

THIS IS AN AUTHOR ACCEPTED MANUSCRIPT OF A PAPER SET TO APPEAR IN THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY. THIS IS NOT THE VERSION OF RECORD AND MAY DIFFER FROM THE FINAL PUBLISHED VERSION.

The activation of contextually predictable words in syntactically illegal positions

Michael G. Cutter¹, Andrea E. Martin^{2,3}, Patrick Sturt⁴

¹ The University of Nottingham

² Max Planck Institute for Psycholinguistics

³ Donders Centre for Cognitive Neuroimaging, Radboud University

⁴ The University of Edinburgh

Word Count: 5141 (Excluding footnotes, figures and tables).

Author Note

Michael G. Cutter, School of Psychology, University of Nottingham, Nottingham, England; Andrea E. Martin, Max Planck Institute for Psycholinguistics and Donders Centre for Cognitive Neuroimaging, Radboud University, Nijmegen, The Netherlands; Patrick Sturt, Department of Psychology, University of Edinburgh, Edinburgh, Scotland.

This research was supported by the Leverhulme Trust Research Project Grant RPG-2017-049, and was completed while the first author was employed at the University of Edinburgh. AEM was supported by the Max Planck Research Group “Language and Computation in Neural Systems” and by the Netherlands Organization for Scientific Research (grant 016.Vidi.188.029)

Correspondence regarding this article should be addressed to Michael G. Cutter, School of Psychology, , University of Nottingham, University Park Campus, Nottingham, NG7 2RD. E-mail: michael.cutter@nottingham.ac.uk.

Our experimental design, sample size, and analysis were pre-registered on the OSF at <https://osf.io/chfub/>. Data files, the R Scripts used in the analysis of our data, and our stimuli are also available on the OSF. We would like to thank Yanxi Lu for assistance with data collection. Our data were presented at a meeting of the Experimental Psychology Society in Bournemouth, UK (July 2019).

Abstract

We present an eye-tracking study testing a hypothesis emerging from several theories of prediction during language processing, whereby predictable words should be skipped more than unpredictable words even in syntactically illegal positions. Participants read sentences in which a target word became predictable by a certain point (e.g. "bone" is 92% predictable given "The dog buried his..."), with the next word actually being an intensifier (e.g. "really"), which a noun cannot follow. The target noun remained predictable to appear later in the sentence. We used the boundary paradigm (Rayner, 1975) to present the predictable noun or an alternative unpredictable noun (e.g. "food") directly after the intensifier, until participants moved beyond the intensifier, at which point the noun changed to a syntactically legal word. Participants also read sentences in which predictable or unpredictable nouns appeared in syntactically legal positions. A Bayesian linear mixed model suggested a 5.7% predictability effect on skipping of nouns in syntactically legal positions, and a 3.1% predictability effect on skipping of nouns in illegal positions. We discuss our findings in relation to theories of lexical prediction during reading.

Keywords; language prediction; eye-movements; word skipping.

Contextual predictability is a key predictor of how easily readers can process a word within a sentence. In electrophysiological work, readers exhibit smaller N400 responses to predictable words (Kutas & Hillyard, 1984); in lexical decision tasks participants respond faster to predictable words (Schwanenflugel & Shoben, 1985); and, in eye movement studies, predictable words are skipped more and fixated for less time (Ehrlich & Rayner, 1981). In the current paper, we outline a hypothesis implicit to several theories of prediction, and test it within an eye-tracking study.

Several theories of prediction assume that readers lexically pre-activate words which are congruent with their current interpretation of the text. Pickering and Gambi (2018) argue that, as people process language, they incrementally update a representation of a speaker's communicative intention with each piece of input, making predictions about upcoming content based on what they believe is being communicated. These predictions lead to pre-activation of information at all linguistic levels, with this pre-activation following the same steps as language production. This includes the pre-activation of a predicted word's form, given enough time. Kuperberg and Jaeger (2016) explain predictability effects within a hierarchical multi-representational generative framework of language comprehension. They propose that as readers process text they incrementally update their beliefs about it at multiple levels, including a message-level representation at the top; beliefs from this level are propagated to lower representational levels, including the level of individual words, resulting in predictive lexical pre-activation. When words predictable from the message-level are encountered, readers have less trouble updating their beliefs, resulting in less processing difficulty. While these approaches differ in other aspects, they both assume readers incrementally construct a representation of a text's meaning, and use this to pre-activate upcoming lexical items.

As mentioned, predictability effects have been observed in eye-movements during reading, with readers skipping predictable words more than unpredictable words (Ehrlich & Rayner, 1981; Fitzsimmons & Drieghe, 2014; Brysbaert, Drieghe, & Vitu, 2005). Even when a predictable word is not skipped, predictability clearly plays a role in the early processing of that word in the parafovea. This was demonstrated by Staub and Goddard (2019), who found that predictability no longer influenced the time spent fixating a word if readers had not received a valid parafoveal preview of that word. Additionally, the fact that readers do not seem to experience a cost of making an incorrect prediction (see Frisson, Harvey, & Staub, 2017; Luke & Christianson, 2016) suggests that these effects are due to the graded pre-activation of potential continuations of a sentence, rather than readers strongly committing to a single sentence continuation. Within the eye-movement literature, predictability effects are attributed to this graded pre-activation shortening early stages of lexical processing, thus affecting saccade programming decisions. For example, in the E-Z Reader model (e.g. Reichle, Warren, & McConnell, 2009), after a reader identifies the fixated word, they begin processing the next word parafoveally while programming a saccade towards it. If what is referred to as the L_1 stage of processing finishes for this word quickly enough then the reader instead programmes a saccade directly to the following word. Crucially, L_1 is shorter for predictable words; consequently, predictable words are skipped more than unpredictable words. This effect of predictability on skipping tends to be of an equivalent size regardless of word length (Rayner, Slattery, Drieghe, & Liversedge, 2011), supporting the notion that these effects are due to an early boost in lexical processing for these words, rather than the completion of lexical identification.

E-Z Reader is fairly agnostic about how predictable words become pre-activated. Equally, theoretical approaches to prediction are more concerned with accounting for findings from ERP studies than the eye-movement literature. However, it is easy to see how

they may be combined. For example in Pickering and Gambi's account the pre-activation of a predicted word's form could explain the shortening of L_1 in E-Z Reader, while in Kuperberg and Jaeger's account greater skipping of predictable words would be due to words consistent with a reader's belief about the text being pre-activated in a graded manner.

An intriguing – and, on the surface, surprising – prediction emerges when considering these theories as mechanisms to explain increased skipping of predictable words. Essentially, within each theory, the pre-activation of specific words should not be limited to syntactic positions within the sentence that these words can appear. Rather, once a word is predictively added to the sentence representation due to the reader's beliefs about the text's intended message, it should remain pre-activated until this belief is disconfirmed. If this is true, readers should skip a contextually predictable word even when it appears in a syntactically illegal position. It is this hypothesis that we test in the current paper. Alternatively, it could be that the pre-activation of a certain word only occurs in positions within the sentence that this word can appear; as such there would be no predictability effects for words appearing in syntactically illegal positions.

We presented participants with sentences in which a specific noun became predictable by a certain point (e.g., *bone* following *The dog buried his...*). In half the sentences this predictable noun or an alternative unpredictable noun (e.g. *food*) appeared at this point. However, in half our sentences the fragment continued with a word that cannot be directly followed by a noun, but that only minimally altered readers' beliefs about the sentence at a message level (e.g. *rather*). Thus, the predictable word should have remained in participants' sentence representation. We used the boundary paradigm (Rayner, 1975) to present a false preview of the word following *rather*, to either be the predictable or unpredictable noun. As the eyes moved beyond *rather* this preview changed to a legal sentence continuation (e.g., *huge*). If predictability exerts its influence on word skipping as hypothesised, readers should

still skip the predictable noun more than the unpredictable noun. If we were to find that the predictability effect is completely absent in syntactically illegal positions it would present a considerable challenge to the theories of predictability discussed above.

Method

Participants

Seventy-eight native English speakers participated.² Furthermore, in our pre-registration we specified a stopping criterion, which was designed to ensure that enough data had been collected to detect a predictability effect of at least 5% in skipping rates. The stopping criterion was found to have been met after 78 participants had been run, and so data collection was terminated. The stopping criterion was based on the 95% credible interval for the estimate of the predictability effect, in the Bayesian analysis described below. The criterion required the difference between the upper and lower limits of the credible interval not to exceed 10%, when transformed from the default logit scale onto percentages. If this criterion were to be met, it would mean that the credible interval would not cross zero, as long as the estimated predictability effect is 5% or more, given that the effect estimate is expected to be at the mid-point of the credible interval.

Apparatus

Eye movements were recorded with an SR-Research Eyelink system running at 1000 hertz. Sentences were displayed on a CRT monitor running at 120 hertz. At this refresh rate display changes took an average of 6.70ms after the eye was detected as having crossed an invisible boundary. Viewing distance was 70cm; 3.4 characters of monospaced font filled 1° of visual angle.

Materials and design

Participants read 56 sentences in a 2X2 design. Our design and sample size were pre-registered (<https://osf.io/chfub/>). The first variable concerned whether our target word appeared in a syntactically legal or illegal position. This variable was between-items, with target words (and sentences) being distinct between these conditions, although well-matched for length, frequency, and predictability. Our second variable was within-items and concerned whether the target noun was predictable or unpredictable. See Figure 1 for an example item in each condition, and <https://osf.io/chfub/> for a full list of stimuli.

Syntactically legal sentences:
 Predictable: The baker put the bread in the hot oven.
 Unpredictable: The baker put the pizza in the hot oven.

Syntactically illegal sentences: 
 Predictable: The dog buried his rather bone poo in the garden.
 Unpredictable: The dog buried his rather food poo in the garden.

Post-boundary:  The dog buried his rather huge poo in the garden.

Figure 1. An example of our stimuli in each condition. For the syntactically illegal sentences the target word would always change to a syntactically valid continuation of the sentence once the eye crossed the invisible boundary, illustrated by the blue line.

Sixty-three English speakers took part in a cloze norming task to assess predictability for 103 items, a majority of which were adapted from Frisson et al. (2017). We designed these stimuli so that an intensifier (e.g. *increasingly*) or co-ordinator following an adjective (e.g. *large but...*) ruled out a noun as the next word in the sentence, once a certain noun had become predictable. Each participant saw ~50% of items in the legal condition, and the remaining items in the illegal condition.³ Each item was rated by at least 31 participants in each condition. For legal items predictability was calculated as the proportion of subjects that provided this word as the next word in the sentence. For illegal items word predictability was calculated as the proportion of subjects that provided this word at some point in their sentence completion.

Based on these ratings we selected 28 sentences to use in the legal condition and 28 to use in the illegal condition. Crucially, our predictable words remained predictable within our sentences even after the introduction of an intensifier or co-ordinator. Stimuli characteristics are shown in Table 1.

These 56 sentences were presented alongside 62 filler items, 42 of which were part of another gaze-contingent experiment. There were also six items provided as practice trials.

Table 1
Characteristics of our Final Stimuli Set.

	Legal		Illegal	
	Predictable	Unpredictable	Predictable	Unpredictable
Mean legal predictability	.89	.01	.91	.01
Mean illegal long-range predictability	-	-	.87	.01
Length	4.89	4.89	4.93	4.93
Log Frequency per million (SUBTLEX-UK)	3.23	3.20	3.23	3.23
Mean log bigram frequency (N-Watch)	2.85	2.96	2.85	2.90
Mean log trigram frequency (N-Watch)	2.11	2.19	2.11	2.07

Legal Predictability refers to the proportion of participants who provided the target word given the sentence up to the word preceding the intensifier/co-ordinator. Illegal long-range predictability refers to the proportion of participants who provided the target word at some point in sentence completions given the sentence up to and including the intensifier/co-ordinator.

Procedure

Participants provided informed written consent upon arrival. A three-point horizontal calibration grid was used, with an acceptance criterion of average error below 0.30 degrees. Each trial was preceded by two drift checks, and a gaze-contingent box at the sentence's first character. Recalibration occurred if either drift check returned a value above 0.40 on two

consecutive trials, or participants failed to trigger the gaze-contingent box. Participants pressed a button after reading each sentence. Yes/no comprehension questions followed 33% of sentences; participants answered using mouse buttons. The experiment lasted ~30 minutes. Our procedure was approved by the University of Edinburgh’s Philosophy, Psychology, and Language Sciences Research Ethics Committee.

Results

Our main analysis was pre-registered (<https://osf.io/chfub/>). Prior to exporting our data from SR-Research DataViewer we removed fixations over 800ms, merged fixations below 80ms with fixations within 0.5 degrees of visual angle, and merged fixations below 40ms with fixations within 1.25 degrees of visual angle. After exporting our data, we removed trials in which pre-target fixations ended in a blink, or in which the display change triggered early, accounting for 14% of data. Accuracy on comprehension questions was high, with participants selecting the correct answer 97% of the time.

Table 2

Mean Skipping Probability for our Target Word in Each Condition

	Legal		Illegal	
	Predictable	Unpredictable	Predictable	Unpredictable
Skipping probability	20.9%	14.8%	18.6%	15.9%

Mean skipping rates are shown in Table 2. We analysed our data using a Bayesian logistic linear-mixed model constructed using the BRMS package (Bürkner, 2017) in R. We used a Bayesian model to allow us to assess to what extent there was a null effect of each contrast, rather than simply test for significant effects (see Nicenboim & Vasishth, 2016 for an accessible introduction to using Bayesian methods in linguistic research). The contrast

structure of our model tested for predictability effects nested under syntactic legality, returning estimates of the main effect of legality, and predictability effects for words in a) legal positions and b) illegal positions. We chose this contrast structure since we were primarily interested in testing whether predictability reliably influences skipping in each sentence type; the nested model allowed us to do this in a single model. While we included a contrast to examine the effect of syntactic legality on skipping, this was only included to account for variance within our model (see Schad, Vasishth, Hohenstein, & Kliegl, 2020), rather than to make any attempt at interpreting this effect. We did not attempt to interpret this effect since our between-items design meant that variables other than syntactic legality could have caused (or suppressed) any observed differences. Our model had weakly informative priors of $Normal(\mu=0, \sigma=10)$ for the intercept and $Normal(0, 1)$ for each fixed effect, and a regularization of 2 on the covariance matrix of random effects. The prior for the intercept reflects the expectation for the grand mean, and, given the logit scale, it corresponds to a distribution with a mean of 50%, and a very wide spread (± 1 standard deviation would cover a range from $<1\%$ to $>99\%$). Thus, it places a negligible restriction on the posterior estimate of the grand mean. The prior for the fixed effects corresponds to a distribution with a zero effect as its mean, and a fairly wide spread---for example, given a mean skipping rate of 50% for one condition, then ± 1 standard deviation would cover a range from 27% to 73%, meaning that we would be 68% certain that the other condition would fall within this range. Given that predictability effects in skipping have typically been found to be in the order of 8% (Brysbaert, Drieghe, & Vitu, 2005), again, this does not place a strong constraint on the posterior estimate of the effect.⁴ It is also worth noting that the posterior estimate will be proportional to the product of the prior and the likelihood of the data, and, with the amount of data we collected the prior will have a minimal influence on the outcome of our analysis. The regularization on the covariance matrix should result in the model returning conservative

estimates for correlations between intercepts and slopes within the model. The models were run with four chains of 5000 iterations, with 1000 iterations being included in the warmup for each chain. The Rhat for all parameters in our model was equal to 1, thus suggesting that the model had successfully converged. The effective sample size for the estimations of the model intercept, the main effect of syntactic legality, the effect of predictability in illegal positions, and the effect of predictability in legal positions was 5352, 9165, 17620, and 16368, respectively.

The population-level effects from our model can be seen in Table 3. Our model showed a clear effect of predictability for words in syntactically legal positions, with $b = -0.50$, 95% CrI = $[-0.83, -0.18]$, $P(b < 0) = 0.999$. The logit estimate for this parameter represents a 5.7% skipping difference, smaller than the average size of this effect (8%; see Brysbaert, Drieghe, & Vitu, 2005). The predictability effect on the skipping of words in illegal positions was also in the predicted direction, and consistent with the majority of the posterior probability mass (see Figure 2), with $b = -0.26$, 95% CrI = $[-0.57, 0.05]$, $P(b < 0) = 0.951$. The mean logit estimate for this parameter represents a relatively small 3.1% skipping difference, and the 95% credible interval included zero. Thus, we cannot claim to have found strong evidence for the effect in the syntactically illegal condition, and the effect is likely very small.

Table 3

Output from our Bayesian mixed effects model. Mean effects estimates (b) and 95% confidence intervals are reported.

	B	L95%	U95%
Intercept	-1.88	-2.15	-1.62
Syntactic legality	0.05	-0.27	0.38
Predictability (legal)	-0.50	-0.83	-0.18
Predictability (illegal)	-0.26	-0.57	0.05




Figure 2. Posterior effects estimates from our Bayesian linear-mixed model. The black line represents an effect of zero. In each posterior distribution the dark grey line represents the mean effect estimate and the shaded grey area represents the 95% credible interval. For completeness we include the estimate for the main effect of syntactic legality here, although do not attempt to interpret this in the main body of our paper.

Supplementary Analysis

The analysis presented above was pre-registered, and as such we consider it the main confirmatory analysis of our data. However, we also report several further analyses, in response to suggestions made by reviewers of a previous version of this paper. These analyses, which are reported below, were not included in the pre-registration, and thus we consider them to be exploratory.

First of all, in addition to our Bayesian linear mixed model a more standard frequentist version of this analysis was requested. Similarly to our Bayesian model, this analysis revealed a non-significant effect of syntactic legality ($b = 0.05$, $SE = 0.15$, $z = 0.35$, $p = 0.73$) and a significant effect of predictability for words in syntactically legal positions ($b = -0.55$, $SE = 0.15$, $z = -3.8$, $p < 0.001$). In contrast to our Bayesian model, this model revealed

a significant effect of the predictability of a word in a syntactically illegal position ($b = -0.32$, $SE = 0.15$, $z = 2.18$, $p = 0.03$). Effects estimates retrieved from the model suggested a 6.6% increase in skipping for predictable relative to unpredictable words in syntactically legal positions and a 3.6% effect in syntactically illegal positions. Thus, in this analysis we do have evidence for an effect of predictability in syntactically illegal positions.

In a further supplementary analysis, we exclusively examined cases in which participants fixated, rather than skipped, the pre-target word. Readers are highly unlikely to skip two consecutive words in the text, meaning that on trials where the pre-target word was not fixated our key dependent variable (i.e. predictability) would have had very little chance of influencing target word skipping. When we re-ran our Bayesian model on only trials where participants fixated the pre-target word there was once again an effect of predictability on the skipping of words in legal positions with $b = -0.59$, 95% CrI = $[-1.02, -0.19]$, $P(b < 0) = 0.997$, with this translating to a skipping effect of 12%. For words in syntactically illegal positions $b = -0.30$, 95% CrI = $[-0.66, 0.06]$, $P(b < 0) = 0.954$, with this translating into an estimated effect on skipping of 4%. Thus, the estimated size of the predictability effect was larger in both legal and illegal positions in this restricted analysis, as would be expected. However, it is worth noting that the 95% credible interval for the effect in illegal positions still included values of 0, meaning that we still do not have strong evidence in favour of our effect in this analysis.

A final supplementary analysis was conducted due to a concern a reviewer had with our stimuli. A small number of target words were repeated throughout our experiment, with about 5% of targets being a repetition of one appearing earlier in the study. It is possible the first instance of a target word may have primed that word the second time it appeared, making it more likely to be skipped. Depending on whether this second instance was the unpredictable word or the predictable word, these repetitions could well have suppressed or

enhanced any skipping differences between these two conditions. To address this, we re-ran our analysis with the second instance of any repetitions removed. In this case, our Bayesian model revealed clear effects of predictability on skipping, in both syntactically legal $b = -0.47$, 95% CrI = $[-0.81, -0.14]$, $P(b < 0) = 0.998$ and illegal positions $b = -0.33$, 95% CrI = $[-0.67, -0.01]$, $P(b < 0) = 0.978$, with these effects estimates translating to skipping effects of 5.5% and 3.7%, respectively.

Discussion

We examined whether participants skipped contextually predictable words more than unpredictable words, both in syntactically legal and illegal positions. Predictability affecting the skipping of words in legal positions is uncontroversial; our data does not contradict this position, with an estimated effect of 5.7%. Whether contextual predictability affects skipping of words in syntactically illegal positions is more novel, and the answer to this question would have interesting implications for several theoretical positions, with a complete absence of an effect potentially being problematic for the theories outlined above. Unfortunately, the pre-registered analysis of our data does not allow us to be overly strong in our conclusions, in that we cannot definitively claim that there are strong predictability effects for words in illegal positions, but we can also certainly not argue in favour of the null hypothesis. Supplementary analyses which were not pre-registered, but which reviewers requested of us, did seem to offer slightly stronger support for our hypothesis, such that our effect was significant in what would be viewed as a more standard frequentist analysis of our data, seemed to grow when we restricted our analysis to cases in which the pre-target word was fixated, and also became reliable when trials potentially confounded by target repetitions were excluded. On this basis, it seems fair to conclude that predictability most likely does affect the skipping of words in illegal positions, but likely to a lesser extent than words in

legal positions. We discuss the implications of our experiment with this interpretation, while acknowledging that future studies may need to further examine this issue.

Our reason for suspecting that contextual predictability may influence the skipping of words even in syntactically illegal positions was grounded in assumptions of theoretical accounts of prediction (e.g. Kuperberg & Jaeger, 2016; Pickering & Gambi, 2018). These theories assume that readers predict words on the basis of their beliefs about the sentence's meaning, with this process resulting in the word being identified faster if it is encountered. Under this assumption, words consistent with people's beliefs about a sentence should be pre-activated for the time between that word becoming predictable and this prediction being confirmed or disconfirmed. As such, the complete absence of a predictability effect in syntactically illegal positions would have been challenging for these theories. Our data suggest that there is a level of pre-activation for predictable words in syntactically illegal positions, thus lending partial support to these accounts. However, our data also suggests this effect is relatively small compared to the effect in legal positions in the current study and the wider literature. Thus, it seems likely that while contextually predictable words are still pre-activated in illegal positions, there is some level of suppression for these words.

It may seem obvious that any predictability effects should be smaller for words in syntactically illegal positions. However, this may not be predicted within the theoretical accounts. Essentially, in our illegal sentences there would have been a longer delay between a word first being added to readers' beliefs about the sentence and that word appearing in the text than in our syntactically legal sentences. This would result in more time for pre-activation of that word to occur; for example, in Pickering and Gambi's (2018) prediction-by-production theory there would be a greater amount of time after readers derived the intention of the text to move through various stages of language production and thus predict the form of the upcoming word, leading to stronger effects. Similarly, within Kuperberg and Jaeger's

theory, once information is inferred at a message-level there should be more time for this to propagate down to lower levels, leading to greater pre-activation of individual words and thus greater predictability effects. Thus, while the (likely) existence of a predictability effect on the skipping of words in illegal positions is in line with these theories, the modest effect size may not be. As these theoretical accounts are developed further they may well need to take account of the current effects, with a mechanism that is able to explain how and why the activation of words that remain contextually predictable is suppressed but not eliminated in certain sentence positions; as it stands, we cannot currently see any mechanism that would account for this. A slight suppression of predictability effects in illegal positions is consistent with a recent theory of language processing based on the idea of hierarchical cue integration, proposed by Martin (2016). In this theory low-level representations cue activation of higher-level representations (e.g. letters cue bigrams), and, crucially, high-level representations cue lower-level representations (e.g. syntactic structure and a situation model can jointly cue word identity). In relation to our study, a particular word may be cued by both the situation model and syntactic structure (alongside a range of low-level representations) in a legal position, but only by the situation model in a syntactically illegal position. Thus, a decrease in predictability effects in syntactically illegal positions would be consistent with this account of language processing.

While we have focussed on how small the effect of predictability is in illegal positions, it should be noted that the effect in legal positions was fairly modest in our study, with an effect of 5.7% compared to an average of 8% (Brysbaert, Drieghe, & Vitu 2005), and maximum of around 14% across the literature (e.g. Fitzsimmons & Drieghe, 2013). This small effect could either be due to certain characteristics of our stimuli or subjects, or it could just be that Brysbaert et al.'s meta-analysis overestimates the size of the effect. Thus, if the predictability effect for words in illegal positions is genuinely smaller than in syntactically

legal positions, in the current study it would have required an exceptionally large sample to determine that an intermediate effect is both a) reliably larger than zero, and b) reliably smaller than 5.7%. Future investigations of the current phenomenon may be able to capitalize upon adapting stimuli from studies that have found larger predictability effects. It may be easier to more definitively establish whether predictability still influences word skipping in illegal positions when the baseline effect is larger. If the effect does still exist in illegal positions but with some level of suppression, then running studies with varying baseline effects may shed light upon the nature of suppression. For example, will the difference always correspond to a proportional influence of syntactic illegality?

One question that readers may have regarding our effects, and the effect of contextual predictability more generally, is the extent to which they are truly due to readers predicting a word on the basis of an event representation of the sentence as opposed to lexical priming between semantically related words. In the former account, predictability effects are due to readers inferring that a *bone* (and not *food*) is a thing that a *dog* would typically *bury*. In the latter account, predictability effects would be due to the words *dog* and *bury* semantically priming *bone* more than *food*. Clearly, we favour the former account, and believe that a subset of data from Brothers and Traxler (2016) may suggest that the latter account is unlikely. Brothers and Traxler (2016) found that a syntactically illegal word was not skipped any more when that word was a repetition of one appearing earlier in the sentence (e.g. *The admiral would not admiral...*) than when it was not (e.g. *the admiral would not surgeon*). This contrasts to a study by Choi and Gordon (2013), who found a repetition priming effect of 5% for the skipping of words in syntactically legal positions. Thus, a word appearing in a syntactically illegal position seems to wipe out lexical priming effects. Arguably, a word seems likely to be more strongly primed by itself than a collection of semantic associates

appearing earlier in the sentence. As such, we favour an alternative locus of our effect, in the form of pre-activation of a word through contextual predictions.

Our study contributes to a growing literature examining how syntactic legality affects word skipping. The picture emerging from this literature suggests that lexical characteristics (e.g. frequency) may still affect word skipping in syntactically illegal positions (e.g. Angele, Laishley, Rayner, & Liversedge, 2014), but that syntactically illegal words may be skipped less than syntactically legal words with similar lexical characteristics (e.g. Brothers & Traxler, 2016; Veldre & Andrews, 2018). Brothers and Traxler argued that their finding was due to an anticipatory mechanism, whereby readers expected a word of a certain syntactic class, leading to pre-activation of words of this class, and thus increased skipping. The fact that contextual predictability may still influence word skipping when a word appears in an illegal position suggests that this anticipatory mechanism does not allow syntactic considerations to completely override other contextual considerations. It should also be noted that the lack of a main effect of syntactic legality on word skipping in our study should not be interpreted as a failure in replicating this prior work. Due to our between-items design there are a number of factors which could have suppressed any syntactic legality effects on word skipping, such as pre-target skipping rates and the location of the final fixation on the pre-target word relative to the start of the target word.

In closing, we tested whether readers are more likely to skip a word which is predictable within their semantic representation of a sentence than one which is not, even when this word appears in an illegal position. Our data suggest that there was indeed a small predictability effect for these words, although further research may be required in order to further determine the size and reliability of this effect, potentially using the posterior estimate from the current paper as a prior in further Bayesian analysis.

References

- Angele, B., Laishley, A. E., Rayner, K., & Liversedge, S. P. (2014). The effect of high- and low-frequency preview and sentential fit on word skipping during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40, 1181-1203.
<http://dx.doi.org/10.1037/a0036396>
- Brothers, T., & Traxler, M.J. (2016). Anticipating syntax during reading: Evidence from the boundary change paradigm. *Journal of Experimental Psychology: Learning Memory and Cognition*, 42, 1894-1906. Doi: 10.1037/xlm0000257
- Brysbaert, M., Drieghe, D., & Vitu, F. (2005). Word skipping: Implications for theories of eye movement control in reading. In G. Underwood (Ed.), *Cognitive Processes in Eye Guidance*, UK: Oxford University Press, pp. 53-77.
- Bürkner, P-C. (2017). *Brms: A package for Bayesian multilevel models using Stan* (R Package Version 2.5.0). Retrieved from <https://cran.r-project.org/web/packages/brms/index.html>
- Christianson, K., & Luke, S. G. (2016). Limits on lexical prediction during reading. *Cognitive Psychology*, 88, 22-60. <https://doi.org/10.1016/j.cogpsych.2016.06.002>
- Ehrlich, S. F., & Rayner, K. (1981). Contextual effects on word perception and eye movements during reading. *Journal of Verbal Learning and Verbal Behaviour*, 20, 641-655. [https://doi.org/10.1016/S0022-5371\(81\)90220-6](https://doi.org/10.1016/S0022-5371(81)90220-6)
- Fitzsimmons, G., & Drieghe, D. (2013). How fast can predictability influence word skipping during reading? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39, 1054-1063. Doi: 10.1037/a0030909

- Frisson, S., Harvey, D. R., & Staub, A. (2017). No prediction error cost in reading: Evidence from eye movements. *Journal of Memory and Language*, 95, 200-214.
<https://doi.org/10.1016/j.jml.2017.04.007>
- Kuperberg, G. R., & Jaeger, T. F. (2016). What do we mean by prediction in language comprehension? *Language, Cognition and Neuroscience*, 31, 32-59.
<https://doi.org/10.1080/23273798.2015.1102299>
- Kutas, M, & Hillyard, S. A. (1984). Brain potentials during reading reflect word expectancy and semantic association. *Nature*, 307, 161-163. <https://doi.org/10.1038/307161a0>
- Martin, A. E. (2016). Language processing as cue integration: Grounding the psychology of language in perception and neurophysiology. *Frontiers in Psychology*.
[doi:10.3389/fpsyg.2016.00120](https://doi.org/10.3389/fpsyg.2016.00120)
- Nicenboim, B., & Vasishth, S. (2016). Statistical methods for linguistic research: Foundation ideas – Part II. *Language and Linguistics Compass*, 10, 591-613.
<https://doi.org/10.1111/lnc3.12207>
- Pickering, M. J., & Gambi, C. (2018). Predicting while comprehending language: A theory and review. *Psychological Bulletin*, 144, 1002-1044.
<http://dx.doi.org/10.1037/bul0000158>
- Rayner, K. (1975). The perceptual span and peripheral cues in reading. *Cognitive Psychology*, 7, 65-81. [https://doi.org/10.1016/0010-0285\(75\)90005-5](https://doi.org/10.1016/0010-0285(75)90005-5)
- Rayner, K., Slattery, T., Drieghe, D., & Liversedge, S. P. (2011). Eye movements and word skipping during reading: Effects of word length and predictability. *Journal of Experimental Psychology: Human Perception and Performance*, 37, 514-528.
<https://doi.org/10.1037/a0020990>

- Reichle, E. D., Warren, T., & McConnell, K. (2009). Using E-Z Reader to model the effects of higher level language processing on eye movements during reading. *Psychonomic Bulletin & Review*, 16, 1-21. <https://doi.org/10.3758/PBR.16.1.1>
- Schad, D. J., Vasishth, S., Hohenstein, S., & Kliegl, R. (2020). How to capitalize on a priori contrast in linear (mixed) model: A tutorial. *Journal of Memory and Language*. Advance online publication. <https://doi.org/10.1016/j.jml.2019.104038>
- Schwanenflugel, P. J., & Shoben, E. J. (1985). The influence of sentence constraint on the scope of facilitation for upcoming words. *Journal of Memory and Language*, 24, 232-252. [https://doi.org/10.1016/0749-596X\(85\)90026-9](https://doi.org/10.1016/0749-596X(85)90026-9)
- Staub, A., & Goddard, K. (2019). The role of preview validity in predictability and frequency effects on eye movements in reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 45, 110-127. <https://doi.org/10.1037/xlm0000561>
- Veldre, A., & Andrews, S. (2018). Beyond cloze probability: Parafoveal processing of semantic and syntactic information during reading. *Journal of Memory and Language*, 100, 1-17. Doi: 10.1016/j.jml.2017.12.002

Footnotes

¹ For brevity we do not discuss other models of eye movement control in depth – however, we do not see that major alternative models would make predictions diverging from those of E-Z Reader in the current study.

² An additional six participants were tested but removed due to poor tracking or noticing more than five display changes. Of the participant's whose data we included 17 noticed at least one display change, with subjects across the whole experiment noticing an average of 0.56 display changes each.

³ In the legal condition participants' instructions were "In the following sentences you simply need to guess the next word, and enter this into the box. There is no need to provide multiple words or finish the sentence. Just make sure that the word fits into the sentence grammatically." For illegal items, instructions were "In the following sentences, you need to complete the sentence sensibly. You can do this in however many words you feel is necessary. Please ensure that the final sentence follows English grammar, and is sensible and complete".

⁴ Results of the analysis were very similar when we adopted the default brms flat priors, which effectively place no constraint at all on the size of the fixed effects.

⁵ The figure of 5% takes into account the effect of counterbalancing on the distribution of stimuli across participants.