

Thematic considerations in the processing of local ambiguities: Evidence from Hebrew

Lola Karsenti, Linguistics Department, Tel Aviv University, IL, lolakarsenti@mail.tau.ac.il

Aya Meltzer-Asscher, Linguistics Department, Tel Aviv University, IL; Sagol School of Neuroscience, Tel Aviv University, IL, ameltzer@tauex.tau.ac.il

In three forced-choice completion experiments in Hebrew, the present study investigates the degree to which attachment decisions in sentences with local ambiguity are driven by the goal to maximize interpretation during incremental processing. In particular, we examine a processing strategy aimed at thematic assignment maximization proposed by Pritchett (1988, 1992): at every point during processing, the Theta Criterion attempts to be satisfied, given the maximal thematic grid of the available verbs. We find a consistent preference for attaching a noun phrase to a preceding verb over attaching it as the subject of a yet-unmentioned verb, in accord with thematic assignment maximization and in line with previous observations. In contrast, when two possible verbs were available before the noun phrase, no consistent attachment preference was observed. This points to a prominent role for thematic assignment in ambiguity resolution. In addition, transitivity bias was found to affect processing choices in the latter case, but not in the former. Finally, we show that local thematic assignment maximization can even override global grammaticality.



1. Introduction

One of the central issues in the research on sentence processing has been parsing choices for locally structurally ambiguous input, as, for example, in (1). (1) illustrates the *object/subject ambiguity*, in which the underlined noun phrase (NP) *the sentence* can be attached in two ways: it can serve as the direct object of the verb in the adverbial clause, as illustrated in (1a); or it can be attached as the subject of the main clause, as illustrated in (1b). Accordingly, this NP is ambiguous with regard to which predicate it receives its thematic role from, i.e., the embedded verb or the main verb.¹

- (1) While Mary was reading the sentence ...
 a. ... she paid attention
 b. ... became clear

Numerous studies (e.g., Adams et al., 1998; Ferreira & Henderson, 1991; Frazier & Rayner, 1982; Staub, 2007) observed a Garden Path (GP) effect in sentences such as (1b), namely, slower reading times at the disambiguating main verb *became*, compared to an unambiguous baseline. This effect suggests that the ambiguous NP is initially attached at the direct object position of the subordinate verb, necessitating reanalysis upon the arrival of the main verb.

An early, influential suggestion made by Frazier (1978, 1987a) to account for the parsing of local ambiguities was that, at points of ambiguity, initial attachment choices are guided by syntactic considerations alone, with no role for lexical, semantic or contextual information. Parsing at this stage is determined by two purely structural principles: *Minimal Attachment*, namely, building the least complex structure possible; and *Late Closure*, namely, attaching incoming material to recent constituents. According to Minimal Attachment and Late Closure, in object/subject ambiguities, such as (1), the NP is attached as a direct object, as this obeys Minimal Attachment (gives rise to the simplest structure) and Late Closure (the NP is attached to a recent constituent). In this article, we do not discuss Minimal Attachment further, but rather focus on Late Closure, which is relevant for the present data.

An alternative to this strategy based on structural principles holds that the processor's aim is to satisfy grammatical relations in such a way as to achieve maximal interpretation during incremental processing (e.g., Crocker, 1992; Frazier & Clifton, 1996; Kamide et al., 2003; Pritchett,

¹ We adopt a view of thematic roles which is fairly standard in theoretical linguistics (e.g., Harley, 2010): thematic roles are verb-specific roles, which are part of the lexical information associated with the verb, and are assigned by it to phrases in the sentence. The roles include Agent, Theme, Goal, Experiencer and several others. They are different than syntactic roles (such as subject and object), in that they are determined not by syntactic position, but by the thematic content of the specific verb, and they carry some semantic information. For example, a subject can be an Agent, a Theme or an Experiencer, depending on the verb in the sentence. Notably, thematic roles do not carry *much* semantic content. For example, an argument which receives the Agent role is interpreted as the initiator of an action, and an argument bearing a Theme role is interpreted as the undergoer of an action, quite broadly.

1988, 1992; Sturt & Crocker, 1997). According to this view, attachment decisions should be defined over thematic structure, rather than constituent structure, with the aim of maximizing thematic role assignment. For instance, a widely observed phenomenon where comprehenders posit a gap for a fronted wh-phrase at the earliest available position (*What did you say that Mary was reading __?*) was first explained by a purely structural principle, i.e., the Active Filler Strategy (de Vincenzi, 1991; Frazier, 1987b). In contrast, Pritchett (1992) and Gibson, Hickok, and Schütze (1994) argued that no special strategy is required; instead, the need to assign thematic roles as soon as possible guides such filler-gap formation. While, in English, these two alternative proposals make similar predictions, Aoshima, Phillips, and Weinberg (2004) showed that the predictions differ in Japanese, a verb-final language. Aoshima et al. (2004) examined sentences with embedded clauses in Japanese, in which a fronted dative wh-phrase can, in principle, be associated either with a gap in an embedded object or main object position (e.g., *which student-dat class teacher-top (__gap1) [principal-nom (__gap2) book-acc read] library-at librarian-dat told*. ‘The class teacher told the librarian at the library which student the principal read a book for.’). The main gap position is the closest to the filler structurally, whereas the embedded position allows the earliest thematic role assignment. The authors found that Japanese speakers preferred to interpret the fronted wh-phrase as if it was displaced from the embedded clause, meaning that the processing of filler-gap dependencies is driven by the need to satisfy thematic requirements, rather than by purely structural principles.

In the current study, we further test the degree to which attachment decisions are guided by the pressure to maximize thematic relations and, hence, interpretation, during processing. In examining *thematic assignment maximization*, we use the implementation of this requirement formulated by Pritchett (1988, 1992). According to Pritchett, parsing performance crucially depends on grammatical competence, meaning that, at every point during processing, a global principle of grammar – the Theta Criterion – drives the parser, through its local application. Accordingly, he proposes the Theta Attachment principle, given in (2).

- (2) Theta Attachment: At every point during processing, the Theta Criterion attempts to be satisfied, given the maximal thematic grid of the processed verbs.

As it turns out, in most structures with local ambiguity examined to date, maximization of thematic assignment as per Theta Attachment (namely, maximal interpretation) is confounded with attaching incoming input to recently built structure (Late Closure). For example, the object reading in object/subject ambiguities, as in (1), will be preferred both according to Late Closure, as the NP is attached to a recent constituent, and according to the principle of maximizing interpretation, as this structure entails assignment of a thematic role to the NP by the adverbial clause verb, and enables the reader to interpret this NP as part of a proposition (‘Mary was reading the sentence’). Thus, previous studies of local ambiguities could not tease apart structural vs. thematic considerations in attachment choices.

However, there are structures with local ambiguity in which the predictions of structural vs. interpretational principles differ. Consider the local ambiguity exemplified in sentences (3–4) (Pritchett, 1992).

- (3) Katrina gave the man who was eating the fudge.
(*Katrina gave the fudge to the man who was eating something.*)
- (4) Katrina gave the man who was eating the fudge the wine.
(*Katrina gave the wine to the man who was eating the fudge.*)

In these sentences, the main verb is a ditransitive verb (here, *gave*), taking both a goal and a theme argument; the goal argument is modified by a subject relative clause which includes an optionally transitive verb (here, *eat*). In such sentences, when the NP following the second verb (*the fudge*) is encountered, it can, in principle, be attached in two ways, as illustrated in (5). It can either be attached to the recent, optionally transitive verb inside the relative clause, with the first verb awaiting its third argument, as in (5a); in this case, another NP has to appear for the sentence to be grammatical, as in (4) above. Alternatively, the ambiguous NP can be attached to the first, ditransitive verb in the main clause, in which case the second verb will be interpreted intransitively, as in (5b). In this case, the sentence is grammatical without additional material. We will call this type of ambiguity *main clause / relative clause (MC/RC) ambiguity*, as the ambiguous NP can be analyzed as either an argument in the main clause, or an argument in the relative clause.²

- (5) a. Katrina gave [the man [who was eating the fudge]]...
b. Katrina gave [the man [who was eating]][the fudge].

For the type of sentence exemplified in (5), models which make predictions based on structural vs. thematic considerations differ. Structural models which highlight locality of attachment, including Late Closure, predict invariable low attachment of the ambiguous NP, as in (5a). In contrast, according to thematic assignment maximization principles, such as Theta Attachment, there is optionality with regard to the attachment of the post-verbal NP. This is so since when this NP is encountered, it can receive a thematic role either from the matrix verb *gave* or from the embedded verb *eat*. Since the NP receives an interpretation in both analyses, no preference is predicted, and a chance distribution of attachment choices is expected.

Note that, at least according to the Theta Attachment principle of Pritchett (1988, 1992), the fact that the Theme role is obligatory for the higher, ditransitive verb (*gave*), whereas it is optional for the lower verb (*eating*), is not relevant at the point of ambiguity. Such global considerations do not have a role in the initial attachment decision. The only relevant factor is maximization

² MC/RC should not be confused with the well-known main verb / reduced relative ambiguity.

of thematic relations, namely, whether the NP can receive a thematic interpretation. This is in contrast with Gibson (1991), according to which assignment of obligatory thematic roles is preferred over assignment of optional ones. In particular, the model presented in Gibson (1991) incorporates both thematic and structural considerations. This model assumes a parallel parser, where structures are evaluated against each other using (a) thematic assignment maximization in accord with the Theta Criterion, specifically, the assignment of *obligatory* thematic roles, and (b) locality, with a preference for local attachments over more distant ones. According to this model, in the MC/RC sentences, the structure with local attachment will satisfy locality, but not thematic assignment, as it leaves the main verb with an unassigned obligatory theta role; in contrast, in the structure with non-local attachment, all obligatory thematic roles are assigned, but locality is not satisfied. Therefore, the structures with local and non-local attachment will have the same weight, predicting, similar to Pritchett's Theta Attachment, a chance distribution of attachment choices.

Importantly, the theories of thematic assignment described above, and, specifically, Theta Attachment, view thematic roles categorically: is there a role that can be assigned? However, other theories hold a probabilistic view of thematic assignment, i.e., they consider how common it is for a particular predicate to have an argument. Therefore, a factor whose effect can be tested here is *verb bias*, namely, the statistical information reflecting verbs' subcategorization preferences.

Note that if verb biases drive attachment choices, then the transitivity bias of the embedded, optionally transitive verb should play a role in both types of ambiguity that we considered, namely, object/subject and MC/RC ambiguities. For object/subject ambiguities, as in (1), if processing is determined by verb biases, then the ambiguous NP will be preferentially attached as object for transitively biased adverbial clause verbs (e.g., *read* or *drink*), whereas it will be preferentially attached as subject if the adverbial clause verb is intransitively biased (e.g., *dance*). For MC/RC ambiguities, as in (3) or (4), the preference for analyzing the ambiguous NP as the object of the optionally transitive verb inside the relative clause should increase for more transitively biased relative clause verbs.

Evidence regarding the effect of verb bias in object/subject ambiguities is limited. In an eye-tracking study, Pickering et al. (2000, Experiment 3) found that participants adopted the object analysis of the ambiguous NP in object/subject ambiguities, even when both prior context and subcategorization frequencies supported the subject analysis (as in *While the pilot was flying the plane/the horse that had arrived stood over by the fence*). Relevant findings are also reported in Itzhak et al. (2010). In this event-related potentials study, participants listened to object/subject ambiguous sentences with no prosodic breaks, or to unambiguous sentences with such breaks. The experiment also manipulated the transitivity bias of the subordinate verb: while all the verbs were optionally transitive, half of the verbs were transitively biased, and the other

half were intransitively biased. Itzhak and colleagues found a larger P600 garden path effect at the disambiguation region for transitively biased, compared to intransitively biased, verbs, suggesting a higher proportion of object attachment in the former. Crucially, though, there was still a P600 effect in sentences with intransitively biased verbs (compared to the unambiguous baseline), which means that the verb's bias did not fully prevent the attachment of the ambiguous NP as an object to the adverbial clause verb. More recently, Huang et al. (2024), in a large-scale self-paced reading study that tested seven structures and included 2000 participants, found no correlation between verb bias and the size of the garden path effect in sentences with the object/subject ambiguity. In particular, in object/subject sentences, every item resulted in a garden path effect that ranged from 59 to 258 ms, but, crucially, this item-level variability was not explained by verb subcategorization bias (cloze-based or corpus-based). These findings all seem to suggest, in line with Pickering et al.'s (2000) conclusion, that while verb bias may have an effect on processing, it is not strong enough to guide initial parsing choices, which are determined by other principles, whether structural (e.g., Late Closure) or interpretational (maximizing thematic assignment).

In contrast to object/subject ambiguities, there are no experimental studies of the MC/RC ambiguity that we know of. Pritchett (1992, p. 114) only reports informal off-line investigations, asking English-speaking participants for their judgments on processing difficulty in sentences such as (3) and (4). For both types of sentences, Pritchett notes that "...hearers appear to judge these sentences either as unprocessable or unproblematic. In other words, some individuals are led down the GP, some aren't". This result aligns with his predictions: when the ambiguous NP is attached to the low verb, a sentence such as (4) will be judged as unproblematic, and a sentence such as (3) will be judged as difficult. When the NP is attached to the high verb, the opposite pattern will emerge. Thus, for each sentence type, both types of judgements should be attested.

In the current study, we examine the MC/RC ambiguity, which enables us to tease apart structural and thematic considerations in attachment choices. We compare the MC/RC ambiguity to the object/subject ambiguity in the same participants and with the same verbs, to see whether there is a dissociation in attachment decisions between the two structures.

1.1 The current study

In this study, we ask whether thematic assignment maximization (implemented as the Theta Attachment principle) guides parsing, and whether this is true regardless of verb bias and the distance between the verb and its arguments. In addition, we ask what happens in cases where thematic assignment maximization does not determine the analysis, because the two attachment choices satisfy thematic relations to the same degree. It is possible that when Theta Attachment is inconclusive, other factors might come into play (Pritchett, 1992). In particular, it is possible that,

in such cases, thematic biases will play a role in the decision. Lastly, we test Pritchett's proposal, according to which thematic assignment maximization operates locally and incrementally, sometimes causing participants to ignore the global obligatoriness of thematic roles.

To test the hypothesis that thematic assignment maximization guides processing, we ran three speeded forced-choice completion experiments in Hebrew, where participants were provided with the locally ambiguous sentence preamble (as in (1)) and had to choose between two completion options (as in (1a) and (1b)). We explored attachment preferences using completion choices, rather than reading times at the disambiguating region, for several reasons. First, reading times at the disambiguation region are extremely variable across items and individuals (e.g., Brothers et al., 2021; Karsenti & Meltzer-Asscher, 2022, and papers cited therein). More importantly, these reading times measure not only *the need* for reanalysis (reflecting initial attachment choices), but also the cost of the reanalysis process itself (Huang et al., 2024). Crucially, this makes it impossible to assess the rates of different attachment choices, especially across structures. If we find an increase in average RT in the disambiguating region of structure A, compared to its baseline, it is impossible to know whether reanalysis with a certain cost happened on all trials, or whether a costlier reanalysis happened on only some portion of the trials, and what this portion was. When comparing two different structures A and B (relative to their baselines), differences in RTs at the disambiguating region are even harder to interpret, since the cost of the reanalysis itself may be different between the two structures. This makes it impossible to know whether reanalysis happened on more trials in one structure than in the other, namely, whether some initial parsing choice was made more often in one structure than in the other.

We, therefore, chose a technique that allowed us to probe attachment preferences directly, namely, RSVP with forced choice completion (e.g., Sikos et al., 2016; Staub, 2009). This method enables us to examine structures that do not cause processing difficulty. Since the sentence fragments examined do not require reanalysis, and since reanalysis is believed to be a last resort mechanism (Fodor & Frazier, 1980; Schneider & Phillips, 2001; Sturt et al., 2001), participants had no reason to modify their initial attachment. This method, therefore, provides a relatively direct window into attachment choices.

2. Experiment 1

The aim of Experiment 1 was to examine whether thematic assignment maximization guides processing, or whether attachment choices are guided by structural considerations (i.e., Late Closure) or by verb bias. To do this, we manipulated *ambiguity type*, contrasting object/subject (obj/subj) ambiguities (6) with main clause/relative clause (MC/RC) ambiguities (7). Two possible completions, each a two-word string in Hebrew, were provided for each fragment. One of the completions indicated local attachment of the ambiguous NP (as object in obj/subj

ambiguities, and as a complement to the relative clause verb in MC/RC ambiguities), and the other indicated non-local attachment (as subject in obj/subj ambiguities, and as a complement to the main clause verb in MC/RC ambiguities). We also varied the transitivity bias of the optionally transitive verb in both ambiguity types. In (6–7), the ambiguous NP is underlined, and the optionally transitive verb is doubly underlined.

- (6) After the guests drank cold water...
- a. the catering arrived (*local attachment*)
 - b. was handed out to everyone (*non-local attachment*)
- (7) The owner gave the guests who drank cold water...
- a. some wine (*local attachment*)
 - b. at the party (*non-local attachment*)

As explained above, Late Closure predicts that completions indicating local attachment will be preferred in both obj/subj and MC/RC ambiguities. Processing guided by verb biases predicts that completion preferences in both ambiguity types will be correlated with the optionally transitive (OpT) verb's transitivity bias. According to the Theta Attachment principle, however, performance will differ between the obj/subj and the MC/RC conditions. In the obj/subj condition, the need to satisfy the Theta Criterion given the maximal thematic grid of the verb will invariably lead to local attachment of the ambiguous NP, regardless of the verb's transitivity bias, as only in this structure will the NP receive a thematic role. In contrast, in the MC/RC condition, either the local or the high attachment option can be chosen, as both satisfy the Theta Criterion to the same degree. Thus, both types of completions may be chosen. In addition, we hypothesized that when optionality in thematic role assignment arises (namely, in the case of MC/RC sentences), transitivity bias will become relevant, affecting processing choices, such that local attachment will be chosen at a higher rate for more transitively biased OpT verbs. The predictions of the different accounts are summarized in **Table 1** below.

Table 1: Predictions of the different parsing strategies.

| | Thematic accounts (Theta Attachment) | Structural strategies (Late Closure) | Verb bias guidance |
|----------|---|---|---|
| Obj/subj | Local attachment, no correlation with verb bias | Local attachment, no correlation with verb bias | Attachment choices correlated with verb bias |
| MC/RC | Chance distribution of attachment choices, possible correlation with verb bias | Local attachment, no correlation with verb bias | Attachment choices correlated with verb bias |

2.1 Methods

2.1.1 Participants

Seventy-two native speakers of Hebrew, students at Tel-Aviv University (mean age = 26.2, range 20-37 years) participated in the experiment, in exchange for financial compensation or course credit. Fourteen participants were excluded, because they did not meet our pre-established criteria, as detailed in 2.2. Two participants were bilingual speakers of Hebrew and either Arabic or Bulgarian, and the rest were monolingual, based on self-report. All participants gave written informed consent. The study was approved by the local ethics committee (Faculty of Humanities, Tel-Aviv University).

2.1.2 Materials

The experiment included 24 Hebrew sentence sets of four conditions, with two experimental conditions, and two control conditions, manipulating ambiguity type (obj/subj / MC/RC). An example set is given in **Table 2**. The experimental obj/subj sentences included an adverbial clause with an optionally transitive verb, followed by an NP. The experimental MC/RC sentences included a ditransitive verb in the main clause, followed by a goal PP which contained a subject relative clause with an optionally transitive verb, and an NP. Importantly, in Hebrew, the word order V PP NP is grammatical.

Similar sentences with intransitive (IN) verbs, where only one completion option was grammatical, served as controls. In the obj/subj control sentences, the adverbial clause contained an intransitive, unaccusative verb, with no thematic role to assign to the subsequent NP, which therefore could only be attached as the subject of the upcoming main verb. In the MC/RC control condition, we did not simply replace the optionally transitive verb with an intransitive one (as we did in the obj/subj control sentences), which would give rise to a sentence like ‘The owner brought [to the guests that departed] cold water’. Rather, we added a PP after the intransitive verb: i.e., ‘The owner brought [to the guests that showered with cold water] orange juice’. We did this to ensure that participants were able to assign a thematic role across a long distance. If participants are able to robustly assign a thematic role across the entire relative clause, which includes a PP, in this control condition, it allows us to maintain that when they assign the thematic role of the main verb to the closer argument in the MC/RC experimental condition, it is not due to difficulty associated with long-distance assignment. Participants who were not accurate on this condition were excluded (see 2.2). In the control conditions, we thus expected high accuracy, manifested, however, at two opposite ends on the scale of preferences in the two conditions: i.e., mostly non-local attachment completions in the obj/subj control and mostly local attachment completions in the MC/RC control.

Table 2: Example set from Experiment 1.

| Condition | Sentence fragment | Completion options (all two-word strings in Hebrew): Local attachment / Non-local attachment |
|---------------------|--|---|
| Obj/subj | <i>axrey she-ha-orxim šatu maim karim</i> after that-the-guests drank water cold 'After the guests drank cold water' | <i>ha-keitering higiya</i> / the-catering arrived (local) |
| Obj/subj Control | <i>axrey she-ha-orxim halxu maim karim</i> after that-the-guests departed water cold 'After the guests departed cold water' | <i>huxzeru la-mekarer</i> was-returned to-the-refrigerator (non-local) |
| MC/RC | <i>baal ha-bait hevi la-orxim še-šatu</i> the owner brought to-the-guests that-drank <i>maim karim</i> water cold 'The owner brought to the guests who drank cold water' | <i>mic tapuzim</i> / orange juice (local) |
| MC/RC Control | <i>baal ha-bait hevi la-orxim še-</i> <i>the owner brought to-the-guests that-</i> <i>hitkalxu be-maim karim</i> <i>showered in-water cold</i> 'The owner brought to the guests who showered with cold water' | <i>emesh ba-xava</i> last-night at-the-farm (non-local) |

For the experimental conditions, we selected 24 optionally transitive verbs whose proportion of occurrence with a direct object ranges from 10% to 92%, based on manual coding of a random sample of 100 occurrences of each verb in Sketch Engine (Hebrew Web 2014 (heTenTen14, Meni/Alon tagged + lempos)).

The experimental and control sentences were intermixed with 24 filler sentences. Twelve filler sentences were similar to the obj/subj condition, but they included one of the following: (i) a transitive verb followed by its direct object, so that the only grammatical completion choice was the local attachment option (as in 'After the student submitted two seminar papers [the test was postponed / yesterday at noon].'); (ii) a transitive verb followed by its direct object and an additional NP, such that the only grammatical completion choice was the non-local attachment option (as in 'After the presenter has finished the economic program the newscast [began immediately / exactly at eight]'); or (iii) an intransitive verb followed by a PP, such that the only grammatical completion choice was the local attachment option (as in 'While the undergraduates were concentrating on the long exam [the supervisor fell asleep / on the large chair].'). The other twelve fillers were similar to the MC/RC conditions. In some

of them, the embedded verb was intransitive, with the post-verbal NP assigned a thematic role by a ditransitive or transitive main verb, so that the completion option was always the non-local attachment choice (as in ‘The instructor corrected for the student who disappeared a number of small mistakes [during the break / easy solutions]’). In others, the main verb was intransitive and the embedded verb, transitive, with the post-verbal NP assigned a thematic role by the embedded verb, such that the correct completion was the local attachment choice (as in ‘The teacher was happy that the pupils solved two difficult exercises [last week / another assignment]’).

Two pre-tests ensured that the two completion options in each structure were equally plausible. In the first pre-test, we compared the plausibility of the two completions in the subj/obj condition, namely, the non-local completion (e.g., ‘After the guests drank, cold water was returned to the refrigerator’) and the local completion (e.g., ‘After the guests drank cold water, the catering arrived’). We used a comma to indicate the intended structure. Twenty-seven Hebrew-speaking volunteers (age 21–43, different from those in other pre-tests or experiments in the study) rated each sentence, along with 12 filler sentences, on a scale of 1–7 on Ibx Farm (Drummond, 2007). The mean plausibility rating was 5.19 for the non-local completion option, and 5.53 for the local completion option. We included in the experiment only sets for which there was no significant difference in plausibility between the two variants. The second pre-test compared the plausibility of high versus low attachment for the MC/RC sentences (e.g., ‘The owner brought to the guests cold water’, ‘The guests drank cold water’). Sixty Hebrew speaking volunteers (age 22–40), different from those of the other pre-tests or experiments in this study, participated in this pre-test. Each sentence was rated on the scale of 1–5, using a Google Forms questionnaire. The mean rating for the high attachment was 4.4, and the mean rating for the low attachment was 4.54. We included in the experiment only sets for which there was no significant difference in plausibility between the high and low attachment.

The 24 sentence sets were assigned to four lists in a Latin Square design. Each participant, thus, read six locally ambiguous obj/subj sentences and six locally ambiguous MC/RC sentences, as well as six control sentences from each structure. The order of presentation was randomized for each participant. The distribution of the 24 experimental sets to lists was done such that each participant read sentences with OpT verbs of very weak (10–21%), weak (35–48%), moderate (49–61%), moderate-strong (61–72%), strong (73–79%) and very strong (80–92%) transitivity bias in each experimental condition.

2.1.3 Procedure

The experiment was web-based, programmed using Ibx Farm. Participants were instructed to read sentences that would appear word-by-word on the screen, and then select the

completion that best continues the sentence, as quickly and naturally as possible. Participants first read two example sentences, and then underwent a practice block of five items. During the experiment, participants pressed a key to begin each trial. Then, the sentence fragment was presented word-by-word at a rate of 400 ms per word with a 100 ms inter-stimulus interval, followed by the two completions, which appeared on the screen simultaneously. The completions remained on the screen until the participant responded. The position of the completions on the screen (left/right) was randomized. The experiment lasted approximately 12 minutes.

2.2 Data analysis

Analyses were conducted using R, version 4.0.2 (R Core Team, 2020-06-22). All data and code can be found on the OSF page for the project at: https://osf.io/uxnmc/?view_only=199ede8bf030465aac9da1187b909b59. Participants were excluded from the analysis if their completion accuracy on items in any of the control conditions (obj/subj and MC/RC) was lower than 85%, resulting in the removal of 14 participants. As explained above, we required high accuracy in the MC/RC control condition to ensure the ability of participants to assign a thematic role across a long distance. For the remaining 58 participants, we performed the analysis only for responses provided within 6000 ms of presentation of the alternatives, affecting 3% of the data. We eliminated responses that took longer than 6 seconds in order to obtain intuitive attachment preferences, and reduce cases of conscious deliberation or reanalysis performed after experiencing a garden path effect. An additional analysis of all data, which showed very similar results to the main analysis, is provided in Appendix A.

The analysis focused only on the two experimental conditions. The completions were coded as zero (0) if they presented non-local attachment, and as one (1) if they presented local attachment. To compare the predictions of the different accounts, we carried out several analyses. First, we tested whether local attachment rates differ in the obj/subj and MC/RC conditions. According to Late Closure, as well as according to guidance by verb bias, local attachment rates should be similar in the two conditions (as locality/verb bias plays a similar role in both), whereas, according to Theta Attachment, attachment should be overwhelmingly local in the obj/subj condition, whereas variability is predicted in the MC/RC condition. In addition, we tested the effect of verb bias in the different ambiguity types. If transitivity biases guide processing, a similar effect of verb bias will appear in both ambiguity types. Likewise, a similar effect of verb bias in both structures is predicted according to Late Closure, as this principle predicts no effect of verb bias in any of the structures. According to Theta Attachment, however, verb bias can affect attachment preferences in MC/RC sentences, where thematic assignment maximization cannot determine the structure, but not in obj/subj sentences, where

Theta Attachment predicts robust local attachment. Therefore, an interaction can be expected here.

A generalized logistic mixed-effects model was fitted, modeling attachment choices with *ambiguity type* (obj/subj / MC/RC, sum-coded (-0.5, +0.5)), *verb bias* (centered), and their interaction as fixed effect predictors, and random effects for participants and items, as in (8).

$$(8) \quad \text{attachment choice} \sim \text{ambiguity type} * \text{verb bias} + (1 + \text{ambiguity type} * \text{verb bias} \mid \text{participant}) + (1 + \text{ambiguity type} * \text{verb bias} \mid \text{item})$$

A follow-up generalized logistic mixed-effects model with nested slopes was fitted to test simple effects of *verb bias* for each of the levels of the factor *ambiguity type*.

We performed both a frequentist and a Bayesian analysis. A generalized logistic mixed-effects model was fitted in the frequentist analysis, using the lme4 package, version 1.1-27.1 (Bates et al., 2021). A maximal model that included all random slopes was not fitted, due to a convergence problem. The final model included random intercepts for participants and items and random slopes for ambiguity type for participants and items. For the Bayesian analysis, a Bernoulli response distribution model was fitted. Bayesian models were created in the Stan computational framework (<http://mc-stan.org/>) accessed with the brms package (Bürkner, 2017, 2018). For the Bayesian analysis, we used mildly informative priors: a standard normal distribution $N(0, 1)$ as the prior distribution for the fixed effects and the standard deviation parameters, a prior of $N(0, 3)$ for the intercept, and the LKJ prior for the correlation matrices of the random effects. Four sampling chains were run with 4000 iterations each, the first 2000 of which were discarded as warm-up. We report percentile-based 95% credible intervals (CrI) for all parameters of interest. Inferences were drawn based on the probability that the parameter of interest is greater (or less) than 0. We concluded that the effect is “significant” if this probability was greater than 0.95. For better interpretability, parameters on the log-odds scale have been back-transformed to the raw proportion scale. In the follow-up Bayesian model, the random effect of nested slopes was not included, due to a convergence problem.

Additionally, for each structure separately, we tested whether attachment choices are consistent with a chance distribution, as, according to Theta Attachment, chance distribution is predicted in MC/RC sentences, but not in obj/subj sentences. To do this, the by-participant data from each structure were analyzed separately, using a binomial test of $H_0: p = 0.5$ against $H_a: p > 0.5$.

2.3 Results

The overall rates of completions (after exclusion of trials with RTs over 6 s) are presented in **Figure 1a** for control conditions, and **Figure 1b** for experimental conditions (for data without exclusion, see **Table A1** in Appendix A).

The control conditions behaved as expected, exhibiting uniform attachment choices in line with the grammatical structure of the sentence, i.e., local attachment in the MC/RC control condition, and non-local attachment in the obj/subj control condition. Mean accuracy in the control conditions for the participants included in the analysis was 95.3%.

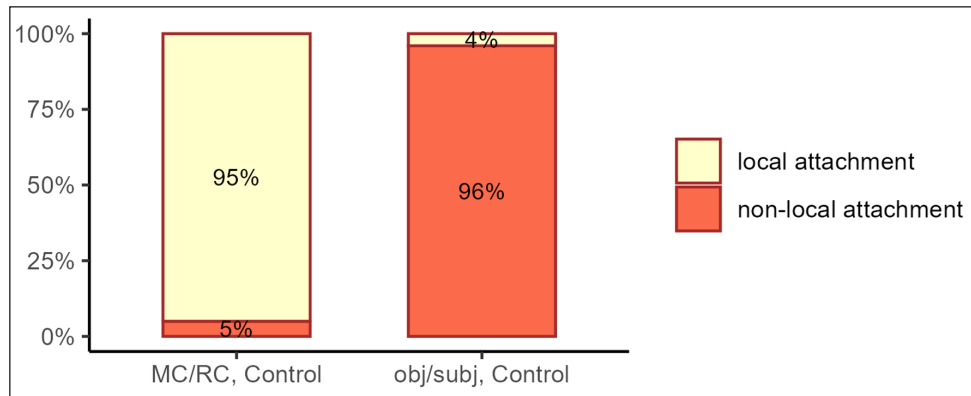


Figure 1a: Rate of completions in control conditions, Experiment 1.

As for the experimental conditions, the obj/subj condition overwhelmingly exhibited local attachment, whereas in the MC/RC experimental condition, local and non-local attachments were chosen at similar rates.




Figure 1b: Rates of completions in experimental conditions, Experiment 1.

The by-item average rates of local attachment in the experimental conditions are presented in **Figure 2**, representing the effect of verb transitivity bias on the rate of choosing local attachment in each sentence type.




Figure 2: The effect of transitivity bias on the rate of choosing local attachment, Experiment 1.

Table 3 and **Table 4** summarize the results of the frequentist and Bayesian analyses, respectively. The frequentist analysis showed an effect for ambiguity type, such that the rates of local attachment in *obj/subj* sentences were significantly higher than in *MC/RC* sentences ($p < .001$). The interaction between ambiguity type and verb bias was not significant. However, the follow-up model with nested slopes showed that in the *obj/subj* sentences, there was no effect of transitivity bias on local attachment rate ($p = .627$), whereas in the *MC/RC* condition, the effect of transitivity bias on local attachment rate was significant ($p = .007$), indicating a higher local attachment rate with more transitively biased OpT verbs.

Table 3: Results of the frequentist analysis, Experiment 1: Estimates, standard-error, t-values and p-values.

| | Estimate | SE | <i>t</i> | <i>p</i> |
|--|----------|-------|----------|-----------|
| Grand Mean (Intercept) | 1.145 | 0.182 | 6.308 | |
| Effect of ambiguity type on completion choice <i>obj/subj</i> > <i>MC/RC</i> | 1.994 | 0.333 | 5.972 | <0.001*** |
| Overall effect of verb transitivity bias on completion choice | 0.012 | 0.006 | 1.856 | 0.064 |
| Interaction between verb bias and ambiguity type | -0.014 | 0.011 | -1.299 | 0.193 |
| Effect of verb bias in the <i>obj/subj</i> condition | 0.005 | 0.010 | 0.485 | 0.627 |
| Effect of verb bias in the <i>MC/RC</i> condition | 0.019 | 0.007 | 2.712 | 0.007** |

The results of the Bayesian analysis were consistent with those of the frequentist analysis. There was an effect of ambiguity type, reflecting a decrease in the proportion of local attachment in the MC/RC condition, relative to the obj/subj condition: the mean of the posterior distribution was 88%, with a CrI of [83, 93], in the obj/subj condition, and 50%, with a CrI of [41, 62], in the MC/RC condition. We consider the effect reliable, as the 95% credible interval for this contrast does not cross zero. In this analysis, the overall effect of verb bias was not reliable, nor was the interaction. However, similarly to the frequentist analysis, the nested analyses showed a reliable effect of verb bias in the MC/RC condition, such that more transitively biased verbs led to more local attachment, but no reliable effect in the obj/subj condition.

Table 4: Results of the Bayesian analysis of question 1, Experiment 1: Posterior means on the log-odds scale (with 95% Bayesian CrI).

| | Log-odds scale posterior mean | [CrI] | Posterior beyond zero |
|---|-------------------------------|-----------------|-----------------------|
| Grand Mean (Intercept) | 1.197 | [0.803, 1.622] | |
| Effect of ambiguity type <i>obj/subj</i> > <i>MC/RC</i> | 2.052 | [1.399, 2.721] | 100% |
| Overall effect of verb transitivity bias | 0.014 | [-0.005, 0.035] | 94% |
| Interaction between verb bias and ambiguity type | -0.02 | [-0.056, 0.015] | 90% |
| Effect of verb bias in the obj/subj condition | 0.005 | [-0.016, 0.028] | 70% |
| Effect of verb bias in the MC/RC condition | 0.022 | [0.004, 0.043] | 99% |

As for the difference from chance in each ambiguity type, the binomial tests showed that in the MC/RC condition, the by-participant average choice of attachment was not significantly different from the null hypothesis of a chance distribution ($t = .7$, $df = 57$, $p = .48$), whereas in the obj/subj condition, the by-participant average choice of attachment was significantly different from the null hypothesis, signifying that the results are inconsistent with a chance distribution ($t = 16$, $df = 57$, $p < .0001$).

2.4 Discussion

Experiment 1 examined attachment preferences for an ambiguous NP during the processing of sentences in which the NP is ambiguous between an object and subject analysis (obj/subj sentences), compared to sentences in which the NP is ambiguous between being analyzed as the object of a matrix verb or as the object of a verb of a relative clause, and can receive a thematic

role in either analysis (MC/RC sentences). In addition, we manipulated the transitivity bias of the optionally transitive verbs directly preceding the ambiguous NP, examining whether this probabilistic information affects attachment preferences in each sentence type.

The results showed that in obj/subj sentences, where only the object attachment choice leads to thematic assignment for the NP, this NP was reliably attached as an object, regardless of the preceding verb's transitivity bias, in accord with the idea that maximal thematic assignment guides processing. In contrast, in MC/RC sentences, where both attachment choices lead to thematic assignment for the ambiguous NP and, therefore, the Theta Criterion could be satisfied in two different ways, this NP was attached locally only in approximately half of the trials, consistent with a chance distribution. In addition, more local attachments were observed for more transitively biased verbs.

The results with regard to the MC/RC structure demonstrate that with carefully normed materials, Pritchett's prediction of a chance distribution of attachment choices when there are two equally good thematic assignment options is borne out. We believe that this pattern was a result of the relatively even distribution of the 24 optionally transitive verbs on the scale of transitivity bias, from 10% to 90%. Namely, Pritchett's prediction does not hold for every verb by itself, but it holds, on average, for verbs with differing transitivity biases.

The present results are inconsistent with a structural approach adopting Late Closure. Specifically, the variability found in attachment choices in the MC/RC condition, both within and between items, provides evidence against this structural strategy, which predicts robust local attachment in this condition. In addition, the results do not support the idea that verb bias guides attachment choices. Although the interaction between verb bias and ambiguity type was not significant, transitivity bias was correlated with rates of selecting a local attachment in the MC/RC condition, whereas a similar effect was absent in the obj/subj condition. The null effect in the obj/subj ambiguity may be due to insufficient power to detect an effect. However, the fact that an effect was observed in the MC/RC condition (as well as the fact that the null result was replicated in Experiment 2) suggests that the effect may, indeed, be null, namely, that verb bias does not guide processing in this case. This result is inconsistent with a strong version of verb bias guidance (e.g., Trueswell et al., 1993) which predicts that in obj/subj sentences as well, attachment will be determined based on the bias of the optionally transitive verb.

The lack of interaction between verb bias and ambiguity type can be related to the fact that the effect of verb bias on attachment choices was positive in both sentence types, even though it was very weak and not significant in the obj/subj sentences, and relatively strong and significant in the MC/RC sentences. In addition, given that correlations in psycholinguistics are never very strong (e.g., the correlation that we found for the MC/RC condition was 0.45), it is not necessarily predicted that a moderate correlation in the MC/RC condition will interact with

a null (or weak) correlation in the obj/subj condition. However, the results are still consistent with the tested hypotheses, namely, that when there is only one theta assigner for the critical NP, the maximal thematic grid of a verb always takes precedence over verb bias in guiding sentence processing, but when there are two possible theta assigners for the critical NP, optionality in theta attachment arises, and then transitivity bias becomes relevant. The results are compatible with the suggestion of Pickering et al. (2000), namely, that while verb bias may have an effect on processing, it is not strong enough to guide initial parsing choices, which are determined by other principles, for example, thematic assignment maximization.

Taken together, the results of Experiment 1 provide evidence for the precedence of thematic considerations during parsing. When there is only one attachment choice which enables thematic assignment, the ambiguous NP invariably attaches in this way, regardless of the verb's transitivity bias, to maximally satisfy the Theta Criterion and, thus, maximize interpretation. However, when there are two attachment choices which enable thematic assignment, and the Theta Criterion can be satisfied in two different ways, either maximizing interpretation of the main clause or the relative clause, other factors – such as transitivity bias – come into play and affect processing choices. The model of Gibson (1991) can also account for the results. According to this model, the structure with local attachment will always be preferred in the obj/subj sentences, because thematic assignment and locality both favor attachment of the ambiguous NP to the preceding verb. However, in the MC/RC sentences, the structures with local attachment and non-local attachment will have the same weight, with the former satisfying locality and the latter satisfying thematic assignment of obligatory thematic roles.

3. Experiment 2

The main goal of Experiment 2 was to test whether the overwhelming local attachment found in Experiment 1 for sentences with object/subject attachment ambiguity can, indeed, be attributed to Theta Attachment alone, or whether some type of locality consideration is taken into account as well. In particular, we asked whether local attachment will be chosen at a lower rate when there is distance between the verb and the subsequent NP, compared to when they are adjacent. Possibly, when the verb is not adjacent to its argument, the role of Theta Attachment decreases. Additionally, if Theta Attachment is, indeed, less decisive in such cases, it is also possible that, in these cases, we will see an effect of transitivity bias, such that local attachment will be chosen at a higher rate for more transitively biased OpT verbs.

To test the effect of distance on thematic assignment maximization, we used Hebrew counterparts of obj/subj ambiguity sentences such as (9), in which an adverb is present between the verb and the NP.

- (9) When the boy was eating at home pizza fell on the floor.

In English, the presence of the adverb *at home* after the verb leads to early closure of the adverbial clause, before the NP *pizza*, because in English, a direct object has to be adjacent to the verb (Chomsky, 1981). Indeed, previous experiments in English that examined sentences with the obj/subj ambiguity used an adverb in baseline sentences on a par with a comma (e.g., Adams et al., 1998; Mitchell, 1987). Unlike English, however, Hebrew allows a word order where an adverb appears between a verb and its direct object. This enables us to explore whether and to what extent the distance between the optionally transitive verb and the ambiguous NP modulates attachment choices.

Distance between verbs and arguments is a relatively robust cause of sentence processing difficulty. This has been proposed in different models, starting with Frazier's (1987a) Late Closure (see also Visibility in Frazier & Clifton, 1998) and was variously attributed to activation decay (e.g., Lewis & Vasishth, 2005), interference by similar items (cue-based retrieval; Lewis & Vasishth, 2005, and subsequent work), or the introduction of new discourse referents (Gibson, 2000; see also Van Dyke & Lewis, 2003). This makes it possible that the parser avoids combining a verb and an argument over a distance, even at the cost of not assigning a thematic role which could otherwise be assigned. Namely, these theories predict lower rates of local attachment when an adverb is present, compared to when it is absent. A similar prediction is derived if parsing decisions are based on statistical information regarding the distribution of adverbs. Although adverbs may appear between the verb and its object in Hebrew, they are rare in this position, and more frequently appear at the end of the verb phrase. Thus, statistically, the probability that the ambiguous NP opens a new clause is larger when it is preceded by an adverb than when it follows the verb directly.³

In contrast, according to thematic assignment maximization approaches, the object attachment is preferable, as it entails assignment of a thematic role, regardless of distance (in Hebrew, where this is grammatically possible). This would result in robust local attachment whether an adverb is present or not. Theories which assume guidance by verb bias predict that the transitivity bias of OpT verbs should play a role both when an adverb appears in the sentence and when it does not, with more local attachment for more transitively biased verbs, as predicted for Experiment 1.

As we noted in Experiment 1, one theory which explicitly includes both thematic considerations and locality is Gibson (1991). According to this theory, in obj/subj sentences, when there is no distance between the verb and the ambiguous NP, the structure with local attachment will be strongly preferred (as already explained in Experiment 1). However, in sentences with increased

³ To verify this generalization, we checked four representative experimental verbs, which are in different places on the transitivity scale, and observed that, indeed, an adverb in Hebrew indicates that the chance of an object is smaller, compared to sentences where an adverb does not appear after the verb (for details, see Appendix B).

distance between the verb and the NP, recency will decrease slightly and, therefore, the structure with local attachment will receive a slightly weaker preference, although it will still be preferred.

In the experiment, we manipulated the distance of the optional direct object from the preceding verb by adding two adverbial phrases between them. We also manipulated the transitivity bias of the verbs, as in Experiment 1. Two possible completions were provided, one indicating local attachment (where *local* means attachment as an object, as in Experiment 1), and the other, nonlocal attachment (namely, attachment as the subject of the upcoming clause). Thus, the experiment included two conditions, both with the same structure (obj/subj ambiguity), differing only in the presence or absence of the adverb (see **Table 5** for a translated example, and **Table 6** for a Hebrew example).

Table 5: Example set, translated from Hebrew.

| Condition | Sentence fragment | Completion options (all two-word strings in Hebrew) |
|-----------|---|--|
| Adverb | When the boy was eating <i>in the kitchen at home</i> hot pizza | Local attachment: the friend called Non-local attachment: fell on the floor |
| No Adverb | When the boy was eating hot pizza | Local attachment: the friend called Non-local attachment: fell on the floor |

3.1 Methods

3.1.1 Participants

Fifty-eight native speakers of Hebrew, students at Tel-Aviv University (mean age = 23.5, range 19-30 years) who did not take part in Experiment 1, participated in the experiment, in exchange for financial compensation or course credit. Fifteen participants were excluded, because they did not meet our pre-established criterion, which was accuracy of at least 85% in filler sentences and in catch trials, which were similar to the MC/RC control condition sentences in Experiment 1. Additionally, four other participants whose feedback showed that they understood the aim of the experiment, due to their previous knowledge of sentence processing, were also excluded. Three participants were bilingual speakers of Hebrew and either English or Russian, and the rest were monolingual, based on self-report. All participants gave written informed consent. The study was approved by the local ethics committee (Faculty of Humanities, Tel-Aviv University).

3.1.2 Materials

The experiment included 24 Hebrew sentence sets of two conditions, manipulating *adverbial intervention*. An example set is given in **Table 6**.

Table 6: Example set from Experiment 2.

| Condition | Sentence | Completion options: Local attachment / Non-local attachment |
|-----------|--|--|
| Adverb | <i>be-zman she-ha-yeled axal ba-mitbax ba-bait</i> at-time that-the-child ate at+ the-kitchen at-home <i>pica xama</i> pizza hot 'When the boy was eating in the kitchen at home hot pizza' | <i>ha-xaver cilcel/</i> the-friend called (local attachment) |
| No-Adverb | <i>be-zman she-ha-yeled axal pica xama</i> at-time that-the-child ate pizza hot 'When the boy was eating hot pizza' | <i>nafla al ha-ricpa</i> fell on the-floor (non- local attachment) |

The *No Adverb* condition sentences were the same as the obj/subj sentences in Experiment 1. To manipulate the distance of the optional direct object from the preceding verb in the *Adverb* condition, we added two adverbial modifiers, each an orthographical word in length, either both temporal, e.g., 'yesterday morning', or both locative, e.g., 'in the kitchen at home'. The experimental sentences were intermixed with 48 filler sentences, six of which were catch trials similar to the MC/RC control sentences from Experiment 1. Twelve fillers were the relative clause fillers from Experiment 1, and the rest (30 fillers) were very similar to them. We included filler sentences and catch trials with a long structure in which only one completion was correct, such as 'The instructor corrected for the student who disappeared a number of small mistakes [during the break / easy solutions]'. We only included participants who mostly answered correctly on these sentence types, to rule out a "guessing" behavior in the Adverb condition, where extra words are added to the sentence.

Sentences were assigned to two lists in a Latin Square design. Each participant, thus, read twenty-four locally ambiguous obj/subj sentences, twelve with, and twelve without, an added adverb. The order of presentation was randomized for each participant.

3.1.3 Procedure

The procedure was similar to that of Experiment 1, with two differences: first, the rate of word presentation was 350 ms + 100 ms ISI, and second, we limited the time for choosing between the two completion options to four seconds. In the instructions, participants were notified that they would have to make their choices within four seconds. Before the experiment began, participants completed a practice block which consisted of four items.

3.2 Data analysis

To test the predictions of the different accounts, we compared local attachment rates in the Adverb condition to those in the No Adverb condition. In addition, we examined whether there is an effect of verb bias of the OpT verbs on attachment choices in each condition.

As in Experiment 1, the completions were coded as zero (0) if they presented non-local attachment, and one (1) if they presented local attachment. A generalized logistic mixed-effects model was fitted, modeling attachment choices with adverbial intervention (Adverb / No Adverb, sum coded), transitivity bias (centered), and their interaction as fixed effect predictors, and random effects for participants and items.

As in the previous experiment, we carried out both a frequentist and a Bayesian analysis and fitted the same models as above. In the frequentist analysis, the maximal models that included all random slopes were not fitted, due to a convergence problem. The final models included random intercepts for participants and items. The Bayesian models included all random intercepts and slopes.

3.3 Results

Eight out of 936 responses (less than 1%) were late responses and, therefore, were not available for analysis. Completion choices and correlation with verb bias are presented in **Figures 3 and 4** (for additional data, see Appendix C).




Figure 3: Rate of completions by condition, Experiment 2.

The results of the frequentist and Bayesian analyses are summarized in **Tables 7 and 8**, respectively. The results show that in both the Adverb and No Adverb conditions, participants principally chose local attachment. However, the difference between Adverb and No Adverb sentences in the rate of selecting a local attachment was significant, with higher rates in the No Adverb sentences, as indicated by both the frequentist and the Bayesian models. In the Bayesian analysis, the mean of the posterior distribution was 90%, with a CrI of [85, 94], in the No Adverb condition, and 80%, with a CrI of [71, 88], in the Adverb condition. The frequentist model also

showed a significant interaction between adverbial intervention and transitivity bias ($p = .04$), such that the insertion of an adverb resulted in a slight, non-significant, positive relation between the rates of local attachment and the verb's transitivity, whereas in sentences without an adverb, the (non-significant) relation was in the opposite direction. However, this interaction did not turn out to be reliable in the Bayesian analysis, as its 95% CrI crossed 0.

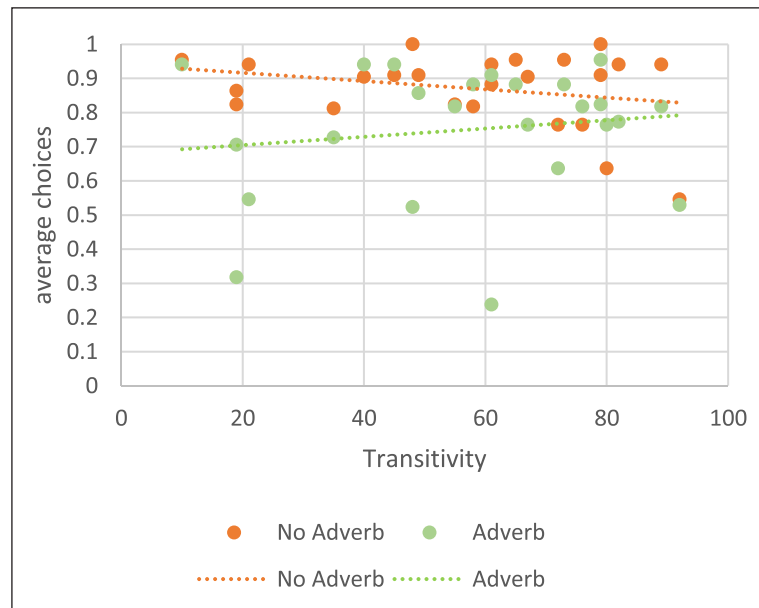


Figure 4: By-item proportion of choices of local attachment in each condition as a function of verb transitivity bias, Experiment 2.

Table 7: Results (on the log-odds scale) of the frequentist analysis of Experiment 2: Estimates, standard error, t-values and p-values.

| | Estimate | SE | t | p |
|---|----------|-------|--------|-----------|
| Grand Mean (Intercept) | 1.715 | 0.195 | 8.791 | |
| Main effect of adverbial intervention <i>Adverb < No-Adverb</i> | -0.815 | 0.233 | -3.504 | <0.001*** |
| Overall effect of verb transitivity bias | -0.003 | 0.008 | -0.382 | 0.702 |
| Interaction between adverbial intervention and transitivity bias | 0.020 | 0.010 | 2.067 | 0.039* |
| Effect of transitivity bias in the No Adverb condition | -0.012 | 0.009 | -1.290 | 0.197 |
| Effect of transitivity bias in the Adverb condition (from the same model after releveling conditions) | 0.008 | 0.008 | 0.939 | 0.348 |

Table 8: Results of the Bayesian analysis of Experiment 2: Posterior means on the log-odds scale (with 95% Bayesian CrI).

| | Posterior mean | [CrI] | Posterior beyond zero |
|--|----------------|----------------|-----------------------|
| Grand Mean (Intercept) | 1.81 | [1.37, 2.26] | |
| Effect of adverbial intervention <i>Adverb < No Adverb</i> | -0.86 | [-0.29, -0.14] | 99% |
| Overall effect of verb transitivity bias | -0.003 | [-0.02, 0.02] | 60% |
| Interaction between adverbial intervention and transitivity | 0.02 | [-0.05, 0.01] | 90% |
| Effect of transitivity in the No Adverb condition | -0.01 | [-0.03, 0.01] | 83% |
| Effect of transitivity in the Adverb condition (from the same model after releveling conditions) | 0.01 | [-0.02, 0.03] | 73% |

3.4 Discussion

Experiment 2 further examined attachment preferences for an ambiguous NP during processing of obj/subj ambiguity sentences. We examined the Theta Attachment strategy, manipulating the distance of the optional direct object from the preceding verb by adding two adverbs, and compared sentences containing an adverb and those without an adverb, to see whether Theta Attachment will still guide processing, even over a distance. We also manipulated the transitivity bias of the optionally transitive verbs, similarly to Experiment 1.

The results showed that participants largely opted for a local attachment. In contrast to experiments in English, where an adverbial phrase was used in baseline sentences on a par with a comma (e.g., Adams et al., 1998; Mitchell, 1987), in more than 70% of trials in Hebrew, distance did not prevent attachment of the ambiguous NP as an object of the preceding verb. However, the effect of distance was statistically significant, such that local attachment was chosen less frequently when an adverb intervened between the verb and the NP.

As in the obj/subj condition in Experiment 1, there was no effect of transitivity bias on attachment choices in Experiment 2, in either the adverb or the no adverb condition. Again, this may be the result of low power, but the fact that Experiments 1 and 3 did show such an effect for MC/RC ambiguities, while the obj/subj ambiguity failed to show this effect across experiments, suggests that there is, indeed, no relation between verb bias and attachment choices in this structure.

The main effect of adverbial intervention is consistent with the predictions of locality approaches, namely, that the proportion of local attachment will be decreased with distance. However, the small difference between the conditions (13%), coupled with the overwhelming choice of object attachment in both conditions, suggest that thematic assignment maximization did play an important role in attachment choices. Thus, the results support some kind of locality or recency preference, which operates together with Theta Attachment to determine attachment choices. It could be thought that this locality preference should be equated with Late Closure. However, Late Closure was conceived as a binary, all-or-none strategy, which cannot be easily applied to graded situations or situations in which multiple constraints are at work (see Frazier & Clifton, 1998, on replacing Late Closure with Visibility). In fact, Frazier and Clifton (1998) note that the effects of distance hold in extreme cases even for unambiguous sentences and, therefore, that it is incorrect to characterize them with an ambiguity resolution principle such as Late Closure. Their idea is that distant nodes can become less visible to the processor because of memory capacity and attentional mechanisms, similarly to the locality theory of Gibson (1998, 2000). Thus, a more cognitively-grounded locality preference seems to us more appropriate than Late Closure as an explanation of these results. The results are well aligned with Gibson (1991), who takes into account both thematic assignment and locality, predicting preference for local attachment both when the verb and NP are adjacent and when they are not, with a slight decrease of this preference in sentences with increased distance between the verb and the subsequent NP.

Notably, the results are not consistent with the predictions of verb bias guidance, because under this strategy, both conditions were predicted to show a positive and significant correlation between transitivity bias and choice of local attachment. However, such correlations were not found, replicating the lack of correlation in the obj/subj sentences in Experiment 1.

4. Experiment 3

The main aim of Experiment 3 was to test a specific prediction made by Pritchett (1988, 1992) that is not necessarily shared by other thematic approaches, namely, that when Theta Attachment allows two attachment options (as in the MC/RC structure), the global obligatoriness of thematic roles is ignored, because thematic assignment maximization operates locally and incrementally.

To test this, in Experiment 3 we focused on the MC/RC ambiguity. The experiment included three conditions, differing only in their embedded verb type. The first two conditions were the same as in Experiment 1, namely, an intransitive relative clause verb with a PP (which served as a control in Experiment 1), and an optionally transitive embedded verb. In this experiment, we added a condition with an *obligatorily transitive* (OblT) embedded verb⁴ (see **Table 9** for a

⁴ By *obligatorily transitive*, we mean extremely highly biased to take a direct object. In our sample, the relevant verbs appear with a direct object in more than 96% of sentences.

translated example, and **Table 10** for a Hebrew example). In the latter condition, both the main verb ('brought' in (9)) and the embedded verb ('received' in (9)) require a direct object. For the sentence to be grammatical, participants must attach 'cold water' locally, as a complement to 'received', and choose the completion 'orange juice', so that this NP can receive the obligatory Theme role of 'brought'. If participants choose 'last night', the resulting sentence is ungrammatical, as one of the verbs remains with an obligatory thematic role unassigned.

Table 9: Example set for Experiment 3, translated from Hebrew.

| Condition | Pre-response context | Completion options (all two-word strings in Hebrew) |
|-----------|---|---|
| IN | The owner brought to the guests that showered with cold water | Local attachment: orange juice |
| OpT | The owner brought to the guests that drank cold water | Non-local attachment: last night at the farm |
| OblT | The owner brought to the guests that received cold water | |

Reminder: in Hebrew, the word order V PP NP is grammatical.

The predictions for the IN and OpT conditions in Experiment 3 are identical to those for Experiment 1. Structural considerations (specifically, Late Closure) predict that the OblT condition should behave very similarly to the OpT condition, with local attachment predicted in both, since the first post-verbal NP will always be attached as an argument of the embedded verb. As explained above, this will eventually lead to a globally grammatical sentence (even if this consideration was not applied locally). Theta Attachment, likewise, predicts similarity between the OpT and OblT conditions, but here, this will be manifested as both local and non-local completions in both conditions, since both present optionality in the attachment of the NP. As explained above, in Pritchett's theory, the obligatoriness of the thematic role should not influence processing at the point of building the initial structure; only local maximization of thematic assignment is considered. Thus, based on Pritchett's proposal, initial attachment should follow a chance distribution of local vs. non-local attachment choices in both the OpT and OblT conditions. Further, if we consider OblT verbs as OpT verbs which are extremely biased to the transitive end of the scale, then our proposal regarding OpT verbs in Experiment 1 is also relevant for OblT verbs, namely, that when Theta Attachment does not decide between the two attachment options, the decision will be influenced by verb bias. Since the average verb bias is larger in the OblT condition than in the OpT condition, rates of local attachment are predicted to be higher in the OblT condition. Crucially, if we do find a substantial proportion of non-local attachments, leading to ungrammatical sentences, this provides strong evidence that participants opt for local thematic assignment maximization, rather than always considering global grammaticality.

4.1 Methods

4.1.1 Participants

Sixty-one native speakers of Hebrew, students at Tel-Aviv University (mean age = 24.7, range 19-33 years) who did not take part in Experiments 1 or 2, participated in Experiment 3, in exchange for financial compensation or course credit. Thirteen participants were excluded, because they did not meet our pre-established criterion, which was the same as in Experiment 1, namely, accuracy of at least 85% in the intransitive condition. Five participants were bilingual speakers of Hebrew and either English, Russian or Spanish, and the rest were monolingual, based on self-report. All participants gave written informed consent. The study was approved by the local ethics committee (Faculty of Humanities, Tel-Aviv University).

4.1.2 Materials and procedure

The experiment included 24 Hebrew sentence sets of three conditions, manipulating verb type (OpT/OblT/IN). An example set is given in **Table 10**.

Table 10: Example set from Experiment 3.

| Condition | Sentence fragment | Completion options (Local attachment / Non-local attachment) |
|-----------|--|--|
| IN | <i>baal ha-bait hevi la-orxim še-hitkalxu</i> the owner brought to-the-guests that-showered <i>im maim karim</i> with water cold 'The owner brought to the guests that showered with cold water' | |
| OpT | <i>baal ha-bait hevi la-orxim še-šatu maim</i> the owner brought to-the-guests that-drank water <i>karim</i> cold 'The owner brought to the guests that drank cold water' | <i>mic tapuzim /</i> orange juice (local attachment) |
| OblT | <i>baal ha-bait hevi la-orxim še-kiblu</i> the owner brought to-the-guests that-received <i>maim karim</i> water cold 'The owner brought to the guests that received cold water' | <i>emesh ba-xava</i> last-night at-the-farm (non-local attachment) |

For the new OblT condition, we selected 19 obligatorily transitive verbs whose proportion of occurrence with a direct object was at least 96%, based on manual coding of a random sample of 100 occurrences of each verb in a Google search. One verb was used in two sets, with

different agreement features, i.e., masculine and feminine; one other verb was used in five sets, with varying agreement features.⁵ OpT verbs were selected according to the same criteria as in Experiment 1. Out of 24 sets, Experiment 3 had 8 sets not included in Experiment 1, and 3 different OpT verbs.

A pre-test was carried out to verify that the different attachment options were equally plausible. Sixty Hebrew speaking volunteers (age 22-40), different from those in the other experiments of this study, participated in this pre-test.⁶ Three sentences were composed for each set, e.g., for the set demonstrated in **Table 10**, the sentences were: ‘The owner brought the guests cold water’, ‘The guests received cold water’, ‘The guests drank cold water’. The sentences were assigned to three lists in a Latin Square design, and their order was randomized for each participant. Each sentence was presented separately on the computer screen, and participants were asked to rate them on a scale of 1-5, using Google Forms. We compared the responses for each set with t-tests (two comparisons for each set: main clause vs. relative clause attachment with the OpT verb and main clause vs. relative clause attachment with the ObIT verb). Plausibility did not differ between the attachment options of each set.

The experimental sentences were intermixed with 24 filler sentences, all of which were unambiguous but contained relative clauses similar to the experimental sentences. Twelve fillers were the fillers from Experiment 1, and an additional twelve fillers were the fillers from Experiment 2. Sentences were assigned to three lists in a Latin Square design. The order of presentation was randomized for each participant.

The procedure was similar to that of Experiment 1. However, we ran the Experiment via Zoom. The experiment was run with open cameras (both the participant’s and the experimenter’s), in order to emulate the laboratory conditions as much as possible.

4.2 Data analysis

The exclusion criterion was the same as in Experiment 1. After exclusion, for the remaining 48 participants, we performed the analysis with a 6800 ms cutoff, which affected 3% of the data. An additional analysis of all data, with very similar results, is provided in Appendix C. Data from one sentence set was excluded from the analysis, due to a typo.

As in the previous experiments, the completions were coded as zero (0) if they presented non-local attachment, and one (1) if they presented local attachment. We compared local attachment rates in the ObIT condition to those in the OpT condition. In contrast to Experiment 1, the

⁵ The repeated verb was ‘received’. Repetition across sets was required in order to create plausible sentences.

⁶ This pre-test is the same one described in Experiment 1. Namely, the pre-test included a third condition, in addition to the two described above.

condition with the intransitive verb was included in the data analysis in Experiment 3.⁷ We also compared the OblT condition to the IN condition, as we wanted to test whether participants always chose the local attachment (opting for a grammatical sentence) in the obligatorily transitive condition, just as they did in the intransitive condition. To do this, a generalized logistic mixed-effects model was fitted, modeling attachment choices with *Verb type* (OpT / OblT / IN, treatment coded) as a fixed effect predictor, and random effects for participants and items. We changed the reference level to test all three contrasts. As in Experiments 1 and 2, we carried out both frequentist and Bayesian analyses. In the frequentist analysis, a maximal model that included all random slopes was not fitted, due to a convergence problem. The final model we fitted included random intercepts for participants and items. The Bayesian model was fitted as maximal and included all random slopes.

For the OpT and OblT conditions, we also asked whether attachment choices were consistent with a chance distribution. Finally, for the OpT sentences, we asked whether there was an effect of transitivity bias on attachment choices, as in Experiment 1. These questions were examined in the same way as in Experiment 1.

4.3 Results

Overall completion rates are illustrated in **Figure 5** (for data without exclusion, see **Table C1** in Appendix C).




Figure 5: Rate of completions by condition, Experiment 3.

The by-item average rates of local attachment in the Optionally Transitive condition are presented in **Figure 6**, representing the effect of verb transitivity bias on the rate of choosing local attachment. For by-item average rates of local attachment in the Obligatory Transitive condition, see **Figure C** in Appendix C.

⁷ Inclusion of the intransitive (control) conditions in the data analysis in Experiment 1 would not have been informative, since the important comparisons were between the two ambiguous conditions, and between each condition and chance. Also, since the control of the obj/subj sentence was one favoring non-local, rather than local, attachment, the comparison between the ambiguous and control obj/subj sentences was meaningless.

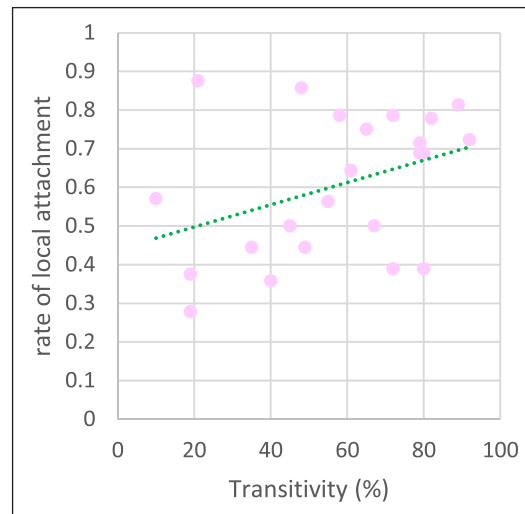


Figure 6: The effect of transitivity bias on the rate of choosing local attachment in the Optionally Transitive condition, Experiment 3.

The results of the frequentist and Bayesian analyses are summarized in **Tables 11** and **12**, respectively. In the frequentist analysis, we observed a main effect of verb type, with significant differences between all three conditions (all p 's < .001), reflecting higher rates of local attachment in the Intransitive condition, compared to both transitive conditions, as well as higher rates of local attachment in the Obligatorily Transitive condition than in the Optionally Transitive condition. The results of the Bayesian analysis were consistent with this, reflecting an increase in proportion of local attachment in the Intransitive condition, relative to both transitive conditions: the mean of the posterior distribution was 63.5%, with CI [54.7%, 72.6%], in the Optionally Transitive condition, 84.1%, with CI [75.6%, 91.3%], in the Obligatorily Transitive condition, and 95.2%, with CI [92.2%, 97.5%], in the Intransitive condition, with none of the credible intervals for the contrasts crossing zero. As for the correlation between the transitivity bias of OpT verbs and attachment preferences in the OpT condition, the results showed that the effect of transitivity bias on local attachment rate was significant in the generalized mixed-effects model ($p = .015$). The result of the Bayesian analysis was consistent with this.

Table 11: Results of the frequentist analysis, Experiment 3: Estimates, standard error, t -values and p -values.

| | Estimate | SE | t | p |
|--|----------|-------|-------|----------|
| OblT condition > OpT | 1.166 | 0.279 | 4.172 | 3.02e-05 |
| IN > OpT | 2.649 | 0.388 | 6.821 | 9.03e-12 |
| IN > OblT | 1.619 | 0.273 | 5.930 | 3.03e-09 |
| Effect of verb bias in the OpT condition | 0.016 | 0.006 | 2.438 | 0.015 |

Table 12: Results of the Bayesian analysis, Experiment 3: Posterior means on the log-odds scale (with 95% Bayesian CrI).

| | Posterior mean | [CrI] | Posterior Beyond zero |
|--|----------------|----------------|-----------------------|
| OblT > OpT | 1.136 | [0.577, 1.740] | 100% |
| IN > OpT | 2.454 | [1.869, 3.088] | 100% |
| IN > OblT | 1.555 | [0.925, 2.232] | 100% |
| Effect of verb bias in the OpT condition | 0.016 | [0.006, 0.027] | 99.8% |

Further, the binomial tests showed that in both the OpT and OblT conditions, the by-participant average choice of attachment was significantly different from the null hypothesis of a chance distribution ($t = 3.260$, $p = .002$ and $t = 10.675$, $p < .001$, correspondingly).

4.4 Discussion

Experiment 3 further examined attachment preferences in sentences with MC/RC ambiguity, testing a proposal made by Pritchett (1988, 1992), namely, that thematic assignment maximization operates locally, without regard to the global obligatoriness of thematic roles. We manipulated the thematic grid of the embedded verb, comparing attachment preferences for ambiguous sentences containing optionally transitive verbs with those for unambiguous sentences containing intransitive verbs, and those for sentences containing obligatorily transitive verbs, which are likewise globally unambiguous.

Similar to Experiment 1, we found that attachment was local in significantly fewer cases in the optionally transitive condition than in the intransitive condition. Interestingly, while local attachment was selected in 53% of the cases in Experiment 1, consistent with a chance distribution, it was selected in 60% of the cases in this experiment, and the responses were not compatible with a chance distribution. In addition, the correlation with transitivity bias in the OpT condition was slightly weaker in Experiment 2, as compared with Experiment 1. These differences may be due to differences in the materials, namely, lack of data for one set which included a typo in Experiment 3, and the addition of three OpT verbs which were included in this experiment and not in the previous one.

For the sentences with obligatorily transitive verbs, we found that while the rate of selecting local attachment was significantly lower, compared with the rate in the intransitive condition (80% compared to 95%), it was significantly higher than that in the OpT condition. The difference in attachment preferences between the OblT and IN sentences demonstrates that attachment decisions were not made based on the Late Closure principle: As the IN sentences included a preposition after the intransitive verb, they always required local attachment; if OblT sentences

followed Late Closure, they would also invariably show local attachment. However, this was not observed.

With regard to Theta Attachment, Pritchett's proposal correctly predicts that there should be a difference in the rates of local attachment between the IN and ObIT conditions: in the latter condition, when the sentence fragment includes two obligatorily transitive verbs, and the Theta Criterion can be satisfied in two different ways, local as well as non-local attachments are considered. This is true even when there is only one globally grammatical option. In fact, in 20% of the trials, participants chose the non-local attachment, giving rise to an ungrammatical sentence. This can be explained if thematic assignment maximization operates locally and incrementally, ignoring the global obligatoriness of thematic roles. The fact that ungrammatical sentences were constructed in 20% of the trials (compared to only 5% in the intransitive case) suggests that participants did locally consider both attachment choices. The higher rates of local attachment in the ObIT condition than in the OpT condition may have two sources. The first is the contribution of verb bias: as shown for OpT verbs in the current experiment, as well as in Experiment 1, more transitively biased verbs in the MC/RC structure lead to more local attachment. As the average verb bias is larger in the ObIT condition than in the OpT condition, more local attachment is predicted in the former. In addition to that, global grammaticality may have had a role, prompting participants towards local attachment (though grammaticality does not have a decisive role, as evidenced by the 20% choices leading to ungrammatical sentences).

While Gibson's (1991) theory, on a par with Theta Attachment, provides an account of the findings in Experiment 1, we rather find more evidence for Theta Attachment in Experiment 3. This is because in Gibson (1991), obligatoriness of thematic roles (and thus, global grammaticality) affects parsing choices: local attachment in the ObIT condition, just as in the OpT condition, has a downside, in that it leaves the obligatory role of the matrix verb unassigned. This contrasts with Pritchett's Theta Attachment, as Theta Attachment holds that such considerations are irrelevant in making the local decision. Therefore, from the theories considered up to now, only Theta Attachment explains why, in 20% of the cases, global grammaticality was not taken into account during attachment decisions.

Another approach which allows parsing decisions which are inconsistent with the global syntactic context is Self-Organized Sentence Processing (Tabor & Hutchins, 2004). In this framework, newly-arrived phrases combine with previous ones in different ways, and the various attachments compete dynamically. Attachments which are inconsistent with the correct global parse are, nonetheless, formed and considered (Tabor et al., 2004). Thus, this view seems compatible with the result we obtained in the ObIT condition. We leave open the question of whether this framework can account for the entirety of our findings.

5. General discussion

The present study investigated the degree to which attachment decisions are guided by thematic considerations. We specifically tested one implementation of this idea, namely, Pritchett's (1988, 1992) principle of Theta Attachment: At every point during comprehension, the parser attempts to maximally satisfy the Theta Criterion. We compared the predictions of this principle with the structural principle of Late Closure, as well as with guidance by verb bias.

Experiment 1 showed that in sentences with object/subject ambiguity, where there is only one attachment choice which enables thematic assignment (10), the ambiguous NP invariably attaches to the preceding verb, maximizing thematic assignment at this point. This result is in line with both Theta Attachment and Late Closure. However, in MC/RC sentences, when both attachment options enable thematic assignment to the ambiguous NP, as in (11), the results show optionality with regard to attachment of this NP, contrary to the predictions of Late Closure, and in line with Pritchett's proposal.

(10) After the guests drank cold water...

(11) The owner brought to the guests that drank cold water...

Additionally, we found that the transitivity bias of the optionally transitive verb was not correlated with attachment choices in sentences like (10), but only in those like (11). We proposed that this is so since in (10), when only one structure enables thematic assignment, this consideration dictates parsing choices. However, in (11), when thematic considerations cannot decide, other factors – i.e., transitivity bias – affect processing choices.

Experiment 2 found that when adverbs were added to sentences with object/subject ambiguity, such that the ambiguous NP was no longer adjacent to the preceding verb, participants still largely opted for local attachment, aiming to maximize thematic assignment, and there was no effect of transitivity bias on attachment choices, replicating Experiment 1. However, the effect of distance was statistically significant, with a higher proportion of local attachment when the NP was adjacent to the preceding verb. Thus, locality and thematic assignment maximization seem to both be at work during structure building.

Experiment 3 further found that when sentences with MC/RC ambiguity, such as (11), included an obligatory transitive subordinate verb, both local and non-local attachments were attested, providing some evidence for Pritchett's assumption that optionality will arise, since the Theta Criterion can be locally satisfied in two different ways, even when there is only one globally grammatical option. However, verb bias (possibly coupled with global grammaticality) influenced parsing choices, such that participants showed a stronger preference for local attachment than in the OpT condition.

Comparing the processing strategies considered in the present study, we can first observe that the structural principle of Late Closure mostly does not account for the results, as it predicts uniform choice of local attachment, regardless of the number of options for thematic assignment, distance, or obligatoriness of thematic roles. Specifically, the variability found in attachment choices in the MC/RC condition, both within and between items, provides evidence against the purely structural approach.

Further, the results show that verb biases do not always guide processing. In the first experiment, the effect of verb bias was qualitatively different between the case with one thematic assignment option and the case with two thematic assignment options, with no effect in the former. This result is hard to explain if verb preferences determine processing choices. In addition, there was no effect of verb bias on attachment choices in Experiment 2.

The experiments suggest a central role for thematic considerations during processing. All three experiments point to the importance of the satisfaction of the Theta Criterion in making parsing decisions, such that structures are selected that maximize thematic assignment. When there is only one structure that maximizes thematic assignment, this structure will be selected in a large majority of the cases, without consideration of the bias of the verb (Experiments 1 and 2), although distance can somewhat reduce this tendency (Experiment 2). In contrast, when both structures satisfy the Theta Criterion to the same degree, as in the MC/RC ambiguity structures in Experiments 1 and 3, other considerations, namely, verb bias, have a much more pronounced effect. Thus, if we compare the effect of transitivity bias when thematic considerations are strong (obj/subj ambiguity, Experiments 1 and 2) vs. weak (MC/RC ambiguity, Experiment 1 and 3), we find a strong effect of bias in the latter case, but no effect in the former. Taken together, the experiments provide evidence for the precedence of thematic assignment and, consequently, of interpretation maximization during processing. Though Pritchett's Theta Attachment does not take into account the influence of distance and verb bias information (or other types of frequency information), its contribution to processing theory is in its implementation of the maximal interpretation hypothesis.

While the current study is the first to investigate MC/RC sentences, much previous work has targeted object/subject ambiguities, and our results largely align with theirs. Most relevant are Pickering et al. (2000, Experiment 3), Itzhak et al. (2010) and Huang et al. (2024), which examined the effect of verb bias in these ambiguities. Notably, Itzhak and colleagues found a correlation of the P600 effect with verb bias in these sentences, though there was processing difficulty regardless of verb bias, whereas Pickering and colleagues and Huang and colleagues did not, similarly to our study. The difference from Itzhak and colleagues may be due to the different measure, or due to Itzhak and colleagues treating verb bias as a binary variable.

While verb bias fails to account for our findings, an open question is whether Surprisal (Hale, 2001; Levy, 2008), which is a more recent probabilistic theory, can account for them. Surprisal

assumes a fully parallel parser which uses the statistics of the language to construct parallel parses of the input, with their weights reflecting the probability of the parse in the language. The focus of Surprisal is on the prediction of localized processing difficulty as reflected in reading times, with prolonged reading times for less predictable input. This is irrelevant to the current study, as our experiments did not measure reading times, but rather attachment choices (but see Huang et al., 2024, for testing the predictions of Surprisal in Garden Path sentences). Surprisal does not assume attachment “choices”, since multiple structures are built at any point. However, it is possible to assume that the preferred completion will be the one which is compatible with the most highly weighted structure to this point. To test whether Surprisal can account for our results, it would be necessary to extract the probabilities of the different continuations in the different ambiguity types from a Hebrew corpus. We leave this for future research. We do note, however, that in the ObIT condition in Experiment 3, readers opted in 20% of trials for completions resulting in globally ungrammatical structures, which have very low frequency in the language. It seems that at least in this case, completion options do not track corpus probabilities.

As the sentences we used, especially in the MC/RC conditions, were very long and complex, another model which might be considered to account for their processing is lossy-context surprisal (Futrell et al., 2020; Hahn, Futrell et al., 2022). This theory adopts Surprisal’s assumptions, but proposes that predictions are based not on a complete, accurate representation of the sentence up to this point, but rather on a lossy, noisy version of the sentence. According to such a model, it can be assumed that readers will sometimes fail to assign the thematic role of the main clause verb in MC/RC sentences not because of optionality, but rather because the memory representation of the first verb is degraded or weak, due to memory capacity and the low probability of this construction. However, we would like to point out that a “forgetting” scenario should lead to the same problem in the MC/RC control condition in Experiment 1 / Intransitive condition in Experiment 3, repeated in (12). Namely, participants would “forget” the ditransitive verb and select the adverb completion. However, participants overwhelmingly selected the NP completion in the experiments (95% in both experiments), showing that they did remember the beginning of the sentence. Future research can evaluate lossy-context surprisal against our data in more detail.

(12) The owner brought to the guests that showered with cold water [orange juice] / [last night at the farm]

Based on the present experiments, we suggest that any processing model should capture the influence of thematic information and the goal of maximizing thematic assignment and interpretation. Other important factors influencing processing, such as thematic biases of verbs and their distance from their argument, have a secondary role, but do not eliminate the importance of Theta Attachment and maximal interpretation.

Appendix A. Analysis of all data from Experiment 1

Table A1 presents the rates of choosing local attachment in Experiment 1 when no trials are excluded based on long RTs.

Table A1: Raw data: the rate of choosing local attachment (SD).

| Condition | Mean choice of local attachment (SD) |
|-------------------|--------------------------------------|
| Obj/subj | 0.845 (0.363) |
| Obj/subj, control | 0.0402(0.197) |
| MC/RC | 0.523(0.500) |
| MC/RC, control | 0.948(0.222) |

Tables A2 and **A3** present the results of the frequentist analysis and Bayesian analyses, respectively, of all data in Experiment 1 (without trial exclusion based on RT).

Table A2: Results of the frequentist analysis, Experiment 1, raw data.

| | Estimate | SE | <i>t</i> | <i>p</i> |
|---|----------|-------|----------|------------|
| Grand Mean (Intercept) | 0.982 | 0.162 | 6.061 | |
| Effect of ambiguity type <i>Obj/subj</i> > <i>MC/RC</i> | 1.765 | 0.232 | 7.599 | < 0.001*** |
| Overall effect of verb transitivity bias | 0.012 | 0.006 | 2.074 | 0.038* |
| Interaction between verb bias and ambiguity type | -0.012 | 0.009 | -1.415 | 0.157 |
| Effect of verb bias in the <i>obj/subj</i> condition | 0.006 | 0.008 | 0.685 | 0.492 |
| Effect of verb bias in the <i>MC/RC</i> condition | 0.018 | 0.007 | 2.719 | 0.007** |

Table A3: Results of the Bayesian analysis, Experiment1, raw data.

| | log odds scale | |
|---|----------------|-----------------|
| | posterior mean | [CrI] |
| Grand Mean (Intercept) | 1.025 | [0.66, 1.422] |
| Effect of ambiguity type <i>obj/subj</i> > <i>MC/RC</i> | 1.866 | [1.312, 2.44] |
| Overall effect of verb transitivity bias | 0.014 | [-0.001, 0.032] |
| Interaction between verb bias and ambiguity type | -0.019 | [-0.05, 0.011] |
| Effect of verb bias in the <i>obj/subj</i> condition | 0.008 | [-0.012, 0.029] |
| Effect of verb bias in the <i>MC/RC</i> condition | 0.022 | [0.005, 0.041] |

Appendix B. Distribution of direct objects with and without preceding adverbs

To verify whether in Hebrew obj/subj sentences, the probability that the ambiguous NP opens a new clause is larger when it is preceded by an adverb than when it follows the verb directly, we manually coded a random sample of 100 occurrences of four verbs in Sketch Engine (Hebrew Web 2014 (heTenTen14, Meni/Alon tagged + lempos)). We selected four representative experimental verbs, which are in different places on the transitivity scale. We counted occurrences of a direct object when an adverb was present after the verb, compared to when there was no adverb. The results showed that, indeed, occurrence of a post-verbal adverb indicates that the chance of a direct object is smaller, compared to when there is no adverb.

Table B1: Direct object NP occurrence when an adverb is present after the verb compared to when it is absent.

| Verb | Total % occurrence with direct object | Total % of occurrence with adverb | % occurrence with direct object when adverb does not appear after verb | % occurrence with direct object when adverb appears after the verb |
|-----------------------------|---------------------------------------|-----------------------------------|--|--|
| רקד dance | 16 | 36 | 21 | 8.6 |
| בישל cook | 46 | 38 | 54 | 12 |
| אכל eat | 69 | 14 | 74 | 58 |
| ניתח operate/ analyze | 89 | 15 | 91 | 62 |

Appendix C. Analysis of all data from Experiment 3

Table C1 presents rates of choosing local attachment in Experiment 3 when no trials are excluded based on long RTs.

Table C1: Rate of choosing local attachment in Experiment 3 when no trials are excluded based on long RTs.

| Condition | Choice of local attachment (SD) |
|------------------------------|---------------------------------|
| Optionally transitive (OpT) | 0.589 (0.490) |
| Obligatory transitive (ObIT) | 0.795 (0.404) |
| Intransitive | 0.943 (0.232) |

Tables C2 and **C3** present the results of the frequentist analysis and Bayesian analysis, respectively, of all data in Experiment 3 (without trial exclusion based on RT).

Table C2: Results of the frequentist analysis, Experiment 3, raw data.

| | Estimate | SE | <i>t</i> | <i>p</i> |
|--|----------|--------|----------|----------|
| OblT > OpT | 1.0216 | 0.1785 | 5.723 | 1.05e-08 |
| IN > OpT | 2.6272 | 0.2617 | 10.040 | < 2e-16 |
| IN > OblT | 1.6056 | 0.2678 | 5.995 | 2.03e-09 |
| Effect of verb bias in the OpT condition | 0.015 | 0.007 | 2.274 | 0.023* |

Table C3: Results of the Bayesian analysis, Experiment 3, raw data.

| | Posterior mean | [CrI] |
|--|----------------|----------------|
| OblT > OpT | 1.136 | [0.576, 1.70] |
| IN > OpT | 2.433 | [1.869, 3.023] |
| IN > OblT | 1.507 | [0.873, 2.147] |
| Effect of verb bias in the OpT condition | 0.015 | [0.004, 0.026] |

Figure C below adds an information concerning the distribution of the by-item rates in the OblT condition, for trials whose RTs < 6800 ms.

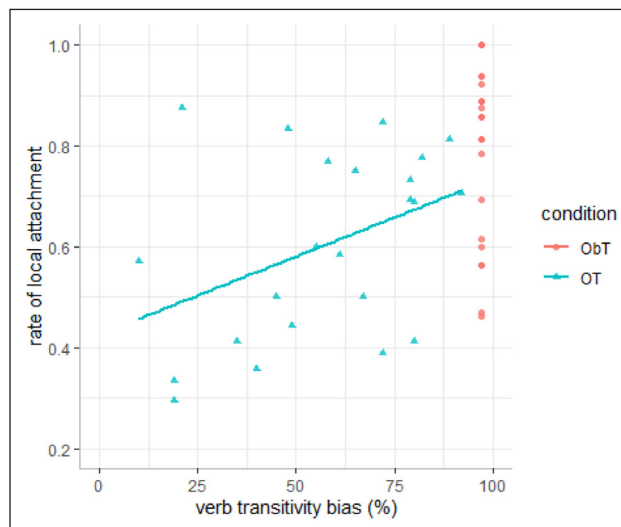


Figure C: The rate of choosing local attachment as a function of the transitivity bias in the Optionally Transitive and Obligatory Transitive conditions, Experiment 3.

Data accessibility statement

All data and R code are available in the following OSF online repository: https://osf.io/uxnmc/?view_only=199ede8bf030465aac9da1187b909b59.

Ethics and consent

Ethical approval for the experiments of this study was granted by the Institutional Ethics Committee at Tel Aviv University; Reference numbers 0000032-3, 0000032-4 and 0000032-5.

Acknowledgements

This research was funded by the Israel Science Foundation, grant 525/22 (PI: Aya Meltzer-Asscher). We thank audiences at AMLaP 27 and HSP 36, and at the workshop of the Sentence Processing and LISC labs at Tel-Aviv University, for their valuable feedback. Lastly, we thank the editor and three reviewers for insightful comments on earlier versions of this article, which have helped us improve it substantially.

Competing interests

The authors have no competing interests to declare.

Authors' contributions

Lola Karsenti: conceptualisation, methodology, software, investigation, data curation, formal analysis, visualization, writing.

Aya Meltzer-Asscher: conceptualisation, methodology, investigation, funding acquisition, supervision, writing, review & editing.

ORCID IDs

Lola Karsenti: orcid.org/0009-0004-7016-4665

Aya Meltzer-Asscher: orcid.org/0000-0003-4886-0029

References

- Adams, B., Clifton, C., & Mitchell, D. (1998). Lexical guidance in sentence processing? *Psychonomic Bulletin & Review*, 5(2), 265–270. <https://doi.org/10.3758/BF03212949>
- Aoshima, S., Phillips, C., & Weinberg, A. (2004). Processing filler-gap dependencies in a head-final language. *Journal of Memory and Language*, 51, 23–54. <https://doi.org/10.1016/j.jml.2004.03.001>

- Bates D., Mächler M., Bolker B., & Walker S. (2021). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Brothers, T., Hoversten, L. J., & Traxler, M. J. (2021). Bilinguals on the garden-path: Individual differences in syntactic ambiguity resolution. *Bilingualism: Language and Cognition*, 24(4), 612–627. <https://doi.org/10.1017/S1366728920000711>
- Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80(1), 1–28. <https://doi.org/10.18637/jss.v080.i01>
- Bürkner, P. C. (2018). Advanced Bayesian multilevel modeling with the R package brms. *arXiv preprint arXiv:1705.11123*. <https://doi.org/10.32614/RJ-2018-017>
- Chomsky, N. (1981). *Lectures on government and binding*. Foris. <https://doi.org/10.1515/9783110884166>
- Crocker, M. W. (1992). *A logical model of competence and performance in the human sentence processor* [Unpublished doctoral dissertation]. University of Edinburgh.
- de Vincenzi, M. (1991). *Syntactic parsing strategies in Italian*. Kluwer Academic Publishers. <https://doi.org/10.1007/978-94-011-3184-1>
- Drummond, A. (2007–2021). Ibex: Internet-based experiments. <https://ibex.spellout.net>
- Ferreira, F., & Henderson, J. M. (1991). Recovery from misanalyses of garden-path sentences. *Journal of Memory and Language*, 30, 725–745. [https://doi.org/10.1016/0749-596X\(91\)90034-H](https://doi.org/10.1016/0749-596X(91)90034-H)
- Fodor, J. D., & Frazier, L. (1980). Is the human sentence parsing mechanism an ATN? *Cognition*, 8, 417–459. [https://doi.org/10.1016/0010-0277\(80\)90003-7](https://doi.org/10.1016/0010-0277(80)90003-7)
- Frazier, L. (1978). *On comprehending sentences: Syntactic parsing strategies* [Doctoral dissertation]. University of Connecticut. <https://digitalcommons.lib.uconn.edu/dissertations/AAI7914150>
- Frazier, L. (1987a). Sentence processing: A tutorial review. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 559–586). Lawrence Erlbaum Associates, Inc.
- Frazier, L. (1987b). Syntactic processing: Evidence from Dutch. *Natural Language and Linguistic Theory*, 5, 519–559. <https://doi.org/10.1007/BF00138988>
- Frazier, L., & Clifton, C. (1996). *Construal*. MIT Press.
- Frazier, L., & Clifton, C. (1998). Sentence reanalysis and visibility. In J. D. Fodor & F. Ferreira (Eds.), *Reanalysis in sentence processing*. Kluwer. https://doi.org/10.1007/978-94-015-9070-9_5
- Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, 14(2), 178–210. [https://doi.org/10.1016/0010-0285\(82\)90008-1](https://doi.org/10.1016/0010-0285(82)90008-1)
- Futrell, R., Gibson, E., & Levy, R. P. (2020). Lossy-context surprisal: An information-theoretic model of memory effects in sentence processing. *Cognitive Science*, 44(3), e12814. <https://doi.org/10.1111/cogs.12814>
- 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: Papers from the first Mind Articulation Project Symposium* (pp. 95–126). MIT Press. <https://doi.org/10.7551/mitpress/3654.003.0008>

- Gibson, E., Hickok, G., & Schütze, C. T. (1994). Processing empty categories: A parallel approach. *Journal of Psycholinguistic Research*, 23, 381–405. <https://doi.org/10.1007/BF02143946>
- Gibson, E. A. F. (1991). *A computational theory of human linguistic processing: Memory limitations and processing breakdown* [Doctoral dissertation]. Carnegie Mellon University.
- Hahn, M., Futrell, R., Levy, R., & Gibson, E. (2022). A resource-rational model of human processing of recursive linguistic structure. *Proceedings of the National Academy of Sciences of the United States of America*, 119(43), 1–9. <https://doi.org/10.1073/pnas.2122602119>
- Hale, J. (2001). A probabilistic Earley parser as a psycholinguistic model. In *Second meeting of the North American Chapter of the Association for Computational Linguistics*. <https://doi.org/10.3115/1073336.1073357>
- Harley, H. (2010). Thematic roles. In P. Hogan (Ed.), *The Cambridge encyclopedia of the language sciences* (pp. 861–862). Cambridge University Press.
- Huang, K. J., Arehalli, S., Kugemoto, M., Muxica, C., Prasad, G., Dillon, B., & Linzen, T. (2024). Large-scale benchmark yields no evidence that language model surprisal explains syntactic disambiguation difficulty. *Journal of Memory and Language*, 137. <https://doi.org/10.1016/j.jml.2024.104510>
- Itzhak, I., Pauker, E., Drury, J. E., Baum, S. R., & Steinhauer, K. (2010). Event-related potentials show online influence of lexical biases on prosodic processing. *NeuroReport*, 21(1), 8–13. <https://doi.org/10.1097/WNR.0b013e328330251d>
- Kamide, Y., Altmann, G. T. M., & Haywood, S. L. (2003). The time-course of prediction in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language*, 49(1), 133–156. [https://doi.org/10.1016/S0749-596X\(03\)00023-8](https://doi.org/10.1016/S0749-596X(03)00023-8)
- Karsenti, L., & Meltzer-Asscher, A. (2022). Prediction of successful reanalysis based on eye-blink rate and reading times in sentences with local ambiguity. *Language and Cognition*, 14(4), 552–574. <https://doi.org/10.1017/langcog.2022.17>
- Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, 106(3), 1126–1177. <https://doi.org/10.1016/j.cognition.2007.05.006>
- Lewis, R. L., & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, 29(3), 375–419. https://doi.org/10.1207/s15516709cog0000_25
- Mitchell, D. C. (1987). Lexical guidance in human parsing: Locus and processing characteristics. In M. Coltheart (Ed.), *Attention and performance XII* (pp. 601–618). Erlbaum, Inc.
- Pickering, M. J., Traxler, M. J., & Crocker, M. W. (2000). Ambiguity resolution in sentence processing: Evidence against frequency-based accounts. *Journal of Memory and Language*, 43, 447–475. <https://doi.org/10.1006/jmla.2000.2708>
- Pritchett, B. L. (1988). Garden path phenomena and the grammatical basis of language processing. *Language*, 64, 539–576. <https://doi.org/10.2307/414532>
- Pritchett, B. L. (1992). *Grammatical competence and parsing performance*. University of Chicago Press.

- R Core Team. (2020). R: A language and environment for statistical computing (Version 4.0.0) [Software]. R Foundation for Statistical Computing. <https://www.R-project.org/>
- Schneider, D., & Phillips, C. (2001). Grammatical search and reanalysis. *Journal of Memory and Language*, 45(2), 308–336. <https://doi.org/10.1006/jmla.2001.2777>
- Sikos L., Duffield, C. J., & Kim, A. E. (2016). Grammatical predictions reveal influences of semantic attraction in online sentence comprehension: Evidence from speeded forced-choice sentence continuations. *Language, Cognition and Neuroscience*, 31(8), 1055–1073. <https://doi.org/10.1080/23273798.2016.1186808>
- Staub, A. (2007). The parser doesn't ignore intransitivity, after all. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(3), 550–569. <https://doi.org/10.1037/0278-7393.33.3.550>
- Staub, A. (2009). On the interpretation of the number attraction effect: Response time evidence. *Journal of Memory and Language*, 60(2), 308–327. <https://doi.org/10.1016/j.jml.2008.11.002>
- Sturt, P., & Crocker, M. W. (1997). Thematic monotonicity. *Journal of Psycholinguistic Research*, 26(3), 297–322. <https://doi.org/10.1023/A:1025076608204>
- Sturt, P., Pickering, M. J., Scheepers, C., & Crocker, M. W. (2001). The preservation of structure in language comprehension: Is reanalysis the last resort? *Journal of Memory and Language*, 45(2), 283–307. <https://doi.org/10.1006/jmla.2001.2776>
- Tabor, W., Galantucci, B., & Richardson, D. (2004). Effects of merely local syntactic coherence on sentence processing. *Journal of Memory and Language*, 50(4), 355–370. <https://doi.org/10.1016/j.jml.2004.01.001>
- Tabor, W., & Hutchins, S. (2004). Evidence for self-organized sentence processing: Digging in effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(2), 431–450. <https://doi.org/10.1037/0278-7393.30.2.431>
- Trueswell, J. C., Tanenhaus, M. K., & Kello, C. (1993). Verb-specific constraints in sentence processing: Separating effects of lexical preference from garden paths. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19, 528–553. <https://doi.org/10.1037/0278-7393.19.3.528>
- Van Dyke, J. A., & Lewis, R. L. (2003). Distinguishing effects of structure and decay on attachment and repair: A cue-based parsing account of recovery from misanalyzed ambiguities. *Journal of Memory and Language*, 49(3), 285–316. [https://doi.org/10.1016/S0749-596X\(03\)00081-0](https://doi.org/10.1016/S0749-596X(03)00081-0)

