Are words pre-activated probabilistically during sentence comprehension? Evidence from new data and a Bayesian random-effects meta-analysis using publicly available data

Bruno Nicenboim^{1,2}, Shravan Vasishth², & Frank Rösler³

¹ Tilburg University

² University of Potsdam

³ University of Hamburg

Abstract

Several studies (e.g., Wicha, Moreno, & Kutas, 2003; DeLong et al., 2005) have shown that readers use information from the sentential context to predict nouns (or some of their features), and that predictability effects can be inferred from the EEG signal in determiners or adjectives appearing *before* the predicted noun. While these findings provide evidence for the pre-activation proposal, recent replication attempts together with inconsistencies in the results from the literature cast doubt on the robustness of this phenomenon. Our study presents the first attempt to use the effect of gender on predictability in German to study the pre-activation hypothesis, capitalizing on the fact that all German nouns have a gender and that their preceding determiners can show an unambiguous gender marking when the noun phrase has accusative case. Despite having a relatively large sample size (of 120 subjects), both our preregistered and exploratory analyses failed to yield conclusive evidence for or against an effect of pre-activation. The sign of the effect is, however, in the expected direction: the more unexpected the gender of the determiner, the larger the negativity. The recent, inconclusive replication attempts by Nieuwland et al. (2018) and others also show effects with signs in the expected direction. We conducted a Bayesian random-effects meta-analysis using our data and the publicly available data from these recent replication attempts. Our meta-analysis shows a relatively clear but very small effect that is consistent with the pre-activation account and demonstrates a very important advantage of the Bayesian data analysis methodology: we can incrementally accumulate evidence to obtain increasingly precise estimates of the effect of interest.

Keywords: ERP, pre-activation, predictions, grammatical gender, Bayesian meta-analysis

Word count: 10714

1 Introduction

Consider the following sentence, with two possible continuations:

- (1) The grandmother wants to pick mushrooms in the forest. She takes ...
 - a. [HIGH-PROB.] *a basket*.
 - b. [LOW-PROB.] *a cloth*.

It is well known that the predictability of a word affects sentence comprehension. Words with low predictability are read slower than words with high predictability (e.g., Ehrlich & Rayner, 1981). In the EEG literature, it has been shown that words with low-predictability are accompanied by a relative negativity that peaks around 300-500 ms after word onset over central parietal scalp sites in comparison with high-predictable words, this difference between conditions is usually referred to as an N400 effect, (first noticed in Kutas & Hillyard, 1980, for semantic anomalies; and in Kutas & Hillyard, 1984, for low predictable words; for a review, see Kutas & Federmeier, 2011).¹ In (1), for example, the continuation “*a basket*” has higher predictability than the continuation “*a cloth*”, and thus we would expect a more negative signal in “*a cloth*” in (1b) in comparison with “*a basket*” in (1a).

There is no doubt that language comprehension is at least minimally predictive: The human parsing system seems to be a left-corner parser (Resnik, 1992) and the context clearly influences the state of the system (see the discussion in Kuperberg & Jaeger, 2016). However, the fact that a word is predictable does not necessarily entail that the lexical item and/or its semantic (and maybe phonological or grammatical) features are always going to be automatically predicted: The N400 effects seen when words with different predictability are compared (e.g., “*basket*” vs. “*cloth*” in (1a) and (1b)) could be due to differences in the difficulty in integrating their meaning with the semantic context rather than differences in preactivation: Any factors that facilitate pre-activation could correspondingly be argued to make semantic integration easier (Van Berkum, Brown, Zwitserlood, Kooijman, & Hagoort, 2005).

There are, in fact, two views regarding the relationship between predictability and the amplitude of the EEG signal at the N400 spatiotemporal window: (i) the access view (Kutas & Federmeier, 2000; Kutas & Federmeier, 2011) which can be linked to the idea

¹However, earlier and later effects have also been reported depending on the modality of the task.

For partial support of this research, we thank the Volkswagen Foundation through grant 89953 to Shravan Vasishth and the Deutsche Forschungsgemeinschaft for two grants: VA 482/8-1 to Shravan Vasishth and Frank Roesler, which funded Bruno Nicenboim, and the Project number 317633480 – SFB 1287, Project Q. Thanks to Arnold Kochari and Monique Flecken for their readiness to share their data and answer questions about their analysis, thanks to Mante Nieuwland and colleagues, Aine Ito and colleagues for sharing their datasets. In addition, thanks to Marte Otten Sussex and Jos van Berkum for sharing stimuli and datasets (even if we couldn’t use it).

Correspondence concerning this article should be addressed to Bruno Nicenboim, Tilburg University, Department of Cognitive Science and Artificial Intelligence, Tilburg University, 5000 LE, Tilburg, The Netherlands. E-mail: bruno.nicenboim@uni-potsdam.de

of prediction error (Rabovsky & McRae, 2014; Kuperberg & Jaeger, 2016), and (ii) the semantic integration view (e.g., Brown & Hagoort, 1993). These views, however, are not mutually exclusive, and both access/prediction error and integration may play a role (Hagoort, Baggio, & Willems, 2009; Nieuwland et al., 2019), but they can be disentangled. We briefly describe these two views below:

(i) Access view Before a word w is read, its context triggers the pre-activation of semantic features in memory (e.g., Kutas & Federmeier, 2000; Rabovsky & McRae, 2014; Kuperberg & Jaeger, 2016) or the entire lexical item (e.g., DeLong, Urbach, & Kutas, 2005). Reading the word w triggers a memory access. The amplitude of the N400 evoked is a function of the mismatch between the pre-activation due to the context and the activation due to word w . While the access view is compatible with both lexical pre-activation and the pre-activation of features, recent computational modeling work shows that most effects reported in the literature are more easily accounted for by the idea of pre-activation of (semantic) features (Rabovsky & McRae, 2014).

(ii) Semantic integration view After the word w is read and retrieved from memory, it needs to be integrated into the partial semantic representation that has been built so far. The N400 is a function of the difficulty of integrating a word in its global and local context (Van Berkum, Hagoort, & Brown, 1999). The reason why predictable words are easier to integrate is that they are also more plausible in the given context and fit the reader/listener’s world knowledge (Baggio & Hagoort, 2011; Hagoort, Hald, Bastiaansen, & Petersson, 2004).

1.1 Evidence for the access view

One of the pieces of evidence that appeared to settle the debate in favor of the access view comes from a series of experiments that manipulated the gender of a determiner or adjective (in Spanish: Wicha, Moreno, & Kutas, 2003; Wicha, Bates, Moreno, & Kutas, 2003; Wicha, Moreno, & Kutas, 2004; and in Dutch: Van Berkum et al., 2005; Otten, Nieuwland, & Van Berkum, 2007; Otten & Van Berkum, 2008; Otten & Van Berkum, 2009) or the phonological form of a determiner (DeLong et al., 2005) that appeared prior to a predicted noun. Although these studies varied considerably regarding methodology (see Table 1), what they had in common was that they all intended to provide evidence in favor of the access view.

Wicha, Moreno, and Kutas (2003) used ERP to examine the role of grammatical gender and semantic congruity in Spanish. In their experiment, native Spanish speakers read sentences in which a drawing depicting a target noun was either congruent or incongruent with sentence meaning, and either agreed or disagreed in gender with the preceding article. The most novel finding in the study was that when determiners with gender markings that were unexpected based on the context up to that point were compared with expected determiners as in (2b) vs (2a), an N400 effect was found (notice that BASKET was presented as a picture in (2)). It is important to notice that a determiner of either gender is syntactically, semantically, and pragmatically correct at this point, since Red Riding Hood in sentence (2) could have been carrying food in a *sack_{masc}* (“*un costal*”) or in a *basket_{fem}* (“*una canasta*”), though *basket* was more expected according to an offline cloze task. It can

be argued that studies using this type of gender manipulation provide evidence for the access view and against semantic integration, because gender information is largely arbitrary, and synonymous words that fit into the context equally well can have different genders.

- (2) a. Caperucita Roja llevaba la comida para su abuela en una
 Little Red Riding Hood carried the food for her grandmother in a_[fem]
 CANASTA muy bonita...
 BASKET_[fem] very pretty...
- b. Caperucita Roja llevaba la comida para su abuela en un
 Little Red Riding Hood carried the food for her grandmother in a_[masc]
 CANASTA muy bonita...
 BASKET_[fem] very pretty...

However, other experiments using gender in Spanish and Dutch did not show a consistent pattern of results at the prenominal determiner or adjective; see Table 1 for a summary of published results. In all these studies, although the ERP elicited by the unexpected noun consistently showed an N400, no consistent ERP effect was found at the determiner or adjective, not even when the experiment was repeated within the same laboratory. As discussed below, either a negativity or a positivity was found at the determiner/adjective.

Spanish shows a more consistent pattern of results than Dutch. However, while both Wicha, Bates, et al. (2003) and Wicha, Moreno, and Kutas (2003) found an N400 effect for determiners with an unpredictable gender in comparison with a predictable one, Wicha et al. (2004) reported a late positivity. While there certainly are differences between the studies (see Table 1), there is no clear explanation for the difference in the patterns. Interestingly, Foucart, Martin, Moreno, and Costa (2014) did find an N400 effect for differences in predictability of determiners in Spanish sentences presented to Spanish native and non-native speakers with a design similar to Wicha et al. (2004), where a positivity was reported. One important difference between the studies is that in contrast to the studies led by Wicha, Foucart et al. (2014) showed no sentences with agreement or semantic violations. Similar results to Foucart et al. (2014) were found by Martin, Branzi, and Bar (2018) with only native speakers.

Dutch, in contrast, shows more mixed results. While Van Berkum et al. (2005) detected a very early positivity at unpredictable marked adjectives in comparison with predictable ones (from 50ms after the gender inflection), Otten et al. (2007) detected a negativity (300-600ms after the gender inflection) with a very similar design (and auditory presentation in both experiments). Furthermore, in a similar experiment with visual presentation, Otten and Van Berkum (2008) showed a very late negativity: about 900 ms after onset of the critical adjective. Otten and Van Berkum (2009) examined predictable vs. unpredictable determiners, which in contrast to adjectives that are inflected, have a different lexical form (*het* vs *de*). However, Otten and Van Berkum (2009) detected a negativity for unpredictable determiners in comparison with predictable determiners with visual presentation, similarly to what was detected in Otten et al. (2007) (with adjectives and auditory presentation). Kochari and Flecken (2019) attempted a replication of Otten and Van Berkum (2009) with a larger sample size ($N = 70$ vs. $N = 38$), and while they failed

to find a significant effect at the gender-marked determiner, they did observe a pattern consistent with the original experiment. Recently, Fleur, Flecken, Rommers, and Nieuwland (2019) manipulated gender expectancy and whether the expected noun phrase would be definite or not in Dutch, and detected an N400-like effect for the gender manipulation.

DeLong et al. (2005) capitalized on the fact that the form of an indefinite determiner in English depends on the phonological form of the following word (“*an*” precedes nouns beginning with vowel sounds, whereas “*a*” precedes nouns beginning with consonant sounds). In their experiment, participants were presented with sentences ending with a predictable noun phrase, such as “*a kite*” in (3a), or an unpredictable one, such as “*an airplane*” in (3b).

- (3) a. The day was breezy so the boy went outside to fly a kite.
- b. The day was breezy so the boy went outside to fly an airplane.

Since the determiners preceding the critical noun have identical minimal meaning, there should be no difference between the fit of “*a*” and “*an*” to the semantic context (i.e., both determiners should incur the same integration costs). However, DeLong et al. (2005) showed that the amplitude of the negativity was smaller with increasing cloze probability of the word, both at the noun and, critically, at the determiner. This is evidence for the access view.

However, recent replication attempts (Ito, Martin, & Nieuwland, 2017a; Nieuwland et al., 2018) have cast doubt on the robustness of the original results of DeLong et al. (2005).²

Ito et al. (2017a) used materials and a procedure similar to those of DeLong et al. (2005), but failed to find a significant N400 effect at the articles in two experiments with different presentation rates (reported for native English speakers in Ito et al., 2017b), but unexpected articles did elicit numerically larger N400 amplitudes than expected ones.

It might be that the differences between DeLong et al. (2005), and Ito et al. (2017a) have to do with low statistical power: The original study was conducted with 32 participants and the attempted replications with 23 participants. It is likely that both studies had relatively low power. It is well-known (Gelman & Carlin, 2014) that when power is low, statistically significant results will be overestimates, Type-M(agnitude) errors, that will be difficult to replicate (Vasishth, Mertzen, Jäger, & Gelman, 2018). However, Nieuwland et al. (2018) attempted a large-scale direct replication (nine laboratories with a total of 334 participants) of DeLong et al. (2005). They also failed to find the statistically significant effect claimed in the original study with the same statistical analysis, although the sign of the effect was in the expected direction in a linear mixed model analysis.

²Martin et al. (2013) also conducted a similar study; however, Martin et al. (2013) used an average reference instead of linked mastoids as is common in studies investigating the N400 making the comparison with the previous findings less straightforward (see Nieuwland et al., 2018; Ito, Martin, & Nieuwland, 2017b). A reviewer points out that it is not too difficult to compare results from different references, although admittedly the average reference, which reduces the data rank, is slightly more challenging than the other unipolar references, which are just affine transformations.

Nieuwland et al. (2018) noticed a possible reason for the failure to find evidence for the access view in the DeLong et al. manipulation: The implicit assumption of the DeLong et al. study was that people expect the initial phoneme of the predicted noun immediately after the determiner. However, unexpected determiners do not necessarily disconfirm the upcoming noun, they might be indicating that another word would come first (e.g., “*an old kite*”). Nieuwland and colleagues argue that given that corpus data (the Corpus of Contemporary American English and the British National Corpus) show that “*a*”/“*an*” is followed by a noun only one third of the time, participants might only occasionally be actively predicting a determiner matching the initial phoneme of the noun.

Table 1
List of EEG studies with a manipulation before the predicted noun.

Study	Year	N _{subj}	Lang.	Manip.	ROI	Pres.	ISI	Others	Results
DeLong, Urbach, and Kutas	2005	32	English	a/an	det.	200	300		Neg 200-500 CP(R)
Ito, Martin, and Nieuwland	2017a	23	English	a/an	det.	300	200		ns (Neg 300-500)
Ito, Martin, and Nieuwland	2017a	23	English	a/an	det.	500	200		ns (Neg 300-500)
Martin et al.	2013	19	English	a/an	det.	200	300		Neg 250-400 ***
Nieuwland et al.	2018	356	English	a/an	det.	200	300		ns (Neg 200-500)
Wicha, Bates, et al.	2003	31	Spanish	gender	det.	audio	-	drawing after the ROI, ungrammatical fillers	Neg 300-500ms FM
Wicha, Moreno, and Kutas	2003	28	Spanish	gender	det.	300	200	drawing after the ROI, ungrammatical fillers	Neg 300-600 W
Wicha, Moreno, and Kutas	2004	28	Spanish	gender	det.	300	200	ungrammatical fillers	Pos 500-700 W (L)
Foucart et al.	2014	17	Spanish	gender	det.	200	500	ungrammatical fillers	Neg 300-500 CP **
Martin, Branzi, and Bar	2018	20	Spanish	gender	det.	200	500	while tong tapping	Neg 300-500
Martin, Branzi, and Bar	2018	20	Spanish	gender	det.	200	500	while listening to own voice saying /ta/	Neg 300-500
Van Berkum et al.	2005	24	Dutch	gender	adj.	audio	-		Pos 50-250 *
Otten et al.	2007	32	Dutch	gender	adj.	audio	-		Neg 300-600 RF
Otten and Van Berkum	2008	24	Dutch	gender	adj.	346	106		Neg 900-1100 (F)
Otten and Van Berkum	2009	38	Dutch	gender	det.	376	106		Neg 200-600 (R)
Kochari and Flecken	2019	70	Dutch	gender	det.	358	400		ns (Neg 200-600 R)

Note. In each study, the EEG signal is analyzed from the onset of the region of interest (except if indicated with (*), where it is measured from the inflection); the results are based only on native speakers (except when indicated with (**), where 29 bilinguals were included in the analysis); the reference for the EEG is the average of left and right mastoids (except when indicated with (***) where the reference was a global average); the experiment is conducted in serial visual presentation (except if it is indicated with 'audio'); and the critical region is always presented alone. The following abbreviations are used; Pres.: presentation duration, ISI: interstimulus interval, ROI: region of interest, F: frontal, C: central, M: medial, P: parietal, W: widely distributed, R: right, L: left, Neg: negative, Pos: positive, ns: non significant.

1.2 Some potential reasons for the inconsistencies in the literature

In the above-mentioned studies, even if there is a difference in the N400 spatiotemporal window at the determiner as a function of gender-marking, it might be the case that it is too small and/or too noisy to be reliably detected. We identify two closely related issues that may be contributing to the inconsistencies across studies.

The first issue is that the statistical comparison across conditions for the determiner or adjective is performed across different words with different lengths and frequencies. In English, “a” is shorter and more frequent than “an”. In languages with gender-marked determiners/adjectives, some genders are more frequent than others. For example, in Dutch, common-gender words are much more frequent than neuter-gender words (Van Berkum, 1996). Even though comparisons across different words are fairly common in the ERP literature on predictions, comparing different words might be problematic when one is dealing with small and subtle effects. One countermeasure adopted to mitigate the effect of comparing different words is to counterbalance the appearance of the words in high and low predictability conditions. However, different words contribute to their own idiosyncratic EEG signal and this may increase the variability of the signal, making it more difficult to find the expected N400 or any other effect. The problem of small effects is coupled with the fact that in many of the studies of the prediction literature, ERP analysis is performed after visual inspection of the data, and thus it is unclear to which extent the results from the studies hinge on the specific choices that were made during the analysis (see also Luck & Gaspelin, 2017; Nieuwland, 2019).

A second issue is that in all the previously mentioned studies, predictability is manipulated by making one continuation highly predictable and the other unpredictable. For example, given a context sentence like “*The day was breezy so the boy went outside to fly...*”, “*a kite*” is highly predictable because it has a high cloze probability, and “*an airplane*” is unpredictable because it has a low cloze probability. However, in the case of the unpredictable “*an airplane*”, not only does the comprehender have to process an unpredictable continuation, but in addition the comprehender experiences a violation of a highly expected continuation (“*a kite*”) in a *constraining context*. The context “*The day was breezy so the boy went outside to fly...*” is constraining because it has a potential *best completion* with high cloze probability, regardless of the cloze probability of the actual continuation. In terms of information theory, a constraining context has lower entropy (or in other words, lower uncertainty) than a less constraining one, because in constraining contexts the cloze probability mass is concentrated in one (or few) responses.³ The low-cloze coupled with the violation of a highly expected continuation (i.e., a low-cloze word in a highly constraining context) may lead to a bi-phasic response: an N400 followed by a late positivity. In fact, Van Petten and Luka (2012) show that biphasic ERP responses (negativities followed by frontal positivities) are fairly common when low vs. high-cloze words are examined in high-constraint contexts, and it has been hypothesized that the late positivity is a response to a violation of a strong expectation (Federmeier, Wlotko, De Ochoa-Dewald, & Kutas, 2007;

³For example, a constraining context might have nine completions of a certain kind (e.g., “*kite*”) and another different one (e.g., “*plane*”), then its entropy would be $H_{constraining} = -\sum_i P_i \log_2 P_i = -(0.9 \log_2(0.9) + 0.1 \log_2(0.1)) = .47$; in contrast, a non-constraining context might have ten different completions, and thus its entropy would be $H_{non-constraining} = -10(0.1 \log_2(0.1)) = 3.3$.

Van Petten & Luka, 2012). Although the late positivity has been studied only for nouns and not for the unexpected gender (or phonological marking) of the determiner or adjective, it might be the case that the effect of the low cloze probability and the violation of a strong expectation are confounded at the determiner/adjective.

A potential solution to this confound is to introduce two conditions that appear in a non-constraining context. Consider for example the design of the studies led by Otten (Otten et al., 2007; Otten & Van Berkum, 2008; Otten & Van Berkum, 2009), where a non-constraining context was also included. This is exemplified with (4) and (5) from Otten et al. (2007) translated from Dutch to English, where the two Dutch “genders”, neuter and common, are indicated with *neut* and *com*.

- (4) Constraining context:
Anne had finally found a quiet place for studying. She sat down and grabbed a ...
 - a. ...*big_{neut}* and pretty well-thumbed book_{*neut*} out of her bag.
 - b. ...*big_{com}* and pretty well-thumbed novel_{*com*} out of her bag.
- (5) Non-constraining context:
After studying Anne had found a quiet place in the park. She sat down and grabbed a ...
 - a. ...*big_{neut}* and pretty well-thumbed book_{*neut*} out of her bag.
 - b. ...*big_{com}* and pretty well-thumbed novel_{*com*} out of her bag.

In this experimental design, in the constraining context (4) there are a highly predictable adjective *big_{neut}*, and an unpredictable adjective *big_{com}* that also involves a violation of a strong expectation. A non-constraining context (5) is added to this design; here, the same adjectives *big_{neut}* and *big_{com}* appear as continuations, but both are unexpected. Thus, the non-constraining context is intended to serve as a baseline for the constraining context: the prediction effect at the adjective would be expressed as an interaction between constraint (constraining vs. non-constraining) and completion (*big_{neut}* vs. *big_{com}*). That is, the focus is on the difference [(4a)– (4b)] – [(5a)– (5b)]. This is because no difference is expected due to predictability between the baseline continuations *big_{neut}* and *big_{com}* in the non-constraining context (5), but a difference is expected due to predictability between the two continuations in the constraining context (4).

We argue that modifying the design to have a constraining and non-constraining context and examining the interaction in the resulting 2×2 design may not be enough to rule out an overlap between negativity and late positivity (since this has been reported for nouns at least; Van Petten & Luka, 2012). This is so because the first term [(4a)– (4b)] might be eliciting both a negativity and a late positivity, while the second term [(5a)– (5b)] is not expected to differ (at least not as much as the first set of sentences), and thus is not acting as a baseline. We suggest that the relevant comparisons would be (5a) vs. (4a) where only an N400 effect is expected, and (5b) vs. (4b) where only a late positivity but not an N400 might appear.

Given these concerns, in our experiment, we follow the 2×2 design by Otten and colleagues, including a constraining and a non-constraining context and two possible con-

tinuations. But instead of looking for an interaction in the 2×2 design, we planned a nested comparison (Schad, Hohenstein, Vasishth, & Kliegl, 2018) within the two levels of the factor constraint: we compare (5a) vs. (4a), where we expect an N400 effect, and (5b) vs. (4b), where we expect a late positivity.

1.3 The current experiment

Our EEG experiment takes advantage of the fact that all German nouns have a gender (i.e., masculine, feminine, or neutral) and that the determiners that precede the nouns can show an unambiguous gender marking when the noun phrase has accusative case.⁴ This allows us to create stimuli such as (6)-(7), where the critical region is the determiner (in bold) preceding the hypothesized (un)predicted target noun. We preregistered the following design, and two models for the determiner region and the noun region at predetermined time windows and electrodes (see <https://osf.io/vdx8g/>). For the determiners, one model examines a potential N400 effect modulated by the gender probability on the determiners of (6a) and (7a), and another model focuses on a potential late positivity modulated by the constraint of the context on the determiners of (6b) and (7b). As we discussed in the introduction, both our models use a different approach than the one commonly used in the studies reviewed in Table 1 (and their replications). In these earlier studies, the context was kept constant (i.e., constraining) and the critical determiner and noun were manipulated, changing the cloze probability. In our study, the context changes (i.e., constraining vs. non-constraining context) and the target word is kept constant. We detail our preregistered models in section 2.5.1. In an exploratory analysis, we investigate a more traditional design that keeps the context constant (i.e., fix the context to constraining).

(6) Constraining context:

Die Großmutter möchte im Wald Pilze sammeln.
The grandmother wants in.the forest mushrooms gather.

‘The grandmother wants to pick mushrooms in the forest.’

a. High-probability continuation:

Sie nimmt **einen** großen Korb mit und macht sich auf den
She takes **a.masc.acc** big.masc.acc basket with and makes herself on the
Weg.
way

‘She takes a big basket and starts on her way.’

b. Low-probability continuation:

⁴This is strictly true only for “*ein*”, “*einen*”, “*eine*” (a.neut, a.masc, a.fem) since otherwise feminine determiners might be ambiguous with plural (of any gender). For example, “*meinen*” (my.masc.sg) and “*mein*” (my.neut.sg) are unambiguously gender marked, while “*meine*” is ambiguous between feminine singular and the plural of any gender. In order to avoid this ambiguity, we used “*ein/e/n*” for all conditions in items where the gender of one of the critical nouns was feminine (since “a” cannot be plural). We used possessive pronouns, such as “*mein/en*” only in cases where no condition had a feminine noun.

Sie nimmt **ein** großes Tuch mit und packt es in einen Korb.
 She takes **a.masc.acc** big.masc.acc cloth with and packs it in a.masc basket

‘She takes a big cloth and packs it in a basket.’

(7) Non-constraining context:

Die Großmutter möchte heute keine Pilze im Wald sammeln.
 The grandmother wants today no mushrooms in.the forest gather.

‘Today the grandmother doesn’t want to pick mushrooms in the forest.’

a. Low-probability continuation:

Sie nimmt **einen** großen Korb mit und bereitet ein
 She takes **a.masc.acc** big.masc.acc basket with and prepares a.neut
 Picknick vor.
 picnic for

‘She takes a big basket and prepares a picnic.’

b. Low-probability continuation:

Sie nimmt **ein** großes Tuch mit und packt es in den
 She takes **a.masc.acc** big.masc.acc cloth with and packs it in the
 Picknickkorb.
 picnic.basket

‘She takes a big cloth and packs it in the picnic basket.’

2 Methods

2.1 Participants

For the EEG experiment, we preregistered a sample size of 100 participants: 50 university students and 50 non-students. Due to time constraints, before the first submission of the manuscript, we only recruited 89 participants: 54 university students and 35 non-students. No power analysis to determine the required sample size was performed before the initial data collection. Instead, in order to maximize the chances to detect an effect, we aimed to collect data from a sample size much larger than in typical language ERP studies ($N=30$), given the time and resources available. Because of a bug in the presentation script, one of the four experimental lists was not presented. To balance the design, we recruited 31 more participants (27 students and 4 non-students) at a later stage. This means that in total, 120 participants completed the EEG experiment.

Informed consent from the participants was obtained before the session. The study was approved by Research Ethics Committee of the faculty of Psychology and Human Movement of the University of Hamburg.

Although psycholinguistic studies have mainly utilized the college “subject-pool” population, extending our research to the community-based sample is important because it is more representative of the population at large. We excluded non-native German speakers,

left-handed subjects, subjects with dreadlocks, subjects with self-reported neuronal abnormalities and/or language impairments, and participants of the related sentence completion tasks detailed below. All participants had normal or corrected-to-normal vision.

Students were recruited from the Vasishth Lab subject pool at the University of Potsdam. Non-university students were recruited using fliers placed on public transportation hubs and local retail, and by the referral of other recruited participants and of the university and administrative staff. We excluded non-university student participants below the age of 18 and above the age of 35 so that they match the age of the student population.

2.2 Materials

2.2.1 Sentence completion tasks. We first conducted a pilot sentence completion task to create the potential experimental items. Next, we conducted a second sentence completion task to establish the predictability of the gender of the target nouns before the gender-marked determiner was read for all the experimental items. Finally, we conducted a third sentence completion task to establish the predictability of the nouns conditional on their previous (gender agreeing) determiner and adjective. Participants of the sentence completion tasks were recruited from the same subject pool used for the EEG experiment; these participants had not taken part in the other tasks mentioned above. We detail the tasks below.

2.2.1.1 Pilot task. We originally created 400 potentially constraining and non-constraining contexts together with sentences truncated before a potential accusative determiner. Some of these items were based on the stimuli of experiment 1B of Otten and Van Berkum (2008). We presented these items in Latin-square design to 40 participants (so that each participants only saw the constraining or the non-constraining context version of each item) using Linger (<https://tedlab.mit.edu/~dr/Linger/>).

2.2.1.2 Sentence completion task for the gender probability and constraint. We refined the truncated sentences based on the results of the pilot task, and we discarded problematic items. We then built 295 potentially constraining and non-constraining contexts together with sentences truncated before a potential accusative determiner. These sentences were presented in Latin-square design to 65 native speaker participants using Linger.

Based on the results of this task, we selected 283 constraining and non-constraining contexts, and we built items similar to (6)–(7). We present a summary of the mean and 90% quantiles for the cloze probability of the genders in the stimuli and the best completion in Table 2. The best completion is the continuation that has the highest cloze probability for a given context. We show gender characteristics of the items in Table 3.

2.2.1.3 Sentence completion task for the noun probability and constraint. We ran an additional sentence completion task where we truncated the experimental items before the noun to calculate the probability of each noun conditional on the complete context. We presented these items in Latin-square design to 83 participants using Linger. In this way, we were able to calculate the predictability of the nouns conditional on the context that includes the gender marked determiner and adjective.⁵ These cloze probabilities were

⁵One sentence had a typo and was removed afterwards.

Table 2

Mean and 90% quantiles of the gender probability for each condition for sentences truncated before the determiner. The column called gender cloze probability refers to the cloze probability of the gender of the actual noun phrase that completed the sentence in the experimental items. The column called gender best completion shows the cloze probability for the most likely gender as continuation for a given context

Constraint level	Completion	Gender cloze probability	Gender best completion
Constraining	a	0.80, [0.57, 0.97]	0.80, [0.58, 0.97]
Constraining	b	0.08, [0.00, 0.21]	0.80, [0.58, 0.97]
Non-constraining	a	0.49, [0.21, 0.78]	0.56, [0.37, 0.78]
Non-constraining	b	0.20, [0.03, 0.41]	0.56, [0.37, 0.78]

Table 3

Proportions of genders of the determiners for each condition.

constraint	completion	fem	masc	neut
Constraining	a	0.34	0.40	0.26
Constraining	b	0.35	0.40	0.25
Non-constraining	a	0.34	0.40	0.26
Non-constraining	b	0.35	0.40	0.25

used for the analysis at the noun region.

Table 4

Mean and 90% quantiles of the noun probability for each condition for sentences truncated before the noun. The column called Noun cloze probability refers to the cloze probability of the actual noun that completed the sentence in the experimental items. The column called Noun best completion shows the cloze probability for the most likely noun continuation for a given context

Constraint level	Completion	Noun cloze probability	Noun best completion
Constraining	a	0.73, [0.35, 0.95]	0.78, [0.50, 0.95]
Constraining	b	0.14, [0.00, 0.47]	0.38, [0.16, 0.67]
Non-constraining	a	0.40, [0.05, 0.79]	0.49, [0.23, 0.80]
Non-constraining	b	0.17, [0.00, 0.50]	0.35, [0.15, 0.60]

2.3 Procedure

Participants were tested in a single session, seating approximately 60 cm from a 56 cm presentation screen. Before starting the EEG experiment, participants performed a stop signal task (Lappin & Eriksen, 1966; Logan & Cowan, 1984) that closely follows the design of Verbruggen, Logan, and Stevens (2008). The experiment was run using OpenSesame (Mathôt, Schreij, & Theeuwes, 2012) and its results will be part of a future exploratory analysis of the data.

Following the stop-signal task, participants performed an EEG experiment run in OpenSesame (Mathôt et al., 2012). Participants were instructed to read sentences appearing in serial visual presentation for comprehension and answer yes/no questions (that followed a quarter of the trials) by pressing “f” or “j” in a keyboard.

Following the electrode cap preparation, the experiment started with five practice trials. The stimuli of the experiment were presented using a Latin-square design. At the beginning of each trial, a fixation cross appeared for 500 ms followed by a blank screen for 500 ms. After that, the context was presented broken in short phrases (e.g., *Die Großmutter|möchte|heute|keine Pilze|im Wald|sammeln.*). Each phrase appeared in the center of the screen for $190 \text{ ms} \times \text{number of words} + 20 \text{ ms} \times \text{number of characters}$ with a minimum of 250 ms, followed by a blank screen for 250 ms. Then the continuation appeared word by word. Each word appeared in the center of the screen for $190 \text{ ms} + 20 \text{ ms} \times \text{number of characters}$ (with a minimum of 250 ms) followed by a blank screen for 250 ms. Although this means that target determiner presentation times differed for the different gender conditions, as we explain later, the preregistered comparisons examined the same determiner (determiner in completions (a) in one model, and determiners in completions (b) in another model).

2.4 EEG recording and data processing

2.4.1 Recording. EEG was recorded from 32 scalp sites by means of AgAgCl active electrodes mounted in an elastic electrode cap at the standard 10-20 system (Jasper, 1958) using an ANT Neuro amplifier manufactured by TMSi. As is standard in ERP research, eye movements and blinks were monitored with bipolar electrodes placed at the outer canthus of each eye (horizontal EOG) and above and below the right eye (vertical EOG). EEG and EOG were recorded with a sampling rate of 512 Hz and a low-pass filter of 138 Hz. Recordings were initially referenced to the left mastoid and re-referenced offline to the average of the left and right mastoid channels.

2.4.2 Preprocessing. The data was pre-processed with the *R* (R Core Team, 2018) package *eeguna* (Nicenboim, 2018). The EEG data was filtered using a zero-phase band-pass finite impulse response (FIR) filter with pass band-edge frequencies of 0.1 and 30 Hz. The width of the transition band for low and high edges was 0.10 and 7.50 Hz respectively. This filter was adapted from the default filters in the python package MNE (Gramfort et al., 2013, p. v 0.0.17.1). Eye movements were corrected using independent components analysis (ICA: Jung et al., 2000) with the deflation-based FastICA algorithm (Hyvärinen & Oja, 1997; Nordhausen, Ilmonen, Mandal, Oja, & Ollila, 2011; Miettinen, Nordhausen, Oja, & Taskinen, 2014). After this, we rejected segments containing a voltage difference of over $100 \mu V$ in a time window of 150 ms or containing a voltage step of over $50 \mu V/ms$. The corrected signal was segmented and then baseline-corrected relative to a 100 ms interval preceding the stimulus. The preprocessing of the data was done as preregistered, with two exceptions: we did not downsample the EEG data to 500 Hz, and, as suggested by a reviewer, we did not use a notch filter.

Data from five participants were excluded from the analysis due to technical problems with the recording. The average comprehension question accuracy of the participants who were included in the analysis was 90% (range 77-100). We lost data from 0.51% trials due

to technical problems, and we rejected 5% of the trials due to large differences in voltage. The data, code, and scripts of this study can be found in <https://osf.io/w85gc/>.

2.5 Statistical analysis

2.5.1 Preregistered data analysis. Taking our literature review as a starting point, the pre-registered analyses focus on two ERP effects that have been associated with predictions: (i) the so-called N400 effect, that is, a negativity with a centroparietal distribution in the electrode sites Cz, CP1, CP2, P3, Pz, P4, POz (DeLong, Quante, & Kutas, 2014) within a time window of 250–600 ms, and (ii) a late positivity with a frontal distribution in the electrode sites F7, F3, Fz, F4, F8, FC5, FC1, FC2, FC6, C3, C4 (DeLong et al., 2014) within a time window of 500–900 ms with respect to the onset of the critical region. Given the large variability on the time windows from the literature (see Table 1), we chose a rather wide time window for both effects. Notice that even though there is an overlap in the time windows, we expect the effects in different models (see also section 1.2).

For the data analysis, we define the ERP amplitudes as the average scalp potential over the pre-specified time windows and electrode sites *for each trial* (as was done in, for example, Frank, Otten, Galli, & Vigliocco, 2015; Nieuwland et al., 2018). This allows us to use a Bayesian linear mixed model of the single-trial data and to simultaneously model variance associated with each subject and with each item.

We use a Bayesian data analysis approach implemented in the probabilistic programming language *Stan* (Stan Development Team, 2018b) using the model wrapper package *brms* (Bürkner, 2017; Bürkner, 2018) in *R* (R Core Team, 2018).⁶ An important motivation for using the Bayesian approach is that it is easy to fit fully hierarchical models with fully specified variance-covariance matrices for all random effects, which provide the most conservative estimates of uncertainty (Schielzeth & Forstmeier, 2009). We use regularizing priors for all our parameters, except in the parameters of interest of the models used for the Bayes factor calculation; see below. Regularizing priors are minimally informative and have the objective of yielding stable inferences (Chung, Gelman, Rabe-Hesketh, Liu, & Dorie, 2013; Gelman, Jakulin, Pittau, & Su, 2008; Gelman, Simpson, & Betancourt, 2017). We fit the models with four chains and 3000 iterations, with 1000 of which were the burn-in or warm-up phase. In order to assess convergence, we verified that there were no divergent transitions, that the $\hat{R}$ s (the between- to within-chain variances) were close to one, that the

⁶The complete list of software that we used is the following: *R* (Version 3.6.3; R Core Team, 2018) and the *R*-packages *bayesplot* (Version 1.7.1; Gabry & Mahr, 2018), *bindrcpp* (Müller, 2018), *brms* (Version 2.12.0; Bürkner, 2017; Bürkner, 2018), *cleanNLP* (Arnold, 2017), *cowplot* (Version 1.0.0; Wilke, 2019), *dplyr* (Version 0.8.5; Wickham, François, Henry, & Müller, 2018), *eeguana* (Version 0.1.4.9000; Nicenboim, 2018), *ggplot2* (Version 3.3.0; Wickham, 2016), *ggridges* (Version 0.5.2; Wilke, 2018), *lme4* (Version 1.1.21; D. Bates, Mächler, Bolker, & Walker, 2015), *MASS* (Version 7.3.51.5; Venables & Ripley, 2002), *Matrix* (Version 1.2.18; D. Bates & Maechler, 2018), *osfr* (Version 0.2.8; Wolen & Hartgerink, 2019), *papaja* (Version 0.1.0.9942; Aust & Barth, 2018), *polspline* (Version 1.1.17; Kooperberg, 2018), *purrr* (Version 0.3.3; Henry & Wickham, 2018), *Rcpp* (Version 1.0.4; Eddelbuettel & François, 2011; Eddelbuettel & Balamuta, 2017), *readr* (Version 1.3.1; Wickham, Hester, & François, 2018), *readxl* (Version 1.3.1; Wickham & Bryan, 2018), *rmarkdown* (Version 2.1; Xie, Allaire, & Golemund, 2018), *rstan* (Version 2.19.3; Stan Development Team, 2018a), *StanHeaders* (Version 2.21.0.1; Stan Development Team, 2018c), *stringr* (Version 1.4.0; Wickham, 2018), *tidyr* (Version 1.0.2; Wickham & Henry, 2018), and *udpipe* (Version 0.8.3; Wijffels, 2019).

number of effective sample size were at least 10% of the number of post-warmup samples (i.e. at least 300 out of 3000), and we visually inspected the chains. All these convergence checks seemed adequate for the models reported in this paper, but some variance parameters had relatively low effective sample size (the lowest one was 488); see the complete models in <https://osf.io/w85gc/>. Nicenboim and Vasishth (2016) discuss the Bayesian approach in detail in the context of linguistic and psycholinguistic research.

We report mean estimates and 95% quantile-based Bayesian credible intervals. A 95% Bayesian credible interval has the following interpretation: it is an interval containing the true value with 95% probability (see, for example, Jaynes & Kempthorne, 1976; Morey, Hoekstra, Rouder, Lee, & Wagenmakers, 2016).

2.5.2 Additional (non preregistered) inference criteria. In the pre-registration, we intended to carry out inference using the posterior probability of the estimate being positive given the data ($P(\beta > 0 | data)$). However, we realized later that this quantity can be confused with the evidence against $\beta = 0$. Given that the research question is often framed in terms of the presence or absence of an effect, we use the Bayes factor to quantify the extent to which the data support a null model (M0), which assumes that a particular parameter of interest, β , is 0, over an alternative model (M1), which assumes that the parameter is not zero. For example, if we want to test the effect of predictability, our alternative model would contain the parameter representing the effect of predictability, and the null model would not have this parameter.

The Bayes factor quantifies the weight of the evidence in favor of the null model M0 relative to some alternative model. The Bayes factor is similar to the likelihood ratio test, with the difference that the likelihood is integrated over the parameter space defined by the priors, instead of being maximized. In Bayes factors computations, the priors can be very influential; vague or uninformative priors on the parameter of interest can strongly bias the Bayes factor in favor of the null model (Lee & Wagenmakers, 2014). Consequently, for each model, we carry out a sensitivity analysis, showing a range of Bayes factors keeping the model specification constant but defining increasingly informative priors for the parameter of interest (e.g., the effect of predictability on the amplitude of the N400). Defining increasingly informative priors amounts to defining smaller standard deviations for the relevant prior. The different Bayes factors BF_{01} represent the evidence against the alternative model (or in favor of the null), contingent on the magnitude of the effect assumed a priori via the priors. In the Bayes factor plots, we indicate what we consider well-calibrated priors, namely, priors that are agnostic regarding the direction of the effect, but relatively informative regarding the size of the possible effect. For these well-calibrated priors we used $Normal(0, SD)$ with standard deviation SD between 0.10–0.50. These standard deviations are based on the N400 estimates at the noun region in Nieuwland et al. (2018).

Bayes factor values for the comparison BF_{01} that are larger than one favor the null model, and values below one favor the alternative model. Values close to one indicate that the evidence is inconclusive (see Wagenmakers, Lodewyckx, Kuriyal, & Grasman, 2010). In order to use the Bayes factor to make decisions about the evidence for or against a particular hypothesis, Jeffreys (1939/1998) has suggested the following approximate criteria; these should not be regarded as hard and fast rules with clear boundaries. Here, we assume that

we are comparing two models, labeled M0 and M1.

- $BF_{01} > 100$: Decisive evidence for Model M0.
- $BF_{01} = 32 - 100$: Very strong evidence for Model M0.
- $BF_{01} = 10 - 32$: Strong evidence for Model M0.
- $BF_{01} = 3 - 10$: Substantial evidence for Model M0.
- $BF_{01} = 2 - 3$: Not worth more than a bare mention.

To calculate the Bayes factor, we used bridge sampling (Bennett, 1976; Meng & Wong, 1996; Gronau et al., 2017) with eight chains and 30000 iterations, 1000 of which were the warm-up phase.

2.5.2.1 Effect of the predictions at the determiner. We fit a linear mixed model to the average amplitude of the N400 spatiotemporal window with respect to the determiner using the predictor *gender predictability* and the nuisance predictor *gender*. We used the data from the same determiners pooled from items with a constraining context and a high probability gender, and from items with a non-constraining context and a low probability gender: “*einen*” in examples (6a) and (7a). *Gender predictability* (**pred** in the model) is the scaled and smoothed probability transformed to $\log_2$ (log to the base two) of gender cloze probability. We use additive smoothing (also called Laplace or Lidstone smoothing; Lidstone, 1920; Chen & Goodman, 1999) with pseudocounts set to one. This means that the smoothed probability is calculated as the number of responses with a given gender plus one divided by the total number of responses plus two. The effect of predictability is then assumed to be continuous and non-linear. The estimate represents the change in the amplitude for a change in one point in the non-linear scale: To make it concrete, it can be understood as the average change in amplitude when one compares a determiner with gender probability of 0.74 with a determiner with a gender probability of 0.26, but also when one compares the change between 0.26 and 0.09. Notice that we use gender probability rather than the determiner probability based on the idea that the pre-activation might be at the feature level and not lexical (Rabovsky & McRae, 2014). However, we cannot disentangle these two possibilities because of the high correlation between determiner cloze probability and gender probability: 0.77 (but see Fleur et al., 2019, for an experimental design where word probability and gender probability are disentangled). The motivation for using log-transformed cloze probability is two-fold: First, there is evidence that at least for reading time experiments, the use of log-transformed probability shows a better fit to the data than other transformations, including no transformation (Smith & Levy, 2013). Second, log-transformed probabilities have been shown to be useful accounting for reading time data (among others: Levy, 2008; Hale, 2001) and EEG/ERP data (Frank et al., 2015; Delaney-Busch, Morgan, Lau, & Kuperberg, 2017). Indeed, model comparison using the EEG data of the noun region of Nieuwland et al. (2018) shows support for the use of log-transformed cloze-probabilities, see Appendix A. *Gender* is a Helmert coded predictor (Schad et al., 2018) that compares *masculine* vs. *feminine*, and *masculine* together with *feminine* vs. *neuter* (**gender**). We did not have specific predictions for the effect of the gender predictor; we chose Helmert coding here because it gives us orthogonal comparisons. For this model and the following ones, we specify the following weakly informative (regularizing) priors in **brms**:

```
reg_priors <- c(prior(normal(0, 10), class = Intercept),
               prior(normal(0, 1), class = b),
               prior(normal(0, 10), class = sd),
               prior(normal(0,10), class = sigma),
               prior(lkj(2), class = cor))
```

We specify the model in `brms` as follows:

```
brm(A_N400 ~ pred + gender +
     (1 | item) + (pred + gender | subj),
    prior = reg_priors,
    data_a, ...)
```

In the code above, `...` refers to some other parameters that are needed for sampling.

2.5.2.2 Effect of the disconfirmation of strong predictions at the determiner. We fit a linear mixed model similar to the previous one to the average amplitude of the late positivity with respect to the determiner. We use the predictor *gender context constraint*, (`constr`) and the nuisance predictors *gender predictability* (`pred`) and *gender*. We used the data from the same low-probability-gender determiners pooled from items with a constraining context and from items with a non-constraining context: “*ein*” in examples (6b) and (7b). *Gender context constraint* (`constr`) is operationalized as the best gender completion, that is the highest cloze probability of a gender of a given context, and it is calculated as the highest number of responses for a given gender divided by the total number of responses and then scaled. In contrast to predictability, and since there are no data in favor (or against) log-transforming best completion values, we left them in the raw scale. We specify this model in `brms` as follows:

```
brm(A_PNP ~ constr + pred + gender +
     (1 | item) + (constr + pred + gender | subj),
    prior = reg_priors,
    data_b, ...)
```

2.5.2.3 Effect of the predictions at the noun. We fit a linear mixed model similar to the one presented in 2.5.2.1 using noun cloze probability instead of gender probabilities. The linear mixed model is fit to the average amplitude of the N400 spatiotemporal window with respect to the target noun using the predictor *noun predictability* (`pred_noun`), which is the scaled, smoothed probability of the noun cloze transformed to log to the base two, taking into account the context that includes the gender information. We used the data from the same nouns pooled from conditions (a): “*Korb*”, “basket”, in examples (6a) and (7a). We specify this model in `brms` as follows:

```
brm(A_N400 ~ pred_noun +
     (1 | item) + (pred_noun | subj),
    prior = reg_priors,
    data_a_noun, ...)
```

2.5.2.4 Effect of the disconfirmation of “strong” predictions at the noun. For completeness, we present here one model that was preregistered but that might be problematic because of the restricted range of noun constraint values in our stimuli; see Table 4. This is the case, because noun cloze probability and best completions are updated

given the entire context where the noun appears, which includes gender information. For this model, we fit a linear mixed model to the average amplitude of the late positivity with respect to the target noun using the predictors *context constraint* and *noun predictability*. We used the data from the same nouns pooled from conditions (b): “*Tuch*”, “cloth”, in examples (6b) and (7b). *Context constraint* (`constr_noun`) is operationalized as the best noun completion, that is the highest cloze probability of a noun of a given context, and it is calculated as the highest number of responses for a given noun divided by the total number of responses and then scaled. We specify this model in `brms` as follows:

```
brm(A_PNP ~ constr_noun + pred_noun +
      (1 | item) + (constr_noun + pred_noun | subj),
  prior = reg_priors,
  data_b_noun, ...)
```

3 Results

3.1 Results of the preregistered analysis at the determiner

Results were inconclusive for the preregistered models at the determiner region. For the determiners pooled from conditions (a) of items with constraining context (i.e., high probability gender) and of items with non-constraining contexts (i.e., low probability gender), the estimate for the effect of predictability on the N400 spatiotemporal window was $\hat{\beta} = 0.17\mu V$, 95% CrI = $[-0.17, 0.51]$; the signal was more negative for more unexpected genders, although the estimate has a large degree of uncertainty. For determiners pooled from conditions (b) of items with constraining context (i.e., low probability gender) and of items with non-constraining contexts (i.e., low probability gender), the estimate for the effect of the constraint of the context on the late positivity was $\hat{\beta} = 0.03\mu V$, 95% CrI = $[-0.17, 0.23]$; the signal was more positive when the context was more constraining (controlling for gender probability), although the estimate has again a large degree of uncertainty. Analyses using the Bayes factor show that for well-calibrated priors ($Normal(0, SD)$ with SD between 0.10–0.50) that represent realistic effect sizes, the evidence is between inconclusive and weakly supporting the null; see Figure 1. By-condition averaged ERP waveforms on determiners for both preregistered models are shown in Figures 2, 3 and in the Appendix C (Figures C1 and C2); scalp topographies are shown in Figure 4.

3.2 Results of the preregistered analysis at the noun

Results were also inconclusive for the preregistered models at the noun region. First, there is no clear N400 effect of predictability at the noun (6a-7a): $\hat{\beta} = 0.05\mu V$, 95% CrI = $[-0.2, 0.3]$; the signal was, as expected, more negative for more unexpected nouns, although the estimate has again a large degree of uncertainty. Analyses using the Bayes factor show that for well-calibrated priors, the evidence is between inconclusive and weakly supporting the null; see Figure 5. By-condition averaged ERP waveforms on nouns following a constraining context are shown in Figure 6 and in the Appendix C (Figure C3); the scalp topography of the difference between conditions is shown in Figure 7(A). We speculate that this could be due to the fact that the *non-constraining* context was in fact somewhat

Figure 1. Figure (A) shows the posterior distributions of the effect of predictability on the N400 spatiotemporal window for the same determiners pooled from items with a constraining context and a high probability gender, and from items with a non-constraining context and a low probability gender; see examples (6a) and (7a). Figure (C) shows the posterior distributions of the effect of the constraint of the context for the late positivity spatiotemporal window for the same determiners with low probability gender pooled from items with a constraining context and a non-constraining context; see examples (6b) and (7b). The horizontal bars show 95% credible intervals and the dots show the posterior means. Figures (B) and (D) show the Bayes factors in favor of the null (BF_{01}) for different standard deviations (SD) of the normally distributed prior centered at zero for the effect of predictability and constraint respectively. The grey area indicates well-calibrated priors: $Normal(0, SD)$ with SD between 0.10–0.50. A Bayes factor below one indicates evidence against the null.

Figure 2. ERP for the grand average across N400 spatial window (Cz, CP1, CP2, P3, Pz, P4, POz) elicited by the presentation of the same determiners pooled from items with a constraining context and a high probability gender, and from items with a non-constraining context and a low probability gender; see examples (6a) and (7a). The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).

constraining after reading the determiner and adjective. Both types of contexts use similar words and, for most experimental items, the main difference between the contexts is that either an event happens or does not happen (e.g., “*the grandmother wanting to pick up mushrooms in the forest*” vs. “*the grandmother not wanting to pick up mushrooms in the forest*”), or the temporal order of events (e.g., “*someone wanting to buy sport shoes*” vs. “*someone having already bought sport shoes*”). For this reason, both type of contexts are likely to be semantically associated to the critical noun. Kuperberg (2016) argues that in the presence of multiple competing cues, those combinations that provide more reliable evidence for a prediction will outweigh other combination of cues, including the full set of contextual cues provided by the semantic and syntactic structure. This entails that comprehenders may temporarily ignore the syntactic structure of the non-constraining (or somewhat constraining) context reducing the mismatch of the low cloze probability noun that follows them. In fact, there is evidence that under some configurations, the N400 amplitudes only depend on semantic associations between the target word and the context without regard to the full

Figure 3. ERP for the grand average across late positivity spatial window (F7, F3, Fz, F4, F8, FC5, FC1, FC2, FC6, C3, C4) elicited by the presentation of the same determiners pooled from items with a constraining context and a high probability gender, and from items with a non-constraining context and a low probability gender; see examples (6b) and (7b). The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).

syntactic and semantic structure (e.g., in negation: Fischler, Bloom, Childers, Roucos, & Perry Jr., 1983; in atypical thematic roles: Chow, Smith, Lau, & Phillips, 2016), and that negative quantifiers (such as *few*) pose a particular challenge to incremental comprehension in plausible but low-cloze continuations (Nieuwland, 2016).

Second, results from the linear mixed model investigating the effect of constraint at the noun region for the late positivity show puzzling results: The model shows an effect of constraint in an unexpected direction, as the constraint of the context increases the average signal is more negative: $\hat{\beta} = -0.27\mu V$, 95% CrI = $[-0.53, 0.0032]$. By-condition averaged ERP waveforms on nouns following a constraining context are shown in Figure 8 and in the Appendix C (C4); the scalp topography of the difference between conditions is shown in Figure 7(b). Analysis with the Bayes Factor shows weak to no support against the null model; see Figures 5(C) and (D). In addition, the more predictable the presented noun, the more positive the average signal: $\hat{\beta} = 0.089\mu V$, 95% CrI = $[-0.17, 0.35]$. As we mentioned

Figure 4. Figure (A) shows a topographic plot with the difference between ERP amplitudes in the time-window of 250–600 ms following the presentation of the same determiners pooled from items with a constraining context and a high probability gender, and from items with a non-constraining context and a low probability gender; see examples (6a) and (7a). Figure (B) shows a topographic plot with the difference between ERP amplitudes in the time-window of 500–900 ms following the presentation of the same determiners with low probability gender pooled from items with a constraining context and a non-constraining context; see examples (6b) and (7b).

before, one possibility is that this manipulation did not work as we expected because of the restricted range of noun constraint values in our stimuli; see Table 4. Once the gender of the upcoming noun is known, the predictability of potential upcoming nouns changes, and that means that the stimuli were not built appropriately. We avoid overinterpreting the effect of the constraint of the context at the nouns.

4 Exploratory analysis

While the preregistered analysis at the determiner showed no evidence for an effect, it could be the case that the effect of the manipulation was not strong enough. One reason to suspect this is that the *non-constraining contexts* were at least slightly constraining and their completions not entirely unexpected. Indeed, the difference between the mean cloze probabilities of the gender of the determiners in conditions (a) following a constraining context vs. the ones in condition (a) following a non-constraining context is 0.31. In contrast, the difference between the mean cloze probability of the gender of the determiner following constraining contexts in conditions (a) vs. (b) is 0.72. While there is a lot of variability among the items, and the model included cloze probability as a continuous predictor (and not as a factorial comparison), it might be the case that a larger range of cloze probabilities would show an effect. Another potential issue is that to keep identical target regions in each item, we rely on varying contexts. This means that the brain activity before the determiner is not identical –whether or not there is prediction– and this might increase the noise in the signal.

Pre-registered models for the noun region

Posterior distribution

Bayes factors

Figure 5. Figure (A) shows the posterior distributions of the effect of predictability on the N400 spatiotemporal window for the same nouns pooled from items with a constraining context and a high probability noun, and from items with a non-constraining context and a low probability noun; see examples 6a and 7a. Figure (C) shows the posterior distributions of the effect of the constraint of the context for the late positivity spatiotemporal window for the same nouns with low probability pooled from items with a constraining context and a non-constraining context; see examples 6b and 7b. The horizontal bars show 95% credible intervals and the dots show the posterior means. Figures (B) and (D) show the Bayes factors in favor of the null (BF01) for different standard deviations (SD) of the normally distributed prior centered at zero for the effect of predictability and constraint respectively. The grey area indicates well-calibrated priors: $Normal(0, SD)$ with SD between 0.10–0.50. A Bayes factor below one indicates evidence against the null.

Figure 6. ERP for the grand average across N400 spatial window (Cz, CP1, CP2, P3, Pz, P4, POz) elicited by the presentation of the same nouns pooled from items with a constraining context and a high probability gender, and from items with a non-constraining context and a low probability gender; see examples (6a) and (7a). The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).

Given that a comparison holding the (constraining) context fixed and varying the completion is the most common in the literature and the one present in all the studies of Table 1, we also explored the effect of predictability under a constraining context using gender probability as a predictor (and gender as a nuisance predictor as before) on the N400 spatiotemporal window. While we will be comparing different determiners in this model, some of the idiosyncratic differences between determiners should be captured by the **gender** predictor and its interaction with predictability.

We specify this model in `brms` as follows:

```
brm(A_N400 ~ pred * gender +
    (1 | item) + (pred * gender | subj),
    prior = reg_priors,
    data_Cab, ...)
```

Figure 7. Figure (A) shows a topographic plot with the difference between ERP amplitudes in the time-window of 250–600 ms following the presentation of the same critical nouns pooled from items with a constraining context and a high probability cloze, and from items with a non-constraining context and a low probability cloze; see examples (6a) and (7a). Figure (B) shows a topographic plot with the difference between ERP amplitudes in the time-window of 500–900 ms following the presentation of the same critical nouns with low probability cloze pooled from items with a constraining context and a non-constraining context; see examples (6b) and (7b).

Another possibility is to treat each context as an item and explore the effect of predictability for our entire experimental stimuli. This would be the same model as before, but replacing `item` for `context` and run on the entire data for the determiner region:

```
brm(A_N400 ~ pred * gender +
      (1 | context) + (pred * gender | subj),
  prior = reg_priors,
  data, ...)
```

Furthermore, as a sanity check that our experiment *can* elicit an N400 effect, we fit another model using only data from the target nouns appearing in the constraining context. We specify the model in `brms` as follows:

```
brm(A_N400 ~ pred_noun +
      (1 | item) + (pred_noun | subj),
  prior = reg_priors,
  data_noun_Cab, ...)
```

We included two more exploratory analysis focusing on the adjective region and focusing on the determiner but replacing gender probability with determiner probability in the Appendix B.

Figure 8. ERP for the grand average across late positivity spatial window (F7, F3, Fz, F4, F8, FC5, FC1, FC2, FC6, C3, C4) elicited by the presentation of the same nouns pooled from items with a constraining context and a high probability gender, and from items with a non-constraining context and a low probability gender; see examples (6b) and (7b). The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).

5 Results of the exploratory analysis

At the determiner region, the results were still inconclusive for the exploratory models. The estimate for the effect of predictability on the N400 spatiotemporal window for the determiners pooled from conditions (a) and (b) of items with constraining context was $\hat{\beta} = 0.11\mu V$, 95% CrI = $[-0.0069, 0.23]$, and the effect of predictability on the N400 spatiotemporal window for the determiners pooled from all the conditions treating context as item was $\hat{\beta} = 0.12\mu V$, 95% CrI = $[0.015, 0.22]$; the signal was more negative for more unexpected genders, although the estimate has a large degree of uncertainty in both cases. Analysis using the Bayes factor for both models shows that the evidence supporting the null is inconclusive and slightly favors the null for a priori larger effect sizes (Figures 9A and C); notice that this is the case even when the 95% credible interval of the model that pools all the items does not overlap with zero (see Rouder, Haaf, & Vandekerckhove, 2018, for a discussion on the difference between estimation and Bayes factor approaches). By-condition

Exploratory models for the determiner region

Posterior distribution

Bayes factors

Figure 9. Figure (A) shows the posterior distributions of the effect of predictability on the N400 spatiotemporal window for the determiners with high and low probability genders from items with a constraining context; see examples 6a and 6b. Figure (C) shows the posterior distributions of the effect of predictability on the N400 spatiotemporal window for determiners of the four conditions pooled together. The horizontal bars show 95% credible intervals and the dots show the posterior means. Figures (B) and (D) show the Bayes factors in favor of the null (BF01) for different standard deviations (SD) of the normally distributed prior centered at zero for the effect of predictability for both models. The grey area indicates well-calibrated priors: $Normal(0, SD)$ with SD between 0.10–0.50. A Bayes factor below one indicates evidence against the null.

averaged ERP waveforms on determiners following a constraining context are shown in Figures C5 and 10, and the scalp topography of the difference between conditions is shown in Figure 11.

By contrast, at the noun region there is a clear N400 effect: $\hat{\beta} = 0.57\mu V$, 95% CrI = [0.42, 0.72]. Analyses using the Bayes factor show that for well-calibrated priors that represent realistic effect sizes, there is very strong evidence for an effect; see Figure 12. By-condition averaged ERP waveforms on nouns following a constraining context are shown in Figures C6 and 13, and the scalp topography of the difference between conditions is shown in Figure 14.

Figure 10. ERP for the grand average across N400 spatial window (Cz, CP1, CP2, P3, Pz, P4, POz) elicited by the presentation of the critical determiners with high and low probability genders from items with a constraining context; see examples (6a) and (6b). The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).

6 Discussion

The results from our preregistered and exploratory models show no clear support for the access view that posits that before a word is read, its context triggers the pre-activation of semantic features (and maybe phonological and grammatical features as well). We found

Figure 11. The topographic plot shows the difference between ERP amplitudes in the time-window 250–600 ms following the presentation of the determiners with high and low probability genders from items with a constraining context; see examples (6a) and (6b).

no evidence for an N400 effect or a late positivity due to differences in predictability or the level of constraint of the context respectively. Thus, from our experiment alone, we should conclude that there is no evidence for the pre-activation of features of expected nouns, measured on gender-marked articles in German. Given that results from recent studies are either inconclusive (Kochari & Flecken, 2019; Ito et al., 2017a), or weakly supporting the null (Nieuwland et al., 2018), *if* there is such pre-activation, its effect is *too* small to be detected reliably by a single experiment, even if the sample size is very large by EEG standards (e.g., 334 participants in Nieuwland et al., 2018). Furthermore, a self-paced reading study with a relatively large-sample size (120 participants and 60 items) has shown that unexpected gender-marking in determiners in Spanish do not lead to a slowdown in comparison to expected ones, and thus *if* pre-activation of gender marking occurs it may have no measurable impact in reading times (Guerra, Nicenboim, & Helo, 2018).

However, as in recent replication attempts (Nieuwland et al., 2018; Ito et al., 2017a; Kochari & Flecken, 2019; Kochari & Flecken, 2019), the patterns we found are in the expected direction in terms of polarity: the more unexpected the gender of the determiner, the larger the negativity of the average amplitude from the typical electrodes and time window of the N400 effect. It is therefore unclear from a qualitative literature review alone how strong the evidence is in favor (or against) pre-activation.

In order to quantitatively synthesize the results of the aforementioned studies and the current one, we conducted a Bayesian random effects meta-analysis. A random effects meta-analysis assumes that there is a unique underlying effect (i.e., the effect of pre-activation on the N400 spatiotemporal window) to be estimated from all the studies, while at the same time it allows for the estimates in the individual studies to vary because of differences between the specific studies and because of sampling variability.

It would also be possible to conduct a frequentist random effects meta-analysis instead

Exploratory model for the noun region

Posterior distribution

Bayes factors

Figure 12. Figure (A) shows the posterior distribution of the effect of predictability on the N400 spatiotemporal window for the nouns with high and low probability genders from items with a constraining context; see examples 6a and 6b. The horizontal bar show the 95% credible interval and the dot show the posterior mean. Figure (B) shows the Bayes factors in favor of the null (BF01) for standard deviations (SD) of the normally distributed prior centered at zero for the effect of predictability. The grey area indicates well-calibrated priors: $Normal(0, SD)$ with SD between 0.10–0.50. A Bayes factor below one indicates evidence against the null.

of a Bayesian one to synthesize the evidence of multiple studies. A frequentist meta-analysis would certainly be an improvement over examining whether the p -value of each study is below .05. However, the end result of a frequentist meta-analysis is a maximum likelihood estimate with its confidence interval and/or a p -value. Reliance on p -values has been criticized on several grounds: It has been argued that p -values promote the dichotomization of evidence into the artificial categories of “statistically significant” and “not statistically significant” which encourages erroneous scientific reasoning, and that, in addition, p -values are usually misinterpreted (among other Cumming, 2012; Vasishth & Nicenboim, 2016; McShane, Gal, Gelman, Robert, & Tackett, 2019). Even though some frequentist approaches do discourage examining whether p -values are below .05 or not (e.g., the new statistics of Cumming, 2012; the continuously cumulating meta-analytic of Braver, Thoemmes, & Rosenthal, 2014) and instead advocate employing only continuous descriptive criteria such as effect size estimates and their confidence intervals, confidence intervals are also often misinterpreted (Hoekstra, Morey, Rouder, & Wagenmakers, 2014). Morey et al. (2016) identify several common misconceptions regarding confidence intervals. In particular, the following two misconceptions are properties of *Bayesian* credible intervals rather than of *fre-*

Figure 13. ERP for the grand average across N400 spatial window (Cz, CP1, CP2, P3, Pz, P4, POz) elicited by the presentation of the critical nouns with high and low probability genders from items with a constraining context; see examples (6a) and (6b). The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).

quentist confidence intervals: (i) confidence intervals cannot be used to assess the certainty with which a parameter is in a particular range, and (ii) there is no necessary connection between the precision of an estimate and the size of a confidence interval. In contrast, a Bayesian meta-analysis inherits all the desirable characteristics of the Bayesian framework, such as the possibility of yielding (i) credible intervals to quantify the uncertainty (given the data and model) over a parameter and (ii) Bayes factors to quantify the weight of the evidence in favor of the null model relative to a specific alternative model.

7 Meta-analysis

We conducted a Bayesian random effects meta-analysis that allows for heterogeneity in the different studies. That is, each individual study is adjusted based on its estimated effect and its estimated standard error (see also Vasishth, Chen, Li, & Guo, 2013; Nicenboim, Roettger, & Vasishth, 2018; Jäger, Engelmann, & Vasishth, 2017). Given that there might be differences between studies manipulating the phonological features of the determiner (Nieuwland et al., 2018; Ito et al., 2017a) and gender marking (Kochari & Flecken,

Figure 14. The topographic plot shows the difference between ERP amplitudes in the time-window 250–600 ms following the presentation of the critical with high and low cloze probability from items with a constraining context; see examples (6a) and (6b).

2019, and the current study), we included a sum-coded predictor that indicates the type of manipulation (1 for phonological features and -1 for gender marking).

To anticipate our conclusion, even though the studies, analyzed individually, were “inconclusive”, as we show below, the meta-analysis does present evidence in favor of an effect consistent with the pre-activation account, albeit a very small one.

7.1 Estimates of the individual studies

For the meta-analysis, we included the publicly available data from Nieuwland et al. (2018), Ito et al. (2017a), and Kochari and Flecken (2019), and a subset of our own data (only data from items with constraining contexts). We reanalyzed the publicly available data using Bayesian linear mixed models that were as similar as possible: For all the studies, the predictor labeled predictability is smoothed, transformed to log base two, and then scaled using the same mean and standard deviation as in our data.⁷ This means that for all the studies the estimate of the effect of predictability represents the change in the average EEG amplitude for a change in one point in the same logarithmic scale. It is, for example, the average change in amplitude when one compares a determiner with cloze probability of 0.74 with a determiner with a cloze of 0.26. In the case of Nieuwland et al. (2018) and Ito et al. (2017a), predictability refers to the log cloze probability of the article form (“a” vs. “an”). In contrast, in the case of Kochari and Flecken (2019) (and our own study), predictability refers to the log cloze probability of the gender of the upcoming noun. For the cloze task data of Kochari and Flecken (2019), we determined the gender of the noun phrases using *frog* (Desmet & Hoste, 2014) and *udpipe* (Wijffels, 2019). Since the studies included had

⁷For each item i the predictor is calculated as follows: $pred_i = (\log_2((n_i + 1)/(N_i + 2)) - M_N)/SD_N$; where n_i is the number of answers matching a given determiner for the item i in the cloze task, N is the total number of answers given for item i , and M_N and SD_N are the mean and standard deviation of the log cloze probabilities in our study.

only two conditions with a constraining context, for the estimate of our own study, we also use data only from items with constraining contexts, that is, the estimate from our exploratory analysis ($\hat{\beta} = 0.11\mu V$, 95% CrI = $[-0.0069, 0.23]$).

For Ito et al. (2017a), we used their data from native speakers averaged across the central parietal electrodes (C1, Cz, C2, CP1, CPz, CP2, P1, Pz, P2, PO3, POz, and PO4) in a time window of 250–400 ms after the determiner for each trial (<https://osf.io/5m4br>). We fit the data of their two experiments with a single model, where we assume heterogeneity of variance across the experiments that were run with different modes of presentation (i.e., both experiments can have different variances). The model is specified in `brms` as follows:

```
var_priors <- c(prior(normal(0, 10), class = Intercept, dpar = sigma),
               prior(normal(0, 1), class = b, dpar = sigma))

brm(bf(A_N400 ~ (pred * det) * experiment +
              (1 | item) + (pred * det | subj),
              sigma ~ experiment),
    prior = c(reg_priors, var_priors) ...)
```

For Nieuwland et al. (2018), we used their data, which were averaged across the central parietal electrodes (C1, Cz, C2, CP1, CPz, CP2, P1, Pz, P2, PO3, POz, and PO4) in a time window of 200–500 ms after the determiner (<https://osf.io/e624a>). We fit all the different labs that conducted the same study with a single model where we assume heterogeneity of variance across the different labs (i.e., different labs can have different variances). We specify the model in `brms` as follows:

```
brm(bf(A_N400 ~ (pred * det) * experiment +
              (1 | item) + (pred * det | subj),
              sigma ~ experiment),
    prior = c(reg_priors, var_priors), ...)
```

For Kochari and Flecken (2019), we used their pre-processed data (osf.io/ptqwm), and averaged the signal from the central parietal electrodes (Cz, CP1, CP2, P3, Pz, P4, and POz) in the same time window as in our study, 250–600 ms after the determiner. We specify the model in `brms` as follows:

```
brm(A_N400 ~ pred * det +
      (1 | item) + (pred * det | subj),
    prior = reg_priors, ...)
```

7.2 Results of the meta-analysis

The meta-analysis shows a relatively clear effect of predictability: $\hat{\beta} = 0.11\mu V$, 95% CrI = $[0.05, 0.16]$, and no clear suggestion of a difference between studies that manipulated gender and phonological features: $\hat{\beta} = 0.0088\mu V$, 95% CrI = $[-0.046, 0.064]$. Figure 15 shows the 95% credible intervals of the meta-analytic estimate (on the top), and of the non-pooled and partially-pooled estimates of the original studies, that is, the 95% credible intervals estimated either without taking into account the other studies or as part of the random-effects meta-analysis. The Bayes factor analysis shows that, for well-calibrated

Figure 15. Forest plot of the estimates of the effect of predictability at a determiner prior to a noun; a positive effect is consistent with the pre-activation account (i.e., a more expected determiner leads to a reduction of the N400 effect). Horizontal lines represent 95% credible intervals. The cross at the top of the plot represents the mean of the meta-analytic estimate, blue circles are the estimates reconstructed from the individual studies, and black triangles are the shrinkage estimates of the individual studies delivered by the random-effects meta-analysis (i.e., partially pooled estimates).

(i.e., informative) priors, there is evidence for an overall effect in the expected direction: the more unexpected the determiner, the larger the average amplitude in the N400, for the given time window and electrodes. As Figure 16 shows, if we assume a priori that the effect is relatively small (but not too small, since the priors with a SD lower than 0.1 show between a weak support for the alternative and some support for the null), the data furnishes strong evidence for an effect. Conversely, if we assume a priori that the effect could be large, the evidence is weaker.

8 General Discussion

We conducted a study investigating pre-activation capitalizing on the fact that all German nouns have gender marking (masculine, feminine, or neutral) and that their preceding determiners can show an unambiguous gender marking when the noun phrase has accusative case. We used a relatively large sample size ($N = 120$). However, both our preregistered and our exploratory analyses failed to yield conclusive evidence for an effect

Figure 16. Figure (A) shows the posterior distribution of the meta-analytic estimate of the effect of predictability at a determiner prior to a noun; a positive effect is consistent with the pre-activation account (more expected determiners lead to a reduction of the N400 effect). Figure (B) shows the Bayes factors in favor of the null (BF01) for different standard deviations (SD) of normally distributed priors centered in zero for the effect of predictability. The grey area indicates well-calibrated priors: $Normal(0, SD)$ with SD between 0.10–0.50. A Bayes factor below one indicates evidence against the null.

of pre-activation (or for the absence of it). However, the sign of the effect is in the expected direction: the more unexpected the gender of the determiner, the larger the negativity on average amplitude of the N400 spatiotemporal window. This seems to be also the case for the recent replication attempts (Nieuwland et al., 2018; Ito et al., 2017a; Kochari & Flecken, 2019).

We see two possible explanations for the small N400 effect at the determiner:

1. Under the view that the N400 is indexing the prediction error for only *semantic* features as suggested by Rabovsky and McRae (2014); Kuperberg (2016), we would interpret the N400 effect at the determiner as follows. Prior to the determiner, the context pre-activates semantic features to different levels and, incidentally, the noun with the highest cloze probability (e.g., *basket.masc*) is the word with features that overlap the most with the highest pre-activated features in memory. Seeing a determiner with a grammatical gender (or a phonological feature in English) that is incompatible (e.g., *a.fem*) with the word with the highest cloze probability changes the context for the upcoming noun, and thus provokes an update of the activation of semantic features in memory. The difference between the previous levels of activation

and the updated ones would correspond to the N400 effect.

2. Another alternative that may explain the very small effect is that the prediction error is not limited to semantic features but extends to grammatical and phonological features. A very small effect at the determiner may be indicating that the mismatch at the grammatical gender (or phonological) features may not be affecting other semantic features that are also activated: For example, an unexpected gender in German at the determiner can still lead to a different noun with a similar meaning (i.e., similar semantic features) as the most predictable word (e.g., “*Couch*” and “*Sofa*” in German have different genders). Even though the N400 has been linked to access to *semantic* memory, there is evidence of purely phonological features playing a role: words that do not rhyme with a prime word produce greater negativity in the N400 time range than rhyming words (Praamstra & Stegeman, 1993; Perrin & García-Larrea, 2003). Furthermore, although the N400 has been generally characterized as purely semantic, there is evidence that the mismatch in grammatical features (such as argument structure) can elicit an N400 effect (Chow et al., 2016); the N400 has also been found when syntactic reanalysis occurs (Bornkessel, McElree, Schlesewsky, & Friederici, 2004). A theoretical advantage of assuming that feature activation is not restricted to semantic features is that the extension of features allows for a single mechanism responsible for the access to memory for prediction generation and for the creation of dependencies with previous words in the sentence (as in Lewis & Vasishth, 2005; Nicenboim & Vasishth, 2018; Engelmann, Jäger, & Vasishth, 2019; Jäger, Mertzen, Van Dyke, & Vasishth, 2019; Vasishth, Nicenboim, Engelmann, & Burchert, 2019, 11).

Interestingly, preliminary work of Fleur et al. (2019) points in the direction of a possible dissociation between the prediction of gender and phonological features (namely the lexical form of the determiner). However, without formal computational modeling it is not trivial to distinguish whether the N400 at the determiner can be explained with the grammatical or phonological context updating the activation of semantic features (Rabovsky & McRae, 2014; Kuperberg, 2016), or with the direct change in the activation of semantic, grammatical, or phonological features (in line, to some extent, with Chow et al., 2016).

To sum up, our meta-analysis shows a relatively clear and very small effect in the direction consistent with the prediction error account (or access view) that assumes that before a word is read, its context leads to the probabilistic pre-activation of semantic (and maybe phonological and grammatical) features in memory. It is noteworthy that almost every experiment in our meta-analysis had a 95% credible interval overlapping zero, and that only very weak conclusions could be deduced from the individual studies in isolation. Given the small effect present in this type of experimental manipulation, it seems unfeasible that a single study could provide closure on this topic. We were only able to gather more conclusive evidence once we quantitatively synthesized these same studies. This highlights the fact that science is and should be a cumulative enterprise, which benefits from data and materials that are publicly available (also see Nicenboim et al., 2018).

8.1 Potential concerns and limitations

There are some potential concerns and some limitations in our meta-analysis. We address these issues below:

8.1.1 Previous studies had much larger effects of the prenominal manipulation. A potential objection to our conclusion of a small effect due to the prenominal manipulation across the studies of the meta-analysis is that it is hard to reconcile our results with studies that did show larger N400 effects with similar manipulations: the original study in English with the a/an manipulation (DeLong et al., 2005), some studies manipulating gender in Dutch (Otten et al., 2007; Otten & Van Berkum, 2009; and recently Fleur et al., 2019), and most studies manipulating gender in Spanish (Wicha, Bates, et al., 2003; Wicha, Moreno, & Kutas, 2003; Foucart et al., 2014; and Martin et al., 2018; but cf. Wicha et al., 2004). For English and Dutch (except maybe Fleur et al., 2019, see below), it is likely that the larger effects are overestimates due to Type-M(agnitude) errors (Gelman & Carlin, 2014). This might be the case because the original studies had low power to detect small effects, and thus only overestimates could be detected and yield significant results (Vasishth et al., 2018). The fact that studies attempting to replicate these findings failed to find significant effects (and had much smaller estimates) supports this idea (e.g., Kochari & Flecken, 2019; Nieuwland et al., 2018). Results of Fleur et al. (2019) might be different since in their experiment the larger N400 effect might have been driven by both the mismatch of gender and lexical features (or lexical form).

Studies with Spanish speakers seem to be relatively consistently showing large effects. Possibly, the effects are genuinely stronger for the gender manipulation in Spanish. One explanation is that due to the transparency of gender marking and the reliability of gender cues in Spanish, readers (and listeners) can rely more strongly on the gender feature assigning more weight to it. With some exceptions, there are clear phonological patterns mostly associated with either feminine or masculine nouns (Beatty-Martínez & Dussias, 2019), and the identity of the determiner reliably indicates the gender of the following noun (except for the cases where the following noun starts with /a/; see Eddington & Hualde, 2015). The transparency of gender marking and the reliability of gender cues in Spanish seem to affect the acquisition of gender, making it easier to acquire in comparison to gender in German and Dutch languages (Eichler, Jansen, & Müller, 2012; Audring, 2014). The general reliability of gender cues and the transparency of gender marking might be relevant to both their acquisition and their weight in predictive processes (as reflected in the amplitude of the EEG signal in the N400 spatiotemporal window when there is a mismatch). This could explain the large effects for Spanish and the small effect found in our experiment in German. Nouns in German are not as clearly gender-marked as in Spanish: There are semantic, phonological, pseudosuffix, as well as suffix regularities on the basis of which gender can be assigned to German nouns, but there is a group of nouns that does not have any gender-marking regularities (Köpcke, 1982; Hohlfeld, 2006). Furthermore, while the critical regions in our experimental stimuli are singular accusative indefinite determiners and these are unambiguously gender marked, this is not the case for the complete paradigm. The complete determiner paradigm in German displays a high degree of homophony since gender marking in German is intertwined with case and number (Eichler et al., 2012). A more complete meta-analysis with studies in different languages that vary in the regularity

of their gender-marking might elucidate this point.

8.1.2 The results are very sensitive to the prior used in the Bayes factor analysis. Regarding the sensitivity analysis for the Bayes factor, a potential objection was pointed out by a reviewer: doesn't varying the priors amount to simply fishing for an effect under different assumptions, analogous to frequentist analyses that use terms like "marginally significant" to claim significance? The goal behind varying the priors, however, is to provide a full picture of what the researcher should believe *under different assumptions*; i.e., under different priors. Our meta-analysis showed that only if we assume a priori that the effect is small do we have relatively strong evidence for an effect. Using informative vs. uninformative priors to understand the range of conclusions one can draw under different assumptions is one of the strengths of the Bayesian framework.

8.1.3 The meta-analysis should be treated as exploratory. Whereas the meta-analysis suggests evidence in favor of a very small effect and provides no evidence for systematic differences between the "a/an" manipulation and the gender manipulation in Dutch or German, the case is by no means closed. The current meta-analysis was based on only the few studies with publicly available data (Ito et al., 2017a; Nieuwland et al., 2018; Kochari & Flecken, 2019) and our own exploratory results, and only one covariate (manipulation type) was investigated. Thus future work can still inform new and more comprehensive meta-analyses. It is certainly possible that our estimates are biased due to the studies included being not representative. However, a meta-analysis is certainly an improvement over the conventional approach, which just counts the number of significant and non-significant effects across studies in order to decide whether an effect is present or not (see Nicenboim et al., 2018). Even though the only compelling solution for the pervasive problem of bias is the use of large-scale preregistered replication attempts, a meta-analysis provides a valuable starting point and recommendations for future studies with respect to sample size and potential covariates (Van Elk et al., 2015). Ideally a large number of studies should be investigated, and studies could be clustered according to their language and paradigm leading to a precise estimate of both the commonality between the studies and the effect of different factors.

8.2 Conclusion

In closing, this study makes several novel contributions. First, this work is to our knowledge the first attempt at investigating the effect of gender on predictability in German. Second, the meta-analysis demonstrates a very important advantage of the Bayesian data analysis methodology: we can incrementally accumulate evidence to obtain increasingly precise estimates of the effect of interest. The approach we present, of assimilating existing evidence into the data analysis, has wide applicability in psycholinguistic research and has the potential to significantly advance our understanding of different phenomena of interest. Finally, although we cannot resolve the question of whether the a/an vs. gender manipulation at the determiner shows any difference in pre-activation, the evidence available so far indicates that regardless of the type of manipulation there is a small pre-activation effect at the determiner.

References

- Arnold, T. B. (2017). A tidy data model for natural language processing using cleanNLP. *R Journal*, 9(2).
- Audring, J. (2014). Gender as a complex feature. *Language Sciences*. Exploring Grammatical Gender, 43, 5–17. doi:10/ggks4g
- Aust, F., & Barth, M. (2018). *papaja: Create APA manuscripts with R Markdown*. R package version 0.1.0.9842.
- Baggio, G., & Hagoort, P. (2011). The balance between memory and unification in semantics: A dynamic account of the N400. *Language and Cognitive Processes*, 26(9), 1338–1367. doi:10.1080/01690965.2010.542671
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. doi:10.18637/jss.v067.i01
- Bates, D., & Maechler, M. (2018). *Matrix: Sparse and dense matrix classes and methods*. R package version 1.2-15.
- Beatty-Martínez, A. L., & Dussias, P. E. (2019). Revisiting Masculine and Feminine Grammatical Gender in Spanish: Linguistic, Psycholinguistic, and Neurolinguistic Evidence. *Frontiers in Psychology*, 10. doi:10/ggk2qc
- Bennett, C. H. (1976). Efficient estimation of free energy differences from Monte Carlo data. *Journal of Computational Physics*, 22(2), 245–268. doi:10.1016/0021-9991(76)90078-4
- Bornkessel, I., McElree, B., Schlesewsky, M., & Friederici, A. D. (2004). Multi-dimensional contributions to garden path strength: Dissociating phrase structure from case marking. *Journal of Memory and Language*, 51(4), 495–522.
- Braver, S. L., Thoenes, F. J., & Rosenthal, R. (2014). Continuously cumulating meta-analysis and replicability. *Perspectives on Psychological Science*, 9(3), 333–342. PMID: 26173268. doi:10.1177/1745691614529796
- Brown, C., & Hagoort, P. (1993). The Processing Nature of the N400: Evidence from Masked Priming. *Journal of Cognitive Neuroscience*, 5(1), 34–44. doi:10.1162/jocn.1993.5.1.34
- Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80(1), 1–28. doi:10.18637/jss.v080.i01
- Bürkner, P.-C. (2018). Advanced Bayesian multilevel modeling with the R package brms. *The R Journal*, 10(1), 395–411.
- Chen, S. F., & Goodman, J. (1999). An empirical study of smoothing techniques for language modeling. *Computer Speech & Language*, 13(4), 359–394. doi:https://doi.org/10.1006/csla.1999.0128
- Chow, W.-Y., Smith, C., Lau, E., & Phillips, C. (2016). A “bag-of-arguments” mechanism for initial verb predictions. *Language, Cognition and Neuroscience*, 31(5), 577–596. doi:10.1080/23273798.2015.1066832
- Chung, Y., Gelman, A., Rabe-Hesketh, S., Liu, J., & Dorie, V. (2013). Weakly informative prior for point estimation of covariance matrices in hierarchical models. *Manuscript submitted for publication*.
- Cumming, G. (2012). *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis*. Routledge.
- Delaney-Busch, N., Morgan, E., Lau, E. F., & Kuperberg, G. R. (2017). Comprehenders rationally adapt semantic predictions to the statistics of the local environment: A

- Bayesian model of trial-by-trial N400 amplitudes. In *Proceedings of the 39th Annual conference of the Cognitive Science Society, London, UK*.
- DeLong, K. A., Quante, L., & Kutas, M. (2014). Predictability, plausibility, and two late ERP positivities during written sentence comprehension. *Neuropsychologia*, *61*, 150–162. doi:10.1016/j.neuropsychologia.2014.06.016
- DeLong, K. A., Urbach, T. P., & Kutas, M. (2005). Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience*, *8*(8), 1117–1121. doi:10.1038/nrn1504
- Desmet, B., & Hoste, V. (2014). Fine-grained Dutch named entity recognition. *Language Resources and Evaluation*, *48*(2), 307–343. doi:10.1007/s10579-013-9255-y
- Eddelbuettel, D., & Balamuta, J. J. (2017). Extending R with C++: A Brief Introduction to Rcpp. *PeerJ Preprints*, *5*, e3188v1. doi:10.7287/peerj.preprints.3188v1
- Eddelbuettel, D., & François, R. (2011). Rcpp: Seamless R and C++ integration. *Journal of Statistical Software*, *40*(8), 1–18. doi:10.18637/jss.v040.i08
- Eddington, D., & Hualde, J. I. (2015). El abundante agua fría: Hermaphroditic Spanish Nouns. *Studies in Hispanic and Lusophone Linguistics*, *1*(1), 5–32. doi:10/ggk2qg
- Ehrlich, S. F., & Rayner, K. (1981). Contextual effects on word perception and eye movements during reading. *Journal of Verbal Learning and Verbal Behavior*, *20*(6), 641–655. doi:10.1016/S0022-5371(81)90220-6
- Eichler, N., Jansen, V., & Müller, N. (2012). Gender acquisition in bilingual children: French–German, Italian–German, Spanish–German and Italian–French: *International Journal of Bilingualism*. doi:10/ggks4w
- Engelmann, F., Jäger, L. A., & Vasishth, S. (2019). The effect of prominence and cue association in retrieval processes: A computational account. *Cognitive Science*. In Press.
- Federmeier, K. D., Wlotko, E. W., De Ochoa-Dewald, E., & Kutas, M. (2007). Multiple effects of sentential constraint on word processing. *Brain Research*, *1146*, 75–84. doi:10.1016/j.brainres.2006.06.101
- Fischler, I., Bloom, P. A., Childers, D. G., Roucos, S. E., & Perry Jr., N. W. (1983). Brain potentials related to stages of sentence verification. *Psychophysiology*, *20*(4), 400–409. doi:10.1111/j.1469-8986.1983.tb00920.x
- Fleur, D., Flecken, M., Rommers, J., & Nieuwland, M. S. (2019). Definitely saw it coming? An ERP study on the role of article gender and definiteness in predictive processing. *bioRxiv*. doi:10.1101/563783
- Foucart, A., Martin, C. D., Moreno, E. M., & Costa, A. (2014). Can bilinguals see it coming? Word anticipation in L2 sentence reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*(5), 1461–1469. doi:10.1037/a0036756
- Frank, S. L., Otten, L. J., Galli, G., & Vigliocco, G. (2015). The ERP response to the amount of information conveyed by words in sentences. *Brain and Language*, *140*, 1–11. doi:10.1016/j.bandl.2014.10.006
- Gabry, J., & Mahr, T. (2018). *Bayesplot: Plotting for bayesian models*. R package version 1.6.0.
- Gelman, A., & Carlin, J. (2014). Beyond power calculations assessing Type S (sign) and Type M (magnitude) errors. *Perspectives on Psychological Science*, *9*(6), 641–651.

- Gelman, A., Jakulin, A., Pittau, M. G., & Su, Y.-S. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 1360–1383.
- Gelman, A., Simpson, D., & Betancourt, M. (2017). The Prior Can Often Only Be Understood in the Context of the Likelihood. *Entropy*, 19, 555. doi:10.3390/e19100555
- Gramfort, A., Luessi, M., Larson, E., Engemann, D., Strohmeier, D., Brodbeck, C., ... Hämäläinen, M. (2013). Meg and eeg data analysis with mne-python. *Frontiers in Neuroscience*, 7, 267. doi:10.3389/fnins.2013.00267
- Gronau, Q. F., Sarafoglou, A., Matzke, D., Ly, A., Boehm, U., Marsman, M., ... Steingroever, H. (2017). A tutorial on bridge sampling. *Journal of Mathematical Psychology*, 81, 80–97. doi:https://doi.org/10.1016/j.jmp.2017.09.005
- Guerra, E., Nicenboim, B., & Helo, A. V. (2018). A crack in the crystal ball: Evidence against pre-activation of gender features in sentence comprehension. In *Architectures and mechanisms for language processing (AMLaP)*, Berlin, Germany.
- Hagoort, P., Baggio, G., & Willems, R. M. (2009). *Semantic unification*. MIT Press.
- Hagoort, P., Hald, L., Bastiaansen, M., & Petersson, K. M. (2004). Integration of Word Meaning and World Knowledge in Language Comprehension. *Science*, 304(5669), 438–441. doi:10.1126/science.1095455
- Hale, J. (2001). A probabilistic earley parser as a psycholinguistic model. In *Second meeting of the North American Chapter of the Association for Computational Linguistics on Language technologies 2001 - NAACL '01* (pp. 1–8). Second meeting of the North American Chapter of the Association for Computational Linguistics. doi:10.3115/1073336.1073357
- Henry, L., & Wickham, H. (2018). *Purrr: Functional programming tools*. R package version 0.2.5.
- Hoekstra, R., Morey, R. D., Rouder, J. N., & Wagenmakers, E.-J. (2014). Robust misinterpretation of confidence intervals. *Psychonomic bulletin & review*, 21(5), 1157–1164.
- Hohlfeld, A. (2006). Accessing grammatical gender in German: The impact of gender-marking regularities. *Applied Psycholinguistics*, 27(2), 127–142. doi:10/c8qvvd
- Hyvärinen, A., & Oja, E. (1997). A Fast Fixed-Point Algorithm for Independent Component Analysis. *Neural Computation*, 9(7), 1483–1492. doi:10.1162/neco.1997.9.7.1483
- Ito, A., Martin, A. E., & Nieuwland, M. S. (2017a). How robust are prediction effects in language comprehension? Failure to replicate article-elicited N400 effects. *Language, Cognition and Neuroscience*, 32(8), 954–965. doi:10.1080/23273798.2016.1242761
- Ito, A., Martin, A. E., & Nieuwland, M. S. (2017b). Why the A/AN prediction effect may be hard to replicate: A rebuttal to DeLong, Urbach, and Kutas (2017). *Language, Cognition and Neuroscience*, 32(8), 974–983. doi:10.1080/23273798.2017.1323112
- Jäger, L. A., Engelmann, F., & Vasishth, S. (2017). Similarity-based interference in sentence comprehension: Literature review and Bayesian meta-analysis. *Journal of Memory and Language*, 94, 316–339. doi:https://doi.org/10.1016/j.jml.2017.01.004
- Jäger, L. A., Mertzen, D., Van Dyke, J. A., & Vasishth, S. (2019). Interference patterns in subject-verb agreement and reflexives revisited: A large-sample study. *Journal of Memory and Language*. In Press. doi:https://doi.org/10.1016/j.jml.2019.104063
- Jasper, H. H. (1958). The 10/20 international electrode system. *EEG and Clinical Neurophysiology*, 10, 371–375.

- Jaynes, E. T., & Kempthorne, O. (1976). Confidence intervals vs. Bayesian intervals. In W. L. Harper & C. A. Hooker (Eds.), *Foundations of probability theory, statistical inference, and statistical theories of science* (Vol. 6b, pp. 175–257). The University of Western Ontario Series in Philosophy of Science. doi:10.1007/978-94-010-1436-6_6
- Jeffreys, H. (1939/1998). *The theory of probability*. Oxford University Press.
- Jung, T.-P., Makeig, S., Humphries, C., Lee, T.-W., Mckeown, M. J., Iragui, V., & Sejnowski, T. J. (2000). Removing electroencephalographic artifacts by blind source separation. *Psychophysiology*, 37(2), 163–178.
- Kochari, A. R., & Flecken, M. (2019). Lexical prediction in language comprehension: A replication study of grammatical gender effects in Dutch. *Language, Cognition and Neuroscience*, 34(2), 239–253. doi:10.1080/23273798.2018.1524500
- Kooperberg, C. (2018). *Polspline: Polynomial spline routines*. R package version 1.1.13.
- Köpcke, K.-M. (1982). *Untersuchungen zum genussystem der deutschen gegenwartssprache*.
- Kuperberg, G. R. (2016). Separate streams or probabilistic inference? What the N400 can tell us about the comprehension of events. *Language, Cognition and Neuroscience*, 31(5), 602–616. doi:10.1080/23273798.2015.1130233
- Kuperberg, G. R., & Jaeger, T. F. (2016). What do we mean by prediction in language comprehension? *Language, Cognition and Neuroscience*, 31(1), 32–59. doi:10.1080/23273798.2015.1102299
- Kutas, M., & Federmeier, K. D. (2000). Electrophysiology reveals semantic memory use in language comprehension. *Trends in Cognitive Sciences*, 4(12), 463–470. doi:10.1016/S1364-6613(00)01560-6
- Kutas, M., & Federmeier, K. D. (2011). Thirty Years and Counting: Finding Meaning in the N400 Component of the Event-Related Brain Potential (ERP). *Annual Review of Psychology*, 62(1), 621–647. doi:10.1146/annurev.psych.093008.131123
- Kutas, M., & Hillyard, S. A. (1980). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science*, 207(4427), 203–205. doi:10.1126/science.7350657
- Kutas, M., & Hillyard, S. A. (1984). Brain potentials during reading reflect word expectancy and semantic association. *Nature*, 307(5947), 161–163. doi:10.1038/307161a0
- Lappin, J. S., & Eriksen, C. W. (1966). Use of a delayed signal to stop a visual reaction-time response. *Journal of Experimental Psychology*, 72(6), 805–811. doi:10.1037/h0021266
- Lee, M. D., & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge University Press.
- Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, 106(3), 1126–1177. doi:10.1016/j.cognition.2007.05.006
- Lewis, R. L., & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, 29(3), 375–419. doi:10.1207/s15516709cog0000_25
- Lidstone, G. J. (1920). Note on the general case of the Bayes-Laplace formula for inductive or a posteriori probabilities. *Transactions of the Faculty of Actuaries*, 8(182-192), 13.
- Logan, G. D., & Cowan, W. B. (1984). On the Ability to Inhibit Thought and Action: A Theory of an Act of Control. *Psychological Review*, 91(3), 295–327.
- Luck, S. J., & Gaspelin, N. (2017). How to get statistically significant effects in any ERP experiment (and why you shouldn't). *Psychophysiology*, 54(1), 146–157. doi:10.1111/psyp.12639

- Martin, C. D., Branzi, F. M., & Bar, M. (2018). Prediction is Production: The missing link between language production and comprehension. *Scientific Reports*, 8(1). doi:10.1038/s41598-018-19499-4
- Martin, C. D., Thierry, G., Kuipers, J.-R., Boutonnet, B., Foucart, A., & Costa, A. (2013). Bilinguals reading in their second language do not predict upcoming words as native readers do. *Journal of Memory and Language*, 69(4), 574–588. doi:10.1016/j.jml.2013.08.001
- Mathôt, S., Schreij, D., & Theeuwes, J. (2012). OpenSesame: An open-source, graphical experiment builder for the social sciences. *Behavior Research Methods*, 44(2), 314–324. doi:10.3758/s13428-011-0168-7
- McShane, B. B., Gal, D., Gelman, A., Robert, C., & Tackett, J. L. (2019). Abandon statistical significance. *The American Statistician*, 73(sup1), 235–245. doi:10.1080/00031305.2018.1527253
- Meng, X.-l., & Wong, W. H. (1996). Simulating ratios of normalizing constants via a simple identity: A theoretical exploration. *Statistica Sinica*, 831–860.
- Miettinen, J., Nordhausen, K., Oja, H., & Taskinen, S. (2014). Deflation-Based FastICA With Adaptive Choices of Nonlinearities. *IEEE Transactions on Signal Processing*, 62(21), 5716–5724. doi:10.1109/TSP.2014.2356442
- Morey, R. D., Hoekstra, R., Rouder, J. N., Lee, M. D., & Wagenmakers, E.-J. (2016). The fallacy of placing confidence in confidence intervals. *Psychonomic bulletin & review*, 23(1), 103–123.
- Müller, K. (2018). *Bindrcpp: An ‘Rcpp’ interface to active bindings*. R package version 0.2.2.
- Nicenboim, B. (2018). *eeguana: A package for manipulating EEG data in R*.
- Nicenboim, B., Roettger, T. B., & Vasishth, S. (2018). Using meta-analysis for evidence synthesis: The case of incomplete neutralization in German. *Journal of Phonetics*, 70, 39–55. doi:10.1016/j.wocn.2018.06.001
- Nicenboim, B., & Vasishth, S. (2016). Statistical methods for linguistic research: Foundational Ideas - Part II. *Language and Linguistics Compass*, 10(11), 591–613. doi:10.1111/lnc3.12207
- Nicenboim, B., & Vasishth, S. (2018). Models of retrieval in sentence comprehension: A computational evaluation using Bayesian hierarchical modeling. *Journal of Memory and Language*, 99, 1–34. doi:10.1016/j.jml.2017.08.004
- Nieuwland, M. S. (2016). Quantification, prediction, and the online impact of sentence truth-value: Evidence from event-related potentials. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(2), 316–334. doi:10.1037/xlm0000173
- Nieuwland, M. S. (2019). Do ‘early’ brain responses reveal word form prediction during language comprehension? A critical review. *Neuroscience & Biobehavioral Reviews*, 96, 367–400. doi:10.1016/j.neubiorev.2018.11.019
- Nieuwland, M. S., Barr, D. J., Bartolozzi, F., Busch-Moreno, S., Darley, E., Donaldson, D. I., ... Von Grebmer Zu Wolfsturn, S. (2019). Dissociable effects of prediction and integration during language comprehension: Evidence from a large-scale study using brain potentials. *bioRxiv*. doi:10.1101/267815

- Nieuwland, M. S., Politzer-Ahles, S., Heyselaar, E., Segaert, K., Darley, E., Kazanina, N., ... Huettig, F. (2018). Large-scale replication study reveals a limit on probabilistic prediction in language comprehension. *eLife*, 7. doi:10.7554/eLife.33468
- Nordhausen, K., Ilmonen, P., Mandal, A., Oja, H., & Ollila, E. (2011). Deflation-based FastICA reloaded. In *2011 19th European Signal Processing Conference* (pp. 1854–1858). 2011 19th European Signal Processing Conference.
- Otten, M., Nieuwland, M. S., & Van Berkum, J. J. A. (2007). Great expectations: Specific lexical anticipation influences the processing of spoken language. *BMC Neuroscience*, 8(1). doi:10.1186/1471-2202-8-89
- Otten, M., & Van Berkum, J. J. A. (2008). Discourse-Based Word Anticipation During Language Processing: Prediction or Priming? *Discourse Processes*, 45(6), 464–496. doi:10.1080/01638530802356463
- Otten, M., & Van Berkum, J. J. A. (2009). Does working memory capacity affect the ability to predict upcoming words in discourse? *Brain Research*, 1291, 92–101. doi:10.1016/j.brainres.2009.07.042
- Perrin, F., & García-Larrea, L. (2003). Modulation of the N400 potential during auditory phonological/semantic interaction. *Cognitive Brain Research*, 17(1), 36–47. doi:https://doi.org/10.1016/S0926-6410(03)00078-8
- Politzer-Ahles, S. (2017). An extension of within-subject confidence intervals to models with crossed random effects. doi:10.20982/tqmp.13.1.p075
- Praamstra, P., & Stegeman, D. F. (1993). Phonological effects on the auditory N400 event-related brain potential. *Cognitive Brain Research*, 1(2), 73–86. doi:https://doi.org/10.1016/0926-6410(93)90013-U
- R Core Team. (2018). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria.
- Rabovsky, M., & McRae, K. (2014). Simulating the N400 ERP component as semantic network error: Insights from a feature-based connectionist attractor model of word meaning. *Cognition*, 132(1), 68–89. doi:10.1016/j.cognition.2014.03.010
- Resnik, P. (1992). Left-corner parsing and psychological plausibility. In *Proceedings of the 14th Conference on Computational Linguistics-Volume 1* (pp. 191–197). Association for Computational Linguistics.
- Rouder, J. N., Haaf, J. M., & Vandekerckhove, J. (2018). Bayesian inference for psychology, part IV: Parameter estimation and Bayes factors. *Psychonomic Bulletin & Review*, 25(1), 102–113. doi:10.3758/s13423-017-1420-7
- Schad, D. J., Hohenstein, S., Vasishth, S., & Kliegl, R. (2018). How to capitalize on a priori contrasts in linear (mixed) models: A tutorial. *arXiv preprint arXiv:1807.10451*.
- Schielzeth, H., & Forstmeier, W. (2009). Conclusions beyond support: Overconfident estimates in mixed models. *Behavioral Ecology*, 20(2), 416–420. doi:10.1093/beheco/arn145
- Smith, N. J., & Levy, R. (2013). The effect of word predictability on reading time is logarithmic. *Cognition*, 128(3), 302–319. doi:10.1016/j.cognition.2013.02.013
- Stan Development Team. (2018a). RStan: The R interface to Stan. R package version 2.18.2.
- Stan Development Team. (2018b). Stan: A C++ library for probability and sampling, version 2.18.2.

- Stan Development Team. (2018c). StanHeaders: Headers for the R interface to Stan. R package version 2.18.0.
- Van Berkum, J. J. A. (1996). *The psycholinguistics of grammatical gender: Studies in language comprehension and production*. OCLC: 246993702. Nijmegen : Max Planck Instituut voor Psycholinguïstiek.
- Van Berkum, J. J. A., Brown, C. M., Zwitserlood, P., Kooijman, V., & Hagoort, P. (2005). Anticipating Upcoming Words in Discourse: Evidence From ERPs and Reading Times. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(3), 443–467. doi:10.1037/0278-7393.31.3.443
- Van Berkum, J. J. A., Hagoort, P., & Brown, C. M. (1999). Semantic Integration in Sentences and Discourse: Evidence from the N400. *Journal of Cognitive Neuroscience*, 11(6), 657–671. doi:10.1162/089892999563724
- Van Elk, M., Matzke, D., Gronau, Q., Guang, M., Vandekerckhove, J., & Wagenmakers, E.-J. (2015). Meta-analyses are no substitute for registered replications: A skeptical perspective on religious priming. *Frontiers in Psychology*, 6. doi:10/f7rjq
- Van Petten, C., & Luka, B. J. (2012). Prediction during language comprehension: Benefits, costs, and ERP components. *International Journal of Psychophysiology*, 83(2), 176–190. doi:10.1016/j.ijpsycho.2011.09.015
- Vasishth, S., Chen, Z., Li, Q., & Guo, G. (2013). Processing Chinese relative clauses: Evidence for the subject-relative advantage. *PLoS ONE*, 8(10), 1–14.
- Vasishth, S., Mertzen, D., Jäger, L. A., & Gelman, A. (2018). The statistical significance filter leads to overoptimistic expectations of replicability. *Journal of Memory and Language*, 103, 151–175. doi:https://doi.org/10.1016/j.jml.2018.07.004
- Vasishth, S., & Nicenboim, B. (2016). Statistical methods for linguistic research: Foundational ideas - Part I. *Language and Linguistics Compass*, 10(8), 349–369. doi:10.1111/lnc3.12201
- Vasishth, S., Nicenboim, B., Engelmann, F., & Burchert, F. (2019). Computational models of retrieval processes in sentence processing. *Trends in Cognitive Sciences*, 23, 49–?? doi:https://doi.org/10.1016/j.tics.2019.09.003
- Vehtari, A., Gelman, A., & Gabry, J. (2015). Pareto Smoothed Importance Sampling.
- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5), 1413–1432. doi:10.1007/s11222-016-9696-4
- Venables, W. N., & Ripley, B. D. (2002). *Modern applied statistics with S* (Fourth). ISBN 0-387-95457-0. New York: Springer.
- Verbruggen, F., Logan, G. D., & Stevens, M. A. (2008). STOP-IT: Windows executable software for the stop-signal paradigm. *Behavior Research Methods*, 40(2), 479–483. doi:10.3758/BRM.40.2.479
- Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H., & Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the Savage–Dickey method. *Cognitive Psychology*, 60(3), 158–189. doi:10.1016/j.cogpsych.2009.12.001
- Wicha, N. Y. Y., Bates, E. A., Moreno, E. M., & Kutas, M. (2003). Potato not Pope: Human brain potentials to gender expectation and agreement in Spanish spoken sentences. *Neuroscience Letters*, 346(3), 165–168. doi:10.1016/S0304-3940(03)00599-8

- Wicha, N. Y. Y., Moreno, E. M., & Kutas, M. (2003). Expecting Gender: An Event Related Brain Potential Study on the Role of Grammatical Gender in Comprehending a Line Drawing Within a Written Sentence in Spanish. *Cortex*, 39(3), 483–508. doi:10.1016/S0010-9452(08)70260-0
- Wicha, N. Y. Y., Moreno, E. M., & Kutas, M. (2004). Anticipating Words and Their Gender: An Event-related Brain Potential Study of Semantic Integration, Gender Expectancy, and Gender Agreement in Spanish Sentence Reading. *Journal of Cognitive Neuroscience*, 16(7), 1272–1288. doi:10.1162/0898929041920487
- Wickham, H. (2016). *Ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York.
- Wickham, H. (2018). *Stringr: Simple, consistent wrappers for common string operations*. <http://stringr.tidyverse.org>, <https://github.com/tidyverse/stringr>.
- Wickham, H., & Bryan, J. (2018). *Readxl: Read excel files*. R package version 1.2.0.
- Wickham, H., François, R., Henry, L., & Müller, K. (2018). *Dplyr: A grammar of data manipulation*. R package version 0.7.8.
- Wickham, H., & Henry, L. (2018). *Tidyr: Easily tidy data with 'spread()' and 'gather()' functions*. R package version 0.8.2.
- Wickham, H., Hester, J., & François, R. (2018). *Readr: Read rectangular text data*. R package version 1.2.1.
- Wijffels, J. (2019). *Udpipe: Tokenization, parts of speech tagging, lemmatization and dependency parsing with the UDPipe NLP toolkit*. R package version 0.8.3.
- Wilke, C. O. (2018). *Ggridges: Ridgeline plots in 'ggplot2'*. R package version 0.5.1.
- Wilke, C. O. (2019). *Cowplot: Streamlined plot theme and plot annotations for 'ggplot2'*. R package version 0.9.4.
- Wolen, A., & Hartgerink, C. (2019). *Osfr: R interface to OSF*. <http://centerforopenscience.github.io/osfr>.
- Xie, Y., Allaire, J., & Golemund, G. (2018). *R markdown: The definitive guide*. ISBN 9781138359338. Boca Raton, Florida: Chapman and Hall/CRC.

Table A1

Model comparison between different measures of predictability.

Models	$\widehat{elpd}$	Difference in $\widehat{elpd}$	Difference SE
log cloze prob.	-93401	0.00	0.00
cloze prob.	-93408	-6.85	3.59
factorial predictor	-93414	-12.76	6.84

Note. The table is ordered by the expected log-predictive density ($\widehat{elpd}$) score of the models, with a higher score indicating better predictive accuracy. The highest scored model is used as a baseline for the difference in $\widehat{elpd}$ and the difference standard error (SE).

Appendix A

Model comparison between different measures of predictability as predictors for the average amplitude of the N400 spatiotemporal window.

We compare here the differences in predictive accuracy between smoothed cloze probabilities, log-transformed cloze probabilities, and factorial (1 for high cloze and -1 for low cloze) predictors using the noun region of Nieuwland et al. (2018). Since this dataset shows a clear effect of predictability, the dataset might be useful to compare different measures of predictability. We compare the out-of-sample predictive accuracy of the different models using Pareto smoothed importance sampling approximation to leave-one-out cross validation implemented in the package `loo` (Vehtari, Gelman, & Gabry, 2017; Vehtari, Gelman, & Gabry, 2015).

Table A1 shows the results of the comparison: There is a modest advantage in predictive accuracy by using log-transformed cloze probability in comparison with both non-transformed cloze scores and a factorial predictor.

Appendix B

Additional exploratory analyses

For completeness, we report here additional exploratory analyses that might be of interest.

Model for the adjective region with a constraining context. For this model, we analyzed the adjective region under a constraining context using gender probability as a predictor (and gender as a nuisance predictor as before) on the average amplitude of the N400 spatiotemporal window. We specify the previous model in `brms` as follows:

```
brm(A_N400 ~ pred * gender +
      (1 | item) + (pred * gender | subj),
  prior = reg_priors,
  data_Cab_adjectives, ...)
```

The estimate for the effect of predictability on the N400 for the adjectives pooled from conditions (a) and (b) of items with constraining context was $\hat{\beta} = 0.039\mu V$, 95% CrI = $[-0.077, 0.16]$.

Model for the determiner region with a constraining context using the cloze probability of the determiner as a predictor. For this model, we analyzed the determiner region under a constraining context using the cloze probability of the determiner, rather than its gender probability, as a predictor (and gender as a nuisance predictor as before) on the average amplitude of the N400 spatiotemporal window. We specify the previous model in `brms` as follows:

```
brm(A_N400 ~ pred_det * gender +
      (1 | item) + (pred_det * gender | subj),
  prior = reg_priors,
  data_Cab_adjectives, ...)
```

The estimate for the effect of predictability on the N400 for the determiners pooled from conditions (a) and (b) of items with constraining context using the cloze probability of the determiner as a predictor was $\hat{\beta} = 0.084\mu V$, 95% CrI = $[-0.043, 0.21]$.

Appendix C

Additional figures

For completeness, we report here additional figures.

ERPs elicited at the determiners (conditions a) in the N400 spatiotemporal window

Figure C1. ERPs elicited by the presentation of the same determiners pooled from items with a constraining context and a high probability gender, and from items with a non-constraining context and a low probability gender; see examples (6a) and (7a). The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).

**ERPs elicited at the determiners (conditions b)
in the late positivity spatiotemporal window**

Figure C2. ERPs elicited by the presentation of the same determiners with low probability pooled from items with a constraining context and a non-constraining context; see examples (6b) and (7b). The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).

ERPs elicited at the nouns (conditions a) in the N400 spatiotemporal window

Figure C3. ERPs elicited by the presentation of the same nouns pooled from items with a constraining context and a high probability noun, and from items with a non-constraining context and a low probability noun; see examples 6a and 7a. The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).

Figure C4. ERPs elicited by the presentation of the same critical nouns with low cloze probability pooled from items with a constraining context and a non-constraining context; see examples (6b) and (7b). The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).

ERPs elicited at the determiners (constraining context) in the N400 spatiotemporal window

Figure C5. ERPs elicited by the presentation of the critical determiners with high and low probability genders from items with a constraining context; see examples (6a) and (6b). The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).

ERPs elicited at the nouns (constraining context) in the N400 spatiotemporal window

Figure C6. ERPs elicited by the presentation of the critical nouns with high and low cloze probability from items with a constraining context; see examples (6a) and (6b). The x-axis indicates time in seconds and the light grey square represents the pre-specified time window used in the analysis. The y-axis indicates voltage in microvolts, note that negative polarity is plotted downwards. The dark grey areas around the traces represent 95% confidence intervals based on cross random effects; see Politzer-Ahles (2017).