

Semantic Diversity, Frequency and the Development of Lexical Quality in Children's Word

Reading

Yaling Hsiao and Kate Nation

Department of Experimental Psychology

University of Oxford, UK

Address for correspondence: Yaling Hsiao (yaling.hsiao@psy.ox.ac.uk) or Kate Nation (kate.nation@psy.ox.ac.uk), Department of Experimental Psychology, University of Oxford, Oxford, UK.

Abstract

Frequency exerts a powerful influence on lexical processing but it is possible that at least part of its effect is caused by high frequency words being experienced in more diverse contexts over an individual's language experience. To capture this variability, we applied Latent Semantic Analysis on a 35-million-word corpus of texts written for children, deriving a measure of semantic diversity that quantifies the similarity of all the contexts a word appears in. Across three experiments with 6-13-year-old children involving reading aloud and lexical decision, we found a main effect of semantic diversity: high diversity words were responded to faster and read more accurately than low diversity words. Frequency, document count and age of acquisition were also significant predictors of reading behavior. These findings demonstrate that contextual variability contributes to word learning and the development of lexical quality, beyond the effect of frequency.

Keywords

semantic diversity; contextual diversity; frequency effect; lexical quality; reading development

Highlights

- We examined the influence of semantic diversity on children's word reading
- Semantic diversity was calculated from a corpus of material written for children
- Semantic diversity predicted children's word naming and lexical decision
- The effect of semantic diversity was independent of frequency, document count and age of acquisition

- Findings suggest contextual variability is important in shaping lexical development

Semantic Diversity, Frequency and the Development of Lexical Quality in Children's Word Reading

Children who read more are better at reading words than children who read less. Estimates of print exposure – a proxy for how much an individual has read – account for unique variance in reading development and are associated with individual differences in orthographic and phonological processing (e.g. Cunningham & Stanovich, 1990; Mol & Bus, 2011). Why might this be? Reading is a skill and like all skills, practice is critical to becoming expert and for word-level reading, practice may be important in at least two distinct ways. First, reading practice allows basic skills to be honed and fine-tuned, promoting the development of reading fluency. In addition, however, an avid reader is likely to experience more words and a larger range of different language contexts than a less prolific reader. In this paper, we investigate whether this broad lexical experience influences the ease with which children read words.

Reading experience provides exposure to individual words, cumulatively adding to both type and token frequency. Word frequency is an item level variable that typically reflects the number of times a word appears in a corpus. In adults, the frequency effect is robust across a range of lexical tasks involving written words, with high frequency words enjoying a processing advantage (Balota & Chumbley, 1984; Brysbaert, Warriner, & Kuperman, 2014; Monsell, 1991; Rayner & Duffy, 1986) (for review, see Brysbaert, Mandera, & Keuleers, 2018). In children too, estimates of word frequency influence how long it takes to read a word, or make a lexical decision to it (Joseph, Nation, & Livversedge, 2013; Schmalz, Marinus, & Castles, 2013). Models of skilled word recognition vary in how they handle the frequency effect (e.g. M. Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001; Murray & Forster, 2004; Norris, 2006; Plaut,

McClelland, Seidenberg, & Patterson, 1996) but centre on the idea that variations in frequency reflect differences in experience, with repeated exposure to a word influencing its accessibility, either by changing the recognition threshold (e.g. McClelland & Rumelhart, 1981), or the weights between nodes in distributed models (e.g. Seidenberg & McClelland, 1989).

Surprisingly, given the ubiquity of the frequency effect, there has been little discussion of how reading experience shapes lexical representations through development so as to influence reading behaviour.

In line with these theoretical accounts, frequency might influence the development of lexical quality via the principle of repetition. On this view, words higher in frequency (and therefore experienced more often) become more strongly represented in memory over time and leading them to be processed more efficiently than words lower in frequency (and therefore experienced less often). A quite different theoretical account is that frequency influences lexical processing via the quality or contextual nature of encounters with each word, not just the number of encounters. In natural reading, words are rarely encountered in isolation: they occur in sentences, paragraphs and texts. Does the linguistic context in which a word has occurred across previous encounters matter for how that word is subsequently processed? The lexical legacy hypothesis (Nation, 2017) suggests it might, as it is the substrate from which knowledge about a word builds. On this view, reading experience provides many different contexts and episodes which over time result in a complex database about a word, its connections to other words and its lexical history within an individual's experience. In turn, these rich and diverse encounters bring about local variation in lexical quality (Perfetti, 2007) at the word level: a legacy that is measurable during word reading behaviour, even as skilled readers process words in a laboratory task such as lexical decision.

Several lines of evidence suggest that variation in contextual experience with words influences how easily they are processed in tasks such as lexical decision. McDonald and Shilcock (2001) devised a new variable, *contextual distinctiveness*, which captured the local lexical environment in which words co-occur across a corpus. They defined contextual distinctiveness as the relative entropy between the posterior distribution (the distribution of words occurring in a ten-word window around a target word) and the prior distribution (the expected distribution of those words when the target word is not taken into account). Contextual distinctiveness was a better predictor of lexical decision latencies than word frequency, supporting the view that distributional statistics that reflect the lexical environment a word has been experienced in have a role to play in theoretical accounts of visual word recognition (see Baayen, 2010 for extended discussion).

A number of more recent studies (for review, see Jones, Dye, & Johns, 2017) offer support to the general approach introduced by McDonald and Shilcock (2001). Adelman, Brown and Quesada (2006) found that contextual diversity (indexed in their study as the number of unique documents a word appears in across a corpus) not only predicted lexical decision and naming latency, it eliminated any effect of word frequency. Variations in document count are also associated with word reading in sentence processing, as revealed by analysis of eye movements (Plummer, Perea, & Rayner, 2014). One issue with these findings is that contextual diversity as indexed by document count is highly correlated with word frequency and it might be that it is simply a better measure of frequency than frequency itself (e.g. Brysbaert & New, 2009). It is also important to note that document count does not take the content of the contexts into account. This is an important point to consider, if we are to understand the theoretical reasons for why document count might influence lexical processing, beyond frequency.

Semantic diversity (or its reciprocal, *semantic distinctiveness*) is a variable that explicitly captures the similarity in content of different contextual experiences of a word (e.g. Hoffman, Lambon Ralph, & Rogers, 2013; Johns, Gruenenfelder, Pisoni, & Jones, 2012; Jones et al., 2017; Jones, Johns, & Recchia, 2012). It is similar in spirit to McDonald and Shilcock's (2001) contextual distinctiveness variable, but rather than calculate contextual similarity via local lexical co-occurrence, similarity is quantified using Latent Semantic Analysis techniques. This produces a metric that captures context-dependent variation in similarity in a graded fashion. Semantic diversity predicts lexical decision and naming latency in adults, outperforming both word frequency and document count (Johns, Dye, & Jones, 2014; Johns et al., 2012; Johns, Dye, & Jones, 2016), with words that appear in more semantically diverse contexts being processed more easily than less semantically diverse words. In this way, semantic diversity behaves similarly to polysemy, and indeed, the processing advantage for polysemous words in lexical decision might be related to the fact that polysemous words tend to be more semantically diverse (Azuma & Van Orden, 1997; Pexman, Hargreaves, Siakaluk, Bodner, & Pope, 2008; Rodd, Gaskell, & Marslen-Wilson, 2002).

Clearly, semantic diversity is associated with how easily skilled adult readers make a lexical decision response. Its basis, however, must stem from reading and language experience. Retuning to the lexical legacy hypothesis (Nation, 2017), a word's semantic diversity at any point in time can be thought of as the product of an individual's contextual experiences with that word and the opportunities for learning that are afforded by those experiences, culminating in variations in lexical quality that in turn govern item-level variation in lexical processing. This description chimes with Adelman et al.'s conclusion that "learning based models of reading cannot accommodate [Adelman et al.'s] results unless they are modified so that learning

mechanisms are sensitive to context, not frequency” (2006, p. 822). While there have been computational implementations of word learning from contextual experience (Hoffman & Woollams, 2015; Johns et al., 2014), it is striking that there is little relevant evidence from studies of children’s reading. Given that semantic diversity is a variable that has its roots in learning and experience, developmental data are important.

There is, however, evidence to show document count influences on children’s word reading. Perea, Soares and Comesaña (2013) calculated a contextual diversity variable from a children’s reading corpus, using document count as their metric. They chose 60 words that varied orthogonally in contextual diversity and frequency and asked 4th Grade Portuguese children to make a lexical decision to them. Words high in contextual diversity were processed more quickly than frequency-matched low diversity words. There was, however, no effect of frequency: latencies to high vs. low frequency words did not differ when the two sets of words were matched for contextual diversity. These findings replicated in a second experiment that adopted a regression design, using a different sample of children and a different set of items.

By 4th grade then, it seems that children are sensitive to contextual factors afforded by reading experience. One limitation to these findings is that diversity was instantiated using document count, rather than semantic diversity. As noted above, document count is highly correlated with word frequency and does not capture the similarity of content between contexts. Sample size was small, with fewer than 30 children in each experiment, and Perea et al. only sampled 60 items. The children’s corpus was also small in terms of the number and range of documents, comprising 3.2 million words taken from 171 elementary textbooks.

In our paper, we aimed to build on these findings. Most importantly, we devised a measure of semantic diversity to capture the semantic similarity in content across contexts using

the *Oxford Children's Corpus* – a developmental corpus of children's written language. Based on evidence from adults, we predicted that this would be less tightly bound with word frequency than document count, and associated with variation in children's lexical decision and naming. We also asked whether frequency influences children's lexical processing. In adults, evidence suggests that word frequency matters less, once semantic diversity is taken into account. In children, however, frequency might be more important as repetition might be critical in a developing system. With this in mind, we also measured the children's reading level and extended the age range, allowing us to investigate whether semantic diversity and frequency have different effects at different levels of proficiency. Finally, we also considered age of acquisition, defined as the approximate age at which a word is learned. Like frequency, age of acquisition influences children's lexical processing (V. Coltheart, Laxon, & Keating, 1988). Its association with semantic diversity has not been investigated in relation to word reading, although there is evidence that in infancy, words that are acquired early are experienced in more diverse contexts and as a result become more associated with other known words (Hills, Maouene, Riordan, & Smith, 2010).

We begin by introducing the *Oxford Children's Corpus* and describing how semantic diversity was calculated before examining how variations in semantic diversity are associated with children's lexical processing in three different datasets.

Calculating Semantic Diversity

The Oxford Children's Corpus (OCC) is a dynamic and growing corpus, initiated in 2006 by Oxford University Press to guide the preparation of dictionaries for children. The version used in this paper contained over 35 million words and 12,000 documents, targeted at children

aged 5-16 years old. Unlike some other children's corpora, the OCC is not restricted to curriculum materials and structured reading schemes. It also contains classic and modern children's fiction, non-fiction, textbooks, websites and magazines. Thus, it broadly samples a wide range of the type of written materials children encounter during their reading experience.

As part of our pre-processing procedure, some documents were removed (such as webpages from museum websites that contained very short texts documenting information about exhibits, e.g., dimensions of exhibits and historical period). We lemmatized the corpus so that inflectional endings of words were removed to return the base forms, known as lemmas. We also excluded blank spaces and words with little semantic content, such as function words (e.g., prepositions, determiners), digits and punctuation marks. Following Hoffman et al. (2013), we subdivided longer documents into 1000-word contexts. This created 20,411 unique contexts. We then excluded words that occurred fewer than 50 times in the corpus, or appeared in fewer than 40 contexts. This left around 12 million word tokens and around 12,500 unique word types. Latent Semantic Analysis was applied to the matrix representing the frequency of each word in all contexts. The values in the matrix underwent log transformation, which then were divided by the entropy of that word in the corpus. This standard log entropy weighting procedure reduces the influence of very high frequency words. Singular Value Decomposition was then used to provide a solution with approximately 300 dimensions to produce a set of vectors for each word, and a set of vectors for each context in the corpus. For a given word, the mean was calculated by averaging the cosines of all pairwise combinations of context vectors the word appeared in, representing the average similarity between any two contexts containing the word. This value was log transformed and the sign reversed to indicate the diversity, rather than the

distinctiveness, of the contexts. The resulting value represents the semantic diversity of each word.

It is worth noting two procedural differences between our calculational approach and that used by Hoffman et al. (2013). As the performance of semantic vector models is not altered following lemmatization, or by the exclusion of function words (Bullinaria & Levy, 2012), we used lemmas rather than word forms as the target unit (on the basis that the underlying semantic representation of words with the same root is similar) and excluded function words (to conserve computational resources). To verify that these procedural changes did not change the nature of the semantic diversity metric, we also calculated semantic diversity without lemmatization and without removing function words (i.e., following the same method as Hoffman et al.). This correlated with our original estimate of semantic diversity, $r=.9$; we retained the original estimates for the analyses reported throughout this paper. We also checked reliability by randomly splitting the corpus in half and calculating semantic diversity across each sub-corpus separately. The split-half correlation was high, $r=.93$, indicating high reliability.

Across all of the lemmas in the OCC, the average semantic diversity value was 1.91 ($SD=0.29$, $range=0.07-2.48$). Its distribution is shown in Figure 1. The mean is a little higher than the mean calculated by Hoffman et al. from a large adult corpus of written language ($M=1.51$, $SD=0.37$, $range=0.08-2.41$). Most words in the OCC received a semantic diversity score of 1.5-2.5. Words at the lowest end of the diversity distribution tend to be more specific and restrictive in meaning, including, for example, proper nouns from children's shows or video games. At the other end of the distribution, high diversity words tend to be higher in frequency and more variable in meaning.




Figure.1 Distribution of words in the Oxford Children's Corpus as a function of semantic diversity

To understand how semantic diversity relates to other lexical properties, we conducted a series of pairwise correlations, summarised in Figure 2. Semantic diversity correlated with log word frequency ($r = .22$) and log document count ($r = .38$), both taken from the OCC. This is not surprising given the procedure that derived the semantic diversity metric was based on word frequency and document count information. However, the correlations are modest, unlike the very high correlation between OCC frequency and OCC document count ($r = .94$). These findings mirror the high correlations seen in Hoffman et al.'s norms from adults. Children's semantic diversity correlated with number of senses ($r = .28$, $N = 8897$; taken from Wordsmyth

Dictionary). Moving to other semantic related properties, there was no correlation with concreteness or imageability ($r = .04$ and $.07$, respectively) (for concreteness, $N = 10238$ from Brysbaert et al., 2014; for imageability, $N = 2199$ from Schock, Cortese, & Khanna, 2012). There was a moderate negative correlation between semantic diversity and age of acquisition ($r = -.40$, $N = 9504$, from Kuperman, Stadthagen-Gonzalez, & Brysbaert, 2012), showing that high diversity words tend to be acquired earlier in life. Finally, our developmental measure of semantic diversity correlated with Hoffman et al.'s measure of semantic diversity from adult corpora, $r = .35$. This modest correlation suggests that while children's reading experiences are different to those accumulated by adulthood, words that occur in diverse contexts in children's reading materials also appear in diverse contexts more generally.




Figure 2. Pairwise correlations summarising the lexical properties of words in the Oxford Children's Corpus (OCC). See text for number of items in each pairwise correlation (reliability of measures: children's semantic diversity $r=.93$ (see text); Hoffman et al.'s semantic diversity $r=.90$; AoA $r=.84$; concreteness $r=.92$; see original sources for further details. Reliability estimates not available for number of senses and imageability).

Having established semantic diversity values for a large set of words, we investigated through a series of three experiments whether semantic diversity is implicated in children's word reading. This allowed us to test the hypothesis that variation in contextual experience (as indexed by semantic diversity) influences the development of lexical quality, beyond the influence of

frequency. If the number of encounters is most critical, as argued by theories that account for frequency in terms of repeated encounters, any contextual diversity associated with those encounters should not matter. Evidence that semantic diversity influences children's reading behaviour beyond frequency would challenge the theoretical assumption that the frequency effect stems from repeated exposure to word tokens during reading experience. Instead, the data would support the view that orthographic representations are shaped by the linguistic context in which words are experienced. By developing and using a reliable measure of semantic diversity that takes semantic content into account (and is quite distinct from frequency itself, cf. document count), we were well placed to test the hypothesis that variation in previous contextual experience with a word influences how easily that word is subsequently processed. Experiment 1 was small in scale, manipulating semantic diversity in categorical design and assessing its influence on lexical decision and reading aloud performance. Experiment 2 also assessed lexical decision and reading aloud but sampled a larger number of words, using a continuous design. Experiment 3 reports secondary analysis of an existing dataset with a large number of observations per word and asks whether semantic diversity influences reading aloud.

EXPERIMENT 1

This experiment replicated the design of Perea et al.'s (2013) investigation of document count and its influence on children's lexical decision but instead used semantic diversity as the variable of interest. This allowed us to compare children's processing of high vs. low semantically diverse words, controlling for the effects of frequency and document count.

Method

Participants

Thirty-five 8-11 year old children ($M=9.2$ years, $SD=0.9$) recruited from classrooms in the UK participated in this experiment. All spoke English as their first language and all scored within normal range on a standardised test of word reading fluency, the *Test of Word Reading Efficiency*, in which children read aloud as many words (or nonwords) as possible in 45 seconds (Torgesen, Wagner, & Rashotte, 1999) (TOWRE, $M=115$, $SD=14.03$). Thirty-seven children were initially recruited but two failed to score within normal range on the TOWRE (operationalized as a standard score below 86) and were therefore excluded from the experiment. Ethical approval was provided by the University of Oxford's Research Ethics Committee.

Materials

Sixty words were selected from the corpus, 30 higher in semantic diversity ($M=2.20$, $SD=0.09$) and 30 lower in semantic diversity ($M=2.02$, $SD=0.06$) ($t(58)=8.78$, $p<.0001$). Following the recommendations of van Heuven, Mandera, Keuleer, and Brysbaert (2014), we used the Zipf scale to estimate frequency, operationalized as $\log_{10}(\text{frequency-per-million-word})+3$. This normalises the distribution and aides interpretation: similar to a Likert scale, values of 1-3 represent low frequency words and values of 4-7 represent high frequency words. Frequency varied within each set, but overall frequency and document count were matched across the high vs. low diversity words, shown in Table 1 (frequency: $t(58)=0.65$, $p=.52$; document count: $t(58)=0.43$, $p=.67$). It is worth noting that the frequency (and document count) range for the selected words in this experiment was small, with most words being relatively high in frequency. Word length varied between 5 and 9 letters ($M=6.4$ letters) and did not differ across sets ($t(58)=-0.68$, $p=.50$). Sixty pseudowords were created using *Wuggy*, a pseudoword generator (Keuleers & Brysbaert, 2010, Version 0.2.0b2). For the lexical decision task, two lists were constructed, each containing 15 high semantic diversity words and 15 low diversity words,

along with 30 pseudowords. Each child completed lexical decision for one list only; the other list was used for the naming (reading aloud) task, with lists counterbalanced across children. See Appendix A for all words used in this experiment along with their semantic diversity scores.

Procedure

The experiment took place in a single session in a quiet area close to the child's classroom. After completing the TOWRE, children completed either the lexical decision task or the naming task first, with order counterbalanced across lists and children.

The lexical decision task was presented on a laptop, using the E-prime software (Version 2.0) (Schneider, Eschman, & Zuccolotto, 2012). Following Perea et al. (2013), we used a go/no-go procedure; Moret-Tatay and Perea (2011) recommended this procedure as it produces faster response times and lower error rates in children than the traditional yes/no lexical decision task. Each trial began with a central fixation cross, presented for 500ms. This was followed by a word or pseudoword. Children were asked to press a key on the response box as quickly as possible if they thought the string was a word. They were asked to refrain from responding if they thought the string was a “made-up” word. The letter string stayed on the screen until a keypress was made, or until 2500 ms if no response was made. The next trial was then initiated. Each child saw 30 words and 30 pseudowords.

For the naming task, each word was printed on a card and children were asked to read each word aloud. There was no time constraint. Each child read aloud 30 words, different to those seen in the lexical decision task.

Results

Descriptive data are shown in Table 1. Raw data for this and all three experiments can be found in the Supplementary Materials. For lexical decision, performance was high with accuracy

averaging 98%. Reaction times were faster for words in the high semantic diversity condition (773 ms) than the low diversity condition (810 ms). This difference was significant by item ($t(58) = -2.33, p = .02$) but not by children ($t(34) = -1.16, p = .25$). For naming, performance was at ceiling and was not analysed further.

Table 1. Lexical statistics and summary of lexical decision and naming performance in Experiment 1

	High Semantic Diversity		Low Semantic Diversity	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Lexical Statistics				
Semantic diversity	2.20	0.09	2.02	0.06
Frequency ^{Zipf}	5.50	0.16	5.44	0.14
Document count ^{log}	3.38	0.17	3.36	0.17
Length ^{letter}	6.30	0.70	6.43	0.80
Lexical decision				
Accuracy (% correct)	97.54	0.16	99.04	0.10
RT (ms)	773	272	810	298
Naming (% correct)	99.04	0.10	99.62	0.06

Following Perea et al., we used multiple regression to investigate lexical decision latencies further (to correct word trials only), and to consider the effects of semantic diversity, frequency, document count and reading proficiency as continuous factors. Because of the high correlation between frequency and document count ($r = .94$), we ran two otherwise identical models, one including frequency and the other including document count.

For all analyses reported in this paper, RTs were transformed into inverse RTs ($-1000/\text{RT}$), following Brysbaert and Stevens (2018). The data were transformed back to raw RTs for ease interpretation, as shown in Figure 3. Note that while the data are plotted categorically

(using the median split of each variable), they were analysed continuously. The variables were entered centred to reduce multicollinearity and scaled for comparability. In the regression model that included word frequency, there were significant effects of semantic diversity ($b = -0.02$, $SE = 0.01$, $t = -2.18$, $p = .045$), proficiency ($b = -0.08$, $SE = 0.01$, $t = -7.04$, $p < .0001$) and frequency ($b = -0.03$, $SE = 0.01$, $t = -2.25$, $p = .025$). RTs were shorter as words increased in semantic diversity and frequency, and for children with higher levels of reading proficiency. In the model where document count replaced frequency, both semantic diversity ($b = -0.02$, $SE = 0.01$, $t = -1.98$, $p = .048$) and proficiency ($b = -0.08$, $SE = 0.01$, $t = -7.03$, $p < .0001$) predicted RT, but document count did not ($b = -0.005$, $SE = 0.01$, $t = -0.41$, $p = .68$).




Figure 3. Mean (SE) lexical decision RT as a function of semantic diversity, frequency, document count and proficiency (plotted as median-split for each variable but analysed continuously).

Discussion

These results provide initial support for the hypothesis that semantic diversity is associated with lexical decision in children. Words high in semantic diversity were responded to

37 ms faster than words low in diversity. This effect maintained when frequency was controlled, consistent with semantic diversity operating in a different manner to frequency. Frequency was more closely associated with performance than document count and the effect of frequency remained when semantic diversity was included in the model. Contextual diversity, as indexed by document count, did not. This is in contrast to Perea et al. (2013) who found that document count outperformed frequency. In our data, document count and semantic diversity were not interchangeable and while semantic diversity was associated with lexical decision performance, document count was not. In addition, proficiency was also significant suggesting that the level of reading ability children brought to the task was important.

One limitation of this experiment was that items were chosen to vary in semantic diversity categorically and the number of words was small, totalling only 60. It is also important to note the relatively restricted range of values, not just of semantic diversity, but also of frequency and document count too. Another limitation is that we did not consider age of acquisition (AoA). It is well-established that the age at which words are acquired during childhood influences how easily they are named and recognised in lexical decision by children and adults (e.g. V. Coltheart et al., 1988; Monaghan, Chang, & Welbourne, 2017; Morrison & Ellis, 1995). There is also evidence that contextual variability has an intimate relationship with the order in which words are acquired in infancy, such that early acquired words have more semantic connections with other known words than those acquired later (Hills, et al., 2010). In line with this, the correlation between semantic diversity and AoA across our corpus is $-.40$ (larger than the correlation between semantic diversity and frequency) indicating that early-acquired words are more semantically diverse. It is thus important to establish whether semantic diversity is in fact an age of acquisition effect in disguise. Although there was no difference in

AoA between the high vs. low semantic diversity words used in this experiment ($t(58) = -0.12$, $p = .90$), when AoA was added to the regression model, it predicted lexical decision performance ($b = 0.05$, $SE = 0.01$, $t = 3.89$, $p = .0001$). Importantly however, the effect of semantic diversity remained significant ($b = -0.02$, $SE = 0.01$, $t = -2.18$, $p = .03$). Taken together, the findings of Experiment 1 suggest that there is an effect of semantic diversity in children's reading which cannot be explained in terms of frequency, document count or AoA.

EXPERIMENT 2

To lessen the risk of selection biases associated with categorical designs, our second experiment included 300 words, chosen to vary more continuously across a range of semantic diversity values. As before, children completed lexical decision and naming tasks and we examined how semantic diversity, frequency (or document count) and proficiency influenced performance in this larger and more powerful experiment. We also included AoA to test further whether the semantic diversity is distinct from AoA, as suggested by the post hoc analysis of data from Experiment 1.

Method

Participants

One hundred and twenty children, recruited from Year 3 to Year 6 classrooms in five UK schools, participated in this experiment. All spoke English as their first language. Six were subsequently excluded due to poor performance on the TOWRE. This left 114 children in the final sample (M age 9.9 years, $SD = 1.1$ years, range = 7.62-11.53 years). Ethical approval was provided by the University of Oxford's Research Ethics Committee.

Materials

Three hundred words were selected to encompass a range of semantic diversity scores ($M=1.97$, $SD=0.25$, $Range=1.60-2.36$), listed at Appendix B. Words were initially sampled randomly. We then checked them for profanity and replaced unsuitable items with other randomly selected words. The 300 words were split into 10 lists, each containing 30 words. Word length and frequency were balanced across lists. Word length averaged 6.8 letters ($SD= 2.16$; range 3-15 letters). The words also varied in zipf frequency ($M=4.67$, $SD=0.53$, range=3.87-6.47), log document count ($M=2.62$, $SD= 0.47$, range=1.79-4.07) and AoA ($M= 7.78$, $SD=2.17$, range= 2.78-13.56). The intercorrelations between item-level properties of the selected words are shown in Table 2. Three hundred pseudowords were generated using Wuggy, as per Experiment 1. These were dispersed into the 10 word lists such that each list contained 30 words and 30 pseudowords.

Table 2. Intercorrelations between lexical variables for the items used in Experiment 2

	2.	3.	4.	5.
1. Semantic diversity	.34***	.37***	-.06	-.32***
2. Frequency ^{Zipf}		.98***	-.39***	-.54***
3. Document count ^{log}			-.36***	-.50***
4. Length ^{letter}				-.36
5. AoA				

* $p < .05$. ** $p < .01$. *** $p < .001$. $N=300$ for all variables except AoA, for which $N=270$.

Procedure

The procedure was identical to Experiment 1. Each child was allocated to one of the ten lists for lexical decision and a different list for naming. Order of task was counterbalanced across

children. Some lists had fewer observations than the others but each word was sampled at least 10 times across the experiment for both lexical decision and naming.

Results

We used linear mixed effects models to examine the factors affecting lexical decision and naming, namely semantic diversity, zipf frequency (or log document count; as in Experiment 1, separate models tested the effects of frequency and document count given their high inter-correlation), AoA and each child's level of reading proficiency. Length in letters was also added as a fixed factor. All variables were again centred and scaled. Given larger amount of data, we were able to fit the data using linear mixed effects models (lme4 package by Bates, Mächler, Bolker, & Walker, 2015) to assess generalizability of the effects across participants and items. Throughout our paper, interaction terms were decided based on model comparison: if two models did not differ in terms of goodness-of-fit, the simpler model was preferred. In this experiment, only main effects were included in each model. Statistical significance was considered based on the criteria of $t > 2$.

Model criticism was performed after the models had been fitted to remove influential outliers. This approach is recommended for mixed effects models over aggressive a priori screening, such as removing observations that are two standard deviations away from the mean (Baayen & Milin, 2010). We removed data points with standardized residuals of the models exceeding 2.5 standard deviations that did not follow a normal distribution. Full details of the analyses, including the random variance components, fixed effects, and the variance explained by each model (we used the correlation between the fitted and the observed values as R^2 is not available for mixed effects models) are provided in the Supplementary Materials, along with the raw data.

Lexical Decision: Accuracy

Unlike Experiment 1, children were not at ceiling on the lexical decision task ($M = 85\%$, $SD = 0.13$) meaning we were able to analyse accuracy as well as RT. A generalized linear mixed effect model with the binomial link function revealed significant main effects of semantic diversity ($b = 0.62$, $SE = 0.16$, $z = 3.90$, $p < .0001$), word frequency ($b = 1.21$, $SE = 0.22$, $z = 5.64$, $p < .0001$) and AoA ($b = -1.06$, $SE = 0.18$, $z = -5.80$, $p < .0001$). Words that were more semantically diverse, higher in frequency and lower in AoA were responded to more accurately. Proficiency was significant too ($b = 0.84$, $SE = 0.20$, $z = 4.29$, $p < .0001$) with those children scoring well on the TOWRE performing more accurately, not surprisingly. Word length was not significant ($b = 0.30$, $SE = 0.18$, $z = 1.74$, $p = 0.08$). The only significant interaction was proficiency * length ($b = 0.32$, $SE = 0.10$, $z = 3.21$, $p = .001$), suggesting that accuracy increased in line with proficiency, when length was held constant. All remaining two-way interactions were not significant.

The results were similar in a second model when document count replaced word frequency, with significant main effects of semantic diversity ($b = 0.60$, $SE = 0.17$, $z = 3.58$, $p = .0003$) and document count ($b = 1.09$, $SE = 0.20$, $z = 5.37$, $p < .0001$) and AoA ($b = -1.07$, $SE = 0.18$, $z = -5.93$, $p < .0001$). Proficiency ($b = 0.84$, $SE = 0.19$, $z = 4.38$, $p < .0001$) was significant in this analysis but word length was not ($p = .07$). As before, the only significant interaction was between proficiency and length ($b = 0.32$, $SE = 0.10$, $z = 3.29$, $p = .001$); all other two-way interactions were not significant.

Lexical Decision: RTs

Turning to latency, mean RT to correct trials was 960 ms ($SD = 386$ ms). A linear mixed effects model including frequency showed significant results for all main effects, with faster responses being associated with higher frequency and greater semantic diversity, lower AoA,

shorter length and more proficient readers: semantic diversity ($b=-0.02$, $SE=0.01$, $t=-2.20$), frequency ($b=-0.09$, $SE=0.01$, $t=-7.27$), AoA ($b=0.09$, $SE=0.01$, $t=8.37$), word length ($b=0.05$, $SE=0.01$, $t=4.54$) and proficiency ($b=-0.06$, $SE=0.02$, $t=-3.09$). All two-way interactions were included in the analysis. The significant interaction between AoA and semantic diversity ($b=0.02$, $SE=0.01$, $t=2.42$) indicated that there was a larger effect of AoA as semantic diversity increased. Also significant was the interaction between AoA and frequency ($b=-0.02$, $SE=0.01$, $t=-2.26$), indicating that the effect of AoA was smaller for higher frequency words. Other interactions were not significant.

Replacing frequency with document count produced a similar pattern of results with significant main effects of semantic diversity ($b=-0.02$, $SE=0.01$, $t=-2.17$), document count ($b=-0.08$, $SE=0.01$, $t=-6.72$), length ($b=0.05$, $SE=0.01$, $t=4.81$), AoA ($b=0.09$, $SE=0.01$, $t=9.07$), and proficiency ($b=-0.06$, $SE=0.02$, $t=-3.09$). There was only one significant interaction, between AoA and semantic diversity ($b=0.02$, $SE=0.01$, $t=2.35$).

Naming Accuracy

Reading aloud accuracy was high at 90% ($SD=0.12$). A generalized linear mixed effects model with binomial link function produced main effects of semantic diversity ($b=0.44$, $SE=0.21$, $z=2.13$, $p=.03$), word frequency ($b=0.77$, $SE=0.31$, $z=2.51$, $p=.01$), AoA ($b=-1.26$, $SE=0.25$, $z=-5.11$, $p<.0001$) and proficiency ($b=1.39$, $SE=0.27$, $z=5.09$, $p<.0001$). Thus, words, higher in diversity, higher in frequency, and lower in AoA were read more accurately and not surprisingly, children with higher levels of reading proficiency read words more accurately overall. Word length was not significant ($p=.08$) and there were no significant interactions. Replacing frequency with document count produced similar results with significant main effects of semantic diversity ($b=0.46$, $SE=0.21$, $z=2.20$, $p=.03$), document count ($b=0.74$, $SE=0.29$,

$z=2.55$, $p=.01$), AoA ($b=-1.27$, $SE=0.24$, $z=-5.21$, $p<.0001$) and proficiency ($b=1.40$, $SE=0.27$, $z=5.15$, $p<.0001$). Once again, there was no effect of word length ($p=.06$) and no significant interactions.

Discussion

The findings of Experiment 2 support those reported in Experiment 1. Semantic diversity was associated with lexical decision latency beyond the contribution of frequency, document count and length. Lexical decision performance was below ceiling, allowing us to analyse children's accuracy. Here again we found an influence of semantic diversity, which mirrored the RT data in pattern: greater accuracy aligned with higher levels of semantic diversity, independent of frequency or document count. Similarly, diversity, frequency, document count and length all influenced reading aloud. Early acquired words enjoyed a lexical processing advantage, but this effect of AoA was clearly separate to the effect of semantic diversity in both naming and lexical decision.

EXPERIMENT 3

Experiment 2 included a large number of items ($N=300$) but the number of observations for each item in each task was low (about 10). To assess the reproducibility of our findings, we took the opportunity to investigate whether semantic diversity was associated with reading aloud via secondary analysis of an existing dataset. The data formed part of the standardisation of the *Diagnostic Test of Word Reading Processes* (DTWRP; Forum for Research in Language and

Literacy, 2012)¹. This is a measure of single word reading where children read aloud 30 regular words, 30 exception words and 30 nonsense words (with regularity defined as adhering to grapheme-phoneme correspondence rules, Rastle & Coltheart, 1999). The dataset comprised 350 children and totalled around 29,500 observations. Of the 60 words appearing in the test, only 46 could be assigned a semantic diversity value. The remaining words were either too low in frequency (defined as occurring fewer than 50 times or in fewer than 40 contexts in the corpus) or too low in semantic content (i.e. they were function words) for us to establish a semantic diversity value.

Method

Participants

The DTWRP dataset contains data from 350 6-13 year old children ($M = 9.8$ years, $SD = 1.7$ years). They were recruited from Year 2 to Year 7 classrooms in the UK. All were reported to speak English as their first language.

Materials

Our analysis included 46 words from the DTWRP (see Table 3), 23 deemed by the test authors to be regular and 23 exception. Length ranged between 2 to 13 letters. As the words come from a closed test that is used clinically, we are not able to make them freely available here. The data were collected in accordance with the test manual instructions, with each word being read aloud, one at a time and without time restriction. Responses were scored as correct or incorrect.

¹ Thanks to Dr. Jessie Ricketts for making these secondary analyses possible. We gratefully acknowledge Jessie and other members of the *Forum for Research in Literacy and Language* for sharing their DTWRP data, and the children's TOWRE scores.

Results

Although our focus was not with the contrast between regular and exception words, we include it here, given the words were initially selected by the test authors on that basis. Table 3 shows reading aloud performance for each set of words, along with item characteristics and intercorrelations between the lexical variables in shown in Table 4. The two word sets did not differ in frequency ($t(44)=-1.45$, $p=.15$), document count ($t(44)=-1.60$, $p=.12$) length ($t(44)=1.20$, $p=.23$) or AoA ($t(44)= 1.45$, $p=.15$). There was however a difference in semantic diversity ($t(44)=-2.28$, $p=.03$) as the exception words were more diverse than the regular words.

We modelled the data as per Experiment 2, assessing the effects of semantic diversity, zipf frequency (or document count), AoA and length. Regularity was also included as a main effect. TOWRE was included as an estimate of reading proficiency. Interaction terms were semantic diversity with proficiency, AoA and regularity. Full details are provided in the Supplementary Materials, along with the raw data.

Table 3. Item characteristics and naming accuracy for the words included in Experiment 3

	Regular Words (n=23)		Exception Words (n=23)	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Lexical Statistics				
Semantic diversity	1.96	0.22	2.08	0.14
Frequency ^{Zipf}	4.88	0.76	5.26	0.78
Document count ^{log}	2.70	0.72	2.99	0.64
Length ^{letter}	7.22	2.28	6.43	2.13
Age of acquisition	7.36	2.72	6.16	2.60
Naming				
N observations	7774	14.96	7682	
Accuracy (%)	86	0.14	81	0.20

Table 4. Intercorrelations between lexical variables for the items used in Experiment 3

	2.	3.	4.	5.
1. Semantic diversity	.55***	.59***	-.41**	-.56***
2. Frequency ^{Zipf}		.98***	-.74***	-.63***
3. Document count ^{log}			-.71***	-.61***
4. Length ^{letter}				-.65***
5. AoA				

* $p < .05$. ** $p < .01$. *** $p < .001$. $N = 46$.

In the model with frequency as a predictor, all main effects except for frequency ($p = .90$) were significant. Regular words were read more accurately than exception words ($b = 1.77$, $SE = 0.45$, $z = 3.90$, $p < .0001$). Words higher in semantic diversity were read more accurately than those lower in diversity ($b = 0.72$, $SE = 0.34$, $z = 2.12$, $p = .03$) and length also mattered, with short words at an advantage ($b = -0.89$, $SE = 0.36$, $z = -2.347$, $p = .01$). AoA was significant too ($b = -1.22$, $SE = 0.41$, $z = -2.95$, $p = .003$) showing that words acquired earlier were read more accurately. There was also a main effect of proficiency: as expected following Experiments 1 and 2, children with higher levels of reading skill had better reading accuracy overall ($b = 1.65$, $SE = 0.14$, $z = 11.58$, $p < .0001$). None of the interactions were significant but two are perhaps worth noting, given $p = .05$. The interaction between AoA and semantic diversity ($b = -0.39$, $SE = 0.20$, $z = -1.94$, $p = .05$) suggests that later acquired words show a smaller increase in accuracy with each unit of increase in semantic diversity. The interaction between proficiency and semantic diversity ($b = 0.09$, $SE = 0.05$, $z = 1.95$, $p = .05$) suggests that better readers showed a bigger effect of semantic diversity. There was no interaction between semantic diversity and regularity ($p = .67$).

Replacing frequency with document count produced the same pattern of results with significant main effects of spelling regularity, ($b = 1.78$, $SE = 0.45$, $z = 3.92$, $p < .0001$), semantic

diversity ($b=0.78$, $SE=0.34$, $z=2.30$, $p=.02$), word length ($b=-0.97$, $SE=0.36$, $z=-2.68$, $p=.007$), AoA ($b=-1.26$, $SE=0.40$, $z=-3.14$, $p=.002$), and proficiency ($b=1.64$, $SE=0.14$, $z=11.57$, $p<.0001$). As for frequency, there was no effect of document count ($p=.66$). The pattern of interactions was similar too but this time the association between AoA and semantic diversity was statistically significant ($b=-0.41$, $SE=0.20$, $z=-2.05$, $p=.04$). The interaction between proficiency and semantic diversity was once again marginal ($b=0.09$, $SE=0.05$, $z=1.95$, $p=.05$) and there was no interaction between semantic diversity and regularity ($p=.70$).

Discussion

This secondary analysis of an existing dataset collected for other purposes replicated the effect of semantic diversity seen in our earlier experiments. We found that items that were more diverse were pronounced more accurately than less diverse words. Like in Experiment 2, there was also an independent effect of AoA: words acquired earlier in life were read more accurately than later acquired words, and there was also a main effect of children's reading proficiency. Unlike Experiment 2, there was no evidence of an independent effect of frequency, either for word frequency itself, or for document count. We note however that the item set is small and while there was variation in frequency, there were relatively few very low frequency words. The item set manipulated spelling regularity and this too predicted naming accuracy with regular words being read more accurately. This experiment also demonstrates the utility of repurposing existing datasets to replicate effects of interest in data collected on items selected for other purposes.

General Discussion

The aim of this study was to investigate whether semantic diversity influences children's word reading. We used a large corpus of reading materials written for children to calculate semantic diversity – a metric that captures variation in the similarity of contexts each word appears in, across the corpus. In three experiments involving different children, different words and two different tasks, children's reading behaviour was sensitive to semantic diversity. Generally, words higher in semantic diversity were processed more quickly and more accurately than words lower in semantic diversity. Importantly, this effect was not a frequency effect in disguise: while frequency and document count were both associated with children's reading, semantic diversity maintained an independent influence. Similarly, semantic diversity could not be explained by variation in age of acquisition. While AoA contributed to children's reading performance, an independent effect of semantic diversity maintained across all analyses.

Our findings replicated the facilitative effect of semantic diversity found in lexical decision and naming in adults (Hoffman & Woollams, 2015; Johns et al., 2016). One way to explain the positive effect semantic diversity is to draw a parallel with the effect of polysemy, namely the finding that highly polysemous words are processed more quickly in lexical decision than non-polysemous words, provided the senses are related in meaning (e.g. Rodd et al., 2002 and for semantic richness effects more generally, Goh, Yap, Lau, Ng, & Tan, 2016; Rodd, Gaskell, & Marslen-Wilson, 2004; Yap, Lim, & Pexman, 2015; Yap, Tan, Pexman, & Hargreaves, 2011). This has been attributed to polysemous words having richer semantic representations, reflecting the fact that they are mapped to a number of distinct (but overlapping) meanings. Greater variability results in rapid semantic activation from orthographic input,

allowing lexical decision responses to be made more quickly (e.g., Rodd, Gaskell, & Marslen-Wilson, 2004). Within this type of theoretical framework, semantic diversity reflects gradations in semantic representation, based on the notion that variation in the meaning of a word is a continuous property of variation in the context in which the word is used (Hoffman & Woollams, 2015; Hoffman et al., 2013). However, it is worth noting that not all of the semantic richness measures are compatible with one another. For example, Hoffman et al.'s semantic diversity measure calculated on the *British National Corpus* showed a positive correlation with number of senses but correlated negatively with concreteness and imageability. Thus, high diversity words tend to be more polysemous and more abstract in meaning than words lower in diversity.

A complementary model that captures the effects of semantic diversity is the semantic distinctiveness model (Johns et al., 2015; Jones et al., 2012; Jones et al., 2017). Central to this account of semantic memory is the notion of updating representations in response to contextual change. The model works by predicting the meaning of a word from its memory representation based on previous encounters in context. If a word is encountered in a similar context to what's predicted, little changes; if however the new context is different to that expected based on previous encounters, it is encoded more strongly and the memory representation of the word is updated. The result of this is that words experienced in more variable contexts are less associated with a particular context and are thus more strongly weighted in memory. In contrast, less diverse words are more predictable from a given context and are therefore less strongly weighted in memory. In this way, changes in semantic context over time influence the strength or nature of the lexical representation in memory; in turn, these become associated with item-level differences in lexical processing.

Taking a developmental perspective, Nation (2017) proposed that once basic skills are in place, the emergent consequence of variation in reading experience is variation in lexical quality. Our finding that semantic diversity influences children's word reading adds support to this view. We also observed a robust frequency effect in Experiment 2, when the item set contained a good range of variation in frequency (cf. Experiments 3). This is in contrast to experiments with adults where the frequency effect is reduced or even eliminated once semantic or contextual diversity is considered (Adelman et al., 2006; Jones et al., 2012). For children then, the principle of repetition is important for explaining variation in word reading: repeated exposure to a word, even in redundant contexts, influences how easily that word is read. This makes sense, given a certain level of frequency is needed to build sensitivity to diversity (and see Monaghan, Chang, Welbourne, & Brysbaert (2017) for complementary data from a computational modeling perspective).

Although not the focus of our paper, AoA had an independent influence on reading behavior. Words that are acquired early tend to be higher in cumulative frequency, and early on at least, they are experienced in more diverse contexts and are more semantically associated with other known words (Hills et al., 2010). In our experiments, variance associated with AoA is clearly separable from what was captured by semantic diversity but nevertheless, AoA reflects differences in lexical experience that shape the developing system in a way that over time becomes reflected in processing differences in word reading and lexical decision (Monaghan, et al., 2017; Nation, 2017).

Our findings make clear that theoretical accounts of visual word recognition need to take both development and contextual experience seriously. Alongside the item-level main effects, we also saw an effect of reading proficiency across all three experiments. Not surprisingly, better

readers out-performed less-skilled readers. Proficiency also modulated some item-level effects, with better readers tending to show larger effects of length and diversity. It is important to interpret these findings cautiously, given our age range (and thus reading range) is quite restricted, sampling children through mid-childhood only. Effects within childhood or within a restricted range of skills level may vary to those seen in the population at large. In a large-scale study across the lifespan (535 participants aged 8-83 years), Davies, Arnell, Birchenough, Grimmond & Houlson (2017) found that the size of the frequency effect reduced with increasing age and increasing levels of reading skill. This is consistent with a smaller frequency effect in adults (and computational models) at higher levels of print exposure (Chateau & Jared, 2000; Monaghan et al., 2017). Johns, Sheppard, Jones and Taler (2016) reported a larger effect of semantic diversity on lexical decision in older adults (65+ years) than younger adults (18-30 years), perhaps reflecting differences in language experience and lexical knowledge.

While these findings echo our own conclusion that the interplay between learning processes and experience plays a critical role in establishing the word recognition system, they also highlight the need for more lifespan data so that we can model the effects of psycholinguistic variables in interaction with age and individual differences. More data are also needed to allow a range of psycholinguistic variables (derived from development corpora) to be explored in the same developmental dataset, akin to megastudies with adults (Balota et al., 2007; Brysbaert, Stevens, Mander, & Keuleers, 2016; Keuleers, Diependaele, & Brysbaert, 2010; Keuleers, Lacey, Rastle, & Brysbaert, 2012). As noted earlier, highly diverse words tend to have more senses; they are also acquired earlier in life (Hills, et al., 2010). A large developmental dataset will enable our understanding of how these and other variables interact and how they fall

out of contextual experience, furthering our insight into how learning promotes the development of structure in the reading system (Davies et al., 2017; Monaghan, Chang, & Welbourne, 2017)

Another future direction includes establishing a causal relationship between semantic diversity, children's reading experience and lexical processing. Two recent experiments highlight the utility of this approach. Rosa, Tapia and Perea (2017) used a word learning paradigm which manipulated contextual variability. Children encountered new words either in the context of short stories, expository texts or math exercises, or across all three types of context. Words experienced in more varied contexts were learned better. Joseph and Nation (2018) also examined children's word learning. New words were encountered either in semantically diverse or semantically uniform contexts; word learning was indexed via changes in fixation duration as a function of exposure through the course of learning, and via a series of post-tests. There was no overall effect of diversity, possibly because 10 exposures to a new word is insufficient. However, semantic diversity did interact with reading comprehension skill suggesting that better readers were more sensitive and adaptive to changing contexts. This speculation calls for future studies to model learning more effectively, using more exposure sessions that carefully manipulate and control lexical statistics, while taking account of individual differences.

A final future direction concerns how to capture lexical experience in a way that is psychological valid. We operationalized semantic diversity by representing individual words and documents in a co-occurrence matrix and applying Latent Semantic Analysis procedures, using all the texts in the Oxford Children's Corpus. While capturing an important aspect of lexical experience that is clearly associated with children's processing of written words, this "count" representation of text experience is not psychologically plausible as it does not accommodate the

dynamics of incremental learning. People do not form lexical representations like our model, where the entire semantic space is represented in one instance, based on all of the language statistics collected in a corpus. An alternative approach with potentially greater psychological validity draws on “predict” models (Hollis, 2017; Mander, Keuleers, & Brysbaert, 2017). For example, models such as Continuous Bag of Words (CBOW) and Skip-gram receive input in sequentially within a small window span and either learn to predict the target word based on the surrounding context (CBOW), or learn to predict the context given the target word (skip-gram). Such models have the potential to capture learning in a way that distributed semantic models based on Latent Semantic Analysis cannot. An important future direction will be for neural network models to be trained on input from a children’s corpus of written language in sequenced fashion (e.g. starting with texts targeted to the youngest audience).

In conclusion, by developing and using a semantic diversity measure based on a large developmental corpus of children’s reading experience, we have established that semantic diversity is an important determiner of children’s lexical decision and reading aloud performance. This effect is independent of the effects of frequency, document count and age of acquisition, although those metrics also influence reading behavior, alongside semantic diversity itself. Our findings demonstrate that semantic variability in the content of contexts a word has been experienced influences how easily it is subsequently processed. This observation cannot be explained by frequency, nor by age of acquisition. One possibility is that experiencing a word across changing contexts updates its memory representation (Jones et al., 2017). In turn, this brings about differences in lexical quality as reading and contextual experience accrues and with time becomes reflected in differences in lexical processing, even in simple single-word tasks such as lexical decision and reading aloud (Nation, 2017). Our findings emphasize the need for

models of skilled word recognition to consider learning, and the context in which learning takes place.

Acknowledgements

We would like to thank Oxford University Press for allowing us access to the Oxford Children's Corpus, with particular thanks to Vineeta Gupta and Nilanjana Banerji for their support and collaboration. We are most grateful to Jessie Ricketts, Morag Stuart and Jackie Masterson who kindly shared data from the standardization sample of the *Diagnostic Test of Word Reading Processes* (Forum for Research in Literacy & Language, 2012) for the secondary data analyses reported in Experiment 3. Phoebe Rogers, Helen Norris and Megan Bird provided outstanding research assistance; we also thank Bram Vanderkerckhove for his work on the corpus in the early stages of the project, Paul Hoffman for his advice on calculating semantic diversity and Victoria Murphy and Stephen Pulman for constructive discussion. This research was funded by a Leverhulme Trust Research Project Grant (reference: RPG-2015-070).

References

- Adelman, J. S., Brown, G. D. A., & Quesada, J. F. (2006). Contextual diversity not frequency determines word naming and lexical decision times. *Psychological Science, 17*, 814–823.
- Azuma, T., & Van Orden, G. C. (1997). Why SAFE is better than FAST: The relatedness of a word's meanings affects lexical decision times. *Journal of Memory and Language, 36*, 484–504. <https://doi.org/10.1006/jmla.1997.2502>
- Baayen, R. H. (2010). Demythologizing the word frequency effect: A discriminative learning perspective. *The Mental Lexicon, 5*, 436–461. <https://doi.org/10.1075/ml.5.3.10baa>
- Baayen, R. H., & Milin, P. (2010). Analyzing Reaction Times. *International Journal of Psychological Research, 3*, 12–28. <https://doi.org/10.1287/mksc.12.4.395>
- Balota, D. A., & Chumbley, J. I. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. *Journal of Experimental Psychology: Human Perception and Performance, 10*, 340–357. <https://doi.org/10.1037/0096-1523.10.3.340>
- Balota, D. A., Yap, M. J., Cortese, M. J., Hutchison, K. A., Kessler, B., Loftis, B., ... Treiman, R. (2007). The English Lexicon Project. *Behavior Research Methods, 39*, 445–459.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using {lme4}. *Journal of Statistical Software, 67*, 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Brysbaert, M., Mandera, P., & Keuleers, E. (2018). The Word Frequency Effect in Word Processing: An Updated Review. *Current Directions in Psychological Science, 27*, 45–50. <https://doi.org/10.1177/0963721417727521>
- Brysbaert, M., & New, B. (2009). Moving beyond Kučera and Francis: A critical evaluation of

current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior Research Methods*, 41, 977–990.

<https://doi.org/10.3758/BRM.41.4.977>

Brysbaert, M., & Stevens, M. (2018). Power Analysis and Effect Size in Mixed Effects Models: A Tutorial. *Journal of Cognition*, 1, 1–20. <https://doi.org/10.5334/joc.10>

Brysbaert, M., Stevens, M., Mandera, P., & Keuleers, E. (2016). The impact of word prevalence on lexical decision times: Evidence from the Dutch lexicon project 2. *Journal of Experimental Psychology: Human Perception and Performance*, 42, 441–458.

<https://doi.org/10.1037/xhp0000159>

Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46, 904–911.

<https://doi.org/10.3758/s13428-013-0403-5>

Bullinaria, J. A., & Levy, J. P. (2012). Extracting semantic representations from word co-occurrence statistics: Stop-lists, stemming, and SVD. *Behavior Research Methods*, 44, 890–907. <https://doi.org/10.3758/s13428-011-0183-8>

Chateau, D., & Jared, D. (2000). Exposure to print and word recognition processes. *Memory & Cognition*, 28, 143–153. <https://doi.org/10.3758/BF03211582>

Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). DRC: a dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, 108, 204–256. <https://doi.org/10.1037/0033-295x.108.1.204>

Coltheart, V., Laxon, V. J., & Keating, C. (1988). Effects of word imageability and age of acquisition on children's reading. *British Journal of Psychology*, 79, 1–12.

Cunningham, A. E., & Stanovich, K. E. (1990). Assessing Print Exposure and Orthographic

- Processing Skill in Children: A Quick Measure of Reading Experience. *Journal of Educational Psychology*, 82, 733–740. <https://doi.org/10.1037/0022-0663.82.4.733>
- Davies, R., Arnell, R., Birchenough, J. M. H., Grimmond, D., & Houlson, S. (2017). Reading through the lifespan: Individual differences in psycholinguistic effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43, 1298–1338. <https://doi.org/10.1037/xlm0000366>
- Forum for Research in Language and Literacy. (2012). Diagnostic Test of Word Reading Processes.
- Goh, W. D., Yap, M. J., Lau, M. C., Ng, M. M. R., & Tan, L. C. (2016). Semantic richness effects in spoken word recognition: A lexical decision and semantic categorization megastudy. *Frontiers in Psychology*, 7, 1–10. <https://doi.org/10.3389/fpsyg.2016.00976>
- Hills, T. T., Maouene, J., Riordan, B., & Smith, L. B. (2010). The Associative Structure of Language: Contextual Diversity in Early Word Learning, 63, 259–273. <https://doi.org/10.1016/j.jml.2010.06.002>.The
- Hoffman, P., Lambon Ralph, M. A., & Rogers, T. T. (2013). Semantic diversity: A measure of semantic ambiguity based on variability in the contextual usage of words. *Behavior Research Methods*, 45, 718–730. <https://doi.org/10.3758/s13428-012-0278-x>
- Hoffman, P., & Woollams, A. M. (2015). Opposing effects of semantic diversity in lexical and semantic relatedness decisions. *Journal of Experimental Psychology: Human Perception and Performance*, 41, 385–402. <https://doi.org/10.1037/a0038995>
- Hollis, G. (2017). Estimating the average need of semantic knowledge from distributional semantic models. *Memory and Cognition*, 45, 1350–1370. <https://doi.org/10.3758/s13421-017-0732-1>

- Johns, B. T., Dye, M., & Jones, M. N. (2014). The influence of contextual diversity on word learning. *Proceedings of the 35th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society, 242–247. <https://doi.org/10.3758/s13423-015-0980-7>
- Johns, B. T., Dye, M., & Jones, M. N. (2016). The influence of contextual diversity on word learning. *Psychonomic Bulletin & Review*, 23, 1214–1220. <https://doi.org/10.3758/s13423-015-0980-7>
- Johns, B. T., Gruenenfelder, T. M., Pisoni, D. B., & Jones, M. N. (2012). Effects of word frequency, contextual diversity, and semantic distinctiveness on spoken word recognition. *The Journal of the Acoustical Society of America*, 132, EL74-EL80. <https://doi.org/10.1121/1.4731641>
- Johns, B. T., Sheppard, C. L., Jones, M. N., & Taler, V. (2016). The role of semantic diversity in word recognition across aging and bilingualism. *Frontiers in Psychology*, 7, 1–11. <https://doi.org/10.3389/fpsyg.2016.00703>
- Jones, M. N., Dye, M., & Johns, B. T. (2017). Context as an Organizing Principle of the Lexicon. In *Psychology of Learning and Motivation* (Vol. 67, pp. 239–283). <https://doi.org/10.1016/bs.plm.2017.03.008>
- Jones, M. N., Johns, B. T., & Recchia, G. (2012). The role of semantic diversity in lexical organization. *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale*, 66, 115–124. <https://doi.org/10.1037/a0026727>
- Joseph, H., & Nation, K. (2018). Examining incidental word learning during reading in children: The role of context. *Journal of Experimental Child Psychology*, 166, 190–211. <https://doi.org/https://doi.org/10.1016/j.jecp.2017.08.010>
- Joseph, H., Nation, K., & Livversedge, S. P. (2013). *Using eye movements to investigate word*

frequency effects in children's sentence reading. School Psychology Review (Vol. 42).

Retrieved from <https://eprints.soton.ac.uk/363349/>

Keuleers, E., & Brysbaert, M. (2010). Wuggy: A multilingual pseudoword generator. *Behavior Research Methods*, 42, 627–633. <https://doi.org/10.3758/BRM.42.3.627>

Keuleers, E., Diependaele, K., & Brysbaert, M. (2010). Practice Effects in Large-Scale Visual Word Recognition Studies: A Lexical Decision Study on 14,000 Dutch Mono- and Disyllabic Words and Nonwords. *Frontiers in Psychology*, 1, 174. <https://doi.org/10.3389/fpsyg.2010.00174>

Keuleers, E., Lacey, P., Rastle, K., & Brysbaert, M. (2012). The British Lexicon Project: lexical decision data for 28,730 monosyllabic and disyllabic English words. *Behavior Research Methods*, 44, 287–304. <https://doi.org/10.3758/s13428-011-0118-4>

Kuperman, V., Stadthagen-Gonzalez, H., & Brysbaert, M. (2012). Age-of-acquisition ratings for 30,000 English words. *Behavior Research Methods*, 44, 978–990. <https://doi.org/10.3758/s13428-012-0210-4>

Mandera, P., Keuleers, E., & Brysbaert, M. (2017). Explaining human performance in psycholinguistic tasks with models of semantic similarity based on prediction and counting: A review and empirical validation. *Journal of Memory and Language*, 92, 57–78. <https://doi.org/10.1016/j.jml.2016.04.001>

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, 88, 175–407.

McDonald, S. A., & Shillcock, R. C. (2001). Rethinking the word frequency effect: the neglected role of distributional information in lexical processing. *Language and Speech*, 44, 295–323. <https://doi.org/10.1177/00238309010440030101>

- Mol, S. E., & Bus, A. G. (2011). To read or not to read: A meta-analysis of print exposure from infancy to early adulthood. *Psychological Bulletin*, 137, 267–296.
<https://doi.org/10.1037/a0021890>
- Monaghan, P., Chang, Y.-N., & Welbourne, S. (2017). Different processes for reading words learned before and after onset of literacy. *Proceedings of the 39th Cognitive Science Society Conference*(June).
- Monaghan, P., Chang, Y. N., Welbourne, S., & Brysbaert, M. (2017). Exploring the relations between word frequency, language exposure, and bilingualism in a computational model of reading. *Journal of Memory and Language*, 93, 1–21.
<https://doi.org/10.1016/j.jml.2016.08.003>
- Monsell, S. (1991). The nature and locus of word frequency effects in reading. In *Basic processes in reading: Visual word recognition*. (pp. 148–197). Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.
- Moret-Tatay, C., & Perea, M. (2011). Is the go/no-go lexical decision task preferable to the yes/no task with developing readers? *Journal of Experimental Child Psychology*, 110, 125–132. <https://doi.org/10.1016/j.jecp.2011.04.005>
- Morrison, C. M., & Ellis, A. W. (1995). Roles of word frequency and age of acquisition in word naming and lexical decision. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. US: American Psychological Association. <https://doi.org/10.1037/0278-7393.21.1.116>
- Murray, W. S., & Forster, K. I. (2004). Serial Mechanisms in Lexical Access: The Rank Hypothesis. *Psychological Review*, 111, 721–756. <https://doi.org/10.1037/0033-295X.111.3.721>

- Nation, K. (2017). Nurturing a lexical legacy: reading experience is critical for the development of word reading skill. *Npj Science of Learning*, 2, 3. <https://doi.org/10.1038/s41539-017-0004-7>
- Norris, D. (2006). The Bayesian Reader: Explaining Word Recognition as an Optimal Bayesian Decision Process. *Psychological Review*, 113, 327–357.
- Perea, M., Soares, A. P., & Comesaña, M. (2013). Contextual diversity is a main determinant of word identification times in young readers. *Journal of Experimental Child Psychology*, 116, 37–44. <https://doi.org/10.1016/j.jecp.2012.10.014>
- Perfetti, C. (2007). Reading Ability: Lexical Quality to Comprehension. *Scientific Studies of Reading*, 11, 357–383. <https://doi.org/10.1080/10888430701530730>
- Pexman, P. M., Hargreaves, I. S., Siakaluk, P. D., Bodner, G. E., & Pope, J. (2008). There are many ways to be rich: Effects of three measures of semantic richness on visual word recognition. *Psychonomic Bulletin and Review*, 15, 161–167. <https://doi.org/10.3758/PBR.15.1.161>
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, 56–115. <https://doi.org/10.1037/0033-295X.103.1.56>
- Plummer, P., Perea, M., & Rayner, K. (2014). The influence of contextual diversity on eye movements in reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40, 275–283. <https://doi.org/10.1037/a0034058>
- Rastle, K., & Coltheart, M. (1999). Serial and strategic effects in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 482–503. <https://doi.org/10.1037/0096-1523.25.2.482>

- Rayner, K., & Duffy, S. (1986). Lexical complexity and fixation times in reading: Effects of word frequency, verb complexity, and lexical ambiguity. *Memory & Cognition*, *14*, 191–201. <https://doi.org/10.3758/BF03197692>
- Rodd, J., Gaskell, G., & Marslen-Wilson, W. (2002). Making Sense of Semantic Ambiguity: Semantic Competition in Lexical Access. *Journal of Memory and Language*, *46*, 245–266. <https://doi.org/10.1006/jmla.2001.2810>
- Rodd, J. M., Gaskell, M. G., & Marslen-Wilson, W. D. (2004). Modelling the effects of semantic ambiguity in word recognition. *Cognitive Science*, *28*, 89–104. <https://doi.org/10.1016/j.cogsci.2003.08.002>
- Rosa, E., Tapia, J. L., & Perea, M. (2017). Contextual diversity facilitates learning new words in the classroom. *PLoS ONE*, *12*, 1–12. <https://doi.org/10.1371/journal.pone.0179004>
- Schmalz, X., Marinus, E., & Castles, A. (2013). Phonological decoding or direct access? Regularity effects in lexical decisions of Grade 3 and 4 children. *Quarterly Journal of Experimental Psychology*, *66*, 338–346. <https://doi.org/10.1080/17470218.2012.711843>
- Schneider, W., Eschman, A., & Zuccolotto, A. (2012). E-Prime User's Guide. Pittsburgh: Psychology Software Tools, Inc.
- Schock, J., Cortese, M. J., & Khanna, M. M. (2012). Imageability estimates for 3,000 disyllabic words. *Behavior Research Methods*, *44*, 374–379. <https://doi.org/10.3758/s13428-011-0162-0>
- Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*. US: American Psychological Association. <https://doi.org/10.1037/0033-295X.96.4.523>
- Torgesen, J., Wagner, R., & Rashotte, C. (1999). Test of Word Reading Efficiency. Austin, TX.

- van Heuven, W. J. B., Mandera, P., Keuleers, E., & Brysbaert, M. (2014). SUBTLEX-UK: a new and improved word frequency database for British English. *Quarterly Journal of Experimental Psychology* (2006), 67, 1176–1190.
<https://doi.org/10.1080/17470218.2013.850521>
- Yap, M. J., Lim, G. Y., & Pexman, P. M. (2015). Semantic richness effects in lexical decision: The role of feedback. *Memory & Cognition*, 43, 1148–1167. <https://doi.org/10.3758/s13421-015-0536-0>
- Yap, M. J., Tan, S. E., Pexman, P. M., & Hargreaves, I. S. (2011). Is more always better? Effects of semantic richness on lexical decision, speeded pronunciation, and semantic classification. *Psychonomic Bulletin & Review*, 18, 742–750. <https://doi.org/10.3758/s13423-011-0092-y>

Appendix A. Materials used in Experiment 1

High Semantic Diversity		Low Semantic Diversity	
Word	SemD	Word	SemD
ancient	2.26	afraid	2.00
appear	2.11	belong	2.10
battle	2.13	breakfast	2.00
building	2.20	castle	1.96
couple	2.11	certain	2.07
danger	2.11	corner	2.04
desert	2.14	dragon	1.91
escape	2.13	figure	2.08
follow	2.16	forest	2.07
giant	2.22	golden	2.00
heavy	2.12	hurry	2.03
human	2.19	journey	2.02
include	2.48	kitchen	2.03
island	2.15	mutter	2.06
machine	2.29	narrow	2.06
manage	2.14	pocket	2.05
member	2.30	present	2.02
message	2.19	quiet	2.04
number	2.34	receive	2.02
parent	2.15	remark	1.88
problem	2.28	repeat	2.01
produce	2.35	servant	1.80
reason	2.11	service	2.09
record	2.30	silence	2.00
secret	2.12	silver	2.03
station	2.15	struggle	2.06
suggest	2.15	subject	2.02
support	2.33	sudden	2.03
weather	2.14	tumble	2.00

weight	2.17	village	2.08
--------	------	---------	------

Appendix B. Materials used in Experiment 2

Word	SemD	Word	SemD	Word	SemD
adventurous	2.12	coach	1.97	exclamation	1.94
advertise	2.07	complicate	2.22	expose	2.06
afford	2.07	comrade	1.76	external	2.16
alert	2.01	condemn	1.80	extraordinarily	1.97
alley	1.91	conduct	1.89	falcon	1.64
anxiously	2.07	confront	1.95	fare	1.87
apartment	1.88	continually	1.90	faster	2.23
applause	1.91	correct	2.22	feathered	1.98
arched	1.99	credit	2.10	fee	2.13
artificial	2.26	crest	1.91	feel	2.11
astonished	1.91	crocodile	2.13	ferry	2.02
attitude	1.94	crutch	1.88	flake	2.01
automatic	2.02	cry	2.00	fleeting	1.94
available	2.33	cue	2.01	flicker	1.99
beaver	1.69	cupboard	2.01	foolishness	1.75
bedside	1.89	cyclone	2.11	foothold	1.95
beforehand	1.83	debate	2.06	forest	2.07
beg	1.91	depressed	1.93	forgive	1.89
bench	2.01	dice	2.01	former	2.07
berry	2.00	disappoint	2.12	found	2.13
bet	2.14	discover	2.20	foundation	1.98
bike	2.00	disgust	2.04	freak	2.03
birth	2.12	dish	2.07	frequent	2.02
bishop	1.97	dismay	1.97	fringe	1.97
bitterness	1.86	dismiss	1.94	gag	1.99
bleached	1.92	distant	2.05	gale	1.91
boiling	2.07	dolphin	1.86	gap	2.09
brave	2.01	doughnut	2.03	gift	2.02
bravery	2.02	dribble	2.17	good	2.18
brother	2.09	drill	2.20	goose	1.97
calendar	2.09	drowsy	1.77	groove	2.14
carbon	1.68	duck	2.10	guilty	2.08
carrier	2.00	dwarf	1.97	gully	1.79
caution	2.01	eagerly	2.00	hardly	2.03
celebration	2.34	earn	2.14	hasten	1.75
cement	2.25	earthly	1.72	hawk	1.95
chart	2.19	electrical	2.09	herb	1.99
chill	1.94	emerge	2.10	highlight	2.18
chime	2.01	enjoyment	1.77	hoarse	1.88

choice	2.14	entertain	1.98	honestly	1.97
circuit	2.23	essential	2.21	hopeless	1.93
classic	2.08	eternity	1.80	horrible	2.03
club	2.21	excitement	2.06	horrid	1.88
horseman	1.81	nagging	1.95	resist	2.13
hue	1.77	notice	2.09	respectful	1.95
humanity	1.78	obedience	1.93	responsible	2.22
hurricane	2.09	oblivious	2.04	restaurant	2.03
hushed	1.92	occupation	1.90	revolt	2.16
incident	2.02	outrageous	2.10	rich	2.11
increase	2.18	outwards	2.04	ride	2.05
incredulously	2.05	overgrown	1.97	right	2.17
independent	2.07	overhear	2.00	rude	1.97
indulge	1.75	oxygen	1.84	rush	2.04
industrial	1.99	partner	2.05	satellite	1.99
influence	2.06	past	2.14	scarf	1.95
inquire	1.74	pear	1.98	scolding	1.72
inspiration	2.11	phase	2.14	scrabble	2.01
intricate	1.95	plane	2.12	scramble	2.04
inventor	1.98	plastic	2.21	sea	2.14
investigation	2.06	plum	2.01	seal	2.13
jest	1.75	polished	2.04	secret	2.12
join	2.30	porcelain	2.04	seemingly	2.05
jolly	1.88	portable	2.22	sentinel	1.90
journey	2.02	positively	1.96	settler	1.82
keeper	1.98	poster	2.17	sharp	2.05
kin	1.84	powder	2.13	shed	2.11
kindly	1.90	predator	1.99	shell	2.13
land	2.19	pregnant	1.86	site	2.36
large	2.23	presently	1.72	slate	1.97
laughing	2.02	previous	2.08	sleep	2.00
leader	2.22	prime	2.23	smile	2.05
lion	2.06	primitive	2.03	snag	1.97
live	2.23	program	2.20	soggy	2.23
lovingly	1.83	property	2.07	sphere	2.10
lunch	2.07	protector	1.97	split	2.27
male	2.18	protest	2.10	spout	2.06
mansion	1.93	pure	2.03	spreading	2.04
margin	1.94	radioactive	1.60	steeple	1.96
maze	2.08	railing	1.97	straighten	2.04
medicine	2.15	readily	1.87	strait	1.92
menagerie	1.79	reading	2.17	stray	1.94

mental	2.03	reality	2.08	strut	2.08
mentor	1.66	recently	2.23	stubbornly	2.08
message	2.19	red	2.15	summon	1.94
moist	2.08	regain	2.05	sustain	1.97
monster	2.01	relieve	1.99	tail	2.10
move	2.22	remains	2.22	taunt	1.96
mumble	1.98	repair	2.24	technician	1.67
telephone	2.21	twinkle	1.99	vulture	2.02
tentatively	1.95	university	2.22	walnut	1.97
thickness	2.10	unlikely	2.16	watery	2.09
tickle	2.08	unmistakably	1.96	weasel	1.89
timid	1.87	unruly	2.05	weird	2.07
tongue	2.03	value	2.10	whipping	2.04
toss	1.98	vegetable	2.14	windowsill	2.09
towel	1.97	vigorously	2.04	withered	1.81
tradition	2.10	village	2.08	woman	2.03
trample	1.94	villager	1.98	wrap	2.04
treaty	1.87	vitamin	1.76	wrinkle	2.03
trifle	1.80	volcanic	2.02	zone	2.18