

wild cards, etc., and – importantly in *c* studies of current change – tree patterns. ICECUP contains a powerful query system, termed *zzy* *n* (FTFs). FTFs are ‘sketches’ of grammatical constructions that can be applied to the corpus to obtain an exhaustive set of matching cases. Figure 2 shows an example of an FTF which matches all instances of a VP followed by a subject complement (CS).⁷ This FTF matches the three nodes highlighted in Figure 1 above.

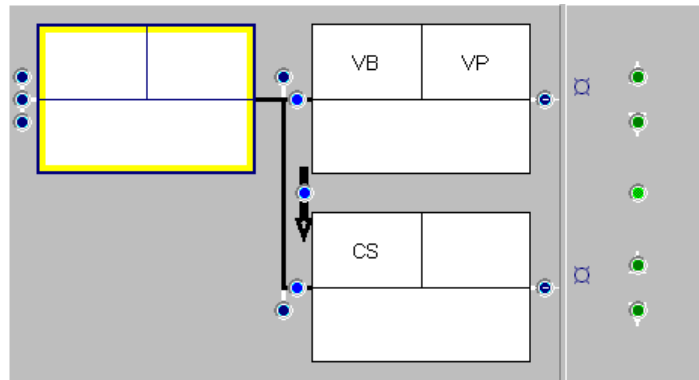


Figure 2: An FTF created with ICECUP, matching the highlighted nodes in Figure 1.

Respecting the fact that linguists disagree about grammar, ICECUP allows users to experiment with the best way of retrieving the grammatical phenomena they are interested in, using the Quirk-style representation in the corpus. The interface is designed to let linguists construct FTFs, apply them to the corpus, identify how they match cases in the corpus, and refine their queries. One can also select part of a tree structure and construct an FTF query from that fragment in order to find how a particular lexical string is analysed, and then seek all similar analyses.

ICECUP offers a range of search tools based around this idea of an abstract ‘FTF’ query, including a lexicon and ‘grammaticon’. DCPSE is an unparalleled resource for linguists interested in short-term changes in spoken English, and in this paper we will demonstrate its value in studies of current change using the examples of the progressive and the *vs.* alternation.⁸

3 Focusing on true alternation: the progressive

For decades, research in the field of sociolinguistics has highlighted the importance of the linguistic variant (see Labov 1969). This impetus has percolated into historical studies of language, but is often overlooked in corpus linguistics. Many studies on current change that have been carried out using corpora have collected frequencies for lexical items or grammatical constructions, but often without considering these frequencies alongside the variants of these patterns as part of a ‘bigger picture’. In the next three sections we look at a number of methodologies for exploring change. First we look at an approach which measures change in the progressive construction using normalised frequency counts. In section 3.2 we then look at a measure which investigates frequency changes as a percentage of the total number of VPs. Section 3.3 considers changes within a set of variants.

3.1 Changes in frequency per million words

Leech (2003) and Smith (2003) both investigate changes in the modal system of English. They carry out a series of independent log-likelihood ‘goodness of fit’ tests for the item,⁹ in this case a modal auxiliary, against the number of words in the corpus, using a method owing to Rayson (2003). This tests whether a perceived difference in a distribution *d* is too large to be explained by accident.

⁷ While the grammar that underlies the ICE-GB parsing (Quirk

When we evaluate rates of progressive VP use, it is more accurate to consider changes in the rate per VP than in the rate per n lexical words. By taking this step we remove this VP density variation, and thereby eliminate the possibility that an observed change could be due to changes in VP density. The revised calculation looks something like the following.

	A: item VP(prog)	B: VP	C: rate (proportion)	D: increase d (LLC = 100%)
LL	2,973	63,314	4.70%	+21.36% \pm 5.46%
AL	3,294	57,801	5.70%	
AL	6,267	121,115	5.17%	

Table 2: Change over time of ‘VP(prog)’ as a proportion of the number of VPs.

Note that we have replaced citations per million words in Column C with the simple proportion (this does not affect the overall calculation). Our results obtain a similar increase (d) to Table 1, but we have eliminated the possibility that variation in VP density accounted for our results.

Changing the baseline frequency from words to an overarching grammatical class (such as VPs) can have a dramatic effect on results. For example, Aarts, Wallis and Bowie (forthcoming) plotted d values for modal auxiliaries can , may , etc. from DCPSE on a per million word and per modal basis and showed that results differed markedly – can rose as a proportion of all modals, but did not change significantly with respect to word frequency; $could$, $should$ and $would$ all fell with respect to word frequency, but this fall could not be distinguished from an overall decline in modal use.

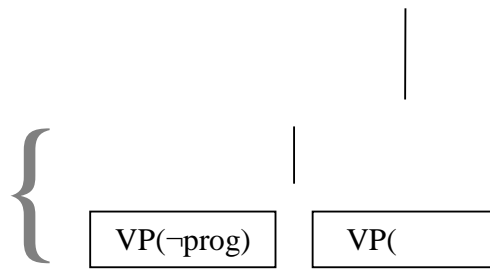
3.3 Changes in one choice out of a set of alternants

Ideally, we wish to evaluate how the progressive changes over time of n con con . The aim should be to focus our experiment on the set of true alternants to which the item in Column A belongs by removing as many distracting factors as possible. In this set of alternants, variation can be hypothesised to take place n members of the set, i.e. such that they compete and substitute for one another over time (Wallis 2003).

A study of modal auxiliaries should ideally therefore distinguish between semantic



their class of ‘progressivisable VPs’ at the point of their first citation, although these novel cases are unlikely to be sufficiently common to make a difference to an experimental outcome.¹¹





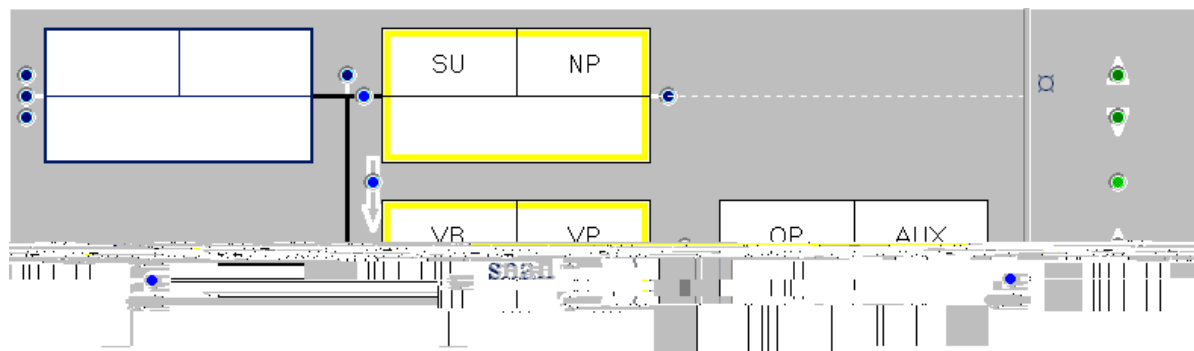


Figure 4: An FTF used to search for *no* after any subject NP.

A second, similar, FTF was used to retrieve instances of *no* and these cases were then subtracted from the results. We exclude all negative cases, including




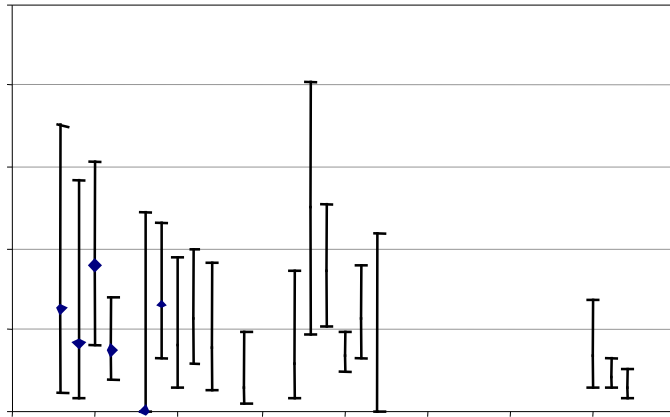


Figure 7a: Declining use of α as a proportion of the set $\{ \alpha, \beta \}$,





(5) Epistemic:

a. $\forall o \exists y f o \quad n y n n \quad o f \underline{L} n \quad o \quad o f \underline{L} y o n \quad c \quad c n$

(Epistemic modals). The fall in χ^2 is therefore not simply attributable to the sharp fall in Epistemic modals from 100 to 28: rather, we have evidence for a shift in use from Epistemic to χ^2 .

Root				χ^2 ()	χ^2 ()	y
LL	33	44	77	0.11	0.08	$d = -12.99\% \pm 38.83\%$
I	22	37	59	0.15	0.10	$\phi = 0.06$
AL	55	81	136	0.26	0.18	$\chi^2 = 0.32ns$

Table 8b: Analysis of change over time for first person declarative Root { , }. The results are not significant and the overall change ϕ is small. Percentage swing d represents the change over time in the proportion of cases of

instances of *o n o* which may alternate with each of the other variants. We use the FTF in Figure 8, again exploiting the parsed corpus. The grammatical annotation of the corpus makes a

Figure 9: Summarising changes for , , and *o n o*, first person positive declarative (non VP-final)

+22.13% ±5.48% (Table 1). Confidence intervals and significance tests are related. Since 22.13 – 5.48 > 0, the change is *n f c n y d f f n f o z o*, i.e. ‘significant’.

By far the most common method for calculating confidence intervals assume that repeated sampling at or around an observation obtains a symmetric, approximately Normally-distributed (‘Gaussian’) interval (Wallis 2009). The formula for the popular Gaussian single-sample interval is simply $z\sqrt{(1 -) n}$, where z is the critical value of the Normal distribution, and n the total number of observations.

However this rough approximation is rather inaccurate when an observation is very skewed (close to 0 or 1) or limited data is available. As approaches either 0 or 1, the confidence interval must tend toward the centre, because not only the v

In summary, employing Wilson intervals and Yates' tests improve precision where conventional χ^2 tests break down: in small, highly skewed datasets. These methods are particularly valuable for corpus linguists, who frequently deal with data of this kind.

A second methodological question

A second methodological question concerns the measurement of Δ , i.e. estimating the size of the change in the rate of 'VP(prog)', the decline of Δ , etc. Statistical significance tells us that the difference is unlikely to be zero (at a given level of confidence, see above). It does not tell us how large this difference actually is.

In the paper we quote the percentage increase (or decrease) d of a variant relative to the first subcorpus of DCPSE (the material from the LLC) with a baseline of 100 percent, and we calculate confidence intervals on d . This approach is relatively intuitive, but it can be misleading – not least because an increase of 20% (say) followed by a decrease of 20% does not bring you back to the start ($1 \times 1.2 \times 0.8 = 0.96$, not 1). It also has the rather unhelpful mathematical property of being unconstrained (it can have any value from minus to plus infinity).

In the statistics literature a number of measures of effect size are occasionally cited. These include the odds ratio, the contingency coefficient C and Yule's Y (Sheskin 1997: 244). A standard measure called Cramér's ϕ can be applied to any rectangular ($r \times c$) χ^2 contingency table. Like Pearson's



Sheskin, D. J. 1997. *Handbook of Computational Nonlinear Analysis*. Boca Raton, FL: CRC Press.