands 1-4 (<0.099 occurrences per million). This overarching picture could help us place the usage-frequency tracing of a specific word over a specific period.

Further, insofar as we are interested in the usage-frequencies of specific words of a given domain or context over a given time-period, we could use Google Ngrams to make up an Orientation Table that is indicative for that domain or context. For instance, if we were interested predominantly in words that are considered explicitly *political* (can be tagged as such), like ‘austerity’, we could use Google Ngrams to make a table of the usage-frequencies of a range of commonly used political words and place ‘austerity’ accordingly. Here is an example of such a made-up Orientation Table for political words according to usage-frequency at the beginning of 2018 by Google Ngrams, and also in terms of OED average usage-frequencies and Bands:

Word	Word occurrence as percentage of corpus at the beginning of 2018 (from Google Ngrams)	OED figures – based on Google Books 1500-2010 and other corpora)	
		Part of speech and Frequency Band	Average Usage - frequency (occurrences per million)
people	0.0720599212	noun, Band 7	600
social	0.0340259715	adjective, Band 7	400
law	0.0239443846	noun, Band 7	300
government	0.0204596273	noun, Band 7	300
rights	0.0147146877	noun (singular), Band 7	300
international	0.0112967235	adjective/noun, Band 7	100
justice	0.0060981896	noun, Band 6	60
democracy	0.0028892535	noun, Band 6	40
capitalism	0.0014383717	noun, Band 6	20
terrorism	0.0010460449	noun, Band 5	8
socialism	0.0005858211	noun, Band 6	10
communism	0.0004231806	noun, Band 5	5
neoliberalism	0.0002656769	noun, Band 4	0.2
populism	0.0001448323	noun, Band 4	0.8
technocratic	0.0000851356	adjective, Band 4	0.8
theocracy	0.0000331292	noun, Band 4	0.4
precariat	0.0000126481	noun, Band 3	0.03

Table 2: Orientation Table for political words in terms of Google Book corpus usage-frequency at the beginning of 2018 and in terms of OED frequency Bands and average word occurrences per million.

Such a purpose-built Orientation Table could be quite useful to place words which we suspect of having become political catchwords by Google Ngram tracings for some time-period around 2018. In other words, contemplating the relationship between the 2018-snapshot figures and the all-time averages could help towards building a kind of mental map in terms of which to read usage-frequency in Google Ngram graphs for specific words.

Such an Orientation Table, however, does not tell us what scale of growth in usage-frequency over a time-period could serve to categorise that word as a *catchword* for that time-period. So, as noted earlier, we also need some way of pegging this. Importantly, the OED

Frequency Bands, while useful for enabling a mental map to place specific words among others, does not have a bearing on which word becomes categorizable as a catchword. The point of catchiness is not how frequently the word is used in relation to a total corpus size at average or at a given time, but about how much and how quickly it catches. So, a catchword can appear as such at any level of usage-frequency, as much in the range of OED’s low usage-frequency Band 2 as in the high usage-frequency Band 5. It is the significant increase in usage-frequency that is relevant, not the OED Band of usage-frequency.

For pegging catchiness, what is relevant here instead are the precise usage-frequency measurements at specific points of time within the time-period in question. In Google Ngrams, that is given as a percentage of the total corpus in 10 decimal points. Going back to ‘austerity’, the Google Ngram graph above has the specific usage frequencies in 2008 and 2018 as follows, showing a steady increase in the interim:

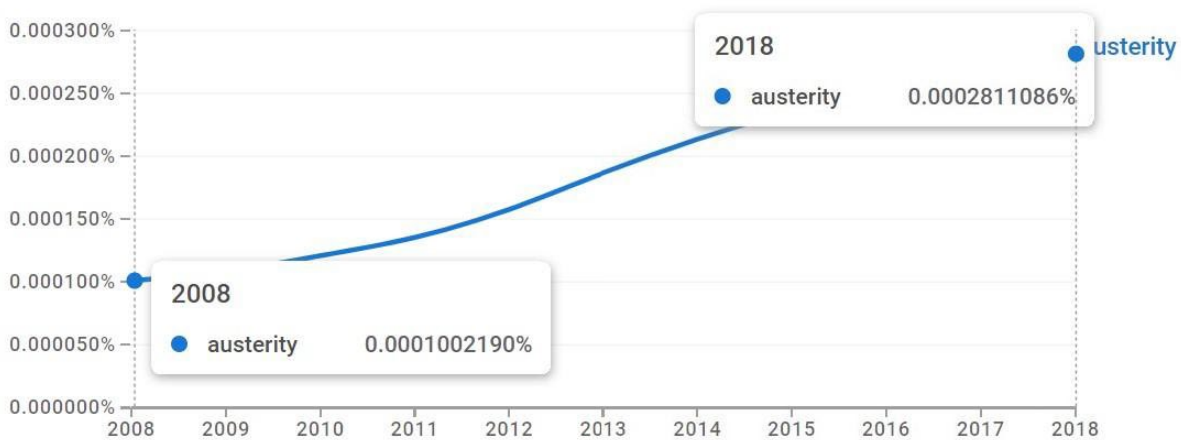


Figure 1: Google Ngram of the word ‘austerity’ in the period 2008-2018 (smoothing 3, English [2019], case-insensitive), with usage-frequencies in 2008 and 2018 indicated.

This shows that from 2008 to 2018 the usage-frequency of the word ‘austerity’ in the Google Book Corpus increased continuously to a multiple of 2.8. This kind of reckoning could serve to fix a scale for pegging which is a ‘catchword’.

To come up with a pegging for the scale at which growth in usage-frequency could be regarded as a ‘catchword,’ there is no straightforward inferential method that we could come up with. So, based on our experience of researching this area, we have gone along with a rule-of-thumb. In general, insofar as Google Ngram indications go, we have it as a rule-of-thumb that: *a popular catchword will, starting from a usage-frequency greater than 0 at a point of time, show a continuous increase in usage-frequency by a multiple of 2 or more by at least the end of 2 years and up to the end of 10 years.* Let us reiterate that this is a rule-of-thumb and not innately meaningful. The boundaries of usage-frequency doubling or more within two years to a decade are based on convenience and experience rather than an inherent rationale; other boundaries may well work better.

The rule-of-thumb may be restated economically as follows. Let us denote the usage-frequency of a word W in year t as $F(t)$. W is a catchword for a time-period of n years, where n ranges from 2 to 10, if $F(t_n) \geq 2 \times F(t_1)$ – where t_n represents the end of n years, and t_1 represents the initial year and $F(t_1)$ is greater than 0.

A very large general-language corpus like Google Books is invaluable for empirical frequency measurements of catchwords, but it is unlikely to be helpful for measuring catchwords which are more extensive than intensive (like *organizational* catchwords) or more intensive than extensive (like *in-group* catchwords). For those, the corpus needs to be delimited by organizational or in-group domains. To track the significant catching on of words and phrases according to domains:

- either the large corpus needs to be annotated so that it can be delimited for those domain-parameters,
- or customised corpora need to be built for those domain-parameters, either by direct text collections from or by capturing text from the digital facilities of the domains in question.

Depending on the resources available, between perceptual frequency estimations and empirical frequency measurements it is possible to identify catchwords for given time-periods reasonably confidently.

Having done that, the significance of these for understanding culture and society entails investigating why they catch. Two approaches to such investigation are considered briefly in the concluding two sections: manual text interpretation and Natural Language Processing (NLP).

Why catchwords catch?: Textual archive and manual text interpretation

Both perceptual frequency estimations and empirical frequency measurements are ultimately sourced from relevant texts. For our purposes, those are texts in which the catchword under investigation is introduced, used, contextualised, defined, conceptualised, renegotiated, etc. Such texts accrue as an archive of news articles, official reports, scholarly publications, fictional narratives, publicity materials and so on with which readers have engaged. These inform their perceptual frequency estimations of catchwords in the first instance. Further, such texts also feature amongst those from which word/discourse corpora and datasets are elicited and made amenable to computational processing -- enabling empirical frequency measurements. A corpus like Google Books obviously links to digitised texts, with their integrity as books or publications in the background. Let us call this base material of texts in collections and datasets focused on a catchword that *catchword's textual archive*.

Manually consulting a catchword's textual archive is the obvious and usual way of trying to work out why that catchword caught on.

A catchword's textual archive allows tracing of the processes and rationales through which the word caught on from one text to others and from one register to others and one region to others and so on. It allows for tracking the pathways through which the word was proposed and contested, adopted and adapted, defined and redefined, standardized and updated, acknowledged and accepted, interrogated and qualified or relegated. In this textualized process, significant insights into *why* certain words caught on at certain social junctures can be found. Those include reasonable insights into *how* the catching on – extensively or intensively or both – has worked.

So far, our own investigations into catchwords have predominantly followed this method. In our catchword's textual archives we look for some of the following:

1. First or initial usage of a catchword, especially noting the connotative nuance that may account for an extant word or neologism to start catching on at a particular juncture.

2. Subsequent definitions of the catchword with its catchy connotations and contexts in view (the performance of ‘defining’ is usually itself rich in contextual nuance, even where definitions appear to be deliberately acontextual).
3. Nudges or adaptations irrespective of definitions in various contexts and registers, with attention to connotative shifts.
4. Reflections on the usage of catchwords as catchwords, especially in terms of scale and application.
5. Opaque usage of a catchword to suggest that it is such, often marked by some sort of emphasiser, such as by italicising or putting in quotation marks.
6. Relationships of such words to other relevant and related words, and/or their placement within a distinctive vocabulary of a collective formation.
7. Observations on catchwords becoming everyday terms and therefore not being catchwords any longer, or coming to be regarded as clichés, or becoming anachronistic, or being replaced by other terms.

Somewhere in a catchword’s textual archive, with such indications we expect to find clues towards explaining why catchwords catch.

Why catchwords catch?: Scaling up textual archive analysis with NLP

Examining a catchword’s textual archive to figure out why the catchword caught on is not straightforward. The manual method only allows for partial reckonings. Importantly, there is a problem of information overload. For instance, consider all the newspaper articles, speeches, books, papers, etc. that constitute the textual archive of ‘austerity’ between 2008 and 2018. This is likely to add up to millions of texts, which is far too many to be analysed through manual reading alone. At this scale, reading needs to be supplemented with quantitative methods of text analysis. Linguistics-related disciplines such as sociolinguistics and corpus linguistics have been developing computer-assisted quantitative text analysis methods since the middle of the 20th century. Developments in natural language processing (NLP), in particular the relatively recent development of neural language models and, even more recently, large, pre-trained neural language models (large language models or LLMs), offer new tools for quantitative text analysis. These hold a lot of potential for examining why catchwords catch and mitigate the problem of information overload.

This section discusses how these developments in NLP can be used to determine why words become catchwords. Before doing this, however, it is worth stepping back a bit to consider the question ‘Why does a term become a catchword?’ in a bit more depth. What kinds of responses would constitute satisfactory answers to this question?

Let us consider a not very informative answer to the question: a term becomes a catchword because lots of people *decide* to start using it in a short period of time. Uninformative as this is, it draws attention to the fact that understanding why terms become catchwords is a matter of understanding the intentions behind a decision to use a term. Some term becomes a catchword when the number of people making the decision to use it has increased significantly. That is indicated when the frequency of the term in its textual archive has increased n times during some time-period. If we knew the motivation behind every single decision to use a term within a textual archive, would we have a complete answer to the question of why the term became a catchword?

The network structure of catchword exposure

The brief answer to that question is: No. Understanding the intentions behind a term’s occurrences is necessary but not sufficient for understanding why it becomes a catchword. The

issue here is that a list of all motivations behind every occurrence of the term is not enough to fully understand the *diffusion* process through which individual decisions accumulate. If the number of people deciding to use a term has increased enough within a time-period for us to consider the term a catchword, it must have ‘diffused’ amongst a population. That is, people must have learnt about the term from seeing others use it, and then have gone on to use it themselves. So, terms becoming catchwords is not just a matter of people deciding to use them, it is also a matter of people becoming *aware* that these terms rise through exposure to what others have said. Understanding this aspect of catchword diffusion requires consideration of the factors that affect exposure.

Sociolinguists have investigated such factors carefully. One sociolinguistic model for linguistic diffusion is the wave model.³¹ According to this, linguistic innovations (neologisms, new connotations, new pronunciations etc.) first spread through the immediate social network of those who know the inventor of the innovation (i.e., the innovator), and then spreads through their own immediate social networks, and so on. The wave model implies a ‘principle of density.’³² That means, if diffusion depends on information travelling between social networks, one would expect diffusion to begin in areas with high population density, where there is a greater prevalence of densely connected social networks, and end in areas with low population density, where there will be a greater prevalence of sparsely connected social networks. This principle of density is formalised by Trudgill (1974),³³ who characterised the capacity for linguistic innovations to flow from one area to another as a function of the populations of both areas and the distance between them.

Salient in this sociolinguistic approach to diffusion is the emphasis placed on interpersonal social network structure. Population density and distance between areas are means of approximating the impact of different social network structures on diffusion: i.e., the densely connected interpersonal networks one might find in populous urban centres vs. the sparsely connected interpersonal networks one might find in less densely populated areas. One can extend the notion of social networks to include not only interpersonal networks but also social media networks (consisting of online ties – followers, friends, etc. – between users) and media consumption networks, consisting of ties between media organisations and consumers of organisations’ outputs (articles, videos etc.). Thinking about media consumption in terms of network structure has become particularly prevalent since social media have become ubiquitous. These provide researchers with data rich enough to represent the relationships between consumers and media organisation relations as node-edge graphs³⁴.

Typically, information flows between interpersonal networks and media consumption networks: e.g., friends recounting to each other and reinterpreting the things they have seen on news sites. It is therefore useful to think of these networks as together constituting a single, overarching exposure network that encompasses both interpersonal and media consumption relations. In fact, all kinds of relations are likely to be part of the network structure of exposure – consumer relations, ownership relations, employer-employee relations and so on. From this perspective, the kinds of information an individual will be exposed to will be fixed by their position within this overarching network. So, which catchwords an individual is aware of will depend upon who their interpersonal connections are and what media content they consume, who their connections are connected with and what content their connections consume, and so on.

The network structure of exposure cannot be reduced to the individual motivations behind decisions to use terms. Consequently, analysis of both the motivations behind the usage of terms and the network structure of exposure is needed for a complete answer to the question of why certain words become catchwords. Neither kind of analysis by itself is sufficient for a complete answer. One might object that there are still other factors to consider. The political economy of media, and factors like market concentration, or the extent to which media are

independent of state interference, or the occurrence of significant events such as pandemics, might all be said to affect the kinds of information people are exposed to. This is undoubtedly true, but such factors will ultimately work through motivations and network structure. If the media ecosystem is oligopolistic, or is under the control of the state, this means a considerable proportion of the information everyone is exposed to originates from a small number of organisations. That further means that those organisations form a core which has a high degree of centrality within the overall network structure of exposure. If some significant event happens that precipitates the diffusion of a catchword, this means that people are using the catchword to talk about the event, returning us again to the issue of the motivations behind decisions to use the catchword.

It is valuable to keep the above in mind. Since textual archives can only tell us about the reasoning behind the decisions to use a word, not the network structure of exposure, the above observations underline the limitations of textual archive analysis as a method for determining why a word becomes a catchword. Linguistic diffusion of lexical items is the central issue of catchwords.

Using LLM driven Word Sense Induction to understand why words become catchwords

Ideally the methods of textual archive analysis scaled up through NLP and social network analysis ought to be fused when considering why a term becomes a catchword. Doubtless this is possible, but for such a fusion to work we need to clarify how inferences from textual archives about the motivations behind decisions to use terms can be scaled up through NLP. This will give some direction on exactly which NLP-based procedures are useful for finding out why terms become catchwords. So, let us return to the issue of how NLP can be used to address the problem of information overload that comes with a catchword's textual archive. The first thing to clarify is how exactly does one figure out intentions from text?

We make inferences about the intentions underlying particular linguistic uses all the time. When someone makes an utterance, whether written or spoken, we recognise the kind of speech-act the utterance is (is the utterance an assertion, a promise, a question, a warning...?). We also grasp how the utterance relates to previous utterances of the text/conversation that forms the overall context of communication, the logical interrelations between the utterance and other utterances, and so on.³⁵ We use all of this information to make inferences about the reasons why the utterance was made. In other words, most of the time, simply understanding an utterance is enough for us to figure out the intention behind the utterance. For instance, understanding that a person *A*'s exclamation of 'Snake!' is a warning, is a reference to a nearby venomous snake, implies propositions like 'If that snake bites you, you will be hospitalised'. This is enough for a person *B* to infer that the intention behind *A*'s exclamation is to draw attention to danger and to prompt *B* to do something about it. This is good for the analyst of catchwords: it appears that much of the information needed to figure out the intention behind an occurrence of a catchword should be contained within the text containing the occurrence. Of course, there are complications to this; for instance, the required information might be spread across many texts, or simply lost. Nevertheless, there is no special method of interpretation required to use this information to figure out intentions. The analyst can rely on their default linguistic competence.

Let us suppose that, based on a usage-frequency analysis using some textual archive, we discovered that 'snake' is a catchword because its usage-frequency tripled within 10 years. If the textual archive contains every single speech act containing the use of 'snake' from all members of the population within those 10 years, and all contextual information needed to correctly interpret every speech act, we would have before us all the information needed to understand what led to 'snake' becoming a catchword. Consideration of this makes clear the

problem of information overload. How can one work through and make generalisations about the millions of speech acts of those 10 years in a reasonable amount of time?

Word Sense Induction (WSI) is an example of a procedure that can address this problem. The purpose of WSI is to automatically isolate the different ways in which a term is used in some collection of texts. For example, a text might contain 10 occurrences of the term ‘bank’, where some of these occurrences discuss financial organisations and others discuss river banks. A good WSI procedure ought to be able to automatically determine which of the 10 occurrences are used to discuss financial organisations and which to discuss river banks. If we can feed the textual archive of ‘snake’ into a WSI algorithm that can automatically isolate the different ways ‘snake’ is used, this will go a long way towards indicating the different reasons why ‘snake’ was used without having to read through the entire textual archive.

Recent developments in NLP concerning pre-trained neural language models or LLMs³⁶ have been particularly effective in realising the goals of WSI. We will not give an extensive, technical explanation of how LLMs are used to conduct WSI here. The reader may directly engage with NLP literature on using LLMs for WSI for that.³⁷ However, there is space within this paper for an intuitive sketch of how all this works.

The central point to grasp is the idea of translating linguistic expressions from their original language into the language of linear algebra. Thus, words, sentences, documents, etc. might be translated from English into vectors and matrices, which are arrays of numbers that can be used to encode information about whatever they are being used to represent. For instance, consider coordinates. The coordinate (3, 2) might represent a location in some space that encodes information about how to reach there, such as, move 3 steps to your right and 2 steps in front of you. Vectors and matrices can be thought of as coordinate-like objects and linear algebra is a mathematical framework for understanding and performing operations on such objects. The coordinate (3, 2) can be thought of as a 2-dimensional vector that can be subjected to the operations of linear algebra. We need not limit ourselves to 2 dimensions. We can also have 3 dimensional vectors: (3, 2, 1) → move 3 steps to your right, 1 step forward, and levitate 1 step up towards the sky. In fact, we can have vectors of whatever number of dimensions we like. In NLP, translated expressions often have hundreds of dimensions.

Matrices generalise the idea of a coordinate even further. Instead of a vector consisting of a single row of numbers, we can have multiple vectors arranged as rows within a single table called a matrix. So, we might have a matrix representing and encoding information about two locations represented by the pair of vectors (3, 2) and (4, 5):

$$\begin{bmatrix} 3 & 2 \\ 4 & 5 \end{bmatrix}$$

Each coordinate vector is now a row in this matrix. As it consists of a pair of 2-dimensional vectors, this matrix is a 2x2 dimensional matrix.

Once linguistic expressions are translated into objects of linear algebra, one can apply to linguistics the linear algebraic tools that are used to compare and transform vectors and matrices in all kinds of ways. For example, in linear algebra one can use measures such as Euclidean distance and cosine similarity to quantify the similarity of a pair of vectors. Translating expressions into linear algebra allows us to use measures like Euclidean distance and cosine similarity to quantify the similarity between linguistic expressions. The process of translating linguistic expressions into linear algebra is called *embedding*. Word occurrences translated into linear algebra are called ‘token embeddings’, translated words ‘word embeddings’, translated sentences ‘sentence embeddings’, and so on.

Of course, none of this is of any use unless there are mechanisms for ensuring that the linear algebraic translations of linguistic expressions really do encode useful information about

them in the same way that (3, 2) encodes useful information about how to arrive at a location. The current paradigm of NLP makes use of artificial neural networks – especially LLMs – to produce useful embeddings.³⁸ Explaining how they do this is beyond the scope of this paper. But it is worth noting that LLMs are capable of being used for a broad variety of tasks, such as sentiment analysis (classifying whether an expression has a positive or negative tone), named entity recognition (classifying whether an expression is a reference to a person, an organisation, a location, etc.), question answering, etc. The most immediately impressive models today are generative models such as ChatGPT, Claude, etc., which can simulate conversations fairly convincingly. These show that artificial neural networks are capable of encoding all sorts of useful linguistic information into linear algebraic objects.

LLMs can be used to conduct WSI. To find out the different ways ‘snake’ is used within its textual archive, we can follow two main steps. First, feed the sentences of all speech acts containing ‘snake’ into an LLM to retrieve token embeddings of each occurrence of ‘snake’. Second, since these token embeddings are vectors, we can make use of linear algebraic methods of quantifying the similarity between vectors and cluster these token embeddings, where embeddings within a cluster are more similar to each other than to embeddings in other clusters. For example, suppose a LLM has given us embeddings for 3 occurrences of ‘snake’, the 2-dimensional vectors (3, 2), (3, 1), (5, 4). Clearly, (3, 2) and (3, 1) are more similar to each other than either (3, 2), (5, 4) or (3, 1), (5, 4). Using Euclidean distance or cosine similarity to measure the similarity between each possible pair of these 3 embeddings would tell us the same thing. So, clustering these 3 embeddings would produce one cluster of two embeddings [(3, 2), (3, 1)] and one cluster of one embedding [(5, 4)]. Since LLMs produce embeddings that capture useful linguistic information, we would expect the two clusters produced here to perhaps correspond to similar ways of using ‘snake’. In this vein, one can imagine that [(3, 2), (3, 1)] corresponds to the pair of expressions ‘That snake is a cobra’ and ‘Cobras are venomous snakes of the genus *naga*’, while [(5, 4)] corresponds to the expression ‘Some snakes use neurotoxic venom’. The greater similarity of (3, 2) and (3, 1) compared to (3, 2), (5, 4) and (3, 1), (5, 4) reflects the similar use of ‘snake’ to talk about cobras. This method of WSI is frequently found in WSI literature (including the examples of LLM-based WSI provided in footnote 33).

Instead of having to read through millions of ‘snake’ speech acts to make generalisations about the main intentions underlying the millions of decisions to use ‘snake’ over 10 years, with WSI we need only to identify the distinctive uses captured by each cluster. Identifying the main uses captured by clustered embeddings is itself a complicated process since each cluster will contain a large quantity of catchword uses. Some manual inspection of clusters will be necessary, but this too can be aided by automatic methods of inspection. For example, one can automatically extract the most distinctive terms of a cluster using measures of distinctiveness (such as term frequency-inverse document frequency)³⁹. These distinctive terms will then give a representative impression of the main uses captured by a cluster in a very compact, concise manner. So, the most distinctive terms of one cluster of ‘snake’ speech acts might be ‘new species’, ‘discover’, ‘venomous’, ‘novel’, and this will tell us that this cluster is concerned with the discovery of a new species of venomous snake.

Limitations

For the sake of clarity, with our example of ‘snake’ we have made the clearly unrealistic assumption of having a textual archive that contains every speech act containing ‘snake’ over a period of 10 years. In the real world there are several limitations that need to be borne in mind. First, one can only extract samples of the speech acts used by very specific subsections of the population that has mentioned some catchword, e.g., the user base of a social media network, the content of newspapers, government publications. Data availability is a key limitation on what kinds of intentions can be understood through methods based on embedding.

There are resources of substantial scale which can be garnered and processed to some degree which we do not detail here. Nevertheless, the case is that one cannot really gain an exhaustive picture of the reasons why a word has become a catchword. At most, one can gain a detailed picture of why a word became a catchword within a range of particular contexts.

The second limitation to bear in mind is that everything has to be written down to be fed into a neural net. There are NLP techniques for automatically converting speech into written strings that can be embedded, but speech recognition is complicated enough to constitute a whole area of study within NLP. In any case, it is not really feasible to go around recording millions of conversations. Furthermore, working with speech data is not just a matter of converting to written form. One also has to consider all those meaningful elements of speech that do not have a written form – significant pauses, tone of voice, glances, etc. A system for transcribing these sorts of elements, as is done in conversation analysis, would be needed. This is, again, a labour-intensive task for large quantities of data. However, in many cases this is not an issue since the original form of many speech acts is written.

Inconclusively

In analysing why certain catchwords catch, much work remains to be done on the precise ways in which manual interpretation and NLP methods bear upon each other. Insofar as conceived here, these appear to be alternative approaches using respectively different scales of textual material to address the same question. Two somewhat different kinds of insight into *catchiness* – possibly complementary insights – may be anticipated. But this distinction is really produced by the methods themselves; this is not a distinction that structures the social and cultural field evinced via catchwords. We have outlined two distinct ways of, so to speak, grasping a coherent process. A satisfactory account of this process, and of the social and cultural field it reflects, would need to bring together these methods in some cohesive way. In principle, NLP methods and manual interpretation could be integrated and somehow merged to produce a closer-to-the-grain account of society and culture at a given juncture. At present, however, this is more a conceivable prospect than a practical undertaking.

In conclusion, here's an impressionistic thought on the substance of this paper. The means that are used to disaggregate, describe, map, rationalise, and otherwise engage systematically with a prevailing cultural and social formation are an aspect of that formation. Raymond Williams's concept of keywords, his reasons for turning to them, and the sources and methods he employed were necessarily implicated in each other – together they were derived from and reflexive of the society and culture he was amidst. What seemed significant and analysable at the time were premised on the prevailing structures of power, political and economic arrangements, institutional orders, technological capacities, ideological contestations, material resources, etc. These are inseparable from the terms 'culture and society'. Those concepts, sources, and methods can be finetuned and stretched for a while, but not indefinitely. When significantly modified or different underpinnings for the terms 'culture and society' settle, new concepts and methods need to be considered. We suggest that the concept of catchwords and the methods outlined here are in synch with the 'culture and society' of our time.

Acknowledgement

The research this paper is based on has not been supported by any funding body. However, an Arts and Humanities Research Council Networks Grant enabled Suman Gupta to organise a series of workshops on the theme *Analysing and Modelling Political Catchwords* in 2023-2024. These offered helpful opportunities for discussing some of the material covered in this paper.

¹ Our use of the phrase ‘collective life’ here is meant to be coterminous with the keywords ‘culture and society’ and allows for a lead up to the latter *as* keywords. The consonant connotations of ‘collective life’ are well captured by Lorens Holm, ‘Foreword: *Collective Life* is a Difficult Term,’ *Architecture and Culture*, 8 no.3–4, 377–381. <https://doi.org/10.1080/20507828.2020.1886509>

² First published in 1976; a later edition is referred here: Raymond Williams, *Keywords: A Vocabulary of Culture and Society* (Oxford: Oxford University Press, 2015).

³ Williams, *Keywords* (2015), vii.

⁴ Williams, *Keywords* (2015), 49.

⁵ The larger project is conceptualised in Raymond Williams, *Culture and Society, 1780-1950* (Harmondsworth: Penguin, 1988 [first edition 1961]); its relation to earlier concepts of culture are outlined usefully in John Eldridge and Lizzie Eldridge, *Raymond Williams: Making Connections* (London: Routledge, 1994), Ch.3.

⁶ The obvious general examples are: Tony Bennett, Lawrence Grossberg, and Meaghan Morris eds., *New Keywords: A Revised Vocabulary of Culture and Society* (Oxford: Blackwell, 2005); and Colin MacCabe and Holly Yanacek eds., *Keywords for Today: A 21st Century Vocabulary* (Oxford: Oxford University Press, 2018). Almost every academic discipline in the Humanities and Social Sciences is blessed with a volume or more of keyword listings with brief essays – several appear in the NYU Press’s *Keywords* series (<https://keywords.nyupress.org/> accessed 20 April 2024); the Penguin/Random House *Keywords* series offers short volumes discussing one keyword each (<https://www.penguinrandomhouse.com/series/DKK/keywords-series/> accessed 20 April 2024).

⁷ ‘Keyword, *N.*’ *Oxford English Dictionary* (OED), Oxford University Press, September 2023, <https://doi.org/10.1093/OED/5928140031> (accessed 20 April 2024).

⁸ ‘Keyword, *N.*’ OED.

⁹ The concept of associative linking was famously proposed in Vannevar Bush. ‘As We May Think’. *The Atlantic*, July 1945. <https://www.theatlantic.com/magazine/archive/1945/07/as-we-may-think/303881/> (accessed 20 April 2023), and has inspired hypertext systems developed later, including the dominant World Wide Web.

¹⁰ ‘Keyword, *N.*’ OED.

¹¹ The early reception is outlined in Fred Inglis, *Raymond Williams* (London: Routledge, 1995), 246-51.

¹² Williams elaborated on this in the Introduction to *Keywords* (2015), xix-xxxi; it is discussed well in Alan O’Connor, *Raymond Williams: Writing, Culture, Politics* (Oxford: Basil Blackwell, 1989), 53-54.

¹³ Williams, *Keywords* (2015), xxxvi.

¹⁴ The issue is considered, with later keyword selection projects in view, by Alan Durant, ‘The significance is in the selection’: identifying contemporary keywords’, *Critical Quarterly* 50, no.1-2 (2008): 122-42.

¹⁵ Stephen Heath, ‘Keywords: abstract, abstraction,’ *Critical Quarterly* 55, no.3 (2013): 6-7.

-
- ¹⁶ Williams, *Keywords* (2015), xxxvi.
- ¹⁷ Williams, *Keywords* (2015), xxxv.
- ¹⁸ Lesley Jeffries and Brian Walker, *Keywords in the Press: The New Labour Years* (London: Bloomsbury, 2018). The methodology is laid out in Ch.1.
- ¹⁹ Jeffries and Walker, *Keywords in the Press*, 23-25.
- ²⁰ Martin Montgomery and Carol Ting, ‘Habeas Corpus? Cultural Keywords, Statistical Keywords, and the Role of a Corpus in their Identification,’ *Critical Quarterly* September 2023, <https://doi.org/10.1111/criq.12734> (accessed 20 April 2024).
- ²¹ Shao Yanjun (邵燕君) and Wang Yusu (王玉王), *破壁书: 网络文化关键词* (Breaking the Wall: Keywords in Chinese Internet Subcultures 破壁书: 网络文化关键词) (Beijing: 生活书店出版有限公司, 2018).
- ²² Wang Yusu 王玉王, ‘Keywords (关键词) and Catchwords (流行语): Notes from China,’ *Analysing Political Catchwords*, March 2024, <https://fass.open.ac.uk/research/projects/analysing-political-catchwords/blogs/keywords-and-catchwords-notes-from-china> (accessed 20 April 2024).
- ²³ “Catchword, *N.*” OED, March 2024, <https://doi.org/10.1093/OED/4520086573> (accessed 21 April 2024).
- ²⁴ Notably, William Safire, *The New Language of Politics: An Anecdotal History of Catchwords, Slogans and Political Usage* (New York: Random House, 1968); Eric Partridge, *A Dictionary of Catch Phrases from the Sixteenth Century to the Present Day* (New York: Stein and Day, 1977 [1985]); Doris Craig, *Catch Phrases, Clichés and Idioms: A Dictionary of Familiar Expressions* (Jefferson NC: McFarland, 1990); Anna Farkas, *The Oxford Dictionary of Catchphrases* (Oxford: Oxford University Press, 2002).
- ²⁵ That growing frequency of usage is in itself socio-politically and culturally indicative is a relatively underexamined idea. Studies of the part high-, medium-, and low-frequency words play in language acquisition and linguistic memory may have some bearing on this argument, but is not taken up here.
- ²⁶ John Algeo, *Fifty Years ‘Among the New Words’: A Dictionary of Neologisms, 1941-1991* (Cambridge: Cambridge University Press, 1992), 2.
- ²⁷ Including our own: Suman Gupta and Ayan-Yue Gupta, “‘Resilience’ as a policy keyword: Arts Council England and austerity,” *Policy Studies* 43, no.2, 2022, <https://doi.org/10.1080/01442872.2019.1645325>; Suman Gupta, *Political Catchphrases and Contemporary History: A Critique of New Normals* (Oxford: Oxford University Press, 2022); Ayan-Yue Gupta, *The Pragmatics of Governmental Discourse: Resilience, Sustainability and Wellbeing* (London: Routledge, 2024).
- ²⁸ Google Ngram Viewer <https://books.google.com/ngrams/>
- ²⁹ See Eitan Adam Pechenick, Christopher M. Danforth, Peter Sheridan Dodds, ‘Characterizing the Google Books Corpus: Strong Limits to Inferences of Socio-Cultural and Linguistic Evolution’, *Plos One* 7 October 2015, <https://doi.org/10.1371/journal.pone.0137041>; Nadja Younes and Ulf-Dietrich Reips, ‘Guideline for improving the reliability of Google Ngram studies: Evidence from religious terms,’ *Plos One PLoS One*. 22 March 2019, <https://doi.org/10.1371/journal.pone.0213554> (accessed 21 April 2024).

-
- ³⁰ See ‘Frequency’, *Oxford English Dictionary* online, <https://www.oed.com/information/understanding-entries/frequency/> (accessed 21 April 2024).
- ³¹ Johannes Schmidt, 1872, *Die Verwandtschaftsverhältnisse Der Indogermanischen Sprachen* (Weimar: Hermann Böhlau, 1872)
- ³² Leonard Bloomfield, *Language* (London: George Allen and Unwin, 1933), 327-328.
- ³³ Peter Trudgill, ‘Linguistic Change and Diffusion: Description and Explanation in Sociolinguistic Dialect Geography’, *Language in Society* 3, no. 2, 1974, <https://doi.org/10.1017/S0047404500004358>, 215-246.
- ³⁴ For examples see Ford Lumban Gaol, Ardian Maulana, Tokuro Matsuo, ‘News consumption patterns on Twitter: fragmentation study on the online news media network’, *Heliyon*, October 2020, <https://doi.org/10.1016/j.heliyon.2020.e05169> (accessed 25th May 2024); Ana Lucía Schmidt, Fabiana Zollo, Michela Del Vicario, Alessandro Bessi, Antonio Scala, Guido Caldarelli, H. Eugene Stanley, Walter Quattrocchi, ‘Anatomy of news consumption on Facebook’, *Proceedings of the National Academy of Sciences*, March 2017, <https://doi.org/10.1073/pnas.1617052114> (accessed 25th May 2024)
- ³⁵ Literature on conversational dynamics and pragmatics is relevant here, see David Lewis, ‘Scorekeeping in a Language Game’, *Journal of Philosophical Logic* 8, no.1, 1979, <https://doi.org/10.1007/BF00258436>; Paolo Labinaz, ‘Brandom’s deontic scorekeeping model and the assertive family’, *Journal of Pragmatics* 128, 2018, <https://doi.org/10.1016/j.pragma.2018.03.006>
- ³⁶ Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, Illia Polosukhin, ‘Attention is All You Need’, *Conference on Neural Information Processing Systems* 30, December 2017, <https://doi.org/10.48550/arXiv.1706.03762>
- ³⁷ See Matan Eyal, Shoval Sadde, Hillel Taub-Tabib, Yoav Goldberg, ‘Large Scale Substitution-based Word Sense Induction’, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics* 1, May 2022, <https://doi.org/10.18653/v1/2022.acl-long.325>; Li Lucy, David Bamman, ‘Characterizing English Variation across Social Media Communities with BERT’, *Transactions of the Association for Computational Linguistics* 9, 2021, https://doi.org/10.1162/tacl_a_00383
- ³⁸ For a popular method of producing token embeddings, see Jacob Devlin, Ming-Wei Chang, Kenton Lee, Kristina Toutanova, ‘BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding’, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* 1, June 2019, <https://doi.org/10.18653/v1/N19-1423>
- ³⁹ Syrielle Montariol, Matej Matinc, Lidia Pivovarov, ‘Scalable and Interpretable Semantic Change Detection’, *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, June 2021, <https://doi.org/10.18653/v1/2021.naacl-main.369>, 4646