Noisy is better than rare: Comprehenders compromise subject-verb agreement to form more probable linguistic structures

Maayan Keshev^{a*} and Aya Meltzer-Asscher^{a,b}

^aLinguistics Department, Tel-Aviv University, Tel-Aviv, Israel; ^bSagol School of Neuroscience, Tel-Aviv University, Tel-Aviv, Israel

*corresponding author maayankeshev@mail.tau.ac.il

Abstract:

Production and perception errors are common in everyday language use. Recent studies suggest that in order to overcome the flawed speech signal, comprehenders engage in rational noisy-channel processing, which can pull their interpretation towards more probable “near-neighbor” analyses, based on the assumption that an error may have occurred in the transmission of the sentence. We investigate this type of processing using subject/object relative clause ambiguity in Hebrew. In four self-paced reading experiments and a sentence completion experiment, we find that during online processing, readers apply elaborate knowledge regarding the distribution of structures in the language, and that they are willing to compromise subject-verb agreement to refrain from (grammatical but) highly improbable structures. The results suggest that the prior probability of alternative analyses modulates the interpretation of agreement.

Keywords: sentence processing; noisy-channel; rational inference; relative clause; subject-verb agreement

1. Introduction

Sentence processing is based on the integration of bottom-up and top-down information. As we read or listen to a sentence, the incoming input is integrated with existing knowledge of different types, e.g. information about lexical items and their properties, about the plausibility of different messages, and about the probability of various linguistic structures. Such probabilistic knowledge can be used, for example, in ambiguity resolution, pulling the interpretation towards the more probable structure (e.g. MacDonald, Pearlmutter, & Seidenberg, 1994; Levy, 2008a). Interestingly, recent studies suggest that comprehenders sometimes utilize similar biases even when the sentence, if taken at face value, is unambiguous, namely corresponds to only one grammatical structure (Levy, Bicknell, Slattery, & Rayner, 2009; Gibson, Bergen, & Piantadosi, 2013). These studies argue that readers take into account the possibility of speech and print errors, as well as perception errors, and maintain some uncertainty with regard to the input. As a result, they may be pulled towards "near-neighbour", extrapolated interpretations. For example, when presented with the sentence "The mother gave the candle the daughter", rather than analysing the sentence faithfully to the input (and arriving at an implausible interpretation), the comprehender may choose to assume that the intended utterance was in fact "The mother gave the candle to the daughter", and that it was corrupted during production or perception, in order to accommodate the more probable message (Gibson et al., 2013).

1.1 Rational noisy-channel processing

In models of rational (Bayesian) noisy-channel processing, the addressee weighs the probability distribution of the different structures against the probability that noise corrupted the input into its current form (Levy, 2008b). More specifically, the addressee ranks the possible underlying structures (marked here as S) of the input (I), deducing the probability that the underlying structure of the input is S_i , by balancing two probability components: the structure's prior probability, $\Pr(S_i)$, and the likelihood of perceiving the current input, given that S_i is the underlying structure (based on a certain noise-model), $\Pr(I/S_i)$. As a result, the highest ranked structure may not be completely faithful to the input. The parser may overlook some properties of the input, or "edit" it in different ways, if this would produce a sentence with higher probability. The existence of such a process can account for comprehenders' ability to recover the intended meaning from noisy utterances. However, it also means that the parser may treat sentences that have only one grammatical reading as ambiguous.

Exemplifying this type of noisy-channel inference, Levy and his colleagues (2009) found, using eye-tracking, that readers look back to reevaluate the spelling of an early word, when it has an orthographic neighbor which grammatically coheres downstream with a structure of higher prior probability. In their experiment, participants regressed to the word *at* in sentences like (1a). This sentence contains a passive reduced relative clause (bracketed), which is a marked and dispreferred structure. However, switching the preposition *at* with its near-neighbor *as* would license the analysis of *tossed* as an active main verb, a much more frequent structure. The rate of regressions to the preposition was significantly higher in (1a) than in (1b), where the verb *thrown* does not allow for an active main verb reading. Moreover, the regression rate in (1a) was higher than in (1c/d), where the preposition (*toward*) had no relevant near-neighbor. This suggests that when there was no possible near-neighbor, the low frequency string was accepted as is. Namely, since the preposition *toward* is too distinct from one which can license a main verb reading, the likelihood of any relevant modification was too low to win over the faithful (though marked) analysis.

- (1) a. The coach smiled at the player [tossed the frisbee].
- b. The coach smiled at the player [thrown the frisbee].
- c. The coach smiled toward the player [tossed the frisbee].
- d. The coach smiled toward the player [thrown the frisbee].

In further support of rational noisy-channel processing, Levy (2011) demonstrated garden path effects which result from compromising faithfulness to the input. The study examined readers' tendency to disregard a comma when a subsequent fronted prepositional phrase could be used as a complement of a preceding verb (in sentences like 2a), in order to avoid the low-frequency locative-inversion structure. In such sentences, adopting a noisy-channel interpretation, namely neglecting the comma, should lead the parser to a dead-end which requires structural reanalysis. At the verb *lurched*, the readers should recognize that a subject was missing, and must reanalyze the structure of the sentence, such that a new clause begins following the comma. In line with this, the results revealed an increase in reading times at the disambiguating verb in (2a) vs. (2b), indicating reanalysis costs which are required if an input-unfaithful analysis was constructed earlier.

- (2) a. As the soldiers marched, toward the tank lurched an injured enemy combatant.
- b. As the soldiers marched, the tank lurched toward an injured enemy combatant.

Another type of support for the noisy-channel processing framework comes from the correlation between the likelihood of edits to the input, and the rate of choosing unfaithful parses (i.e. the number final interpretations which not in like with the literal meaning, across multiple trials and participants). Each possible structure has some likelihood given the input, $\Pr(I/S_i)$. One can experimentally investigate the relation between this likelihood and the probability of choosing an interpretation which that target interpretation by: (i) increasing the general rate of errors within the experiment, with the prediction that when a signal is conveyed over a noisier channel and errors are more likely overall, interpretation will be more susceptible to unfaithful parses; (ii) comparing different types of edits, with the prediction that unfaithful parses will be adopted less often if the required edit is more notable (i.e. when the unfaithful parse has of even lower likelihood).

Using comprehension questions, Gibson and his colleagues (2013) probed the rate of unfaithful interpretations in different settings and structures, with the underlying assumption is that a rise in the rate of unfaithful parses means that $\Pr(S_{noisy}/I)$ is rising relative to $\Pr(S_{faithful}/I)$. They found that an unfaithful, semantically more plausible, reading was chosen at a higher rate when this could be achieved with fewer edits of the input. They also showed that readers engaged in this noise inference at a higher rate when the experiment included a high proportion of sentences with syntactic errors. The rate of unfaithful parses thus increases with the perceived noise rate of the situation. Following these findings, additional studies of noisy channel inference (e.g. Poppels & Levy, 2016; Futrell & Levy, 2017; Gibson, et al., 2017; Ryskin, et al., 2018) have utilized the paradigm developed by Gibson and colleagues (2013), which focuses on semantic plausibility as the prior probability manipulation.

In the current study, we target the prior probability of syntactic structures, and how it affects noisy-channel interpretation. We focus on the prior probability of the faithful analysis, comparing two possible faithful interpretations with the same unfaithful one. Thus, we ask how rare a structure has to be in order for readers to assume, when confronted with it, that it is a noisy version of another structure (i.e. that another structure underlies it and was corrupted by noise). We aim to exhibit that the prior probability of the structure of the faithful parse affects adoption of noisy, unfaithful, interpretation. In addition, we aim to provide evidence for such noisy-channel inferences in incremental processing based on online measures. We believe that this is an important step in assessing the claims of the noisy-channel processing model.

We examine the effect of a structure's prior probability using Hebrew relative clauses which are temporarily ambiguous between a grammatical, but rare, object relative clause and

a subject relative clause containing an agreement mismatch. We will argue that readers may or may not interpret a mismatch in agreement features as evidence for an object relative, depending on the markedness of that object relative structure. Namely, the prior probability of the syntactic structure can modulate how likely readers are to treat the mismatch as a cue for interpretation shift. Before approaching the experimental design, we take a short detour to consider the subject/object relative clause asymmetry and Hebrew word order.

1.2 The subject/object relative clause asymmetry

One of the most well-known findings in the sentence processing literature is the difference in difficulty of two types of relative clauses: Object-extracted (3b) and subject-extracted (3a) relative clauses. Subject relatives (SRs) are preferred over object relatives (ORs): They are read faster (King & Just 1991; Traxler, Morris, & Seely, 2002; Staub, 2010), recalled better (Baird & Koslick, 1974), acquired earlier (de Villiers, et al., 1979; Tavakolian 1981; Diessel & Tomasello 2005), and retained at a higher rate in language disorders (Caplan & Futter, 1986; Grodzinsky, 1989).¹ For Hebrew, this asymmetry was shown primarily in language acquisition research (Arnon, 2010; Friedmann & Novogrodsky 2004; Friedmann, Belletti, & Rizzi, 2009).

(3) a. Here is a filler [which _ has a corresponding gap at the embedded subject position]

b. Here is a filler [which we interpret _ as the object in the bracketed clause]

The ease of processing SRs in comparison to ORs can be attributed to various factors. First, in languages where the filler precedes the embedded clause, and word order adheres to the subject-verb-object scheme (like English and Hebrew), fewer constituents intervene between the filler and the gap in SRs than in ORs. The additional distance in ORs can increase processing load as it requires longer maintenance of the filler (Wanner & Maratsos, 1978; Frazier, 1987a; Just and Carpenter, 1992), or more challenging retrieval at the gap site (Gibson, 2000; Lewis & Vasishth, 2005).

Alternatively, experience-based accounts suggest that the preference for SRs over ORs should be attributed to prior language exposure. SRs are more frequent than ORs (in English - Fox & Thompson, 1990; Real & Christiansen, 2007; Roland, Dick, & Elman, 2007; and in Hebrew - Arnon, 2010). Therefore, the SR reading will be favored in early stages of

¹ Yet it should be mentioned that several factors modulate this subject/object relative clause asymmetry. These factors include the animacy of the filler noun (Gennari & MacDonald, 2008; Traxler, Williams, Blozis, & Morris, 2005), the referential type of intervening elements (Gordon et al., 2004; Real & Christiansen, 2007), the verb type (Gennari & MacDonald, 2009), the relative order of the relative clause and its head (Wagers, Borja, & Chung, 2018), and word order within the relative clause (Levy et al., 2013).

interpretation (Traxler, Morris, & Seely, 2002) and should encompass less surprisal (Levy, 2008a). Moreover, in English (and Hebrew), SRs maintain the surface form of subject-verb-object, and thus align with general word order expectations (MacDonald & Christiansen, 2002), while ORs present a non-canonical order.

Another family of accounts for the SR/OR asymmetry focuses on the syntactic complexity of their final interpretation. Structural theories (Keenan & Comrie, 1977; O'Grady, 1997) propose that SRs are easier to comprehend due to greater accessibility of the subject extraction site: The subject position is hierarchically higher in the sentence structure, and presents a more obligatory constituent. Finally, more than one factor can contribute to the SR preference (see Staub, 2010 and Levy, Fedorenko & Gibson 2013, for evidence for involvement of both memory-based and expectation-based effects; and Pizzaro-Guerva, 2020, for an argument in favor of maintaining structural complexity factors alongside other accounts).

Either of the above hypotheses (or any combination of them) predicts that in Hebrew, like in English and many other languages, SRs would be preferred over ORs as well as over relative clauses with more embedded relativization sites. In the next subsection we describe how the grammatical properties of Hebrew may allow some ambiguity in verb-initial relative clauses. This ambiguity will be utilized in the current study to test whether, and under what circumstances, the preference for SRs can motivate readers to disregard other grammatical information.

1.3 The occasional ambiguity in Hebrew relative clauses

The default word order in Hebrew is SVO (Ravid, 1977; Berman, 1980). However, there are some cases where verb-initial sentences are licensed. First, subject-verb inversion is possible in certain registers (mostly in formal or literary contexts). In such cases (4), the subject would follow the verb, and preferably a clause-initial prepositional phrase (Shlonsky, 1990; Shlonsky & Doron, 1992).

(4) *ba-pšita ha-leylit acra ha-mištara harbe pe'ilim.*

in.the-raid the-nightly detained the-police many activists.

'The police detained many activists in the nightly raid.' (Shlonsky & Doron, 1992)

Second, Hebrew allows various sentence constructions where an overt grammatical subject is lacking (Berman, 1980). The most productive among these subjectless constructions are third-person plural impersonals (5). These constructions allow the speaker not to specify who was

the agent of the event or action, and may compensate for the scarcity of passive constructions in Hebrew (Berman, 1979; Dekel, 2014).

- (5) a. *ganvu et ha-ofana'im šeli etmol.*
 stole-3.PL ACC the-bicycle mine yesterday.
 'Someone stole my bicycle yesterday'.
 b. *be-avar, axlu yoter yerakot.*
 in-past, ate-3.PL more vegetables.
 'In the past, people ate more vegetables'.

These Hebrew structures which allow verb-initial clauses can create local ambiguity between subject and object relatives, as shown in (6). A relative clause starting with a verb could be interpreted either as a canonical subject relative (6a), or as an object relative whose subject was removed: either omitted as an impersonal (6b, only possible when the verb is in the plural, as explained above), or dislocated to the post-verbal position (6c). The relative clauses in (6) include an initial prepositional phrase since the option of a post-verbal subject construction, which is rare to begin with (see section 1.4 below) is further degraded when the verb is not preceded by such a phrase, which supposedly "triggers" the inversion (Shlonsky, 1990; Shlonsky & Doron, 1992).

- (6) *ha-yeladim she-be-sof ha-yom mac'u...*

The-kids that-in-end the-day found.PL

a. Subject relative:

ha-yeladim_i she-be-sof ha-yom, ___i mac'u et ha-derex habayta

The-kids_i that-in-end the-day, ___i found.PL ACC the-way home

'The kids who, at the end of the day, found the way home'

b. Object relative with an impersonal null subject:

ha-yeladim_i she-be-sof ha-yom, Ø mac'u ___i be-ezrat kol toshvey ha-shxuna

The-kids_i that-in-end the-day, Ø found.PL ___i in-help all residents the-neighbourhood

'The kids who, at the end of the day, were found thanks to the neighbourhood residents'

c. Object relative with a post-verbal subject:

ha-yeladim_i she-be-sof ha-yom, mac'u ha-horim ___i yeshenim al ha-sapa

The-kids_i that-in-end the-day, found.PL the-parents ___i sleeping on the-sofa

'The kids who, at the end of the day, the parents found sleeping on the-sofa'

Upon reading the verb, in sentences like (6), the initial interpretation should adhere to the SR/OR asymmetry. Similarly ambiguous relative clauses in Dutch, exhibit an SR preference. Thus, Dutch relative clauses which are fully ambiguous between object and subject relatives are interpreted as SRs over 70% of the time (Frazier, 1987a). Similarly, in German, processing disruption arises when locally ambiguous wh-questions are disambiguated to the object extraction reading (Meng & Bader, 2000). Therefore, in cases like (6), we expect the SR reading to be preferred over either of the OR structures.

1.4 The current study: Prior probability of SR and OR analyses and noisy agreement

In this study, we use relative clauses which are temporarily ambiguous between subject and object relatives in order to test noisy-channel inference. We manipulate the agreement match between the filler and the verb in order to test whether readers are willing to compromise subject-verb agreement (which is required for the SR reading), in order to avoid the less frequent OR structures. Namely, we test whether the parser acts upon a top-down preference for SRs over ORs, even when the verb's features do not licence grammatical subject-verb (i.e. filler-verb) agreement. If an SR analysis is preferred in such cases despite the agreement mismatch, this would present evidence for noisy-channel inference in incremental sentence processing.

Why should it be rational for the parser to compromise the grammaticality of subject-verb agreement? We suggest that agreement is prone to errors, and thus can provide high-likelihood competitors for the faithful parse. Different orthographic, linguistic and psycholinguistic factors may conspire to make agreement errors relatively frequent. First, agreement marking on both verbs and nouns is usually established orthographically by manipulation of only one or two letters. This should make it a likely target of typing errors. Second, the manifestation of agreement features on different constituents varies across languages. Thus, L2 learners may acquire such paradigms inaccurately when the marking is unavailable or different in their native tongue (e.g. Jackson, Mormer, & Brehm, 2018). Utterances of non-native speakers form part of the everyday language input of native speakers, especially in a multi-lingual environment like Israel (Spolsky & Shohamy, 1999), where only 49% of the population speak Hebrew natively (Israel's Central Bureau of Statistics, 2013). Finally, subject-verb agreement is malleable and systematically vulnerable to interference, in both production and comprehension (Bock & Miller, 1991; Eberhard, Cutting, & Bock, 2005; Pearlmutter, Garnsey, & Bock, 1999; Wagers, Lau & Phillips, 2009, among others; for evidence from Hebrew see - Dank & Deutsch, 2010; Dank, Deutsch, & Bock, 2015; Deutsch & Dank, 2009; 2011). Given the considerations

above, in choosing between a faithful, rare structure and a slightly noisy but more common one, a rational inference framework predicts that a noisy SR would be formed. Namely, readers may dismiss the agreement mismatch as an error, to avoid constructing a low probability structure.

In this context, it should also be noted while testing the SR preference in Dutch, Frazier (1987a) found a lingering footprint of the SR bias even when agreement was supposed to disambiguate the structure (7). When the verb did not agree with the filler (pointing to the OR reading, as in 7b), readers still identified the filler as the subject of the relative clause in 31% of the cases. The opposite pattern of errors was much less likely - only 3.7% of the trials were analyzed as an OR when agreement suggested otherwise (in sentences like 7a). This provides evidence that readers may compromise a mismatch between a filler and a verb in order to construct an unmarked SR.

(7) *Ik schreef aan de vriend die mijn tantes ...*

I wrote to the friend who my aunts ...

a. Subject relative:

de vriend₁ [die ___₁ mijn tantes heeft bezocht].

the friend₁ [who ___₁ my aunts has visited].

'the friend who has visited my aunts.'

b. Object relative:

de vriend₁ [die mijn tantes ___₁ hebben bezocht].

the friend₁ [who my aunts ___₁ have visited].

'the friend who my aunts have visited.'

Interestingly, manipulating the availability of the two OR strategies in Hebrew, the post-verbal subject and impersonal subject, will allow us to test the willingness to compromise agreement when different prior probabilities are at stake. Post-verbal subjects are rare and marked, while impersonal null subjects are frequent (though presumably not as frequent as simple SRs). Opting for the noisy SR interpretation could be rational only if the alternative, grammatical OR analysis has a very low prior probability. Thus, we test whether ORs with low and intermediate prior probability produce the same noisy-channel inference.

Full data on the frequency of the different structures in Hebrew are missing, since corpus studies of Hebrew, like experimental ones, are relatively rare. However, there are a few observations on the frequency of post-verbal subject sentences as well as on impersonal null subjects. Maschler (2015) found only 57 tokens of post-verbal subjects in a Hebrew corpus

which comprises approximately 6.5 hours of conversation among 396 different speakers.² In contrast, impersonal null subjects are frequent in different registers and styles of Israeli Hebrew (Taube, 2007), including both written and colloquial language. For example, Berman (2011) identifies such uses "in nearly every transcript" in child-directed speech.

On the model of rational noisy-channel inference, such differences in the prior probability of the pristine OR structure analysis should affect readers' tendency to consider the corrupt SR interpretation and to ignore the agreement mismatch. If readers take these probabilities into account, we should observe a stronger bias for a corrupted SR analysis when the grammatical alternative is an OR with a post-verbal subject. On the other hand, when they have the option of impersonal null subjects, comprehenders might not opt for the corrupted SR interpretation.

Other models of sentence processing do not predict such a dependence on prior probability. One prominent alternative to noisy-channel processing maintains that the parser takes into account only interpretations which are grammatical and faithful to the input. This parsing principle could fit in different frameworks, from serial-modular processing models (Frazier, 1987b) to parallel-interactive processing models (MacDonald, Pearlmutter, & Seidenberg, 1994), and has been implicitly assumed by many (Frazier & Fodor, 1978; Gibson, 1998; Jurafsky, 1996; Hale, 2001). Since in the case of filler-verb mismatch, the SR analysis involves assuming an ungrammatical dependency or a noisy input, such a parser should not even consider this interpretation. Only pristine interpretations are considered, and hence the OR parse should be preferred consistently, regardless of the frequency of the required grammatical structure.

In contrast, other frameworks could maintain that the parser retains a SR reading despite the verb's agreement features, regardless of the frequency of the OR alternative. The "good-enough" processing approach (Ferreira, Bailey, & Ferraro, 2002; Ferreira & Patson, 2007) suggests that when analysing a sentence, in addition to grammatical constraints, comprehenders also use simple heuristics, e.g. an agent-first bias (Ferreira, 2003). In addition, comprehenders may have difficulty in inhibiting their initial interpretations and completing the reanalysis (Christianson et al., 2001). Adopting general heuristics can yield effects similar to those

² It might be questionable to deduce word order preferences within relative clauses from overall frequencies (Levy, Fedorenko, & Gibson, 2013). Yet, in the case of Hebrew post-verbal subjects, it seems that if anything, their proportion in relative clauses might be even smaller than in main clauses. Givón (1976) suggests that the post-verbal subject strategy is suppressed in RCs for complexity considerations. In addition, Maschler (2015) suggests that the discourse function of verb-subject constructions is to introduce a new referent or a shift in the narrative. Backgrounded constituents like relative clauses should provide less motivation for such a discourse mechanism.

predicted by noisy-channel models; however, the frequency of the unfaithful structure should not affect the rate with which it is adopted. Effects are derived from the prominence of general heuristics and difficulty of inhibiting rudimentary interpretations. Thus, to the best of our knowledge, this approach does not suggest a different interpretation strategy for the two OR structures in question. In both cases, readers are equally likely to prefer the SR reading, which aligns best with the agent-first bias, and/or fail to fully reanalyse an initial SR interpretation at the non-agreeing verb. Differences related to the frequency of the OR analysis (i.e. whether it is very rare or moderately infrequent) are not expected to arise as long as the initial and simplest prediction is that of an SR. Yet, the characterisation of "good-enough" processing as a shallow processing procedure in online parsing has been challenged by recent findings (e.g. Slattery et al., 2013; Huang & Ferreira, 2020). See the General Discussion for an interpretation of the results within current formulations of "good-enough" processing, which are based on grammatically proper parsing and subsequent memory fallibility which perturbs offline interpretation.

To test the predictions of the noisy-channel processing model, in Experiment 1, we use cases where the only faithful analysis is that of a OR with the rare post-verbal subject, while in Experiment 2, we test cases where the more common OR interpretation, with an impersonal null subject, is available. Finally, the two structures are directly contrasted in Experiment 3, which also examines comprehenders' interpretations more directly, using a sentence completion paradigm. To disclose the results early, the manipulation of prior probability indeed affects how readers treat the agreement mismatch. Experiments 1A-B exhibit that, faced with the choice between a noisy SR and an OR with a rare word order, readers opt for the SR analysis even though it includes an agreement mismatch. However, when the OR option does not involve a syntactic structure of very low probability (Experiments 2A-B), the grammatical OR analysis is preferred over a noisy SR. This preference pattern is also mirrored in the production data of Experiment 3.

2. Experiments 1A & 1B

As explained above, Hebrew licenses post-verbal subjects (mostly in higher registers) when the clause begins with an adjunct or an adverbial phrase (Shlonsky, 1990; Shlonsky & Doron, 1992). Therefore, when a relative clause (with a singular head) begins with an adverbial, followed by a (singular) verb, the structure is ambiguous between a SR and an OR with a post-verbal subject (we use singular verbs for testing this reading to neutralize the impersonal null subject reading which is licensed with plural verbs). Disambiguation arises at the first NP

following the verb. If this is a definite NP preceded by an accusative case marker (*et*, as in 6a), or an NP that does not agree with the verb, this NP is the object, meaning that the clause is a subject relative clause.³ If the post-verbal NP is a definite NP which is not preceded by an accusative case marker nor by a preposition, and it agrees with the verb (as in 6c), this NP is the subject, meaning that the clause is an object relative clause, with a post-verbal subject.

It is most likely that, faced with this local ambiguity in, comprehenders would initially prefer the SR reading. Therefore, a subject NP following the verb would be inconsistent with the initial parsing preferences and should give rise to a processing difficulty. The two experiments reported in this subsection build on such a reanalysis effect as evidence for the SR interpretation. We test whether readers compromise agreement and endorse the SR structure, by measuring reading times on the post-verbal subject in OR clauses, when the verb matches the filler (confirming the basic SR bias) and when it mismatches it (as in 8).

(8) *ha-yeladim_i she-be-sof ha-yom, maca ha-dod ____i mul ha-televizia*
 The-kids_i that-in-end the-day, found.SG **the-uncle** ____i in.front.of the-TV

If disruption on the post-verbal subject is evident when the verb mismatches the filler, this would suggest that readers overlook the agreement mismatch and form a SR, in line with Rational Noisy-Channel processing. However, if readers take the mismatch as indication that a subject gap is impossible, they should predict a post-verbal subject. In that case, reading times on the post-verbal subject should align with those in a baseline condition including no dependency (see below), as in both cases it should be consistent with the parser's prediction.

We evaluate the processing disruption relative to a baseline condition with no filler-gap dependency, where a post-verbal subject is expected. In such sentences, since there is no option to assume a gap, the absence of a subject before the verb should lead readers to unambiguously predict a post-verbal subject, and such a phrase should not disrupt processing. We chose this baseline in line with the tradition in the ambiguity resolution literature and in contrast to the direct OR-SR comparisons which are often found in the literature on relative clause asymmetry. We believe that this baseline provides a better basis for comparison than a simple SR baseline as it includes the low frequency structure of a post-verbal subject, similarly to the OR structure. Thus, we can avoid possible confounds arising from processing of this infrequent word order.⁴

³ In Hebrew, the distinction between nominative and accusative case is not marked for indefinite noun phrases; thus, a post-verbal indefinite NP will preserve the OR/SR ambiguity, if it agrees with the verb in number, gender and person.

⁴ As Gennari & Macdonald (2008, p. 162-163) argue about using SR as a baseline: "This choice contrasts with typical baseline conditions in ambiguity resolution studies in which unambiguous structures are chosen to have

Moreover, we wanted to allow a complete lexical overlap in the critical word (the nominative subject), as well as the words immediately preceding and following it, across all conditions. This would not have been possible if the verb was followed by its complements, as would occur in a SR baseline sentence.

2.1 Experiment 1A

2.1.1 Methods

Participants

Participants were 52 native speakers of Hebrew (according to self-report) (mean age: 24.18, range: 18-35). Nine participants were bilingual of Hebrew and either Russian or English, and the rest were monolingual. Participants gave informed consent and received payment (20 NIS, approximately 5 US dollars) for their participation. The experiment was approved by the Ethics Committee in Tel Aviv University.

Materials

The experiment included 30 experimental items of three conditions. Materials manipulated the agreement on the filler (*matching* or *mismatching* the following verb, see example set in Table 1). As our baseline, we included an additional condition with an unambiguous post-verbal subject in an adjunct clause (without a filler-gap dependency).

In the Match condition, the filler and verb were both marked for singular. In the Mismatch condition, the verb was still inflected for singular, but the filler was a noun phrase marked for plural. As explained above, to detect reanalysis of an initial SR interpretation, we designed the sentences to include a post-verbal subject, namely a definite NP, not preceded by the accusative case marker, which agreed in features with the preceding verb. In our materials, the object relative was ultimately resolved with an accusative resumptive pronoun (grammatical in Hebrew) within the clausal complement of the verb. Thus, all the sentences in the experiment were eventually grammatical. In addition, to avoid contamination of the critical region (the post-verbal NP) from the verb, we used an adverb as a buffer between the verb and the subject.

Condition	Sentence		
Match	<i>avarnu</i>	<i>leyad ha-pakax_l</i>	<i>še-be-mahalax ha-pkak</i>

similar meanings to those in ambiguous conditions, though the sentences may have minor differences in number of words or word choices. Because subject and object relatives differ in both word order and final sentence meaning, it is possible to attribute differences in comprehension difficulty to either of these factors".

(Exp. 1A&B) we+passed near the-inspector-M.SG₁ that-during the-traffic jam
himšix lixora ha-nahag ha-alim litkof oto₁
 continued-3.M.SG allegedly **the-driver**.M.SG the-violent to-attack him₁
bli šum buša
 without any shame.

'We went past the inspector who the violent driver continued to attack shamelessly, during the traffic jam'.

Mismatch
 (one feature)
 (Exp. 1A&B) *avarnu leyad ha-pakaxim₁ še-be-mahalax ha-pkak*
 we+passed near the-inspectors-M.PL₁ that-during the-traffic jam
himšix lixora ha-nahag ha-alim litkof otam₁
 continued-3.M.SG allegedly **the-driver**.M.SG the-violent to-attack them-M₁
bli šum buša.
 without any shame.

'We went past the inspectors who the violent driver continued to attack shamelessly, during the traffic jam'.

Baseline
 (Exp. 1A&B) *avarnu bezman še-be-mahalax ha-pkak himšix*
 we+passed while that-during the-traffic.jam continued-3.M.SG
lixora ha-nahag ha-alim litkof et ha-pakax bli
 allegedly **the-driver**.M.SG the-violent to-attack ACC the-inspector without
šum buša.
 any shame.

'We went past while the violent driver continued to attack the inspector shamelessly, during the traffic jam'.

Two-feature
 mismatch
 (Exp. 1B
 only) *avarnu leyad ha-pakaxiyot₁ še-be-mahalax ha-pkak*
 we.passed near the-inspectors-F.PL₁ that-during the-traffic jam
himšix lixora ha-nahag ha-alim litkof otan₁
 continued-3.M.SG allegedly **the-driver**.M.SG the-violent to-attack them-F₁
bli šum buša.
 without any shame.

'We went passed the inspectors who the violent driver continued to attack shamelessly, during the traffic jam'.

Table 1. Example set from Experiments 1A & 1B.

The critical NP is marked in bold; ACC = accusative case marker; F = feminine grammatical gender; M = male grammatical gender; SG = singular; PL = plural.

The experimental items were distributed in a Latin square design across three lists. Within each list, the target sentences were combined with 45 grammatical filler sentences, for a total of 75 sentences. The filler items included 10 sentences with no filler-gap dependency, 15 unambiguous subject relatives, and 20 long filler-gap dependencies (10 object relatives and 10 relative clauses with an embedded subject resolution).

Procedure

Sentences were presented on a desktop PC using the Linger software package (Rhode, 2003) in a moving window, non-cumulative self-paced reading display. Sentences were initially masked by dashes with no spaces or punctuation. Participants pressed the space bar to reveal each word. Before they began the experiment, participants underwent a practice block of four items. Seventy-five percent of the trials, in both the practice and experimental blocks, were followed by a 'yes/no' comprehension question. No feedback was provided for correct or incorrect answers. The order of presentation was randomized for each participant. Participants were instructed to read each sentence continuously, at their own pace, such that they understand it, but avoid lingering on words. The experiment took approximately 20 minutes.

Data analysis

Participants were excluded from analysis if their performance on comprehension questions of experimental items was lower than 70% (resulting in the removal of five participants), or if their average RT in experimental items was more than 2.5 SDs above the group's average (resulting in the removal of two additional participants). For the remaining 45 participants, RTs higher than 2.5 SDs above the individual's average RT were trimmed to that cutoff (affecting 2.6% of the data). RTs shorter than 120ms were excluded (affecting 0.03% of the data).

We analyzed RTs from the critical region (the post-verbal subject noun), as well as from the spillover word (the adjective following the post-verbal noun).⁵ We used a treatment

⁵ We did not analyze the verb region, since we believe there is no strong prediction for the RT pattern there. Under rational noisy-channel processing, the agreement mismatch should be detected at the verb. Yet the availability of the more probable near-neighbor pushes the parser to the SR reading. Thus, it could be that the mismatch would

(dummy) coding scheme to detect an increase in RT for each condition relative to the baseline condition. Namely, for every region, two comparisons were made: match vs. baseline and mismatch vs. baseline. In addition, due to an unpredicted contrast between the pattern of the effects at the critical and the spillover region, we fitted another model, using the data of both regions, to test for the main effects over both regions and of the region x condition interaction. Therefore, in this model, treatment (dummy) coding was used for the condition factor and sum coding was used for the region factor.

The data were analyzed in R using mixed-effects linear models in Bayesian analysis with weakly informative priors (Gelman et al., 2014). The Bayesian hierarchical models were fit using Stan (Carpenter et al., 2017), via the brms package (Bürkner, 2017). We first fitted two separate models for the critical and the spillover region. Then we constructed a unified model, including both regions, to verify an unexpected by-region interaction. A maximal random-effect structure was used in this model, with intercepts and slopes for by-item and by-participant effects. We summarize the posterior distribution by presenting its mean along with the 95% credible interval (CrI). We consider the posterior as consistent with an effect if over 95% of the sampled posterior distribution is beyond zero in the predicted direction.

We used a LogNormal likelihood function to model the raw values in milliseconds (ms). The models included a full random effects structure and only the experimental fixed effect. In the text, we report the estimates along with their credible intervals (CrI) transformed back to milliseconds. We used regularizing, weakly informative priors based on Schad, Betancourt, and Vasishth (2020): a standard normal distribution, $N(0, 1)$, was used as the prior distribution for the fixed effects and for the standard deviation parameters. For the intercept, we had a wider prior of $N(0, 10)$. Finally, for the correlation matrices of the random effects, priors were defined using the LKJ prior (Lewandowski, Kurowicka, & Joe, 2009). For each of the models, we sampled from the posterior distribution by running four MCMC chains at 4000 iterations each. The first half of the samples was discarded as warm-up samples. Convergence was checked using the R-hat convergence diagnostic and by visual inspection of the chains (Gelman et al., 2014).

induce a measurable processing disruption, or that it would dissolve without an increase in measured RTs (if self-paced reading times reflect late enough processing stages where agreement is already compromised). Namely, both detecting and not detecting mismatch effects at the verb could be compatible with the noisy channel hypothesis, rendering this comparison uninformative for the purpose of this study.

Additional frequentist analysis revealed similar effect and is available in Appendix A. The code of all analyses and the associated data are available at <https://osf.io/v6y8n/>.

2.1.2 Results

Word-by-word reading times of the different experimental conditions and by-condition numerical means of the tested regions are presented in Figure 1.

Figure 1. Word-by-word RT means (msec) by condition, Experiment 1A. Error bars represent $\pm$ SE.

We observed an unexpected pattern of reanalysis timing when modeling the critical and the spillover regions separately, which can also be seen by visual inspection of the results: processing difficulty in the Match condition only arises at the spillover region, whereas difficulty arises already in the critical region for the Mismatch condition. In order to validate the interaction between region (critical vs. spillover) and condition effects (Match-Baseline contrast and Mismatch-Baseline contrast), we conducted another analysis using data from both regions. Since results of both versions of the analysis were compatible, we report here only the results of the unified model (results of the by-region models are available in Appendix B).

The results are consistent with an increase in RTs at the spillover region of the Match condition relative to the baseline condition: the mean of the posterior distribution was 29ms with a CrI of [11, 47], and over 99% of the posterior distribution was beyond zero. A parallel

effect was not observed in the critical region (posterior mean [CrI]: -7ms [-24, 10], $\Pr(\beta > 0) = 20\%$). The posterior of the main effect for this contrast was not conclusive, with a mean and CrI of 10 [-3, 24] and 94% of the posterior distribution beyond zero. Instead, the posterior distribution was consistent with an interaction between region and condition, for the Match vs. baseline contrast: the mean of the posterior was -17ms with CrI of [-28, -6], and over 99% of the posterior distribution was beyond zero. This suggests a differential pattern in the critical and post-critical position, due to the late timing of the reanalysis cost in the Match condition.

Crucially, the results are consistent with increased RTs in the Mismatch condition relative to baseline. This can be observed in the pairwise comparisons, at the critical region: the mean of the posterior was 17ms with CrI of [-2, 36], and 96% of the posterior distribution was above zero. For parallel contrast in the spillover region, the posterior distribution was inconclusive (posterior mean [CrI]: 11ms [-7, 30], $\Pr(\beta > 0) = 89\%$). Yet the main effect posterior was consistent with an increase in RT the comparison between the Mismatch and Baseline conditions across the two regions (posterior mean [CrI]: 14ms [1, 27], $\Pr(\beta > 0) = 98\%$), and the posterior distribution was not conclusive as to an interaction effect (posterior mean [CrI]: 3ms [-10, 16], $\Pr(\beta > 0) = 67\%$).

2.1.3 Discussion

The results suggest increased processing costs at the disambiguating post-verbal subject both in the Match and Mismatch conditions, relative to an unambiguous subject-inversion structure. We interpret these increased RTs as reanalysis costs, required if readers constructed a SR at the verb. We suggest that readers tend to disregard the mismatch between the filler and the verb and construct a SR in order to avoid the (possible yet rare) subject-verb inversion structure. Note that the tendency to dismiss the agreement mismatch cannot be attributed to a strategy developed throughout the experiment. Since all experimental materials and filler items were grammatical, there should be no reason to expect ungrammaticality as part of the experiment.

It should be noted that while the reanalysis effect in the Match condition was shifted and was found only at the spillover region, in the Mismatch condition, reanalysis costs were evident on the post-verbal subject itself. This could suggest that the detection of the required reanalysis was faster in mismatching condition relative to sentences with a matching verb (further interpretation of this timing difference is discussed 2.2.3). Since this result was unexpected, and in order to validate the reanalysis costs we observed in the Mismatch condition, we conducted another experiment to replicate the findings of this experiment.

2.2 Experiment 1B

In Experiment 1B we aimed to replicate the pattern of results obtained in Experiment 1A. In addition, we wanted to examine how robust the preference for discarding agreement is, by extending our manipulation to more pronounced filler-verb agreement mismatches. We thus added another condition, in which the filler and the verb were differentiated by the two features, with discrepancy in both gender and number (rather than number only).

2.2.1 Methods

Participants

Participants were 68 native speakers of Hebrew (according to self-report) (mean age: 25.48, range: 21-35). Three participants were bilingual of Hebrew and English, and the rest were monolingual. Participants gave informed consent and received payment (20 NIS, approximately 5 US dollars) for their participation. The experiment was approved by the Ethics Committee in Tel Aviv University.

Materials

The experiment included 28 sets out of the sets in Experiment 1A. The experiment comprised the three conditions of Experiment 1A, and an additional condition, the Two-feature mismatch condition, in which the filler and verb differed in both gender and number features: the filler was a feminine plural noun phrase, while the verb was inflected for male singular (see the last condition in Table 1). Items were distributed into four lists in a Latin square design. Filler items were also the same as in Experiment 1A.

Procedure

The procedure was the same as in Experiment 1A.

Data analysis

Data analysis followed the same protocol as in Experiment 1A, only this time we fitted the data directly with the unified model (testing for an interaction between region and condition). Exclusion criteria resulted in the removal of six participants (four for low accuracy in comprehension questions and two for abnormal reading times). For the remaining 60 participants, trimming of high and low RTs affected 2.6% and 0.06% of the data, correspondingly.

2.2.2 Results

Word-by-word reading times of the different experimental conditions and by-condition numerical means of the tested regions are presented in Figure 2.

Figure 2. Word-by-word RT means (msec) by condition, Experiment 1B. Error bars represent $\pm$ SE.

The results of the Bayesian analysis are consistent with a main effect of condition, reflecting a general increase in RT in the Mismatch condition relative to baseline: the posterior had a mean of 15ms with CrI of [-1, 32] and a probability of 97% for being beyond zero. The posterior distributions were not conclusive as to main effects for the Match condition or the Two-feature mismatch, relative to baseline (Match posterior of 2ms [-14, 17] and $\Pr(\beta > 0) = 58\%$; the Two-feature mismatch posterior of 10ms [-6, 26] and $\Pr(\beta > 0) = 88\%$).

In addition, we the results are consistent with an interaction between region and condition for the Match-baseline contrast, but not for the contrasts between the baseline and the Mismatch or the Two-feature mismatch conditions. While the by-region interaction term for the Match-baseline contrast had more than 99% of the posterior being beyond zero (mean of -17ms with CrI of [-30, -3]), the Mismatch-baseline and Two-feature mismatch-baseline by region interactions were mainly centered around zero (Mismatch: mean of 4ms with CrI of [-10, 18]; Two-feature mismatch: mean of -2ms with CrI of [-15, 12]).

We conducted pairwise comparisons, contrasting each condition with the baseline, in each region. We observed an increase in RT at the spillover region of the Match condition relative to baseline: the posterior of this contrast had a mean of 19ms with a CrI of [-2, 41] and 96%

probability of being above zero. The analysis also indicated increased RT at the critical region of the Mismatch condition: the posterior of the contrast also had a mean of 19ms with a CrI of [-2, 41] and 96% probability of being beyond zero. The Two-feature mismatch contrast with the baseline condition was not conclusive: the posterior had a mean and CrI of 11ms [-10, 33], in both regions ($\Pr(\beta > 0) = 85\%$ for the critical region and $\Pr(\beta > 0) = 84\%$ for the spillover region).

Figure 3. The sampled posterior distributions of the pairwise contrasts in Experiments 1A&B, on the log-scale (with 95% CrI shaded in gray).

2.2.3 Discussion

Results of this experiment largely replicated those of Experiment 1A, in that we observed increased RTs in the Mismatch condition (see Figure 3 for a comparison of the effects in both experiments). These results suggest that when the relative clause is ambiguous between a corrupted SR and a OR with a post-verbal subject, readers prefer constructing a noisy SR. In addition, Experiment 1B replicated the temporal pattern observed in Experiment 1A, such that in the Match condition, the reanalysis effect was only observed in the post-critical position, while the effect in the Mismatch condition was detected in the critical region.

We suggest that this pattern may reflect earlier detection of the required reanalysis in the Mismatch condition relative to the Match condition, and/or easier reanalysis in this condition, as a result of the earlier origin of uncertainty. Namely, we propose that readers maintain some

uncertainty regarding the correct structure after encountering the mismatching verb in the Mismatch condition. This in turn facilitates the reanalysis procedure at the post-verbal subject. When cues from the bottom-up input (verb-filler agreement) and top-down priors (SR preference) align, the parser might become more committed to the interpretation. Thus, it would take readers longer to resolve the conflicting evidence and reanalyse the sentence structure, when the ambiguity resolves to the OR interpretation.

Such a strategy is in line with other findings from semantic cues for filler-verb mismatch. Pickering and Traxler (1998, 2003) suggest that a semantic implausibility of the filler-verb association reduces the parser's commitment to the dependency and thus lowers reanalysis costs. This type of uncertainty and lower commitment in cases of mismatch is also in line with the general idea behind the noisy-channel framework, assuming that the parser maintains and continuously updates the uncertainty as to its representations (Levy et al., 2009). This is usually implemented in different studies as uncertainty regarding single word representations (e.g. plural vs. singular word form in this study, and *at* vs. *as* in Levy et al., 2009). However, it could naturally be extended to whole parses while incrementally constructing an interpretation.

For the two-feature mismatch condition, we did not detect processing costs relative to baseline (the posterior had considerable uncertainty below and over zero). Yet there seems to be a large overlap in the posterior distribution for this contrast and the one-feature mismatch (see Figure 3). We therefore believe the results are inconclusive as to whether readers are willing to compromise a double-feature mismatch, or rather take the mismatch as an indication for the rare subject-verb inversion.

It should be noted that in most sets, the additional mismatching feature did not affect the number of orthographic edits which are required for restoring agreement. Future research might examine more closely whether manipulating an abstract feature mismatch, without affecting the likelihood of low-level errors (in visual perception or in typing), can affect readers' parsing strategy. Ryskin and her colleagues (2018) recently suggested that the comprehender's noise model treats some types of errors as analogous and some as categorically distinct, possibly differentiating typing errors and planning errors. If indeed readers make inferences regarding the producer's planning errors, they should consider abstract/linguistic similarity when calculating the likelihood of different errors.

Two possible concerns should be addressed at this point. First, it could be suggested that the increased reading times at the post-verbal subject, which we interpret as reanalysis costs, merely reflect mismatch-detection costs that spilled over from the verb region. However, we believe that this is not likely to be the case. This interpretation of the results wrongfully predicts

that processing costs would be greater when the mismatch is more prominent. Yet, a two-feature mismatch (in Experiment 1B) did not elicit greater processing disruption; If anything, the effect was smaller than that of the one-feature mismatch (in line with the prediction of noisy-channel processing). In addition, we believe that the inclusion of an adverbial buffer in the materials (positioned between the verb and the critical region) makes the spillover interpretation less likely, as it would require the effect to migrate two words downstream. Moreover, in Experiment 3 we will provide additional evidence for a SR preference (from sentence completion data), thus supporting the conclusion that the costs detected in Experiment 1 indeed reflect reanalysis.

Second, reading times data on earlier regions are rather noisy in this experiment, exhibiting unpredicted variability, as can be observed at the pre-critical region in Figure 2. We do not have a compelling account for these fluctuations, nor did we analyze these regions to determine the statistical reliability of these effects. One possible interpretation of this variability could suggest that participants engage in strategic more in this experiment, as the mismatch cues, and specifically the double-mismatch, might make the manipulation more obvious than in Experiment 1.⁶ While we cannot exclude this possibility, we believe that it does not weaken our conclusions. If readers still exhibit the same effects at the critical regions even with more careful, strategic parsing, then the automaticity of noisy-channel inferences is exhibited more strongly.

Overall, we take the results to replicate our findings from Experiment 1A and suggest that readers are willing to compromise agreement in order to construct a structure of higher prior probability. These findings are in line with informal judgments provided in Givón (1976). In his review of the distribution of post-verbal subjects in Hebrew, Givón suggests that post-verbal subjects might be less common in ORs, due to the ambiguity they introduce. He states that when the filler mismatches the verb's features, there seems to be a "tendency" to interpret the sentence as a SR with an agreement error.

These findings suggest that readers consider ungrammatical and input unfaithful interpretations. Thus, they provide evidence against parsing models which generate only pristine analyses. However, these results by themselves cannot support the noisy-channel model over shallow interpretation. It could be suggested that negligence of agreement information does not reflect a rational processing strategy, favouring noisy over rare representations, but a heuristic of "good-enough" processing (Ferreira, Bailey, Ferraro, 2002;

⁶ We thank an anonymous reviewer for pointing this out.

Ferreira & Patson, 2007), which prefers SR clauses across-the-board and does not take into account elaborate frequency information. For example, English passives and object-clefts were suggested by Ferreira (2003) to be systematically misinterpreted, based on good-enough processing and an agent-first heuristic. It is possible that the tendency to interpret the relative clauses of Experiment 1 as SRs is based on shallow processing, general under-specification of agreement information, and/or consistent dismissing of agreement in favor of the short SR dependencies, which conform with an agent-first bias.⁷ In order to test this, in Experiment 2 we use another case where the SR/OR ambiguity might arise, in which the prior probability of the faithful OR analysis is not as low as that of post-verbal subjects.

3. Experiments 2A & 2B

In Experiments 2, we test the SR/OR ambiguity with a different OR alternative, which is not as rare as the subject-verb inversion we used in Experiment 1. Thus, we aim to contrast parsing preferences when different priors are involved. To do this, we use the leeway of impersonal null subjects in Hebrew.

In Hebrew, verbs with plural third person agreement licence empty subjects (Shlonsky, 1997), which are interpreted as referring to an unknown, immaterial or generic entity (see Berman, 2011, for the semantic interpretation of such impersonals). Therefore, in a relative clause with a plural head, a third person plural verb (with no previous subject), would have two possible readings, as in (9a-b). One possible analysis (9a) posits the gap in subject position (i.e. forming an SR). The second (9b) posits it in the object position (i.e. forming an OR) and assumes an impersonal null subject (marked here by $\emptyset$).

(9) *ele ha-yeladim [še-mac'u...*

These the-kids [that-found-3.PL ...

a. Subject relative:

ele ha-yeladim_I [še-____I-mac'u et ha-ofana'im šeli ha-boker].

These the-kids_I [that-____I-found-3.PL ACC the-bicycle mine the-morning].

'These are the kids who found my bicycle this morning.'

b. Object relative:

ele ha-yeladim_I [še- $\emptyset$ -mac'u ____I im ha-ofana'im šeli ha-boker].

These the-kids_I [that-found-3.PL ____I with the-bicycle mine the-morning].

'These are the kids who someone found with my bicycle this morning.'

⁷ In Hebrew, like in English, subject relatives recreate the canonical SVO word order, on the surface, and allow for an agent-verb-theme interpretation. Object relatives, on the other hand, present the object first (as it heads the relative clause). Thus regardless of the subject-verb word order, object relatives present a less canonical word order, which is at odds with heuristics like agent-first.

As in Experiments 1, we test the parser's preference in the face of this ambiguity, manipulating the agreement features of the filler to allow or disallow the subject-verb agreement, required under a SR reading. In Experiments 1, the low frequency of the alternative (object relative with a post-verbal subject) pushed readers towards adopting the SR reading, despite the agreement mismatch. However, we predict that the current experiments, using the plural verb form and therefore introducing the alternative impersonal OR analysis, will produce different results, due to the higher prior probability of this structure. As mentioned in the introduction to this study, while ORs with post-verbal subjects are exceptionally rare, impersonal ORs are quite frequent (although presumably still less probable than simple SRs).

If the prior probability of the alternative structure affects the interpretation of filler-verb agreement mismatch, the results of Experiment 2 might indicate an OR preference instead of the noisy SR analysis we detected in Experiment 1. While SRs should still be favoured over ORs in general, the frequency of ORs with impersonal subjects might not be low enough to tip the balance and make the noisy SR analysis rank higher than the faithful OR analysis. In contrast, if rational noisy-channel processing is not behind the findings of Experiment 1, and agreement is categorically ignored in favour of any probable SR structure, the results of Experiment 2 should be in line with an initial SR interpretation (as in Experiment 1), even though an OR with an impersonal subject is not as rare as an OR with a post-verbal subject.

To assess whether the parser adopts an OR representation, we test for reanalysis costs when a lexically realized argument appears in the direct object position. In addition, as in Experiment 1, we manipulate the agreement of the filler to match or mismatch the (plural) verb. If a SR analysis is constructed for the relative clause (like in the previous experiments), an object NP is predicted and should not cause any difficulty. When the filler matches the verb, the structure is fully ambiguous between a SR and an OR reading, before that NP. If in this case, the object NP will not entail processing costs, this would suggest that readers prefer the SR interpretation, when both options (SR and impersonal OR) are grammatical (exemplifying the SR/OR asymmetry).

As for the case where the verb mismatches the agreement features of the filler, predictions are as follows: If it is the case that agreement information is underspecified or that mismatches are categorically ignored (due to general "good-enough" processing), then an object NP should not cause processing disruption in this condition as well (in line with the SR preference in the match condition). However, the framework of noisy-channel processing predicts a different pattern. Note that here, since the verb is plural, a reading of an OR with an impersonal null subject is possible. This reading is more probable than ORs with post-verbal subjects, but still

arguably less probable than SRs. In this case, since the pristine alternative is not very rare, rational readers should take the mismatch as indication that a subject gap is impossible and construct an OR. If indeed readers construct an OR (using the null subject option), a realized direct object should disrupt processing, as a gap was posited in object position, and reanalysis is required (i.e. the "filled-gap" effect, cf. Stowe, 1986).

Lastly, to detect these possible reanalysis costs at the lexically realized direct object, we compare these relative clauses (with matching or mismatching filler-verb agreement) to a baseline with no filler-gap dependency. This again was done to ensure that the object position is predicted in the baseline condition, in line with Experiment 1 and with previous "filled-gap" studies (e.g. Stowe, 1986; Boland et al., 1995).

3.1 Experiment 2A

3.1.1 Methods

Participants

Participants were 69 native speakers of Hebrew (according to self-report) (mean age: 25.12, range: 18-35). Seven participants were bilingual of Hebrew and either Russian, English, or Spanish and the rest were monolingual. Participants gave informed consent and received payment (20 NIS, approximately 5 US dollars) for their participation. The experiment was approved by the Ethics Committee in Tel Aviv University.

Materials

The experiment included 24 experimental items of four conditions (see example set in Table 2): temporarily ambiguous relative clauses where the filler matched or mismatched the embedded plural verb, a classic "filled-gap" condition, and a baseline condition. The experimental conditions were constructed such that at the verb, the Match and Mismatch conditions can be interpreted as ORs with an impersonal null subject.⁸ Thus, the verb always appeared in the plural form (in the Match condition, the filler and the verb were both plural; in the Mismatch condition, the verb was plural, but the filler was a singular noun phrase). Following the verb, an object NP appeared, marked by the accusative marker. This lexical NP fills the potential object gap position. The dependency grammatically resolved later with a

⁸ As explained above, in the Match condition the resolution is ambiguous between a SR, with grammatical agreement between the filler and the verb (the more likely interpretation), and an analysis with an impersonal null subject and a possessive resumptive pronoun (as in the Mismatch and "filled-gap" conditions).

possessive resumptive pronoun, matching the filler. At the object position, if an OR analysis (with a null subject) was adopted, processing difficulty is predicted, since the gap posited in this position needs to be replaced with a lexical NP, causing a "filled-gap" effect. In the baseline condition, we used an adjunct clause with a realized subject. No relative clause structure was involved, and no gap could be posited in the verb's object position at any point. Thus, the occurrence of a direct object NP should not disrupt processing.

In addition to the Match, Mismatch, and Baseline conditions, we included a "standard" filled-gap effect (FGE) condition. In this condition, the filler was plural, but a first-person pronoun was cliticized to the verb, so a subject gap would not be considered.⁹ Therefore, in this condition, readers are predicted to construct an object relative, and engage in reanalysis upon encountering the verb's object. This condition was included to avoid reliance on possible null effects. That is, if we did not detect any processing difficulty in the Mismatch condition, we would be able to use the contrast with the FGE condition to evaluate the reliability of the results.

Condition	Sentence
Match	<p><i>ha-saparit dibra im ha-me'acvim₁ še- baxru larov</i> the-hairdresser talked with the-designers-M.PL₁ that-chose-3.PL mostly <i>et ha-dugmaniyot ha-rašiyot šelahem₁ ba-rega ha-axaron.</i> ACC the-models the-main of.them-M₁ in.the-moment the-last. 'The hairdresser talked with the designers who usually chose their leading models at the last moment.'</p>
Mismatch	<p><i>ha-saparit dibra im ha-me'acev₁ še- baxru larov</i> the-hairdresser talked with the-designer-M.SP₁ that- chose-3.PL mostly <i>et ha-dugmaniyot ha-rašiyot šelo₁ ba-rega ha-axaron.</i> ACC the-models the-main of.him₁ in.the-moment the-last. 'The hairdresser talked with the designer whose leading models were usually chosen at the last moment.'</p>

⁹ In principle, it would be possible for readers to "edit" the verb form in this condition as well, changing it to the third person plural form in order to adopt a subject relative analysis. However, we assume that readers are much less likely to do so in the FGE condition, relative to the Mismatch condition, since in the latter, a SR can be established by merely assuming misreading of the filler.

Filled-gap effect	<i>ha-saparit dibra im ha-me'acvim₁ še-baxarnu larov</i> the-hairdresser talked with the-designers-M.PL ₁ that-we.chose mostly <i>et ha-dugmaniyot ha-rašiyot šelahem₁ ba-rega ha-axaron.</i> ACC the-models the-main of.them-M ₁ in.the-moment the-last. 'The hairdresser talked with the designer whose leading models we usually chose at the last moment'
Baseline	<i>ha-saparit dibra im ha-me'acvim₁ mikeyvan še-hem₁</i> the-hairdresser talked with the-designers-M.PL ₁ because that-they-M <i>baxru larov et ha-dugmaniyot ha-rašiyot šelahem₁ ba-</i> chose mostly ACC the-models the-main of.them-M ₁ in.the- <i>rega ha-axaron.</i> moment the-last. 'The hairdresser talked with the designers because they usually chose their leading models at the last moment'

Table 2. Example set from Experiment 2A.

The critical NP is marked in bold; ACC = accusative case marker. PL = plural.

The experimental items were distributed in a Latin square design across four lists. Within each list, the target sentences were combined with 51 grammatical filler sentences, for a total of 75 sentences. The filler items included 27 sentences with no relative clauses, 12 unambiguous SRs with a singular head and verb, and 12 unambiguous ORs. It should be noted that in this experiment, the final structure unambiguously contained impersonal null subjects in the Mismatch condition (the Match condition is globally ambiguous between a SR and an OR with an impersonal subject, see footnote 4). Thus, each experimental list included only 6-12 instances of the tested structure, while in Experiments 1A-B all experimental sentences ultimately contained the tested structure (a post-verbal subject). To minimize between-experiments differences in adaptation to the tested structures, we included 10 globally unambiguous impersonal subjects within the filler items of this experiment.

Procedure

The procedure was the same as in the previous experiments.

Data analysis

Data analysis followed the same protocol as in Experiment 1. Since there was no prediction or evidence for a differential pattern across regions (critical and post-critical words), we returned to separate by-region models. Exclusion criteria resulted in removal of nine participants (seven for low accuracy in comprehension questions and two for abnormal reading time averages). For the remaining 60 participants, trimming of high and low RTs affected 2.65% and 0.03% of the data, correspondingly.

3.1.2 Results

Word-by-word reading times of the different experimental conditions and by-condition numerical means of the tested regions are presented in Figure 4.

Figure 4. Word-by-word RT means (msec) by condition, Experiment 2A. Error bars represent $\pm$ SE.

At the critical region (the direct object position), we observed an increase in the Mismatch and in the FGE conditions relative to the baseline condition. The posteriors for these contrasts had high means (80ms and 99ms, correspondingly) with CrI that were far from zero ([52, 110] and [68, 133], correspondingly), and both had over 99% of the distribution above zero. These contrasts extended to the spillover region as well. In this region too, the posterior distributions of the Mismatch-Baseline contrast and the FGE-baseline contrast had high means (62ms and 53ms, correspondingly), CrI that were far from zero ([36, 90] and [30, 76], correspondingly) and over 99% above zero.

In the Match condition, the posterior was inconclusive with regards to increase in RT relative to baseline, both in the critical region (posterior mean and CrI of 16ms [-8, 42]) and in the spillover (posterior mean and CrI of 10ms [-8, 30]). Although relatively high percentage of the posterior distribution was above zero, in the critical region (90%), it did not reach the required level of reliability (95%), thus we failed to detect such contrast. At any rate, the posterior for this contrast is clearly lower than that observed in the Mismatch and FGE conditions, such that credible interval of the Match-Baseline contrast falls clearly outside those of the other contrasts (see Figure 6).

3.1.3 Discussion

In this experiment, we observed increased reading times at the direct object in the Mismatch condition relative to an unambiguous baseline (with no filler-gap dependency), on a par with the FGE condition. We interpret this as evidence that readers constructed an OR in this condition and were therefore required to perform reanalysis at this position. We suggest that when the filler mismatches the verb, and an impersonal null subject analysis is available, readers interpret the input as a faithful OR with an impersonal subject, rather than assume a noisy SR.

As we predicted, the results contrast with the results from Experiments 1, in which a noisy SR analysis was preferred over an OR. This aligns with the hypothesis that the prior probability of the alternative structure affects the interpretation of the agreement mismatch, and that the canonical SR reading is not applied across the board as a heuristic. The faithful analysis in Experiment 2A is of a higher probability than that in Experiments 1, allowing for a higher ranking of this OR analysis. Namely, we suggest that the contrast between the results of Experiments 1 and 2A should be interpreted as evidence for rational noisy-channel processing, and specifically for the effect of prior probability on the interpretation of possibly noisy input.

Note that in principle, in the current experiment readers could resort to the impersonal null subject reading in the Match condition as well, and posit an object gap in order to quickly locate all obligatory arguments of the verb. However, it seems that this strategy was not adopted by the readers in the Match condition, at least not as consistently as in the Mismatch condition. Therefore, it is unlikely that an OR analysis is categorically adopted when the impersonal form is available.

An alternative explanation for the contrast between the results of Experiment 2A and those of Experiment 1 could attribute it to the distance between the filler and the verb. Recall that, in Experiments 1, a temporal or locative adjunct opened the relative clause and served as a buffer

between the filler and the verb. This was done in order to somewhat increase the probability of a post-verbal subject structure (Shlonsky & Doron, 1992). However, it might have also contributed to the decay of the filler's representation and to an increase in the uncertainty regarding its agreement features by the time the verb is processed. The sentences in Experiment 2A lacked this adjunct "buffer" and thus, reliance on the filler's features during dependency formation might have been greater in this experiment. By this logic, agreement features are dismissed when the distance between the filler and the verb is long (as in Experiment 1) but can be weighed higher in shorter dependencies (Experiment 2A). In order to test this hypothesis, we conducted Experiment 2B.

3.2 Experiment 2B

To rule out the possibility that the contrast between Experiments 1 and 2A can be reduced to dependency length, we ran a modified version of Experiment 2A. In Experiment 2B, an adjunct phrase, of the same length as those in Experiments 1A and 1B, was inserted at the beginning of the relative clause.

3.2.1 Methods

Participants

Participants were 51 native speakers of Hebrew (according to self-report) (mean age: 25.96, range: 20-35). One participant was a bilingual speaker of Hebrew and English, and the rest were monolingual. Participants gave informed consent and received payment (20 NIS, approximately 5 US dollars) for their participation. The experiment was approved by the Ethics Committee in Tel Aviv University.

Materials

Materials were based on the 24 sets of Experiment 2A. We added a two-word adjunct at the beginning of the relative clause (in all conditions) and removed the condition of the "standard" filled-gap effect from the design (see example set in Table 3). Filler items were also identical to those used in Experiment 2A, except for a similar adjunct insertion in ten of the sentences.

Condition	Sentence
Match	<i>ha-saparit dibra im ha-me'acvim še-be-tahalix ha-sinun</i> the-hairdresser talked with the-designers that-in-process the-selection <i>baxru larov et ha-dugmaniyot ha-rašiyot šelahem ba-</i>

	<p>chose-3.PL mostly ACC the-models the-main of.them in.the- <i>rega ha-axaron.</i> moment the-last.</p> <p>'The hairdresser talked with the designers who usually chose their leading models at the last moment in the selection process.'</p>
Mismatch	<p><i>ha-saparit dibra im ha-me'acev še-be-tahalix ha-sinun</i> the-hairdresser talked with the-designer that-in-process the-selection <i>baxru-3.PL larov et ha-dugmaniyot ha-rašiyot šelo ba-rega</i> chose-3.PL mostly ACC the-models the-main of.him in.the-moment <i>ha-axaron.</i> the-last.</p> <p>'The hairdresser talked with the designer whose leading models were usually chosen at the last moment in the selection process'</p>
Baseline	<p><i>ha-saparit dibra im ha-me'acvim mikeyvan še-be-tahalix</i> the-hairdresser talked with the-designers because that-in-process <i>ha-sinun hem baxru larov et ha-dugmaniyot ha-rašiyot</i> the-selection they chose-3.PL mostly ACC the-models the-main <i>šelahem ba-rega ha-axaron.</i> of.them in.the-moment the-last.</p> <p>'The hairdresser talked with the designers because they usually chose their leading models at the last moment'</p>

Table 3. Example set from Experiment 2B.

The critical NP is marked in bold; ACC = accusative case marker. PL = plural.

Procedure

The procedure was the same as in the previous experiments.

Data analysis

Data analysis followed the same protocol as in Experiment 2A. Exclusion criteria resulted in removal of six participants (four for low accuracy in comprehension questions and two for abnormal reading time averages). For the remaining 45 participants, trimming of high and low RTs affected 2.57% and 0.02% of the data, correspondingly.

3.2.2 Results

Word-by-word reading times in the different experimental conditions and by-condition numerical means of the tested regions are presented in Figure 5.

Figure 5. Word-by-word RT means (msec) by condition, Experiment 2B. Error bars represent $\pm$ SE.

At the critical region (the direct object position), we observed increased RTs in the Mismatch condition relative to the baseline condition: the posterior distribution of this contrast had a high mean of 65ms with CrI that were far from zero [28, 105]. Accordingly, over 99% of the sampled posterior was above zero. As in Experiment 2A, we failed to detect a contrast between the Match and Baseline conditions. Although 90% of the posterior distribution was above zero, this possible contrast (with a posterior mean of 16ms and CrI of [-9, 41]) was again considerably smaller than that of the Mismatch condition (see Figure 6). Finally, at the spillover region, the posterior distribution has considerable uncertainty on both positive and negative directions, and thus was inconclusive as for any contrasts.

Figure 6. The sampled posterior distributions of the pairwise contrasts in Experiments 2A-B, on the log-scale (with 95% CrI shaded in gray).

3.2.3 Discussion

Experiment 2B replicated the findings of Experiment 2A. We observed increased reading times at the object position (i.e. a "filled-gap" effect) in the Mismatch condition, in line with our prediction. This suggests that the impersonal OR reading is adopted even with prolonged filler-verb distance. In addition, we did not observe a corresponding filled-gap effect when the filler matched the verb's features. This suggests that the parser does not prioritize the OR reading whenever the impersonal null subject is possible, and that the balance between the SR and OR analyses depends also on the mismatch in agreement features.

It might be noted that the filled-gap effect in this experiment produced a numerically smaller estimate relative to that of Experiment 2A (see Figure 6). However, since there is also high overlap between the posterior distributions, this difference is not reliable. Yet, if we are to speculate on the source of this difference, it can be due to lower weighing of agreement information when the distance from the filler is greater. Namely, the certainty regarding the filler's features may decrease with distance, and the readers may dismiss the mismatch more easily.

Another possibly interesting pattern suggested by this data is presented by the considerably larger effects we observed in Experiments 2A-B relative to Experiments 1A-B (see Figures 3 and 6). Reanalysis of at the direct object in Experiment 2 caused a more prominent increase in reading times relative to the reanalysis at the post-verbal subject in Experiment 1. This difference in the magnitude of the effects might suggest that reanalysis of a subject gap is less costly than that of an object gap. It is known that filled-gap effects are harder to find in subject position (Stowe, 1986; Lee, 2004). However, this was previously attributed to the short distance between the initiation of the "active filler" strategy and the filled-gap position, or to the absence of a verb with which the subject gap or argument is integrated. In our case, neither could underlie the contrast in effect magnitude (in all experiments the reanalysis occurs after the verb). It could be that subject gaps are reanalyzed more frequently in everyday communication and thus their reanalysis creates less processing disruption.

Overall, the replication of Experiment 2A suggests that the contrast in readers' interpretation of agreement information in Experiments 1 and 2 cannot be reduced to distance effects only. We therefore interpret the results obtained in this experiment as evidence that the difference in the prior probability of the faithful OR alternative is the most likely source of the contrast between readers' strategies in Experiments 1 and 2. To provide additional support for this claim, in the next section we present our final experiment, which compares the two structures within one experimental design.

4. Experiment 3

In Experiment 3, we wanted to test our interpretation of the results from the previous experiments, and of the contrast between them, within one experiment. We chose to use a different methodology – a sentence completion task – in order to tap into readers' preferences more directly. Experiments 1-2 deduce the preferred parse from reanalysis costs when the sentence disambiguates to the supposedly dispreferred interpretation. Since this is a rather indirect measure, we wanted to corroborate these findings with a task that allows us to get the participants' preferred interpretation without assuming it in the experimental design.

We conducted a sentence completion experiment in which participants were asked to complete a preamble that was presented word-by-word in rapid serial visual presentation. As the preamble, we used sentences similar to those used in Experiments 1-2, truncated after the relative clause verb. The experimental design included four conditions, crossing the number agreement on the filler and the verb (thus including singular verbs as in Experiments 1 and plural verbs as in Experiments 2). We hypothesized that the rate of SR completions will be

lower when the filler-verb mismatch involves a plural verb, thus allowing an impersonal null subject OR, relatively to a mismatch with a singular verb, where the only grammatical and faithful parse involves the rare construction of a post-verbal subject.

Additionally, we manipulated presentation speed of the words in the preamble, to test the effect of perceptual uncertainty. Increasing the perceptual uncertainty by presenting individual words for shorter duration can increase noisy-channel inferences. With short exposure, readers might be less certain of their representation of the filler or the verb, and thus might be more lenient towards compromising agreement. We thus hypothesized that SR completions will be more common in the mismatch conditions when participants read the preamble through a noisier channel – that is with shorter exposure to the individual words.

4.1 Methods

Participants

Participants were 132 native speakers of Hebrew (according to self-report) (mean age: 25.2, range: 18-44). Thirty-two participants were bilingual of Hebrew and either Russian, English, Spanish, French, Farsi or Hindi, and the rest were monolingual. Participants gave informed consent and received payment (20 NIS, approximately 5 US dollars) for their participation. The experiment was approved by the Ethics Committee in Tel Aviv University.

Materials

The experiment included 24 experimental items of four conditions, crossing the form of the verb (singular vs. plural) and the agreement of the filler (matching or mismatching the verb). The sentences presented a beginning of a relative clause and were truncated after the embedded verb (see example set in Table 4). The materials included a clause initial adjunct, as a buffer between the filler and the verb (on a par with Experiments 1A-B and 2B). The experimental items were distributed in a Latin square design along with 12 OR filler sentences (to balance a possible bias in favor of SRs in our materials, originating in the Match conditions).

Condition	Sentence
Match -Verb.SG	<i>xibavnu et ha-talmida še-lamrot ha-xašašot mac'a ...</i> we.liked ACC the-pupil.SG that-despite the-concerns found.SG ...
Match -Verb.PL	<i>xibavnu et ha-talmidot še-lamrot ha-xašašot mac'u ...</i> we.liked ACC the-pupils.PL that-despite the-concerns found.PL ...

Mismatch -Verb.SG	<i>xibavnu et ha-talmidot še-lamrot ha-xašašot mac'a ...</i> we.liked ACC the-pupils.PL that-despite the-concerns found.SG ...
Mismatch -Verb.PL	<i>xibavnu et ha-talmida še-lamrot ha-xašašot mac'u ...</i> we.liked ACC the-pupil.SG that-despite the-concerns found.PL ...

Table 4. Example set from Experiment 3.

Procedure

The experiment was built using Ibex Farm. The preamble of the sentences was presented using rapid visual serial presentation (RSVP) at a rate of 400ms per word. To test the effect of perceptual uncertainty, we manipulated the duration of word presentation, such that items were presented either for 400ms with no inter-stimulus interval, or for 250ms with a 150ms inter-stimulus interval (maintaining the same rate of 400ms per word). Word duration was manipulated within-participant and counterbalanced in a Latin Square design

After the end of the preamble, participants were instructed to complete the sentence fragment in a designated text box. Before they began the experiment, participants underwent a practice block of four items. The order of presentation was randomized for each participant.

Data analysis

Data from 15 participants were excluded as they did not understand the nature of the task (repeating the sentence fragment instead of completing it). Productions of the remaining 117 participants were coded for SR vs. OR completions. For example, a sentence like (9a) would be considered a SR, as it includes indication (via an accusative marked NP) that the object position does not contain a gap. On the other hand, a completion like (9b) is a OR production since it presents a post-verbal subject (see Appendix C for more details regarding the coding scheme).

- (9) *xibavnu et ha-talmidot še-mac'a ...*
we.liked ACC the-students.PL that-found.SG ...
- a. *et ha-xatul bari ve-shalem.*
ACC the-cat healthy and-whole
'We liked the students who found the cat in a good shape.'
- b. *ha-madrixa be-kce ha-maslul.*
the-guide in-edge the-trail.
'We liked the students who the guide found at the end of the trail.'

The data were analyzed in R using a logistic mixed-effects model. To test the effect of verb form (i.e. the possibility of a null subject) on the rate of SR productions we used treatment (dummy) coding for this factor, thus obtaining a measure which compares the singular and plural verb conditions when the filler and the verb matched (regardless of the effects in the Mismatch conditions). Similarly, to test the base rate of the agreement mismatch effect, the agreement match factor was also treatment (dummy) coded, such that the effect represents the contrast between the Match and Mismatch conditions of the singular verb version (independent of the effect of the impersonal null subject option which intervene in the plural verb conditions). Finally, we used sum coding of the word duration factor in order to average the above-mentioned effects over both presentation rates. We conducted a Bayesian analysis using weakly informative priors. The modeling decisions and parameters were on par with the analyses of the previous experiments.

4.2 Results

The distribution of productions across conditions is presented in Figure 7 and 8.

Figure 7. Rate of subject relative completions by condition, Experiment 3.

Figure 8. Overall completion proportions, in Experiment 3 (collapsing over presentation type). Note. OR completion include various types of completions where a simple gap cannot be assumed at the subject position. They are not exclusively limited to direct object extraction and include various gap/resumption positions (e.g. within a preposition, at an embedded subject position, etc).

The results are consistent with an effect of verb form, such that the usage of a plural verb (i.e. allowing impersonals) resulted in a decreased rate of SR productions (posterior mean [CrI]: -1.65 [-2.21, -1.08]). The results are also consistent with an effect of agreement, such that mismatch between the verb and the filler resulted in a decreased rate of SR productions (posterior mean [CrI]: -1.23 [-1.81, -0.65]). Crucially, we observed the predicted two-way interaction. The posterior is consistent with an interaction effect of verb form by filler matching (posterior mean [CrI]: -1.73 [-2.38, -1.09]). This interaction term suggests that the decrease in SR production rate due to agreement mismatch was more prominent when the verb was plural (allowing impersonals). The CrI of all three posteriors did not include zero, and accordingly more than 99% of the sample in these posterior distributions was beyond zero.

The results are also consistent with an interaction between verb form and word duration: 95% of the posterior sample was beyond zero, such that sentences with a plural verb were completed as SR in a higher rate in the speeded presentation conditions (posterior mean [CrI]: -0.33 [-0.73, 0.07]). The effect of word duration and its other interactions were inconclusive. Specifically, the analysis failed to detect the interactions which would indicate an effect of perceptual uncertainty on a noisy SR interpretation of the agreement mismatch. The predicted three-way interaction had only 79% of its posterior distribution beyond zero (posterior mean [CrI]: 0.22 [-0.29, 0.73]), and for the interaction term of word duration and agreement matching only 84% of the posterior distribution beyond zero (posterior mean [CrI]: 0.22 [-0.29, 0.73]).

4.3 Discussion

In this experiment we tested the contrast between the two structures of Experiments 1 and 2 within one experimental design, which probes readers' preferred interpretation more directly, with a sentence completion task. The results suggest that the rate of SR interpretations depends on the availability of the impersonal null subject reading. The effect of verb form suggests an overall decrease in SR completion rate when the verb licenses an impersonal null subject (i.e. when the verb is in its plural form). In addition, we also observed an overall decrease in SR rate for the Mismatch conditions, suggesting that, to some extent, readers use the mismatch itself as evidence against the SR structure. Crucially, the interaction between verb form and agreement matching suggests that the combination of the Mismatch cue and the availability of

an impersonal null subject reading yields a change in strategy, over and above the effect of each of these factors independently.

Both factors, the match between the filler and the verb and the availability of the impersonal reading under the verb form, conspire to increase the probability of an OR structure. Their interaction suggests the way readers interpret the agreement mismatch depends on the prior probability of the grammatical faithful alternative. Specifically, agreement is more likely to be compromised in favor of a SR analysis when the faithful alternative is one of low prior probability, like the post-verbal subject alternative (i.e. with singular verbs), than they are when there is an additional moderately frequent option of impersonal null subject (i.e. with plural verbs). Namely, a corrupted SR analysis is more likely than an OR with a post-verbal subject, but is still less likely than that of an OR with an impersonal null subject.

The presentation rate manipulation did not produce the predicted pattern of results. Since the noisy-channel inference (balancing prior probability and comprising agreement features) is expressed in the current design as an interaction between verb form and agreement mismatch (discussed above), an effect of perceptual uncertainty on the prevalence of noisy-channel inferences should be expressed as a three-way interaction in the data. However, we only observed two way interaction, such that the rate of SR productions increases in conditions with plural verbs and speeded presentation. This two-way interaction cannot be attributed to noise compensation as it affected match and mismatch conditions similarly (see Figure 7). The lack of such a three-way interaction could be due to the low sensitivity of the manipulation. Namely, it could be that the speeded condition was not sufficiently difficult and did not increase readers' uncertainty enough to allow detection of a contrast. Alternatively, the lack of effects could be attributed to other types of uncertainty that speeded presentation might give rise to. It could be that in the speeded presentation, readers are not only uncertain about the perceived word form, but also incorporate higher likelihood of missing a whole word. In the current setting, these two types of uncertainty may pull the interpretation in different directions. If readers assume that they might be missing some of the words, they may interpret the mismatch between the filler and the verb as indication that they missed a word corresponding to the subject. Therefore, while the uncertainty regarding the perceived words may increase the rate of SR completions, the likelihood of missing a whole word may increase the rate of OR completions.

5. General discussion

5.1 Summary

In this study, we investigated the processing of Hebrew relative clauses that are ambiguous between subject relatives (SR) with an agreement error, and input-faithful, grammatical, but less common, object relative (OR) structures. We found that the agreement mismatch between the filler and the verb does not entirely rule out the possibility of a subject gap for comprehenders, but is treated differently in different environments. In Experiment 1, we observed reanalysis costs at the true (post-verbal) subject in the relative clause, suggesting that readers had preferred constructing a SR, dismissing the agreement mismatch, over adopting a rare OR structure with a post-verbal subject. However, in Experiment 2, when another, more common, OR reading was grammatically possible (an impersonal null subject clause), we observed reanalysis costs at the object position of the relative clause, suggesting that an OR was constructed. Importantly, we observed all the effects while using only globally grammatical sentences (in both experimental and filler items), suggesting that the consideration of a possibility of error is not a task-dependent strategy. Finally, Experiment 3 corroborated these findings in a sentence completion task. We observed mostly SR completions when only the rare OR structure (with a post-verbal subject) could restore grammaticality of the sentence, but mostly OR completions when the possibility of an impersonal null subject was introduced.

5.2 The effect of prior probability on the interpretation of agreement mismatches

Human communication is frequently corrupted by noise, originating either in the production system of the interlocutor or in perception or memory errors of the addressee. Yet, addressees often manage to recover the intended meaning from the noisy input with little effort. Inspired by recent models of noise inferences and uncertainty maintenance in sentence processing (Levy, 2008b; Gibson et al., 2013), our study provides evidence for rational noisy-channel inference during incremental processing.

First, we demonstrate the readers treat a string which has only one grammatical analysis as ambiguous. Namely, we find that readers consider ungrammatical representations (in this case mismatching subject-verb agreement) during online processing, even though a semantically plausible and grammatically pristine analysis is available. Second, we show that assuming a corrupted structural dependency depends on the prior probability of the faithful alternative. Readers adopted a corrupted analysis only when the grammatical/fairful one was based on a very infrequent phenomenon in the language (the inverted word order in Experiments 1A-B). When another, relatively probable, structure was grammatically possible (Experiment 2A-B),

the strategy shifted, and interpretation was pulled towards the faithful analysis. This suggests that the above strategy is aimed for forming more probable structures and does not represent a categorical negligence. We would like to stress that the idea of dismissing agreement information, based on rational noisy-channel processing, should not be taken to suggest that agreement is not computed or that the mismatch is not detected. On the contrary, we suggest that although the parser identifies the mismatch, low probability of the grammatical alternative can push the parser to the SR reading.

Overall, our results suggest that readers apply elaborate probabilistic knowledge regarding the distribution of structures in their language during online processing. Thus, the frequency of the alternative syntactic structures affects how readers interpret bottom-up cues like an agreement mismatch and the adoption of noisy, unfaithful, interpretations. We demonstrate this in Hebrew, a language which has not yet been studied in this context. Previous studies of noisy-channel processing have focused on English, which is more limited in terms of manipulating word order, morphological marking and orthographic neighborhood sizes.

5.3 Alternative accounts

Here we address several possible alternatives to our account of the findings, and point out why we believe they cannot explain the entirety of the results we observed, unlike the noisy-channel interpretation, and specifically the hypothesis that the prior probability of alternative analyses affects the interpretation of agreement mismatches.

The costs of filler-gap dependency processing and the SR/OR asymmetry

One alternative account for our results would suggest that they can be explained based only on a general SR preference, or on the costs of infrequent structures. It is widely accepted that the probability of an utterance affects its processing difficulty (e.g. Levy, 2008; Smith & Levy, 2013), and that SR structures are more probable and simpler to process than ORs (e.g. Gennari & MacDonald, 2008; Grodner & Gibson, 2005; King & Just, 1991; Staub, 2010; Traxler, Morris, & Seely, 2002). These assumptions predict that a SR analysis would be preferred over an OR analysis, in line with our results from Experiments 1. Yet they do not account for the contrast between Experiments 1 and 2, or between the singular verb conditions and the plural verb conditions in Experiment 3. In Experiments 1, the agreement mismatch was indeed dismissed in favor of an SR analysis; However, in Experiments 2 the mismatch was taken as evidence for an OR analysis with a null subject, which is still less probable than a SR.

Therefore, the account must include an interaction between the dismissal of agreement and the probability of the alternative structure.

A similar line of thought could suggest that the contrasts which we rely on, namely those between the processing of relative clauses and that of baseline sentences with no dependency, are confounded by the general costs of dependency processing. However, once again this does not seem likely given the results of Experiment 2, where reading times in the match condition aligned with reading times in the baseline condition, despite the existence of a dependency in the former. We maintain that the critical region is distant enough from the initiation of the dependency, and reading times of this region are therefore less affected by general dependency processing costs and more likely to reflect the abandoning of a gap that was previously postulated.

Likelihood of edits

The contrast between the results of Experiments 1 and 2 can, to some extent, be attributed to the likelihood of the corruption, since edits of different types are required for different agreement errors, and some may be more likely than others. In particular, deletion errors are more likely than insertion errors, according to the Bayesian “size principle” (Xu & Tenenbaum, 2007).¹⁰ In line with this, it was observed that readers are more likely to adopt the “corrupted” analysis when it involves assuming deletion rather than the insertion (Gibson et al., 2013; Poppels & Levy, 2016; Ryskin, Futrell, Kiran & Gibson, 2018).

In the context of our study, it should be noted that in Experiments 1A-B, the corrupted SR analysis requires assuming either an insertion error at the filler (the intended filler was singular, and it was corrupted to produce a plural form by adding the plural suffix) or a deletion error at the verb (the intended verb was plural, and the singular form is the result of deletion of the plural morpheme). In contrast, in Experiments 2A-B, the corrupted SR analysis requires assuming the opposite (a deletion error at the filler, editing it from plural to singular, or an insertion error at the verb, editing it from singular to plural).

If one assumes that readers consistently edit the verb’s features, the deletion/insertion contrast may confound our results, as it pulls towards a SR interpretation to a greater extent in Experiments 1A-B (in line with our predictions). However, we believe that this interpretation

¹⁰ A deletion only requires a particular word to be randomly selected from the set of words in the sentence, whereas an insertion requires the selection of a specific word from the producer’s vocabulary. Since vocabulary size is considerably larger than the number of words in a sentence, the Bayesian “size principle” suggests that insertion of a specific word has smaller likelihood than deletion of a specific word.

of our results is not likely, for several reasons. First, we suggest that editing of the filler's features, rather than the verb's, is more likely. Given the non-cumulative presentation of the sentence, we believe that readers maintained higher uncertainty regarding the filler's features (rather than the verb's), as it was no longer available for re-reading at the point of mismatch detection, that is upon encountering the verb. Therefore, readers should be more likely to amend the filler's representation than that of the verb. The likelihood of deletion and insertion errors thus cannot account for our findings. If the filler's features are edited, errors are likely in Experiments 2A-B, thus pulling the results in the direction opposite to our findings and predictions.

Moreover, in Experiment 3 we aimed to better control for this possible confound by using feminine forms only, for which editing of the verb's features – either from singular to plural or from plural to singular - would always require readers to assume a substitution error. This eliminates the possible effect of the deletion/insertion contrast on the analysis of the verb. The fact that the results of Experiments 1 and 2 were still mirrored in Experiment 3 supports our conclusion that the different types of edits cannot account for our findings.

Lastly, the role of the likelihood of insertion vs. deletion errors might not be relevant at all in the case of morphological inflection errors like agreement mismatch.¹¹ The applicability of the size principle depends on the assumed noise model. For error like those we assume in this paper, readers' noise model might involve the probability of selecting from the wrong slot in the morphological paradigm. Under such a (very plausible) noise model, separate letter-deletion versus letter-insertion error types should not come into play. If this type of noise model is indeed the one used by readers, the contrast between the SR/OR preferences in Experiments 1 and 2 could not be attributed to insertion versus deletion of a verbal morpheme.

Avoidance of predicting new material

A different alternative account for our findings may attribute them to a general preference for avoiding prediction of a new discourse referent. In Experiments 1, in order to construct an OR interpretation at the verb, the reader has to predict an upcoming new discourse entity, but this is not the case in Experiments 2. This would produce the higher tendency to adopt SR analyses in Experiments 1 relative to 2.

However, this interpretation is not consistent with results from the match condition of Experiments 2 and 3. In the match condition with a plural filler and a plural verb, namely when

¹¹ We thank Roger Levy for pointing this out to us.

an impersonal null subject is possible, full thematic assignment without need to predict any new discourse entities can be established immediately by assuming an OR reading, assigning the agent role of the verb to the null generic subject, and its theme role to the gap. If readers refrain from predicting new discourse reference, an OR should be consistently constructed in these cases, to avoid prediction of a yet-unknown complement to the verb. However, we did reliably observe adoption of OR analyses in the match condition of Experiments 2A-B. On the contrary, in Experiment 3, we observed mostly SR completions in that condition. Therefore, we argue that this account is not consistent with the overall pattern of results.

Mixed predictions through parallel or stochastic serial processing

Experiment 1 exhibited processing costs of a failed SR prediction, manifested as elevated RTs on a post-verbal subject. Experiment 2 exhibited processing costs of a failed OR prediction, manifested as elevated RTs on a lexical direct object. However, in principle, these results do not necessarily show that readers uniformly predict an SR when the relative head is plural and the verb is singular, and conversely predict an OR when the relative head is singular and the verb is plural. Rather, it is possible that comprehenders entertain a mixture of interpretations including some portion of noisy-channel SR prediction and some portion of input-faithful OR prediction, in both configurations. Such a mixed prediction could be implemented either in parallel processing (Levy, 2008a), where multiple predictions are entertained at once, or in stochastic serial processing (van Gompel, Pickering, & Traxler, 2000), where each trial could reflect a different prediction and the overall proportion of executing one prediction over the other reflects the probabilistic bias. In such cases of a mixed predictions, we should observe some costs for any disambiguation, whether it aligns with a SR reading, an OR reading or neither (relative to an unambiguous baseline). Thus, it could be claimed that comparable mixed predictions cause the processing costs observed in both Experiment 1 and 2, rather than a differential preference for the SR or OR readings (when the alternative is a post-verbal subject or an impersonal, correspondingly).

We believe that a couple of arguments can be made against such interpretation of our data. First, under this interpretation, the mismatch conditions, which lead to such mixed predictions, should exhibit more moderate processing costs relatively to cases where readers entertain a "purer" set of predictions, with clear, more uniform preference for the ultimately unrealized interpretation. In Experiment 1, the Match condition reflects such a case, where a SR should clearly be predicted. In Experiment 2, we had a condition with a classic "filled-gap" effect, which can be viewed similarly, i.e. a case where an OR should be formed. Yet, in both cases

we did not see a clear difference in the effect size between these "uniform prediction" conditions and the mismatch conditions. Moreover, the results of Experiment 3 point toward an evident contrast in the rate of SR and OR preferences for these mismatching relative clauses, depending on the availability of the impersonal form.

Good-enough processing

So far we considered the predictions of "good-enough" processing in the context of incremental heuristic processing. However, current versions of the "good-enough" processing approach suggest that incremental processing adheres to fully grammatical analyses of the input (Slattery et al., 2013), but that an offline interpretation corresponding may be perturbed, possibly due to later memory fallibility (Huang & Ferreira, 2020). Under such an approach, competing analyses, which were disproved during parsing, resurface in late comprehension measures due to difficulty in their inhibition. To the best of our knowledge, this approach does not make clear predictions with regard to reading times measures, as in Experiments 1-2. The grammatical parsing stage, which takes place before memory failure may obscure the correct interpretation, could conform to noisy-channel inference, or to assume strictly grammatical and faithful representations. Therefore, the mechanisms for online processing which the current study argues for are independent of the source of lingering offline misinterpretations.

In addition, we believe that "good-enough" processing cannot offer a full account of the findings in Experiment 3. While "good-enough" processing predicts SR misinterpretations such as those which arise in Experiment 3 (and in previous research on relative clauses in Dutch, Frazier, 1987a), memory fallibility cannot straightforwardly account for difference between the rate of these SR interpretations depending to the available OR alternative. We assume that in all cases, a SR is initially favoured (appears). As a lack of proper inhibition of this SR interpretation in memory should not by itself depend on the frequency of the OR alternative, a more complex process (e.g. noisy-channel inference) should arise at the verb to account for the full pattern.

5.4 Conclusions

The results of our study suggest that the prior probability of alternative analyses modulates the interpretation of agreement. We show that during online processing, readers apply elaborate knowledge regarding the distribution of structures in their language, and they are willing to compromise subject-verb agreement to refrain from (grammatical but) highly improbable structures. We propose (i) that incoming input is integrated with existing knowledge about the

probability of various linguistic structures, pulling the interpretation towards the more probable structure; and (ii) that in line with rational noisy-channel models, the bias towards more probable structures applies not only for ambiguous strings (local or global), but also for cases where only one grammatical reading is available, based on consideration of possible production and perception errors.

Acknowledgements

We would like to thank Roger Levy, Rachel Ryskin, and Jesse Harris for helpful comments and suggestions, as well as the audience at the CUNY 2019 Conference on Human Sentence Processing (Boulder, CO). We also thank Hila Davidovitch for her help with coding the results of Experiment 3, and Radan Nasrallah for her help with collecting data for Experiments 1 and 2.

This work was supported by an Israel Science Foundation grant (number 216/18) and by a German-Israeli Foundation grant (number I-2496-104.4/2018) to Aya Meltzer-Asscher.

Appendix A: Frequentist analyses results

Apart from the Bayesian analyses which appear in the main text, we also performed a frequentist analysis (details and code of the frequentist analysis procedure are available on [osf](#)). Results align with the effects deduced from the Bayesian analysis and are available in the following tables.

	<u>Experiment 1A</u>		<u>Experiment 1B</u>	
	<i>t</i>	<i>p</i>	<i>t</i>	<i>P</i>
<u>Main effects and interactions:</u>				
Main effect of Match - Baseline	3.70	< .001	2.22	.028
Main effect of Mismatch - Baseline	3.85	< .001	2.83	.005
Main effect of Mismatch2 - Baseline	-	-	2.55	.011
Effect of Region	0.48	.631	-2.20	.028
Interaction Match:Region	-2.70	.011	-2.27	.024
Interaction Mismatch:Region	0.62	.538	0.40	.689
Interaction Mismatch2:Region	-	-	-0.37	.715
<u>Pairwise comparisons:</u>				
Match - Baseline, critical region	0.53	> .99	-0.06	> .99
Mismatch - Baseline, critical region	2.72	.026	2.25	.145
Mismatch2 - Baseline, critical region	-	-	1.54	.735
Match - Baseline, spillover region	4.36	< .001	3.16	.009
Mismatch - Baseline, spillover region	1.93	.214	1.68	.559
Mismatch2 - Baseline, spillover region	-	-	2.06	.235

Table A1. Results of the frequentist analysis of Experiment 1: *t*-value and *p*-values (corrected, using Bonferroni correction for four and six comparisons in Exp. 1A and 1B, correspondingly).

	Estimate	SE	<i>t</i>	<i>P</i>
Effect of verb form	-1.950	0.204	-9.57	< .001

Effect of agreement match	-1.731	0.206	-8.41	< .001
Effect of word duration	0.260	0.181	1.44	.151
Interaction of verb form and agreement match	-0.536	0.243	-2.21	.027
Interaction of verb form and word duration	-0.399	0.203	-1.96	.050
Interaction of agreement match and word duration	-0.238	0.205	-1.16	.247
Three-way interaction of verb form, agreement match and word duration	0.313	0.241	1.29	.196

Table A2. Results of the frequentist analysis of Experiment 3: Estimate, standard-error, t -value and p -values.

Appendix B: Results of by-region models for Experiment 1A

It should be noted that Experiment 1A was initially analyzed with two by-region models (one for the critical word and one for the post-critical one). We subsequently explored the interaction between region and condition and conducted the unified analysis (with data from both regions) as reported in the main text. Here we report the results of the two by-region models, which are broadly in line with those of the unified model. In addition to the analyses appearing on the main text we obtained similar effects in a frequentist analysis.

	Posterior mean [CrI]	Pr($\beta > 0$)	<i>t</i>	<i>P</i>
Model 1, critical region:				
Match – Baseline	-7 [-24, 11]	79%	0.62	> .99
Mismatch – Baseline	17 [-3, 38]	95%	2.91	.010
Model 2, spillover region:				
Match – Baseline	28 [11, 46]	> 99%	0.62	> .99
Mismatch – Baseline	11 [-5, 27]	91%	2.91	.010

Table B1. Results of by-region (Bayesian and frequentist) analyses of Experiment 1A: Posterior means of the tested contrasts in milliseconds (with 95% Bayesian credible intervals) and the probability of the parameter being beyond zero, *t*-value and *p*-values (corrected, within each model, using Bonferroni correction for two comparisons).

Appendix C: Coding criteria used in the analysis of Experiment 3

In the coding of production data from Experiment 3, sentences were considered SR completions if they included one of the following:

1. A (definite) NP marked for accusative case (82% of SR completions).
2. An (indefinite) NP which was unmarked for case but mismatched the verb (11% of SR completions).
3. A verbal argument of another category (e.g. a PP) which did not include a resumptive pronoun (obligatory in indirect objects in Hebrew) and was not grammatically compatible with a direct object gap (3% of SR completions).

A sentence was considered as an OR completion if:

1. It included a (definite) NP marked for nominative case (22% of OR completions).
2. It was missing an obligatory direct object (24% of OR completions).
3. It included a corresponding resumptive pronoun which was grammatically inconsistent with another gap in subject position (53% of OR completions).

The full completion dataset (in Hebrew) with coding is available at <https://osf.io/v6y8n/>.

We aimed to refrain from assuming typing errors of the participants (other than those induced by the manipulation) in order to avoid presuming the likelihood of different readings. Therefore, completions which did not match these criteria were coded as ambiguous productions. This included sentences with an ambiguous structure (e.g. with indefinite noun phrases, unmarked for case, that match the verb; with resumptive pronouns in possessive positions, allowing an additional subject gap), and completions which were partial or contained errors unrelated to the filler's agreement features. For grammatically ambiguous productions, an exception was made when the three coders judged that the production was semantically plausible only in one of the readings (these cases constituted 4% from the SR coded data and less than 1% of the OR coded data).

The analysis counted ambiguous productions as non-SR productions along with the OR ones, to keep conservative estimates of SR productions and to allow variance in the Match-Verb.SG condition (where OR productions were extremely sparse). Yet, a similar pattern of results can be observed when excluding the ambiguous responses (see Figure B1).

Figure C1. Rate of subject relative completions by condition in Experiment 3, when excluding all ambiguous and partial responses.

References

- Aoshima, S., Phillips C., & Weinberg A. (2004). Processing filler-gap dependencies in a head-final language. *Journal of Memory and Language*, 51(1), 23-54.
- Arnon, I. (2010). Rethinking child difficulty: The effect of NP type on children's processing of relative clauses in Hebrew. *Journal of Child Language*, 37(1), 27-57.
- Baird, R., & Koslick, J. D. (1974). Recall of grammatical relations within clause-containing sentences. *Journal of Psycholinguistic Research*, 3(2), 165-171.
- Berman, R. A. (1979). Form and function: passives, middles, and impersonals in Modern Hebrew. In *Annual Meeting of the Berkeley Linguistics Society* (Vol. 5, pp. 1-39).
- Berman, R. A. (1980). The case of an (S) VO language: Subjectless constructions in Modern Hebrew. *Language*, 759-776.
- Berman, R. A. (2011). Revisiting impersonal constructions in Modern Hebrew: Discourse-based perspectives. *Impersonal constructions: A cross-linguistic perspective*, 323-356.
- Bock, K., & Miller, C. A. (1991). Broken agreement. *Cognitive psychology*, 23(1), 45-93.
- Boland, J. E., Tanenhaus, M. K., Garnsey, S. M., & Carlson, G. N. (1995). Verb argument structure in parsing and interpretation: Evidence from wh-questions. *Journal of Memory and Language*, 34(6), 774-806.
- Bürkner P. C. (2017). brms: An R Package for Bayesian Multilevel Models using Stan. *Journal of Statistical Software*. 80(1), 1-28.
- Caplan, D., & Futter, C. (1986). Assignment of thematic roles to nouns in sentence comprehension by an agrammatic patient. *Brain and Language*, 27(1), 117-134.
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1), 1-32.
- Dank, M., & Deutsch, A. (2010). The role of morpho-phonological factors in subject–predicate gender agreement in Hebrew. *Language and cognitive processes*, 25(10), 1380-1410.
- Dank, M., Deutsch, A., & Bock, K. (2015). Resolving conflicts in natural and grammatical gender agreement: Evidence from eye movements. *Journal of Psycholinguistic Research*, 44(4), 435-467.
- De Villiers, J. G., Flusberg, H. B. T., Hakuta, K., & Cohen, M. (1979). Children's comprehension of relative clauses. *Journal of psycholinguistic Research*, 8(5), 499-518.
- Dekel, N. (2014). *Colloquial Israeli Hebrew: A corpus-based survey*. Berlin: De Gruyter Mouton.
- Deutsch, A., & Dank, M. (2009). Conflicting cues and competition between notional and grammatical factors in producing number and gender agreement: Evidence from Hebrew. *Journal of Memory and Language*, 60(1), 112-143.
- Deutsch, A., & Dank, M. (2011). Symmetric and asymmetric patterns of attraction errors in producing subject–predicate agreement in Hebrew: An issue of morphological structure. *Language and cognitive processes*, 26(1), 24-46.
- Diessel, H., & Tomasello, M. (2005). A new look at the acquisition of relative clauses. *Language*, 81(4), 882-906.

- Eberhard, K. M., Cutting, J. C., & Bock, K. (2005). Making syntax of sense: number agreement in sentence production. *Psychological Review*, 112(3), 531-559.
- Ferreira, F. (2003). The misinterpretation of noncanonical sentences. *Cognitive psychology*, 47(2), 164-203.
- Ferreira, F., Bailey, K. G., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science*, 11(1), 11-15.
- Ferreira, F., & Patson, N. D. (2007). The 'good enough' approach to language comprehension. *Language and Linguistics Compass*, 1(1-2), 71-83.
- Fox, B. A., & Thompson, S. A. (1990). A discourse explanation of the grammar of relative clauses in English conversation. *Language*, 297-316.
- Frazier, L. (1987a). Syntactic processing: Evidence from Dutch. *Natural Language and Linguistic Theory*, 5(4), 519-559.
- Frazier, L. (1987b). Syntactic processing: Evidence from Dutch. *Natural Language and Linguistic Theory*, 5(4), 519-559.
- Friedmann, N., & Novogrodsky, R. (2004). The acquisition of relative clause comprehension in Hebrew: A study of SLI and normal development. *Journal of Child language*, 31(3), 661.
- Friedmann, N., Belletti, A., & Rizzi, L. (2009). Relativized relatives: Types of intervention in the acquisition of A-bar dependencies. *Lingua*, 119(1), 67-88.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis* (Third ed.). Boca Raton, FL: Chapman and Hall/CRC Press.
- Gennari, S. P., & MacDonald, M. C. (2008). Semantic indeterminacy in object relative clauses. *Journal of memory and language*, 58(2), 161-187.
- Gennari, S. P., & MacDonald, M. C. (2009). Linking production and comprehension processes: The case of relative clauses. *Cognition*, 111(1), 1-23.
- Gibson, E. (1998). Linguistic complexity: Locality of syntactic dependencies. *Cognition*, 68(1), 1-76.
- Gibson, E. (2000). The dependency locality theory: A distance-based theory of linguistic complexity. In A. Marantz, Y. Miyashita, & W. O'Neil (Eds.), *Image, language, brain*, (pp. 95-126). Cambridge, MA: The MIT Press.
- Gibson, E., Bergen, L., & Piantadosi, S. T. (2013). Rational integration of noisy evidence and prior semantic expectations in sentence interpretation. *Proceedings of the National Academy of Sciences*, 110(20), 8051-8056.
- Gibson, E., Tan, C., Futrell, R., Mahowald, K., Konieczny, L., Hemforth, B., & Fedorenko, E. (2017). Don't underestimate the benefits of being misunderstood. *Psychological science*, 28(6), 703-712.
- Givón, T. (1976). On the VS word order in Israeli Hebrew: Pragmatics and typological change. In Cole P. (ed.) *Studies in modern Hebrew syntax and semantics: the transformational-generative approach* (pp. 153-183). Amsterdam: North-Holland Publishing Company.
- Gordon, P. C., Hendrick, R., & Johnson, M. (2004). Effects of noun phrase type on sentence complexity. *Journal of memory and Language*, 51(1), 97-114.

- Grodner, D., & Gibson, E. (2005). Consequences of the serial nature of linguistic input for sentential complexity. *Cognitive science*, 29(2), 261-290.
- Grodzinsky, Y. (1989). Agrammatic comprehension of relative clauses. *Brain and Language*, 37(3), 480-499.
- Hawkins, J. A. (1999). Processing complexity and filler-gap dependencies across grammars. *Language*, 75(2), 244-285.
- Israel's Central Bureau of Statistics (2013). Selected data from the 2011 social survey on mastery of the Hebrew language and usage of languages (in Hebrew). Retrieve from: https://www.cbs.gov.il/he/mediarelease/DocLib/2013/017/19_13_017b.pdf
- Jackson, C. N., Mormer, E., & Brehm, L. (2018). The production of subject-verb agreement among Swedish and Chinese second language speakers of English. *Studies in Second Language Acquisition*, 40(4), 907-921.
- Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, 98, 122-149.
- Keenan, E. L., & Comrie, B. (1977). Noun phrase accessibility and universal grammar. *Linguistic Inquiry*, 8(1), 63-99.
- King, J., & Just, M. A. (1991). Individual differences in syntactic processing: The role of working memory. *Journal of memory and language*, 30(5), 580-602.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: tests in linear mixed effects models. *Journal of Statistical Software*, 82(13).
- Lee, M. W. (2004). Another look at the role of empty categories in sentence processing (and grammar). *Journal of Psycholinguistic Research*, 33(1), 51-73.
- Levy R., (2008a). Expectation-Based Syntactic Comprehension. *Cognition*, 106(3), 1126-1177.
- Levy, R. (2008b). A noisy-channel model of rational human sentence comprehension under uncertain input. In *Proceedings of the conference on empirical methods in natural language processing* (pp. 234-243). Association for Computational Linguistics.
- Levy, R. (2011). Integrating surprisal and uncertain-input models in online sentence comprehension: formal techniques and empirical results. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1* (pp. 1055-1065). Association for Computational Linguistics.
- Levy, R., Bicknell, K., Slattery, T., & Rayner, K. (2009). Eye movement evidence that readers maintain and act on uncertainty about past linguistic input. *Proceedings of the National Academy of Sciences*, 106(50), 21086-21090.
- Levy, R., Fedorenko, E., & Gibson, E. (2013). The syntactic complexity of Russian relative clauses. *Journal of Memory and Language*, 69(4), 461-495.
- Lewandowski, D., Kurowicka, D., & Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of multivariate analysis*, 100(9), 1989-2001.
- Lewis, R. L., & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive science*, 29(3), 375-419.

- MacDonald, M. C., & Christiansen, M. H. (2002). Reassessing working memory: comment on Just and Carpenter (1992) and Waters and Caplan (1996). *Psychological Review*, 109(1), 35-54.
- MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological review*, 101(4), 676-703.
- Maschler, Y. (2015). Word order in time. In Deppermann, A. & Günthner, S. (eds.), *Temporality in interaction* (pp. 201-236). Amsterdam: John Benjamins Publishing Company.
- O'Grady, W. (1997). *Syntactic development*. Chicago: University of Chicago Press.
- Pearlmutter, N. J., Garnsey, S. M., & Bock, K. (1999). Agreement processes in sentence comprehension. *Journal of Memory and language*, 41(3), 427-456.
- Pickering, M. J., & Traxler, M. J. (1998). Plausibility and recovery from garden paths: An eye-tracking study. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(4), 940-961.
- Pickering, M. J., & Traxler, M. J. (2003). Evidence against the use of subcategorisation frequency in the processing of unbounded dependencies. *Language and Cognitive Processes*, 18(4), 469-503.
- Pizarro-Guevara, J. (2020). *When human universal meets language specifics*. Doctoral dissertation, University of California, Santa Cruz.
- Poppels, T., & Levy, R. (2016). Structure-sensitive Noise Inference: Comprehenders Expect Exchange Errors. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (eds.), *Proceedings of the 38th Annual Meeting of the Cognitive Science Society* (pp. 378–383). Austin, TX: Cognitive Science Society.
- Ravid, D., (1977). Aspects of constituent ordering in Modern Israeli Hebrew. *Hebrew Computational Linguistics Bulletin*, 11, 1-45. [In Hebrew]
- Real, F., & Christiansen, M. H. (2007). Processing of relative clauses is made easier by frequency of occurrence. *Journal of memory and language*, 57(1), 1-23.
- Roland, D., Dick, F., & Elman, J. L. (2007). Frequency of basic English grammatical structures: A corpus analysis. *Journal of memory and language*, 57(3), 348-379.
- Ryskin, R., Futrell, R., Kiran, S., & Gibson, E. (2018). Comprehenders model the nature of noise in the environment. *Cognition*, 181, 141-150.
- Schlesewsky, M., Fanselow, G., Kliegl, R., & Krems, J. (2000). The subject preference in the processing of locally ambiguous wh-questions in German. In B. Hemforth, L. & Konieczny (eds.), *German sentence processing* (pp. 65-93). Dordrecht: Kluwer.
- Shlonsky, U. (1990). Pro in Hebrew subject inversion. *Linguistic Inquiry*, 21(2), 263-275.
- Shlonsky, U. (1997). *Clause structure and word order in Hebrew and Arabic: An essay in comparative Semitic syntax*. Oxford: Oxford University Press.
- Shlonsky, U., & Doron, E. (1992). Verb second in Hebrew. In *Proceedings of the West Coast Conference on Formal Linguistics* (Vol. 10, pp. 431-446).
- Smith, N. J., & Levy, R. (2013). The effect of word predictability on reading time is logarithmic. *Cognition*, 128(3), 302-319.

- Spolsky, B., & Shohamy, E. G. (1999). *The languages of Israel: Policy, ideology, and practice*. Clevedon: Multilingual Matters LTD.
- Staub, A. (2010). Eye movements and processing difficulty in object relative clauses. *Cognition*, 116(1), 71-86.
- Tavakolian, S. (1981). The conjoined-clause analysis of relative clauses. In S. Tavakolian (Ed.), *Language acquisition and linguistic theory*, pp. 167-187.
- Taube, D. (2007). Impersonal and passive constructions in contemporary Hebrew. In Bar, T. & Cohen E. (eds.), *Studies in Semitic and General Linguistics in Honor of Gideon Goldenberg* (pp. 277-297). Münster: Ugarit-Verlag.
- Traxler, M. J., Morris, R. K., & Seely, R. E. (2002). Processing subject and object relative clauses: Evidence from eye movements. *Journal of Memory and Language*, 47(1), 69-90.
- Traxler, M. J., & Pickering, M. J. (1996). Plausibility and the processing of unbounded dependencies: An eye-tracking study. *Journal of Memory and Language*, 35(3), 454-475.
- Traxler, M. J., Williams, R. S., Blozis, S. A., & Morris, R. K. (2005). Working memory, animacy, and verb class in the processing of relative clauses. *Journal of Memory and Language*, 53(2), 204-224.
- Wagers, M. W., Borja, M. F., & Chung, S. (2018). Grammatical licensing and relative clause parsing in a flexible word-order language. *Cognition*, 178, 207-221.
- Wagers, M. W., Lau, E. F., & Phillips, C. (2009). Agreement attraction in comprehension: Representations and processes. *Journal of Memory and Language*, 61(2), 206-237.
- Wanner, E., & Maratsos, M. (1978). An ATN approach to comprehension. In M. Halle, J. Bresnan, & G. Miller (Eds.), *Linguistic theory and psychological reality*, pp. 119-161.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological review*, 114(2), 245-272.