Are high-frequency collocations psychologically real? Investigating the thesis of collocational priming

PHILIP DURRANT and ALICE DOHERTY

Abstract

Words which frequently co-occur in language ('collocations') are often thought to be independently stored in speakers' minds. This idea is tested here through experiments investigating the extent to which corpus-identified collocations exhibit mental 'priming' in a group of native speakers. Collocational priming is found to exist. However, in an experiment which aimed to exclude higher-order mental processes, and focus instead on the 'automatic' processes which are thought to best reflect the organisation of the mental lexicon, priming is restricted to collocations which are also psychological associates. While the former finding suggests that collocations found in a large corpus are likely to have psychological reality, the latter suggests that we may need to elaborate our models of how they are represented.

Keywords: *collocation; priming; second language teaching*

1. Introduction

The phenomenon of high frequency collocation has been a focus of research by applied linguists for several decades. This interest has its roots in the work of Firth, who noted a "mutual expectancy" (1968, p. 181) between certain words, such that where we find one, we are likely to find the other. Researchers following Firth's lead have formalised the notion of collocation as "the relationship a lexical item has with items that appear with greater than random probability in its (textual) context" (Hoey, 1991, p. 7). That is, words are collocates of each other if, in a given sample of language, they are found together more often than their individual frequencies would predict (Jones & Sinclair, 1974, p. 19). Words which stand in such a relationship can be said to 'predict' one another in that the presence of one makes the presence of the other more likely than it would otherwise be.

Corpus Linguistics and Linguistic Theory 6–2 (2010), 125–155 1613-7027/10/0006–0125
DOI 10.1515/CLLT.2010.006

AUTHOR'S COPY | AUTORENEXEMPLAR

Though this definition is a textual one, grounded in corpus work, researchers have often moved beyond the text to posit collocation as a psycholinguistic phenomenon. Sinclair, for example, suggested that the existence of collocations in corpora indicates that language users store "a large number of semi-preconstructed phrases that constitute single choices, even though they might appear to be analysable into segments" (Sinclair, 1987, p. 319); Ellis (2001) has claimed that words which are frequently encountered together form 'chunks' in long-term memory; and Hoey (2005) has suggested that high-frequency collocates must mentally 'prime' each other.

Viewing collocation in this way has important implications for second language teaching. The idea that high frequency multi-word forms have some form of independent representation in the minds of competent speakers has motivated models of idiomaticity and fluency (Pawley & Syder, 1983), first and second language acquisition (Wray, 2002; Ellis, 2003; Tomasello, 2003), language processing (Ellis, 2002; Schmitt et al., 2004; Jiang & Nekrasova, 2007), and whole teaching approaches (Nattinger & DeCarrico, 1992; Lewis, 1993).

However, many researchers remain sceptical regarding the psycholinguistic reality of high frequency word combinations. One objection is that the frequency of combinations can be explained in terms of real-world coincidences, rather than requiring any special representation in the mind. This point has been put forcefully by Herbst (1996), who observes that the frequent collocations found in corpora may simply reflect certain extra-linguistic facts about the world. *Dark night*, he comments, is a significant collocation "because nights tend to be dark and not bright" (Herbst, 1996). Similarly Bley-Vroman (2002) has argued that recurrent patterns are merely a product of the use of language to express meaning in context, and do not have any strong direct explanatory force. Newmeyer sums up this position with the comment that frequency-based analysis "is no more defensible as an approach to language and the mind than would be a theory of vision that tries to tell us what we are likely to look at" (Newmeyer, 2003, p. 697).

A second problem with the link from corpus to mind is a pragmatic one: It is unlikely that any existing corpus matches the linguistic experience of any particular speaker, since the content of any given corpus will depend on the purpose for which it was compiled: large corpora such as the British National Corpus attempt to cover the whole range of discourse found in a particular national variety, while others concentrate on the language of a particular writer, or of a defined group such as second language learners or biologists. Therefore even if high frequency of co-occurrence in an individual speaker's language experience did lead to collocations being holistically stored, the inevitable differences between what is in any given corpus and what any given individual has experienced must lead us to question whether the collocations revealed by corpus analysis would be psychologically real for any individual speaker.

From a theoretical standpoint, this is not especially troubling. In a recent formulation of the frequency-representation link, Hoey acknowledges that the body of linguistic experience that provides each language user with their set of mental collocations is "irretrievable, unstudiable, and unique" (2005, p. 14). However, he reasons that although corpora cannot tell us about the primings of any particular speaker, they are likely to be representative of the types of input speakers are likely to have encountered, and so can be used as a testing ground for the validity of claims made about priming (2005, p. 14).

If we have more practical goals, however, this may not be sufficient. For example, a number of researchers (e.g., Ellis et al., 2008; Shin & Nation, 2008; Durrant, 2009) are interested in identifying lists of collocations which language learners need to know. For this, we will want to know more than just whether the collocations found in a corpus are similar *in kind* to those known to individual speakers. We will want to know whether learners need to acquire the specific high frequency collocations identified. Clearly, the set of collocations in any given corpus and the sets of collocations known to each individual speaker will be different. The key question is how much remains constant. If there is a strong overlap between the knowledge of large numbers of speakers and the corpus, then what is found in the latter is likely to be worth learning. If there is a large amount of variation between the knowledge of different individuals and what is found in a corpus, however, they probably are not. Establishing whether the high frequency collocations found in a corpus are psychologically real therefore has clear practical implications.

Although from a practical perspective it is important to determine whether high frequency collocations have psychological reality, little experimental work has been carried out into the psycholinguistic correlates of corpus frequency. The majority of work on the processing of formulaic language has focused on non-propositional (e.g., Van Lancker-Sidtis, 2004) and idiomatic (Gibbs & Nayak, 1989; Gibbs et al., 1989; Titone & Connine, 1999; Peterson et al., 2001) language, rather than frequent collocations. While recent years, have seen a number of studies looking at formulaic language derived from corpora (e.g., Schmitt et al., 2004; Schmitt & Underwood, 2004; Underwood et al., 2004; e.g., Jiang & Nekrasova, 2007; Conklin & Schmitt, 2008), the majority of these have not provided explicit frequency information, and so do not allow us to draw any strong conclusions on the implications of frequency. Exceptions to this rule are studies by McKoon and Ratcliff (1992) and Ellis and his colleagues (2008; 2009), to which we will return below.

Given this lack, Hoey's (2005) model of the relationship between collocations in text and collocations in the mind takes on considerable importance. Hoey describes collocation as "a psychological association between words" which is "evidenced by their occurrence together in corpora more often than is explicable in terms of random distribution" (2005, pp. 3–5). On this model, psychological association is measurable in terms of the psycholinguistic notion of 'priming'. This is the phenomenon, first documented by Meyer and Schvaneveldt (1971), whereby recognition of a word is facilitated by its preceding context, such that a hearer or reader recognises a given word faster if they have previously seen or heard a word which is related to it in some way. Thus, the word *girl* is recognised more quickly when it comes soon after the word *boy* than it does when it follows a semantically unrelated word. In such cases, the context is said to *prime* the target word. The existence of priming between words has been interpreted by psycholinguists as indicating facts about the organisation of the mental lexicon. In particular, it is commonly held that speeded reaction to primed words is a result of neurological activation 'spreading' from the context word to related words (see Neely, 1991 for a review of various theoretical interpretations of the priming effect).

Hoey suggests that speakers produce collocations in a predictable manner because of priming relationships between word pairs. That is, when one part of a collocation is brought to mind (e.g. *rain*), then recall of the other part (e.g. *heavy*) will be prompted. The speaker will therefore be more likely than usual to also produce this collocating form (so completing the collocation *heavy rain*), and is likely to prefer it over other possible synonyms (e.g. *strong*). Hoey argues that such priming is also the basis of our creative language system. The grammatical categories assigned to words, he suggests, are not specified by an independently-existing grammar. Rather, they emerge from lexically-specific patterns of priming. Labels such as 'noun' or 'verb' simply describe sets of words which share certain characteristic, genre independent primings. Nouns, for example, tend – amongst other things – to have *a* and *the* as left-hand collocates, and *of* as a right-hand collocate (2005, p. 154). This view has much in common with the emergentism of usage-based models (Barlow & Kemmer, 2000).

Hoey's suggestion is important because it provides an explicit and readily testable model of how frequency might relate to mental representation. The aim of the present paper is to test his thesis. In the next section, the statistical methods which have been used to quantify frequency of collocation will be briefly reviewed. The following sections then consider how frequency relates to mental priming. In Section 3, the priming paradigm will be described and evidence of the relationship between corpus frequency and priming will be reviewed. We will see that there is at present little evidence for such a link. It will be argued that a key issue here concerns the failure of previous studies to untangle high frequency of co-occurrence from the related phenomenon of psychological association. Section 4 will then describe a series of experiments

AUTHOR'S COPY | AUTORENEXEMPLAR

which test directly the link between co-occurrence in a corpus and mental priming.

2. Frequency-based measures of strength of collocation

The simplest frequency-based method of establishing whether a particular word combination is a collocation is to count the number of times the combination occurs. The problem with this is that, using this definition, many of the strongest collocations in any corpus would simply be those made up of the most frequent words; amongst the strongest collocations in any English corpus, for example, would be *a-the, of-and* and *to-was*. Such combinations are frequent, not because the words stand in any particularly interesting relationship to each other, but simply because they are so common that their regular co-incidence comes about by chance. Moreover, the simple frequency-based approach to collocation also misses word pairs which we might consider collocationally interesting, since strongly associated word pairs composed of words which are individually rare (*zero-sum game, abject poverty*) would not register at all. Corpus linguists have used two main types of methods to improve on raw frequency counts: asymptotic hypothesis tests and mutual information. The two approaches are conceptually different and typically produce rather different types of results.

The main hypothesis testing methods of quantifying collocation are the *z*-score, *t*-score, *chi-squared* and *log-likelihood* tests. These test the null-hypothesis that words appear together no more frequently than we would expect by chance alone. They can therefore be seen as formalisations of Hoey's (1991: 7) definition of collocations as "the relationship a lexical item has with items that appear with greater than random probability in its (textual) context."

The first step in each of these hypothesis testing methods is to calculate how many times we would expect to find a word pair together in a corpus of a certain size by chance alone, given the frequencies of its component words. To calculate this, we first determine how probable it is that any word pair, chosen at random from the corpus, will be the combination we are studying. This is calculated with the formula:

$$P(w1w2) = P(w1) * P(w2)$$

This states that the probability that any randomly selected pair of words will be the combination *w1 w*2 is equal to the probability of *w*1 occurring on its own multiplied by the probability of *w2* occurring on its own. For example, the word *strong* appears in the British National Corpus (BNC) 15,768 times, the word *tea* appears 8,030 times. Since the BNC has a total of 100,467,090 words, we can calculate the probabilities of occurrence of each as follows:

$$P(strong) = \frac{15{,}768}{100{,}467{,}090} = .00016$$

$$P(tea) = \frac{8{,}030}{100{,}467{,}090} = .00008$$

This tells us that if we select any word at random from the BNC, the probability that it will be the word *strong* is 0.00016 and the probability that it will be *tea* is 0.00008. We can then conclude that the probability that any two words, picked at random from the BNC will be the pair *strong tea* is:

$$P(strong\ tea) = 0.00016 * 0.00008 = 1.25\text{e-}08$$

Although this probability is very low, we would still expect *strong* and *tea* to occur together sometimes, simply because the corpus is so large. In fact, *strong tea* can be predicted to occur:

$$1.25\text{e-}08 * 100{,}467{,}090 = 1.26 \text{ times}$$

Since a corpus search shows that strong tea actually occurs 28 times, we can conclude that the pair collocates more frequently than chance. The aim of the hypothesis testing methods is to determine the statistical significance of this apparently greater than chance frequency (Manning and Schütze, 1999: 162–163). To take the most widely-used example, *t-score* is calculated as follows:

$$t\text{-}score = \frac{O - E}{\sqrt{O}}$$

where O is the observed frequency of occurrence of the collocation, and E is the expected frequency of occurrence on the null hypothesis that there is no relationship between the words. Thus, for the pair *strong tea*:

$$t\text{-score} = \frac{28 - 1.26}{\sqrt{28}} = 5.05$$

Some caution is needed in interpreting this formulation, however. The calculation of expected occurrence is based on a model in which words are drawn as if from a hat, entirely at random. As Manning and Schütze note, however, language is far more regular than a "random word generator" (1999: 166). Grammar, semantics, and real-world occurrences all constrain the construction of real language. It is therefore very common for word pairs to co-occur 'more frequently than random', regardless of specifically collocational relations. Given this, levels of statistical significance are not usually thought to constitute useful cut-off points in identifying collocations. Rather, the statistical tests are used to *rank* word pairs according to their relative likelihood of being a collocation (Stubbs 1995: 33; Manning and Schütze 1999: 166).

The second main method for quantifying collocations is the *mutual information* (MI) score, proposed by Church and Hanks (1990). This compares the observed number of occurrences of a word pair with its expected number of occurrences, as follows:

$$MI = \log_2 \frac{O}{E}$$

Thus, for *strong tea*:

$$MI = \log_2 \frac{28}{1.3} = 4.43$$

Mutual information can be conceptualised as a "measure of how much one word tells us about the other" (Manning and Schütze 1999: 178). In other words, when we encounter one part of a word pair which has a high mutual information score, we can predict that the other part of the pair is likely to be nearby. This is importantly different from the hypothesis testing methods. Clear (1993: 279–282) sums up the difference well, noting that, whereas "MI is a measure of *the strength of association between two words*", hypothesis testing methods are measures of "*the confidence with which we can claim there is some association*" (original emphases). This difference has important implications for the types of word pairs retrieved. Clear gives the pair *taste-arbiters* as a typical example of a combination attaining a high *MI* score. Though the pairing is not particularly frequent, it accounts for a high proportion of the occurrences of its component words. In fact, Clear reports that one quarter of all appearances of *arbiters* are within two words of an appearance of *taste*. The two are strongly associated, then, in that where we find *arbiters*, we have a good chance of finding *taste*. However, its relatively low frequency of occurrence reduces the statistical reliability of this pattern – i.e. we cannot be confident that the relationship will be generalisable to other samples of language. A typical example of a pair with a high score on hypothesis testing methods, on the other hand, is *taste-for*. While the association between these words is much weaker than that between *taste* and *arbiters*, the pair occurs with much higher frequency. The connection, though weaker, is therefore more reliable.

It has been acknowledged in the literature that both of these measures of collocation are 'non-directional', in the sense that it makes no difference which member of the pair is taken as node and which as collocate (Stubbs, 1995). Clearly, however, *arbiters* predicts *taste* much more reliably than *taste* predicts *arbiters*. When our aim is to determine how reliably one word primes another, such directionality becomes important. For this reason, directional association measures, such as that suggested by Handl (2008) may be more accurate predictors of priming. However, such measures have not been widely adopted in corpus linguistics, and since our present aim is to determine how well the commonly-used measures of corpus linguists predict priming, we will not use directional measures here. In the next section we review the evidence regarding how well traditional non-directional measures predict mental priming.

3. Evidence on collocational priming

Priming has been researched for several decades, and much is known about its workings. Priming has been shown to exist between words with similar orthographies and phonologies, between words which are related in meaning, and between syntactically congruous words (e.g. determiner – noun) (Balota, 1994: 334–341). Priming is usually investigated by some variation of the two-stage task described by Neely (1991: 265): first, the informant is presented with a single word (the *prime*), to which they are not required to make any overt response. Second, they are presented with a letter string which may or may not be a real word (the *target*), and are required to respond either by making a word/nonword decision (the 'lexical decision task', or 'LDT'), or by saying the target aloud (the 'naming task'). When the target is a word, it is either related or unrelated to the prime. When reaction times or percentage errors are significantly lower for the target word following presentation of the prime the latter is said to prime the former. Priming effects are assumed to reflect the structure of the mental lexicon. Where a priming effect occurs between two words this is considered evidence that there is a structural connection between them.

Much work has been concerned with separating the effects of semantic relations (the relations found between, for example, categories and their exemplars, co-hyponyms, or words with featural similarities between their concepts) from the effects of 'associative' relations (see Lucas 2000; Hutchison 2003 for reviews). Association is the relationship between a word and other words which it 'brings to mind'. This is typically evidenced by studies in which large numbers of subjects are given a prompt word and asked to respond with the first other word that comes to mind. Results from such studies are available as listings of 'normative' associations (e.g. <http://www.eat.rl.ac.uk/>).

The primary aim of this work has been to distinguish between models of the mental lexicon in which word form and word meaning are integrated, such that purely semantic information can influence word recognition, and models in which the two types of information are strictly separated, such that semantic information cannot feed back to aid recognition of form. The former model predicts that pure semantic priming is possible; the latter predicts that it is not (Lucas 2000: 618–619). For present purposes, however, the issue of semantic vs. associative priming is of interest because it has been claimed that associative priming may be a product of the frequent co-occurrence of associated items in text (Charles and Miller 1989).

If this were right, then the existing evidence for association priming would constitute direct experimental support for Hoey's collocational priming thesis. However, though collocation and association are known to be linked, they are far from identical: not all collocations are associates, and not all associates are collocations (Fitzpatrick, 2007). Re-analysis of the prime-target pairs used in studies of associative priming confirms the distinction. In Shelton and Martin's (1992) influential study, for example, pairs in the 'associated' condition have a wide range of *MI*-scores, from the very low (*min* = 0.90) to the moderately high (*max* = 8.85), with a median score which is above the traditional cut-off point of 3, but only modestly so (*Mdn* = 3.98). The same remarks apply to other measures of collocational significance (*t*-score*: min* = 1.81, *max* = 24.34, *Mdn* = 4.87).

A particular problem with using data on association priming as evidence for collocational priming is that those collocations which appear prominently on association norms are likely to be especially salient in a way that other collocations are not. Even if an associative priming study were based entirely on collocating items therefore, the results may be generalisable only to the most salient of collocations. If this is the case, priming may not be a suitable paradigm for understanding collocation in general. Studies of associative priming cannot, therefore, provide adequate support for the thesis of collocational priming.

There are only a handful of studies which have aimed specifically to study collocational priming. Hodgson (1991) found priming for 'phrasal associates' (e.g. 'private-property', 'vacant-building', 'arm-chair') in a lexical decision task, but not in a naming task. However, since he does not give any information regarding how his test items were identified as phrases, and in particular does not provide any frequency data for the items, his results have little to tell us about the relationship between high frequency collocation and priming. Moreover, as Hutchison (2003: 789) points out, Hodgson combines his results for phrasal associates with those from five other types of relations (synonyms, antonyms, conceptual associates, co-ordinates, and super/subordinates), reporting only the overall priming effects for all six types of pairs and not the significance of each type alone. Even if corpus data were available, therefore, the study would not be able to tell us about the effects of collocational priming alone.

Williams (1996) finds significant priming for word pairs which were graded by native speaker informants as 'highly familiar' in conjunctive phrases (i.e. 'X and Y'). As in the Hodgson study, however, insufficient frequency data is provided for the nature of any link between corpus and priming to be stated. Though Williams reports a post-hoc frequency analysis of his items using the one million-word Lancaster-Oslo/Bergen corpus, he is only able to tell us that the pairs were found in the same sentence, on average, 2.94 times (range = 0–11), and within a span of +/−2, on average, 2.06 times (range = 0–6). The small size of the corpus employed and the variability of the results (with some 'collocations' not being attested together at all) mean that it is again difficult to draw any strong conclusions regarding the effects of frequent co-occurrence. Moreover, all of Williams' collocations were attested to be strong associates. It is therefore possible that the priming found is limited to those collocations which are sufficiently salient to be normative associates. A potentially helpful result in this context was the finding that associates that were posited to be collocations produced a greater priming effect (18ms) than did associates which were not collocations (3ms), hinting at an independent effect of collocation over and above that of association alone. However, the difference between the two sets, though large, was not significant. This was because, whereas the eight most strongly associated collocations produced very strong priming effects, the other eight associates showed minimal effects. Williams suggests that priming may be obtainable only for those collocates with "the highest psychological salience" (Williams, 1996: 133). Concluding that collocates which are sufficiently psychologically salient to produce priming are likely also to be strong normative associates, he abandons any attempt to separate the two effects (1996: 134).

Unlike Hodgson and Williams, McKoon and Ratcliff (1992, Experiment 3) used corpus-derived frequency data in compiling collocations for their priming experiment. They compared priming between associated primes with that between collocating pairs. Collocations were identified as such on the basis of *MI* scores calculated from frequencies of occurrence in a 6 million word AP newswire corpus. McKoon and Ratcliff also considered the priming effects of collocations with low *MI* scores. The study found that the greatest facilitation in comparison with an unrelated prime was between associated primes (49ms); significant priming was also found in the collocating condition (21ms); collocations with low *MI* also speeded recognition (17ms) though this improvement on the unrelated condition was not significant.

While acknowledging that the small size of their corpus must bring the reliability of their findings into some doubt, McKoon and Ratcliff (1992: 1164) tentatively conclude that "co-occurrence statistics calculated from large corpora have potential applicability as predictors of priming effects." However, a re-evaluation of their 'high' *MI* collocations against data from the British National Corpus, confirms their worries regarding reliability. Indeed, the collocations in this study are no stronger (and if anything, rather weaker) than those found in Shelton and Martin's study (*MI*: *min* = −3.26; max = 9.64; *Mdn* = 3.43; *t*-score: *min* = −8.58; *max* = 25.86; *Mdn* = 3.90). Moreover, as in Williams' study, many word pairs used in McKoon and Ratcliff's collocating condition were – as the authors acknowledge – also strong associates of the targets. Again then, it is not clear whether the priming effects found are attributable to collocations in general or are restricted to normative associates.

More recently, Ellis et al. (2008) have studied the influence of both frequency of occurrence and mutual information score on priming for the final word of 3, 4, and 5-word sequences. Participants were briefly shown either an incomplete phrase or a series of 'X's. This was then replaced with a single word, which they were told to pronounce into a microphone as quickly as possible. Where the word was the completion of the preceding part-phrase, speed of response amongst native speakers was shown in a multiple regression to be predicted by *MI* score ($\beta = -0.47$), though not by frequency. The same experiment performed with non-native speakers showed no priming effects. This study provides an interesting extension to previous results by studying longer collocations and by separating out the influences of frequency and mutual information scores. However, as with Ratcliff and MacKoon's study, no attempt is made to separate collocational priming from the well-known effect of association.

Finally, Ellis et al. (2009, Experiment 1) used a lexical decision task to determine whether high frequency adverb + adjective and verb + object pairs were recognised more quickly than lower frequency pairings. They found that recognition times were negatively correlated with the frequency of the word pairs, suggesting a processing advantage for collocations. They argue on this basis that the so-called 'semantic priming' seen between associated pairs is in fact due to collocational association. However, as we have already seen, collocations and semantic associates are only partially overlapping sets. As in the other studies discussed here, no attempt is made to separate out associated from non-associated collocations in the stimulus items, so it is not clear that their inference from one to the other is legitimate.

In sum, while some psycholinguists have suggested that collocational priming exists, we do not yet have a clear picture of the phenomenon. Most previous studies have not provided sufficiently robust frequency data to evaluate any claim of a link between priming and frequency of occurrence in a corpus. The few studies which have provided such information suggest that high frequency collocations are psychologically real. However, these results require broader confirmation. Moreover, previous studies have made no attempt to distinguish the effect of collocation from that of psychological association. The studies presented here aim to make up for these shortcomings by exploring the extent to which high frequency collocations identified in the British National Corpus are reflected in the priming of native speakers of British English, regardless of whether they are strong normative associates.

4. Priming studies

4.1 *Introduction*

We have argued that it is important to study empirically the relationship between the frequency of occurrence of collocations in a corpus and their representation in the minds of proficient speakers. We have also noted that Hoey's (2005) theory of collocational priming offers an explicit and testable model of this relationship. However, we have seen that the existing experimental literature does not provide sufficient evidence to support Hoey's theory fully. Though there is good evidence for priming between collocations which also appear on lists of word associations, it is not clear whether this effect can be generalised to high-frequency collocations as a whole, or whether it is restricted to these specially salient items. The present experiments aim to test this.

4.2 *Experiment One*

4.2.1 *Introduction*

The first experiment uses a lexical decision task of the sort described by Neely (1991, see Section 3 above) to determine whether priming can reliably be found between high-frequency collocations, and whether any such effect is independent of whether collocations are psychological associates. Participants are briefly shown a prime word, followed by a second letter string which either is or is not a real word in English. Their task is to decide, as quickly as possible, whether the second string is a real word and to indicate their decision by pressing one of two buttons. When the string is a word, it is, according to corpus data, either a collocate of the prime word or a word which is not attested following the prime in the corpus. Collocational priming will be indicated if participants respond more rapidly to collocate than to non-collocate targets.

Our use of words which are not attested following the prime as control items requires some explanation. This specification leads to comparisons such as:

parish church (collocation) vs. *feature church* (control)
spoken word (collocation) vs. *lower word* (control)

In such comparisons, the contrast is not simply between high frequency collocations and low frequency pairings, but between high frequency collocations and word pairs which are rather difficult to make sense of semantically. It might be felt therefore that a more meaningful control group would be items which are attested in corpora, but with insufficient frequency to qualify as collocations (e.g. *biggest church, stupid word*). However, piloting suggested that such items made poor controls for studying priming, since the well-attested phenomenon of priming between associates was not detectable in comparison with such controls. A search of the literature also indicated that previous evidence for priming between associated words appears to have been based entirely on comparisons of associates with 'incongruous' word pairs. All of the associative priming studies with which we are familiar and which report their method of creating 'unrelated pairs' have followed the lead of Meyer and Schvanevendlt's original paper (1971) by interchanging words from the associated items such that there are "no obvious associations within the resulting pairs" (1971: 228). Most studies do not list the items used in the control condition, but Meyer and Schvaneveldt provide the illustrative examples of BREAD-BUTTER and DOCTOR-NURSE being re-paired as BREAD-DOCTOR and NURSE-BUTTER.

Since such a strong and well-attested phenomenon as associative priming could not be evidenced in comparison with such controls, it was decided to follow the common practice of using non-attested pairs of the sort exemplified above. A disadvantage of this approach is that it raises the possibility that any priming effect may be the result of a difference in reaction to meaningful vs. incongruous word pairs, rather than of any specifically collocational priming effect. To counter this possibility, the incongruous control pairs were compared not only with collocations, but also with attested, but low-frequency combinations. If priming is the result of a contrast between semantically congruous and incongruous pairs, it should also be visible in these comparisons.

The experimental pairings are of four types:

1) *Low frequency combinations*: word combinations which appear in the corpus but are not sufficiently frequent to count as collocations;
2) *Moderate collocations*: moderately frequent collocations which are not strong associates;
3) *Frequent collocations*: very frequent collocations which are not strong associates;
4) *Associated frequent collocations*: very frequent collocations which are strong associates.

As discussed above, the low-frequency combinations were included as a control condition to check whether any priming exhibited for collocations was also true of congruous word combinations in general. Moderate and frequent collocations are separated in order to test whether any priming effect is restricted to the most frequent collocations. The set of associated frequent collocations is included to determine whether there is a difference in priming between associated and non-associated collocations.

4.2.2 *Materials*

16 collocating prime-target pairs were created for each of the three conditions described in Table 1.

Frequency measures were derived from British National Corpus (BNC). To determine whether pairs were strong normative associates or not, two methods were used. First, the Edinburgh Association Thesaurus (EAT) was consulted.[1] Pairs in the moderate and frequent collocation conditions were deemed not to be strong associates only if neither the target word nor any word form derivationally or inflectionally related to the target was listed as an associate of the prime in the Thesaurus; similarly, neither the priming word nor any word form derivationally or inflectionally related to the prime was listed as an associate of the target. Pairs in the associated frequent collocations condition were judged to be strong associates only if the target word was listed as either the first or second strongest associate of the prime in the Thesaurus and had a minimum association score of 10% (i.e. was supplied by at least 10 out of 100 respondents). The mean score of the associations used was 26% (range: 10%–66%). Some of the associated pairs also exhibited 'backward association'. That is, the target also prompted the prime. Seven of the 16 associates were recorded as backward associates, though only two of these had an association score of more than 10% (*match-football*: 12%; *list-shopping*: 17%). The forward and backward association scores for all pairs are listed in the Appendix.

Because the EAT was elicited from a different population from that taking part in the present study (EAT was elicited from British university students between 1968 and 1971), word pairs which are not attested as associates in the EAT may nevertheless be strongly associated for our participants (British university students in 2007). Similarly, some pairs which are prominent on the EAT may not be strong associates for these participants. To take some simple examples, it is unlikely that if the EAT experiment were repeated with our participants, the second most common associate of *politics* would be *Wilson*, or that *mobile* would fail to elicit the response *phone*. Moreover, because the EAT elicited only a single response from each participant, it is possible that some highly salient collocations do not feature on its listings. To ensure that 'non-associated' pairs in the current study were indeed not likely to be strong associates for our participants, and that the putatively associated pairs were, a second test was used in addition to the EAT check. Two groups of 22 subjects were each presented with 40 stimulus words (a different stimulus list for each

Table 1. *Experimental conditions for Experiment 1*

	Examples	Frequency	*MI* score	*t*-score	Psychological associate?
Low frequency combinations	*direct-danger; weak-ground; famous-saying*	2–4 (*Mdn* = 2)	<2 (*Mdn* = 0.91)	<1.5 (*Mdn* = 1.46)	No
Moderate collocations	*greater-concern; likely-effects; recent-figures*	21–65 (*Mdn* = 33)	4–5 (*Mdn* = 4.47)	4–8 (*Mdn* = 5.52)	No
Frequent collocations	*foreign-debt; colour-scheme; mental-picture*	60–357 (*Mdn* = 121.5)	>6 (*Mdn* = 7.65)	>7.5 (*Mdn* = 10.95)	No
Associated frequent collocations	*estate-agent; cutting-edge; card-game*	38–1324 (*Mdn* = 117.5)	>5.5 (*Mdn* = 7.01)	>6 (*Mdn* = 10.63)	Yes

group, giving a total of 80 stimuli) and asked to write down the first three words which came to mind on reading each stimulus. The subjects were taken from the same pool as those participating in the main priming experiment (although none of the participants in this test took part in any of the main priming experiments), and so should provide a good indication of the likely associates of the main study participants. Moreover, by eliciting three associates for each stimulus, we may move a little beyond the very strongest associates. Association scores were calculated by awarding three points to a first choice word, two points to a second choice word, one point to a third choice word, and then finding the average score across all participants. Thus, a word pair could receive an association score of between 0 (not listed by any participant) and 3 (the first word listed by all participants). Pairs in the moderate and frequent collocation conditions were deemed not to be strong associates only if neither the target word nor any word forms derivationally related to the target was supplied as an associate of the prime. Pairs in the associated frequent collocation condition were judged to be strong associates only if the target was supplied as an associate of the prime by at least two respondents. The mean association score was 1.12 (range: 0.27–2.09) and the targets were supplied on average by 9.4 out of 22 respondents (range 2–17).

The 48 target nouns from the three collocation lists were also matched with 48 control primes. The control prime-target pairs were combinations which did not co-occur in the corpus. No very common or very rare words were used as targets or as primes: all words used occurred in the BNC between 3,000 and 30,000 times; placing them well outside the top 300 word forms in the corpus and well within the top 3,500 (Leech et al., 2001). All words were one or two syllables (four to seven letters) in length. Since the target items were the same in both conditions, response times should be identical in each, barring any priming effects.

The collocating and control primes were combined into two counterbalanced lists such that eight collocating pairs from each level were included in each list and targets which were matched with their collocating prime in one list were matched with their control prime in the other. No prime or target word was used more than once in either list. A single set of 48 prime-non-word pairs was also added to both lists. Non-words were items of four to seven letters, generated using the ARC Nonword Database (Rastle et al., 2002). Primes were items which appeared in the BNC between 3,000 and 30,000 times and which were attested to be used as pre-modifiers but which had not been used elsewhere in the experiment. The final materials are shown in the Appendix.

4.2.3 *Participants*

32 students at the University of Nottingham. All were native speakers of British English.

4.2.4 *Procedure*

Participants were randomly assigned in equal numbers to one of two groups. Each group saw one list only in a randomised order. Participants were tested individually in a quiet room. Presentation of the stimuli and recording of the reaction times were controlled by Psychology Software Tools' *E-Prime* software and items were displayed on a CRT monitor. On each trial, a fixation point ("+") was presented, centred on the screen, for 1,500ms. This was replaced with a priming word, which was presented in lowercase letters for 600ms. The prime was then immediately replaced by the target, in uppercase letters. The target stayed on the screen until the participant made a response. Following the response, the screen went blank for 1,000ms before the onset of the next trial. Participants were instructed to press the right button on a button-box if the string was a word and the left button if it was not. They were told to make this decision as quickly as possible. Reaction times were measured from target onset to response. Participants received 10 practice trials, and were allowed a self-timed break before commencing the experiment. Items were presented in two blocks, with a self-paced break between blocks. Each block contained an equal number of items from each level, and the order of presentation of the blocks was counterbalanced between participants.

4.2.5 *Results and discussion*

Reaction times of less that 200ms or more than 3 standard deviations above the participant's mean (2.05% of the total) were replaced with the mean for that participant. Mean accuracy by participant, collapsed across conditions was 96%. There was no effect of condition upon accuracy. Average reaction times for collocations and non-collocations at each of the four levels are shown in Table 2. Reaction times were negatively skewed within conditions—i.e. there was a long 'tail' of participants with reaction times well above the mean. Because of this, times were not normally distributed within conditions and so parametric tests could not be used.

Table 2. *Average reactions times in each condition for Experiment 1*

	collocations	non-collocations	Wilcoxon signed ranks test
Low frequency combinations Median RT (ms)	522	510	$T = 194, p = ns, r = -0.16$
Moderate collocations Median RT (ms)	512	520	$T = 243, p = ns, r = -0.05$
Frequent collocations Median RT (ms)	520	544	$T = 106, p < 0.001, r = -0.37$
Associated frequent collocations Median RT (ms)	507	517	$T = 115, p < 0.001, r = -0.35$

No statistically significant difference was found between collocation and non-collocation conditions in the low frequency or moderate collocation conditions. However, in the frequent and associated frequent collocation conditions, collocating pairs exhibited significant priming in comparison to non-collocations.

These results suggest two main conclusions. First, priming was seen here only between collocations with very high frequencies of occurrence, i.e. those in the frequent (*Mdn MI* = 7.65, *Mdn t*-score = 10.95) and associated frequent collocation conditions (*Mdn MI* = 7.01, *Mdn t*-score = 10.63). More modest collocations, and low frequency word pairs did not demonstrate significant priming relative to implausible controls. The corpus literature (Stubbs, 1995; Hunston, 2002) often cites an *MI* score of 3 and *t*-score of 2 as minimum cut-off points above which collocations are likely to be of linguistic interest. The present results suggest that much higher thresholds are required to identify collocations which are likely to demonstrate priming. Second, priming between high frequency collocations has been shown to exist regardless of whether the pair is likely to be a strong psychological associate. Frequent collocations and associated frequent collocations demonstrated approximately equal levels of priming in terms of their effect sizes (Level 2: $r = -0.37$; Level 3: $r = -0.35$), suggesting that the effect seen here is entirely independent of association.

4.3 *Experiment Two*

4.3.1 *Introduction*

A limitation of the first experiment is that it did not preclude the intervention of 'strategic' processes on the part of subjects. That is to say, it is possible that during the experiment subjects attempted to find relationships between primes and targets and in light of their hypotheses adopted strategies to hasten their responses (e.g. by attempting to guess possible upcoming targets). It has been suggested in the literature that such processes may not reflect the long-lasting organisation of the lexicon, but rather are the ad-hoc products of the experimental task, controlled by higher-order mental faculties (Lucas, 2000: 619).

Two main methods have been used to elicit automatic priming. The first presents items one-by-one and requires participants to make a word/non-word decision for each, so removing the overt pairing between prime and target. McNamara and Altarriba (1988) and Shelton and Martin (1992) both report results suggestive of automatic priming using this technique. A second method is to present primes only very briefly, and preceded and/or followed by a 'pattern mask' (e.g., "######"). Under such conditions, subjects are not usually conscious of the prime so cannot make use of conscious strategies. Several studies have found evidence of apparently automatic priming using such methods (de Groot & Nas, 1991; Sereno, 1991; Perea & Rosa, 2002). Of the two methods, the latter is preferred here since the former, by requiring an explicit response from participants to every word seen, seems likely to encourage participants to process words as isolated items, so undermining any natural collocational processing.

4.3.2 *Materials*

The materials used in this experiment are identical to those in Experiment 1.

4.3.3 *Participants*

38 students at the University of Nottingham. All were native speakers of British English.

4.3.4 *Procedure*

The procedure was the same as that in experiment one except for the timing of exposure to the prime and the addition of a 'pattern mask' before the prime was shown. For each experimental item, after seeing a fixation point ("+") for 1,500ms, participants were shown a patterns mask (i.e., "#######") for 500ms. The priming word then appeared in place of the mask, in lowercase letters, for 60ms only (an exposure of this length is usually too short for participants to be aware of the prime). The prime was then immediately replaced by the target, in upper case letters. The target stayed on the screen until the participant made a response. Following the response, the screen went blank for 1,000ms before the onset of the next trial.

4.3.5 *Results and discussion*

Three participants claimed to be able to read the primes, so their results were excluded from the analysis. As in the previous experiments, reaction times of less that 200ms or more than three standard deviations above the average for a participant (1.83% of the total) were replaced with the participant's mean. Mean accuracy by participant, collapsed across conditions was 97%. There was no effect of condition upon accuracy. Average reaction times for correct responses to targets following collocations and non-collocations at each of the three levels are shown in Table 3. Reaction times were again (for the reasons described above) not normally distributed within levels.

As in Experiment 1, significant facilitation was found for associated frequent collocations; however, no facilitation was observed in any other condition. The lack of facilitation in the low and moderate frequency conditions replicates the findings of Experiment 2. However, the failure to find any priming between frequent collocations is more surprising, and forces us to

Table 3. *Average reaction times in each condition for Experiment 2*

	collocations	non-collocations	Wilcoxon signed ranks test
Low frequency combinations Median RT (ms)	556	581	$T = 236$, $p = ns$, $r = -0.18$
Moderate collocations Median RT (ms)	572	616	$T = 255$, $p = ns$, $r = -0.14$
Frequent collocations Median RT (ms)	563	583	$T = 243$, $p = ns$, $r = -0.17$
Associated frequent collocations Median RT (ms)	548	568	$T = 218$, $p < 0.05$, $r = -0.21$

revise somewhat our conclusion that priming is independent of psychological association. These results seem to suggest that priming between associates may be different in type from that between collocates, with the former controlled by automatic processes, and the latter by strategic.

If this is right, it is of considerable theoretical interest. As was noted above, it has often been suggested that priming between associates is a result of their frequent collocation (e.g., Charles and Miller, 1989). If the two types of word combination elicit different types of priming, however, this view would need to be revised. Some caution is required here, however. Although Level 3 collocations were the only type to pass the $p < 0.05$ threshold, there is, in contrast to Experiment 1, also some suggestion of facilitation at other levels. In all collocation types, median responses were lower in the collocating condition, with effect sizes which are roughly comparable across levels. Moreover, the low frequency and frequent collocation conditions both approached significance in the Wilcoxon signed ranks test (with p values of 0.067 and 0.083 respectively). While robust proof of it has not been found here, therefore, we cannot entirely discount the possibility that the other types of word combination tested may exhibit automatic priming. More research is required to address this issue.

5. General discussion

This paper aimed to determine whether frequency of co-occurrence in a large general purpose corpus is a reliable indicator of psychological priming between words. It was argued that, in the absence of such evidence, it is difficult to judge the validity of frequency-based methods for such tasks as identifying target collocations for language learners to study.

The results presented here appear to indicate that frequency of occurrence does indicate psychological reality, confirming the findings of Ellis and his colleagues (Ellis et al., 2008; 2009) and suggesting that frequency-based methods are a valid method for identifying collocations which are likely to be good targets for language learning. Moreover, Experiment 1 extends previous research by showing that for high-frequency pairs the size and statistical robustness of the priming effect is independent of whether collocations are or are not psychological associates. The relevance of frequency as an indicator of the collocations native speakers are likely to know appears therefore to be generalisable beyond the small group of highly salient items which are found on association lists. As far as we are aware, this is the first experiment to demonstrate that frequency-based collocational priming can exist independently of psychological association.

Experiment 1 also suggested that collocational priming may be restricted to word pairs which score very highly on association measures (*MI* and *t*-score). Priming was only demonstrated here between pairs with *MI* scores of at least 6 and *t*-scores of at least 7.5. In the Applied Linguistic literature, a *MI* score ≥ 3 and a *t*-score ≥ 2 are commonly cited as a useful threshold for identifying genuine collocations (Stubbs, 1995; Hunston, 2002). The current results suggest that, in a search for psychologically real collocations, such guidelines need to be revised upwards.

While these findings confirm the validity of frequency data from a general corpus as a indictor of which word pairings are psychologically real collocations for native speakers, an important caveat must be made regarding the exact psycholinguistic nature of these collocations. Our finding that both associated and non-associated collocations exhibit priming applied only in an experimental set-up which permitted the intervention of 'strategic' (i.e. higher-order) processes. In Experiment 2, which aimed to elicit only the automatic processes which are thought to reflect the long-lasting structure of the mental lexicon, high-frequency collocations did not exhibit a robust priming effect, whereas associated word-pairs did.

This finding suggests the need for further research into the exact nature of collocational priming. While we do not wish to enter into any extensive speculations on the basis of negative results, the contrasting results of Experiments 1 and 2 suggest the possibility of a theoretically interesting divergence in the way associated and non-associated collocations are represented in the mind which may be worthy of further exploration. This divergence suggests that the nature of the mental representation of collocations may not be uniform, but rather that some collocations may be represented in different ways than others. It also suggests that the commonly-repeated assumption that associative priming is a product of collocation may not be tenable, since associated pairs appear to have properties not shared by 'mere' collocations. Future research is needed

AUTHOR'S COPY | AUTORENEXEMPLAR

to pursue this divergence between association and collocation in more detail, and consider the extent to which association and collocation constitute separable priming effects.

A further caveat which must be made is that the present research focuses very much on written, rather than spoken language. This is true both in that the collocations used were archetypically written in nature (a result of the make up of the BNC, 90% of which comprises written language) and in that our procedure employed only written prompts. Further research may therefore wish to consider whether our are generalisable to the spoken medium[2].

As a final comment, it needs to be recognised that Hoey's model of collocational priming is only one possible account of the relationship between frequency of co-occurrence and psychological representation – other possibilities are described by, for example, Ellis (2001) and Schmidt (1992). While Hoey's model calls for experimental paradigms of the sort used here, it is quite possible that other paradigms will give us further insights into the link between frequency and representation. The experimental approach used here has some evident weaknesses. In particular, by presenting collocates as two individual words, divorced from any meaningful context, processing advantages which might be found in more natural situations may have been obscured. Moreover, the artificial nature of the word-recognition task, which does not correspond to any real-life language situation, may also have prevented normal processing effects. Further research may therefore wish to study the processing of collocations in larger contexts and with less intrusive measurement techniques. Such work will require a great deal of theoretical sophistication, however. In particular, once collocations are embedded in larger environments, it will be necessary to take account not only of the mutual predictability of the two words which make up a collocation, but also their probabilistic relationships with other (lexical, grammatical, and discoursal) items in the surrounding context. Such measures have not yet, to our knowledge, been developed.

Acknowledgement

This research was partly supported by a studentship from the Economic and Social Research Council.

Appendix: Experimental items

Condition	Experimental prime	Control prime	Target	Control pair			Experiment pair			Forward association (EAT)	Backward association (EAT)	Association (pilot)[1]	List
				BNC occur-rences	*t*	*MI*	BNC occur-rences	*t*	*MI*				
Low frequency combinations	main	armed	concept	2	0.30	0.34	0	n/a	n/a	0%	0%	N/A	1
Low frequency combinations	direct	market	danger	2	0.93	1.55	0	n/a	n/a	0%	0%	N/A	1
Low frequency combinations	market	direct	experts	3	1.18	1.66	0	n/a	n/a	0%	0%	N/A	1
Low frequency combinations	weak	past	ground	2	1.03	1.87	0	n/a	n/a	0%	0%	N/A	1
Low frequency combinations	fixed	pretty	levels	2	0.88	1.40	0	n/a	n/a	0%	0%	N/A	1
Low frequency combinations	royal	current	lunch	2	0.89	1.42	0	n/a	n/a	0%	0%	N/A	1
Low frequency combinations	huge	famous	powers	2	1.05	1.95	0	n/a	n/a	0%	0%	N/A	1
Low frequency combinations	famous	main	saying	2	0.61	0.81	0	n/a	n/a	0%	0%	N/A	1
Low frequency combinations	simple	worst	access	3	0.88	1.02	0	n/a	n/a	0%	0%	N/A	2
Low frequency combinations	useful	slow	balance	2	0.80	1.20	0	n/a	n/a	0%	0%	N/A	2
Low frequency combinations	front	price	engine	2	0.70	0.99	0	n/a	n/a	0%	0%	N/A	2

AUTHOR'S COPY | AUTORENEXEMPLAR

Condition	Experimental prime	Control prime	Target	Control pair			Experiment pair			Forward association (EAT)	Backward association (EAT)	Association (pilot)[1]	List
				BNC occurrences	*t*	*MI*	BNC occurrences	*t*	*MI*				
Low frequency combinations	easy	spoken	journey	2	0.94	1.57	0	n/a	n/a	0%	0%	N/A	2
Low frequency combinations	central	post	measure	2	0.55	0.71	0	n/a	n/a	0%	0%	N/A	2
Low frequency combinations	strong	single	mixture	2	1.06	1.99	0	n/a	n/a	0%	0%	N/A	2
Low frequency combinations	final	complete	status	4	1.29	1.50	0	n/a	n/a	0%	0%	N/A	2
Low frequency combinations	single	daily	tower	2	0.99	1.75	0	n/a	n/a	0%	0%	N/A	2
Moderate collocation	greater	fixed	concern	40	6.08	4.68	0	n/a	n/a	0%	0%	0 (0)	1
Moderate collocation	subject	likely	content	32	5.41	4.54	0	n/a	n/a	0%	0%	0 (0)	1
Moderate collocation	human	recent	culture	47	6.62	4.85	0	n/a	n/a	0%	0%	0 (0)	1
Moderate collocation	likely	card	effects	40	5.94	4.02	0	n/a	n/a	0%	0%	0 (0)	1
Moderate collocation	recent	cutting	figures	31	5.25	4.14	0	n/a	n/a	0%	0%	0 (0)	1
Moderate collocation	complex	human	series	28	5.04	4.40	0	n/a	n/a	0%	0%	0 (0)	1
Moderate collocation	former	rapid	student	26	4.85	4.35	0	n/a	n/a	0%	0%	0 (0)	1

AUTHOR'S COPY | AUTORENEXEMPLAR

Moderate collocation	special	estate	unit	51	6.81	4.43	0	n/a	n/a	0%	0%	0 (0)	1
Moderate collocation	total	warm	freedom	21	4.36	4.35	0	n/a	n/a	0%	0%	0 (0)	2
Moderate collocation	real	finance	impact	38	5.89	4.51	0	n/a	n/a	0%	0%	0 (0)	2
Moderate collocation	worst	waiting	kind	34	5.63	4.88	0	n/a	n/a	0%	0%	0 (0)	2
Moderate collocation	complete	strong	loss	26	4.82	4.17	0	n/a	n/a	0%	0%	0 (0)	2
Moderate collocation	true	total	owner	24	4.72	4.78	0	n/a	n/a	0%	0%	0 (0)	2
Moderate collocation	lower	narrow	risk	25	4.71	4.13	0	n/a	n/a	0%	0%	0 (0)	2
Moderate collocation	short	real	stay	65	7.77	4.80	0	n/a	n/a	0%	0%	0 (0)	2
Moderate collocation	full	express	text	63	7.67	4.89	0	n/a	n/a	0%	0%	0 (0)	2
Frequent collocation	foreign	colour	debt	226	14.98	8.04	0	n/a	n/a	0%	0%	0 (0)	1
Frequent collocation	past	complex	decade	357	18.85	8.58	0	n/a	n/a	0%	0%	0 (0)	1
Frequent collocation	double	former	doors	115	10.69	8.35	0	n/a	n/a	0%	0%	0 (0)	1
Frequent collocation	stone	greater	floor	67	8.08	6.29	0	n/a	n/a	0%	0%	0 (0)	1
Frequent collocation	rapid	royal	growth	243	15.56	9.07	0	n/a	n/a	0%	0%	0 (0)	1
Frequent collocation	music	special	hall	162	12.59	6.55	0	n/a	n/a	0%	0%	0 (0)	1

Condition	Experimental prime	Control prime	Target	Control pair			Experiment pair			Forward association (EAT)	Backward association (EAT)	Association (pilot)[1]	List
				BNC occur-rences	*t*	*MI*	BNC occur-rences	*t*	*MI*				
Frequent collocation	colour	weak	scheme	126	11.10	6.53	0	n/a	n/a	0%	0%	0 (0)	1
Frequent collocation	armed	music	struggle	117	10.80	9.16	0	n/a	n/a	0%	0%	0 (0)	1
Frequent collocation	finance	close	bill	78	8.72	6.27	0	n/a	n/a	0%	0%	0 (0)	2
Frequent collocation	private	football	homes	77	8.65	6.19	0	n/a	n/a	0%	0%	0 (0)	2
Frequent collocation	price	protest	index	168	12.90	7.63	0	n/a	n/a	0%	0%	0 (0)	2
Frequent collocation	close	front	links	181	13.39	7.66	0	n/a	n/a	0%	0%	0 (0)	2
Frequent collocation	slow	true	motion	107	10.32	8.68	0	n/a	n/a	0%	0%	0 (0)	2
Frequent collocation	mental	shopping	picture	60	7.67	6.63	0	n/a	n/a	0%	0%	0 (0)	2
Frequent collocation	narrow	final	range	93	9.54	6.57	0	n/a	n/a	0%	0%	0 (0)	2
Frequent collocation	warm	full	welcome	180	13.38	8.72	0	n/a	n/a	0%	0%	0 (0)	2
Associated frequent collocation	current	stone	affairs	188	13.64	7.53	0	n/a	n/a	10%	7%	1.55 (13)	1

AUTHOR'S COPY | AUTORENEXEMPLAR

Associated frequent collocation	estate	huge	agent	328	18.10	10.49	0	n/a	n/a	21%	5%	1.23 (10)	1
Associated frequent collocation	parish	feature	church	411	20.23	9.03	0	n/a	n/a	66%	0%	2.09 (17)	1
Associated frequent collocation	cutting	parish	edge	173	13.13	9.27	0	n/a	n/a	11%	0%	1.18 (9)	1
Associated frequent collocation	feature	subject	film	67	8.11	6.75	0	n/a	n/a	31%	0%	1.55 (13)	1
Associated frequent collocation	card	village	game	38	6.03	5.54	0	n/a	n/a	16%	1%	0.27 (2)	1
Associated frequent collocation	pretty	double	girl	87	9.21	6.31	0	n/a	n/a	16%	0%	1.73 (14)	1
Associated frequent collocation	village	foreign	green	105	10.10	6.08	0	n/a	n/a	19%	0%	0.64 (8)	1
Associated frequent collocation	shopping	central	list	144	11.96	8.28	0	n/a	n/a	17%	17%	0.41 (3)	2
Associated frequent collocation	protest	private	march	42	6.40	6.25	0	n/a	n/a	36%	0%	1.00 (10)	2
Associated frequent collocation	football	simple	match	147	12.07	7.89	0	n/a	n/a	14%	12%	0.32 (3)	2

AUTHOR’S COPY | AUTORENEXEMPLAR

Condition	Experimental prime	Control prime	Target	Control pair			Experiment pair			Forward association (EAT)	Backward association (EAT)	Association (pilot)[1]	List
				BNC occur-rences	*t*	*MI*	BNC occur-rences	*t*	*MI*				
Associated frequent collocation	post	useful	office	1324	36.32	9.17	0	n/a	n/a	21%	2%	0.64 (6)	2
Associated frequent collocation	daily	mental	paper	59	7.52	5.54	0	n/a	n/a	23%	0%	0.59 (6)	2
Associated frequent collocation	waiting	short	room	130	11.16	5.58	0	n/a	n/a	29%	0%	1.50 (11)	2
Associated frequent collocation	express	easy	train	38	6.10	6.64	0	n/a	n/a	57%	0%	1.36 (11)	2
Associated frequent collocation	spoken	lower	word	101	9.98	7.26	0	n/a	n/a	36%	2%	1.86 (15)	2

1 The first figure in each cell is the association score (calculated as described above); the figure in brackets is the number of participants to who provided the associate.
2 List = the experimental list on which the target word appears with its experimental prime.

AUTHOR'S COPY | AUTORENEXEMPLAR

Bionotes

Phil Durrant is currently Visiting Assistant Professor in the Graduate School of Education at Bilkent University, where he teaches on the MA TEFL program. He is the author of several papers on the topic of collocation and formulaic language.

Alice Doherty is studying for a PhD in Psycholinguistics at the University of Nottingham. E-mail: durrant.phil@gmail.com; aexad4@nottingham.ac.uk

Notes

1. Cf. <http://www.eat.rl.ac.uk/>
2. We would like to thank an anonymous reviewer for pointing out this limitation.

References

Balota, David A. 1994. Visual word recognition: The journey from features to meaning. In Morton Ann Gernsbacher (ed.), *Handbook of psycholinguistics*, 340–357. San Diego: Academic Press.

Barlow, Michael & Suzanne Kemmer (eds.). 2000. *Usage-based models of language*. Stanford, California: CSLI Publications.

Bley-Vroman, Robert. 2002. Frequency in production, comprehension, and acquisition. *Studies in Second Language Acquisition* 24(2). 209–213.

Charles, Walter G. & George A. Miller. 1989. Contexts of antonymous adjectives. *Applied Psycholinguistics* 10. 357–375.

Church, Kenneth W. & Patrick Hanks. 1990. Word association norms, mutual information, and lexicography. *Computational Linguistics* 16(1). 22–29.

Clear, Jem 1993. From Firth principles: computational tools for the study of collocations. In Mona Baker, Gill Francis & Elena Tognini-Bonelli (eds.), *Text and technology: in honour of John Sinclair*, 271–292. Amsterdam & Philadelphia: John Benjamins.

Conklin, Kathy. & Norbert Schmitt. 2008. Formulaic sequences: Are they processed more quickly than nonformulaic language by native and nonnative speakers. *Applied Linguistics* 29(1). 72–89.

de Groot, Anette M. & Gerard L. Nas. 1991. Lexical representation of cognates and noncognates in compound bilinguals. *Journal of Memory and Language* 30(1). 90–123.

Durrant, Phil. 2009. Investigating the viability of a collocation list for students of English for academic purposes. *English for Specific Purposes* 28(3). 157–179.

Ellis, Nick C. 2001. Memory for language. In Peter Robinson (ed.), *Cognition and second language instruction*, 33–68. Cambridge: Cambridge University Press.

Ellis, Nick C. 2002. Frequency effects in language processing. *Studies in Second Language Acquisition* 24(2). 143–188.

Ellis, Nick C. 2003. Constructions, chunking, and connectionism: the emergence of second language structure. In Catherine J. Doughty & Michael H. Long (eds.), *The handbook of second language acquisition*, 63–103. Oxford: Blackwell.

Ellis, Nick C., Eric Frey, & Isaac Jalkanen. 2009. The psycholinguistic reality of collocation and semantic prosody (1): Lexical access. In Ute Römer & Rainer Schulze (eds.), *Exploring the lexis-grammar interface*, 89–114. Amsterdam & Philadelphia: John Benjamins.

Ellis, Nick C., Rita Simpson-Vlach, & Carson Maynard. 2008. Formulaic language in native and second-language speakers: psycholinguistics, corpus linguistics, and TESOL. *TESOL Quarterly* 41(3). 375–396.

Firth, John R. 1968. A synopsis of linguistic theory, 1930–55. In F. R. Palmer (ed.), *Selected papers of J.R. Firth 1952–1959*, 168–205. Harlow: Longman.

Fitzpatrick, Tess. 2007. Word association patterns: unpacking the assumptions. *International Journal of Applied Linguistics* 17(3). 319–331.

Gibbs, Raymond W., & Nandini. Nayak. 1989. Psycholinguistic studies on the syntactic behavior of idioms. *Cognitive Psychology* 21(1). 100–138.

Gibbs, Raymond W., N. Nayak, & Cooper Cutting. 1989. How to kick the bucket and not decompose: analyzability and idiom processing. *Journal of Memory and Language* 28(5). 576–593.

Handl, Susanne. 2008. Essential collocations for learners of English: The role of collocational direction and weight. In Fanny Meunier & Sylviane Granger (eds.), *Phraseology in foreign language learning and teaching*, 43–66. Amsterdam: John Benjamins.

Herbst, Thomas. 1996. What are collocations: sandy beaches or false teeth? *English Studies* 77(4). 379–393.

Hodgson, J. M. 1991. Informational constraints on pre-lexical priming. *Language and Cognitive Processes* 6(3). 169–205.

Hoey, Michael. 1991. *Patterns of lexis in text*. Oxford: Oxford University Press.

Hoey, Michael. 2005. *Lexical priming: A new theory of words and language*. London: Routledge.

Hunston, Susan. 2002. *Corpora in applied linguistics*. Cambridge: Cambridge University Press.

Hutchison, Keith A. 2003. Is semantic priming due to association strength or feature overlap? A microanalytic review. *Psychonomic Bulletin and Review* 10(4). 785–813.

Jiang, Nan & Tatiana M. Nekrasova, M. 2007. The processing of formulaic sequences by second language speakers. *The Modern Language Journal* 91(3). 433–445.

Jones, Susan & John M. Sinclair. 1974. English lexical collocations. A study in computational linguistics. *Cahiers de lexicologie* 24. 15–61.

Leech, Geoffrey, Paul Rayson, & Andrew Wilson. 2001. *Word frequencies in written and spoken English: based on the British National Corpus*. London: Longman.

Lewis, Michael. 1993. *The lexical approach: The state of ELT and a way forward*. London: Thomson Heinle.

Lucas, Margery. 2000. Semantic priming without association: a meta-analytic review. *Psychonomic Bulletin & Review* 7(4). 618–630.

Manning, Christopher D. & Hinrich Schütze. 1999. *Foundations of statistical natural language processing*. Cambridge, MA: MIT Press.

McKoon, Gail & Roger Ratcliff. 1992. Spreading activation versus compound cue accounts of priming: Mediated priming revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 18(6). 1155–1172.

McNamara, Timothy P. & Jeanette Altarriba. 1988. Depth of spreading activation revisited: semantic mediated priming occurs in lexical decisions. *Journal of Memory and Language* 27(5). 545–559.

Meyer, David E. & Roger W. Schvaneveldt. 1971. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology* 90(2). 227–234.

Nattinger, James R. & Jeanette S. DeCarrico. 1992. *Lexical phrases and language teaching*. Oxford: Oxford University Press.

Neely, James H. 1991. Semantic priming effects in visual word recognition: a selective review of current findings and theories. In Derek Besner & Glyn W. Humphreys (eds.), *Basic processes in reading: visual word recognition*, 264–336. Hillsdale, NJ: Lawrence Erlbaum.

Newmeyer, Frederick. 2003. Grammar is grammar and usage is usage. *Language* 79(4). 682–707.

AUTHOR'S COPY | AUTORENEXEMPLAR

Pawley, Andrew, & Frances H. Syder. 1983. Two puzzles for linguistic theory: Nativelike selection and nativelike fluency. In Jack C. Richards & Richard W. Schmidt (eds.), *Language and communication*, 191–226. New York: Longman.

Perea, Manuel & Eva Rosa. 2002. The effects of associative and semantic priming in the lexical decision task. *Psychological Research* 66(3). 180–194.

Peterson, Robert R., Curt Burgess, Gary S. Dell, & Kathleen M. Eberhard. 2001. Dissociation between syntactic and semantic processing during idiom comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27(5). 1223–1237.

Rastle, Kathleen, Jonathan Harrington, & Max. Coltheart. 2002. 358,534 nonwords: The ARC Nonword Database. *Quarterly Journal of Experimental Psychology* 55(4). 1339–1362.

Schmidt, Robert. 1992. Psychological mechanisms underlying second language fluency. *Studies in Second Language Acquisition* 14(3). 357–385.

Schmitt, Norbert, Sarah Grandage, & Svenja Adolphs. 2004. Are corpus-derived recurrent clusters psycholinguistically valid? In Norbert Schmitt (ed.), *Formulaic sequences: acquisition, processing and use*, 127–151. Amsterdam & Philadelphia: John Benjamins.

Schmitt, Norbert & Geoffrey Underwood. 2004. Exploring the processing of formulaic sequences through a self-paced reading task. In Norbert Schmitt (ed.), *Formulaic sequences: acquisition, processing and use*, 173–189. Amsterdam & Philadelphia: John Benjamins.

Sereno, Joan A. 1991. Graphemic, associative, and syntactic priming effects as a brief stimulus onset asynchrony in lexical decision and naming. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 17(3). 459–477.

Shelton, Jennifer R. & Randi C. Martin. 1992. How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory, and Cognition* 18(6). 272–281.

Shin, Dongkwang & Paul Nation. 2008. Beyond single words: the most frequent collocations in spoken English. *ELT Journal* 62(4). 339–348.

Sinclair, John M. 1987. Collocation: a progress report. In Ross Steele & Terry Threadgold (eds.), *Language topics: Essays in honour of Michael Halliday*, 319–331. Amsterdam & Philadelphia: John Benjamins.

Stubbs, Michael. 1995. Collocations and semantic profiles: on the cause of the trouble with quantitative methods. *Functions of Language* 2(1). 1–33.

Titone, Debra A., & Cynthia M. Connine. 1999. On the compositional and noncompositional nature of idiomatic expressions. *Journal of Pragmatics* 31(12). 1655–1674.

Tomasello, Michael. 2003. *Constructing a language: a usage-based theory of language acquisition*. Cambridge MA, London: Harvard University Press.

Underwood, Geoffrey, Norbert Schmitt, & Adam Galpin. 2004. The eyes have it: an eye-movement study into the processing of formulaic sequences. In Norbert Schmitt (ed.), *Formulaic sequences: acquisition, processing and use*, 153–172. Amsterdam: John Benjamins.

Van Lancker-Sidtis, Diana. 2004. When novel sentences spoken or heard for the first time in the history of the universe are not enough: toward a dual-process model of language. *International Journal of Language and Communication Disorders* 39(1). 1–44.

Williams, John N. 1996. Is automatic priming semantic. *European Journal of Cognitive Psychology* 8(2). 139–151.

Wray, Alison. 2002. *Formulaic language and the lexicon*. Cambridge: Cambridge University Press.