DISCOURSE PROCESSES, *43*(2), 153–172

Negation in Context: A Functional Approach to Suppression

Rachel Giora
Linguistics
Tel Aviv University

Ofer Fein
School of Behavioral Sciences
The Academic College of Tel Aviv Yaffo

Keren Aschkenazi and Inbar Alkabets-Zlozover
Cognitive Studies of Language Use
Tel Aviv University

Three experiments show that, contrary to the current view, comprehenders do not unconditionally deactivate information marked by negation. Instead, they discard negated information when it is functionally motivated. In Experiment 1, comprehenders discarded negated concepts when cued by a topic shift to dampen recently processed information. However, in the presence of a global cue suggesting topic continuity, they retained it, despite a local negation marker that might prompt it. Specifically, when negative statements (*The train to Boston was **no rocket***; Hasson & Glucksberg, 2006) were furnished with relevant (compared to irrelevant) subsequent contexts (*The trip to the city was **fast** though*), incompatible meanings ('fast'), related to the affirmative sense of the negative metaphor (*rocket*), were not suppressed. Instead they were retained and primed related targets (*fast*) appearing in the late context. Experiment 2 showed that preceding contexts had similar effects, inducing retention of probes related to the affirmative meaning of a negated target. Such effects, however, waned after a lengthy delay (Experiment 3).

Correspondence should be addressed to Rachel Giora, Department of Linguistics, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: giorar@post.tau.ac.il

*Hitler cannot **win** the war, he can only **prolong** it!* (Leaflets of the *White Rose*—Resistance in Germany: A call to all Germans! 1943)

Does negation necessarily cue comprehenders to replace a negated concept by an available opposite? Was it the case, then, that when informed by the French official that Arafat was "not dead," comprehenders activated 'alive' while deactivating *dead*? Most probably they did not represent the dying Palestinian Chairman as *alive and kicking*. Along the same lines, is it the case that, when Szymborska (1996) wrote of those she does *not love*, she intended her readers to suppress *love* and activate 'hate'? Apparently not, as her poem discloses:[1]

(1) A "Thank You" Note

There is much I owe
to those I do **not love**.
[…]
My peace be with them
for with them I am free,
and this, love can neither give,
nor know how to take.
[…]
My trips with them always turn out well.
Concerts are heard.
Cathedrals are toured.
Landscapes are distinct.
(Reprinted from *Miracle Fair* by Wislawa Szymborska, p. 24, translated by Joanna Trzeciak. Copyright 2001 by Joanna Trzeciak. With the permission of the publisher, W. W. Norton & Company, Inc.)

The evidence from natural conversations and written texts suggests that speakers do not always want to get across an opposite alterative when using negation. For example, in the following, what is introduced by negation (*not to mention*) is an explicit list of all that the speaker intends the addressee to note rather than ignore:

(2) **Not to mention** *the interminable trip to the hospital* and *the seven days in which he lay dying* while most of his family members were not permitted to come visit, and then *the macabre trip home*, with the corpse in the back of

[1]Throughout the article emphases are added for convenience.

the ambulance, traveling halfway round the West Bank to get around the checkpoints (Levy, 2003).

No wonder comprehenders often perceive negation as affirmation (*this is not a personal attack* = "this is a personal attack"), as the following example demonstrates:

(3) P: … it was very clear. You know.
… She kept saying,
… prefacing everything with,
… you know, this is *not a personal attack*.
… This is *not a personal vendetta*,
B: Yeah, yeah yeah yeah.
Right Right.
P: **Which tells you, that it is**.
B: Yeah.
P: **That's immediately what it said.**
And that's what everybody perceived it.
B: Yeah.
(Du Bois, 2000; as cited in Giora, Balaban, Fein, & Alkabets, 2005, p. 234).

Indeed, information within the scope of negation is often assumed available to both speakers and addressees (see Giora, 2006). In the following example, information within the scope of negation is retained in the mind of the speaker (B). Consequently, it features in her next discourse segment where it is treated as given (as indicated by initial position and intonation):

(4) A: Listen, with your car, you are there, at most, in 5 minutes …
B: Come on … My Daihatsu is **not a jet**. A **fast** car … superb car … But there's a limit … (originally Hebrew, cited in Giora, 2006, p. 999)

That speakers assume that negated information is available to addressees is evidenced by the use of high accessibility markers, such as pronouns and zeros, when referring to this information (Ariel, 1990). In the following, the negated information is considered so highly accessible that it can be referred to by a pronoun (5a) or even be elided (5b; indicated by square brackets for convenience):

(5a) No disabled included in Israeli delegation to United Nations debate on **their** rights (Sinai, 2006).

(5b) We are **not opposed** to the existence of Israel, but [] to its actions (Zandberg, 2006).

Despite such evidence, the received view among linguists and psycholinguists is that negation is a reducing availability operator (for a review, see Giora, 2006). It is perceived as an instruction from a communicator to an addressee to eliminate the negated concept from the mental representation and replace it with an alternative opposite. Such elimination—practically the reduction of the concept's levels of activation to baseline levels or below—is considered obligatory. It should occur late in the comprehension process (about 500 ms after offset of the target constituent) and should result in a focus shift from the negated concept to an emerging opposite.[2] For example, in MacDonald and Just (1989, Experiments 1–2), participants were faster to recognize and name the probe word *bread* following *Every weekend, Mary bakes some bread but no cookies for the children* than following *Every weekend, Mary bakes no bread but only cookies for the children*. Such findings attest to the reduced accessibility of the concepts within the scope of negation. Inconsistently, though, as demonstrated by Experiment 3, when probes were not words appearing in the target sentences but close associates, such as *butter* related to *bread* and tested following both *no bread* and *some bread*, no suppression was demonstrated. Instead, the associated concepts have been found to be primed by the negated concepts. Such findings render the obligatory view of suppression suspect.

It is important to note that MacDonald and Just (1989) also examined early processes. In all their studies, the early lexical access stage was tapped by the use of moving windows (Just, Carpenter, & Woolley, 1982). Results showed that, at this early processing stage, reading times of affirmative (*some bread*) and negative (*no bread*) concepts did not vary. These findings indicate that, initially at least, processing is insensitive to negation, as shown by a great number of studies using different methodologies (Giora et al., 2005; Hasson & Glucksberg, 2006; Kaup, Yaxley, Madden, Zwaan, & Lüdtke, in press; see Giora, 2006, for a review). As opposed to the early reading times, the response and naming times were measured rather late in the comprehension process (a few words following the target word) and partially attested to reduced accessibility of concepts activated initially.

An important piece of research, looking into the time course of processing negation, also suggests that suppression following negation is a default strategy (Hasson & Glucksberg, 2006). In this study, Hasson and Glucksberg examined negated (compared to nonnegated) metaphors, presented out of a specific context.

---

[2]Following Gernsbacher (1990), we take any reduction of initial levels of activation to baseline levels and below to indicate suppression, whether this involves an active effort or an effortless focus shift whereby comprehenders dispose of recently processed information to allocate resources to the initiation of the next concept or substructure. In both cases, there is an elimination of information from the mental representation.

Findings showed that, at relatively short interstimulus intervals (ISIs) such as 150 and 500 ms, only facilitation of incompatible concepts (*fast*) was demonstrated following negation. However, 1,000 ms after offset of negative statements (*The train to Boston was no rocket*), there was no facilitation of either of the related meanings (*fast/slow*; as opposed to preserved accessibility of the appropriate related meaning following affirmative statements). Deactivating incompatible meanings to baseline levels, then, took quite a long time to become effective—between 500 to 1,000 ms following offset of the negated concepts. Such findings are consistent with the view that, given enough processing time and in the absence of a specific context, negation (*no rocket*) deactivates the affirmative meaning of the negated concept (*fast*). These findings, however, do not demonstrate that, after such a long delay, deactivation allows focus to shift from the negated concept to an alternative interpretation (*slow*).

This, however, has been demonstrated by Kaup, Lüdtke, and Zwaan (2006). As predicted by Hasson and Glucksberg (2006), in Kaup et al., concepts within the scope of negation lost accessibility 750 ms following their offset. However, at an ISI of 1,500 ms, they were replaced by alternative opposites. Indeed, after such a lengthy delay, negating concepts, presented in isolated sentences, shifted focus from the negated concept (*The door is not **open***) toward an alternative opposite ('The door is **closed**').

A number of studies have attempted to outline the conditions under which a negation marker would obligatorily suppress information within its scope, regardless of context. For instance, Fillenbaum (1966) showed that dichotomous concepts (*alive/dead*) were relatively receptive to suppression following negation and tended to be remembered in terms of their opposite alternative, referred to as *gist*. Consequently, recall errors following negation of dichotomous concepts (*not alive*) indicated suppression of surface information (*alive*) and memory for the gist antonym ('dead'; for rather limited support for the suppressability of dichotomous concepts, see Paradis & Willners, 2006). In contrast, scalar concepts (*warm/cold*) exhibited a different tendency. When negated, these concepts (*not warm*) showed traces of verbatim memory (*warm*) when misrecalled.

Mayo, Schul, and Burnstein (2004) investigated different aspects of negated concepts. They studied negation effects on bipolar and unipolar descriptions. Bipolar concepts, which are not necessarily dichotomous, have a ready-made antonym (*tidy/messy*) at their disposal; unipolar concepts do not (*adventurous/not adventurous*). Mayo et al. asked their participants to judge whether a description following a negated concept was either congruent or incongruent with a previous description. They found that negative bipolar descriptions gave rise to an alternative antonym. In contrast, unipolar descriptions retained their negated kernel (*adventurous*). Thus, having read a sentence involving a bipolar adjective such as *Tom is not a tidy person*, participants were faster to judge *Tom forgets where he left his car keys* as congruent than to judge *Tom's clothes are folded neatly in his closet* as

incongruent. However, the opposite was true of sentences involving unipolar adjectives. Thus, having read *Roy is not an adventurous person*, participants were slower to judge *Roy is stressed by any change in his life* as congruent than *Roy loves to travel to distant places* as incongruent. Findings in Mayo et al. showed, then, that the availability of a complement concept (*messy*) facilitates rejection of the negated antonym (*tidy* in *not tidy*) and allows a focus shift to a complement concept ('messy').

Most of the research demonstrating suppression effects following negation has investigated sentences in isolation. Such studies cannot address the issue of context effects on the processing of negation. In contrast, results from the few studies looking into negation in context demonstrate that suppression following negation is not obligatory but sensitive to discourse considerations. For instance, Kaup and Zwaan (2003), who introduced a new contextual variable—presence in or absence from the situation model—showed that what eventually counts is not so much the presence or absence of a negation marker but the presence or absence of the negated concept from the situation described. Objects dispelled from the situation lost accessibility; objects existing in the situation gained accessibility, regardless of negation. Thus, 1,500 ms after participants had read target sentences, they were faster to respond to a color probe that was present in the situation than to a color probe that was not, irrespective of negation.

Glenberg, Robertson, Jansen, and Johnson-Glenberg (1999) showed that, in a supportive context, relevant to an oncoming message, oncoming utterances, including negated and nonnegated items, were processed along the same lines. Thus, in a supportive context that mentioned that the choice of color of a new couch was important, the next sentence, which was either *The couch was black* or *The couch wasn't black* took similarly long to read (after correcting for sentence length). However, in a nonsupportive context, sentences involving negation took longer to read than their positive counterparts. Such findings attest to the facilitative effects of relevant prior context on the processing of negation.

More recent research by Lüdtke and Kaup (2006) replicated Glenberg et al.'s (1999) results under more specific and stricter conditions. It showed, first, that early context, featuring an explicit mention of information to be rejected later on in the subsequent discourse segment, facilitated negated information relative to a neutral control. Nonnegated information, however, did not benefit from such a mention. Interestingly, such facilitative effects applied even when prior information was not explicit but only highly suggestive of the concepts to be negated later on. Thus, relevance to a strongly implied prior expectation (e.g., for a kid's dirty shirt in the context of a playground) facilitated rejection of this expectation when it was not met.

So far, then, research into negation effects has adduced evidence supporting a limited view of suppression, depending on the nature of the negated concept, the processing time allowed, the presence or absence of contextual information, and the nature of that information. Although dichotomous and bipolar concepts seem

receptive to suppression effects, unipolar and scalar concepts seem more resistant to such effects, and whereas the absence of a specific context coupled with enough processing time indicates that suppression following negation is obligatory, the presence of a specific context demonstrates sensitivity to discourse factors (such as presence or absence of concepts from the situation described). Other studies showing that context is highly relevant for negation also invite a further look at the functional aspects of suppression.

In this study, we further investigated the effects of context on the retention and suppression of negated information. We first tested these effects with regard to late context—context subsequent to a negative utterance (Experiment 1). We then tested contextual effects with regard to early context—context preceding a negative utterance (Experiments 2 and 3).

THE WHENS AND HOWS OF CONTEXT EFFECTS

Traditionally, research in pragmatics, discourse analysis, and psycholinguistics has examined how interpreting a given discourse segment is affected by early context. For instance, researchers have studied how prior context is instrumental in resolving the ambiguity of a given lexical item, whether via suppression of contextually incompatible meanings (Gernsbacher, 1990) or via narrowing down underspecified senses (Carston, 2002; Frisson & Pickering, 2001). In this study, we focus on the effect of both prior and late context on processes following negation. Departing from recent models of negation (for a review, see Giora, 2006), but following functional approaches to suppression (Gernsbacher, 1990; Giora, 2003, 2006; Giora & Fein, 1999; Keysar, 1994), we propose that suppression following negation is not obligatory (as found for contextless targets by, e.g., Hasson & Glucksberg, 2006; Kaup, 2001; Kaup et al., 2006; MacDonald & Just, 1989). Rather, in the presence of specific contextual information (type of task included), comprehenders are sensitive to global discourse considerations rather than to local cues such as negation.

Based on Gernsbacher's (1990) Structure Building Framework, our experiments tested the view that sensitivity to global discourse considerations, shown to apply to nonnegated information (Gernsbacher, 1990), should also apply to negated concepts and should override sensitivity to local cues such as negation. Thus, although suppression might at times be insensitive to negation, it should always be affected by more general discourse operations such as a change of topic or schema. It will therefore be triggered when an oncoming message signals a shift and consequently a need to initiate a new substructure. According to Gernsbacher, when a new substructure, triggered by a topic or schema shift, is being built, the accessibility of recently processed information is reduced, because cognitive resources are allocated to the initiation of the new substructure. On the basis of the Structure

Building Framework, we predict that information recently comprehended will not be discarded, even when negated, in the case that the oncoming discourse segment does not alert the comprehender to a change of topic or schema. However, if such a change is signaled, recently processed information would be disposed of. Similarly, if information within the scope of negation is relevant to prior context, it will be retained so as to be mappable on the substructure recently being processed, following which its accessibility will be reduced.

Is there a way to predict whether early and late relevant contexts might differ in how they affect the time course of retention and integration of negated information? Our studies have not been designed to resolve this question. However, it should come as no surprise were they to support the view that backward coherence is faster to establish than forward coherence. As shown by Gernsbacher (1990), comprehenders "quickly forget the exact form of recently comprehended information" (p. 72) once this information has been integrated with prior context, so that cognitive resources might be available for the initiation of the next substructure. Research has demonstrated that comprehenders are just as fast at verifying some types of information they assumed through backward inferences as they are at verifying information that was explicitly stated. In contrast, comprehenders are much slower at verifying information assumed through predictive or elaborative inferences (Singer, 1979, 1980; Singer & Ferreira, 1983). As shown by a number of studies (Duffy, 1986; McKoon & Ratcliff, 1986; O'Brien, Shank, Myers, & Rayner, 1988; Potts, Keenan, & Golding, 1988), backward-coherence inferences "are, in general, more likely to be drawn than predictive or elaborative inferences" (Gernsbacher, 1990, p. 80).

In our study, we tested the effects of forward coherence in Experiment 1. We used Hasson and Glucksberg's (2006) English items. Whereas Hasson and Glucksberg showed that, out of a specific context, suppression took place between 500 and 1,000 ms following negation, we found that comprehenders retained such information as long as 1,000 ms, until oncoming messages signal whether it might be usable or dispensable.

The effects of backward coherence were studied in Experiments 2 and 3. In these experiments, we used Giora et al.'s (2005) Hebrew items. Whereas Giora et al. attested to the availability of negated concepts, presented in isolated sentences, as early as 100 ms after offset of the negated concepts, we found that comprehenders retained this information even as long as 750 ms following its offset. However, between 750 and 1,000 ms following its mention, once this information was mapped onto a relevant prior substructure, it was no longer available.

In all, the experiments reported here test the hypothesis that global discourse considerations override local ones, particularly those attributed to negation. Such a view assumes that retention and suppression of information within the scope of negation are not obligatory processes and do not operate unconditionally. Rather,

they are context sensitive: They are suspended until global discourse requirements invite them (for a similar view, see also Garnham, 1992; Giora, 2003, 2006).

EXPERIMENT 1

To examine the possibility that global mechanisms reign supreme and neutralize what seem obligatory local mechanisms, we furnished Hasson and Glucksberg's (2006) negative metaphors with late contexts, which either cohered or did not cohere with these items. We assumed that if the processing of a relevant oncoming message is facilitated by negated information in prior context whereas the processing of an incoherent message is not, this will support the superiority of global (coherence) over local (negation) mechanisms. In other words, if such results are indeed obtained, they will demonstrate that suppression does not operate unconditionally but is rather sensitive to global discourse demands.

Experiment 1 thus aimed to show that related targets (*fast*) would be primed by the negative metaphors (*The train to Boston was* ***no rocket***) in coherent but not in incoherent late contexts. That is, in spite of it being within the scope of negation, the affirmative sense ('fast') of a given negative metaphor (*no rocket*) would be retained and prime a related target (*fast*) in an oncoming coherent but not in an oncoming incoherent string. Whereas a coherent discourse might benefit from information mentioned earlier in the discourse, an incoherent segment involving a change of discourse topic should dampen information mentioned earlier (Gernsbacher, 1990).

Method

Participants. Thirty two undergraduates of the American Program at the Rothberg International School of the Hebrew University of Jerusalem served as volunteer participants; 12 were women and 20 were men, and their ages ranged from 18 to 25. They were all native speakers of American English and had a normal or corrected-to-normal vision.

Materials. Materials comprised 160 strings each containing a target sentence and a late context. In addition, there were 16 yes–no comprehension questions. The 160 strings were made up of (a) the 32 negative metaphors used in Hasson and Glucksberg's (2006) study, which here were furnished with two different late contexts, thus making up 64 experimental items; (b) a set of 64 made-up controls involving the same late contexts as the metaphors in (a) but not the metaphors; and (c) 32 made-up fillers. For the sentences containing negative metaphors, late contexts either cohered (6) or did not cohere (7) with the negative items. That is, they either kept discussing the same discourse topic or shifted to a new one

(see Gernsbacher, 1990; Giora, 1985). The controls had the same late contexts as the experimental items (8–9). However, because they were preceded by sentences that did not contain the negative metaphors, they could control for the possibility that any facilitation to be found in the experimental items might be assignable to the late contexts themselves. Thus, if the controls do not show facilitation of target words, facilitation in the items following the negative metaphors (if there was any) should be viewed as a function of the target sentences only.

All the strings included the same target word (*fast*) in their late context, related to the inappropriate affirmative metaphoric sense ('fast') of the negative metaphors (*no rocket*). This target word was the probe selected by Hasson and Glucksberg (2006) for their lexical decision task. Here it was preceded by five to eight syllables to ensure that readers spent at least 1,000 ms before they encountered it—the only ISI, which, in Hasson and Glucksberg's study, yielded suppression effects. It is important to note that the target word (*fast*) was not always the penultimate or ultimate word in any of the sentence types and was not preceded by a contrast word (*though*):

(6) The train to Boston was no *rocket*. The trip to the city was fast though. (coherent string)
(7) The train to Boston was no *rocket*. The old man in the film spoke fast. (incoherent string)
(8) She poured me a glass of water. The trip to the city was fast though. (coherent string control)
(9) She poured me a glass of water. The old man in the film spoke fast. (incoherent string control)

To control for the relative coherence of the late contexts, 20 North American native speakers of English, 10 women and 10 men, ages 25 to 50, were presented all the items. They rated them on a 7-point coherence scale. Sentences scoring below 3 were classified as incoherent with prior context; sentences scoring above 5 were classified as coherent with prior context. Sentences scoring between 3 and 5 (11 pairs out of the 64) were rewritten and rerated by another 5 native speakers. The next time, they scored either below 3 or above 5.

Four sets of stimuli were prepared so that participants were exposed to only one item of the quartet (6–9). Sentences were displayed in random order.

Procedure. A moving windows procedure was used, which is a self-paced, word-by-word reading task in which each sentence initially appears as a set of dashes to be replaced by words (Just et al., 1982). This procedure allows measuring reading times of each word and further ensures that reading latencies would not be task related nor affected by strategies that readers adopt in order to accelerate their reading pace.

The participants were presented two trial pairs followed by two buffer trials. They were instructed to read the sentences at their own natural pace and answer comprehension questions by pressing the yes–no designated keys on the keyboard. They were tested individually in a well-lit and quiet room. The viewing distance was approximately 40 to 50 cm (between the monitor and the participant). The session took approximately 20 min.

Results and Discussion

Results from four sentence pairs were discarded due to equipment failure. Reading times of target words and of the words preceding these targets were averaged. Reading times of the words preceding the target words were indeed longer than 1,000 ms and thus ensured that comprehenders had enough time to deactivate information when suppression was invited.

As shown in Table 1 and Figure 1, results obtained from participant (t_1) and item (t_2) analyses demonstrate that target words in the coherent strings, in which no shift took place (6), were read faster than targets in the incoherent strings (7), in

TABLE 1
Mean and Standard Deviation Reading Times of Target Words in Each Sentence Type in Experiment 1

Sentence Type	*M*	*SD*
Coherent strings	371	93
Incoherent strings	489	197
Coherent strings controls	457	222
Incoherent strings controls	474	154

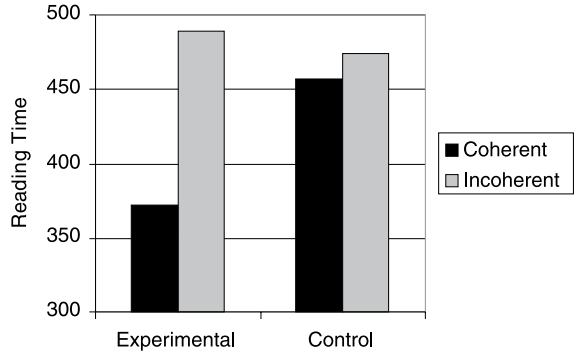

FIGURE 1 Mean reading times of target words in late experimental and control contexts (Experiment 1).

which there was a shift of the discourse topic, $t_1(31) = 3.78$, $p < .0005$, $t_2(27) = 3.94$, $p < .0005$. However, there was no significant difference in the reading times of target words in the control strings (8–9), $t_1(31) < 1$, *ns*, $t_2(27) < 1$, *ns*.

Such findings support the view that suppression is not primarily sensitive to local negation cues. Instead, it is sensitive to global discourse requirements. When a shift of discourse topic occurred, this signaled the need to dampen information recently comprehended and initiate a new substructure (Gernsbacher, 1990). However, when no such shift was signaled, the presumption of coherence cued comprehenders to retain information recently being processed.

It might be argued that it is the mere coherence of the discourse in (6) rather than the affirmative meaning of the negated concept (*rocket*) that could account for the facilitation of the target (*fast*). To control for that possibility, we ran an additional analysis. We compared the amount of facilitation of the target (*fast*) in the coherent (6) versus the control (8) strings to the amount of facilitation of the word preceding the target (*was*) in these two conditions. Recall that both strings shared a late context but differed in an early context. In one (6), both sentences made up a coherent string; in the other (8), they did not. If the amount of facilitation of the target (*fast*) in the coherent versus control strings was greater than the amount of facilitation of the pretarget word (*was*) in these two conditions, this would confirm that the facilitation found for the target was also induced by the negated concept (*rocket*) and did not just benefit from the coherence of the string.

Findings indeed showed that the target word (*fast*) was read significantly faster (371 ms; $SD = 93$) in the coherent string than in its (incoherent) control (457 ms; $SD = 222$), $t_1(31) = 2.76$, $p < .005$, $t_2(27) = 2.47$, $p < .05$. They further showed that the pretarget word (*was*) was also read somewhat faster (353 ms; $SD = 73$) in the coherent discourse than in the (incoherent) control (371 ms; $SD = 92$), $t_1(31) = 1.47$, $p = .078$, $t_2(27) = 1.42$, $p = .084$. However, an analysis of variance with word location (target or pretarget) and sentence type (coherent or control) as factors resulted in a Location × Type interaction, $F_1(1, 31) = 6.55$, $p < .05$, $F_2(1, 27) = 3.37$, $p = .077$, suggesting that the target (*fast*) was facilitated to a greater extent than the pretarget word. The only way to explain its greater facilitation would be to assume that it was primed by the affirmative meaning of the negated concept (*rocket*).

Taken together, these findings demonstrate that suppression did not apply unconditionally. Instead, when the next discourse segment made recently processed information seem useless for future use, this information was dampened. However, when, in light of oncoming messages, information just processed was deemed useful for future purposes, this information was retained for future use, irrespective of negation. Indeed, although the negative metaphors (*no rocket*) in the coherent items primed the target words (*fast*) appearing in the next sentence, they did not affect such facilitation when the next sentence did not cohere with information recently processed. The fact that the controls exhibited no facilitation differences ne-

gates the possibility that this facilitation might have been affected by the late contexts themselves.

In all, these results support the view that suppression is not obligatory but can be suspended until the processor encounters information inviting it. In the absence of such information and in the presence of a global cue to the contrary, suppression is not triggered, despite local cues such as negation that might prompt it.

To further investigate the effects of contextual information on the processing of negation, we conducted Experiments 2 and 3. One of the goals of Experiments 2 and 3 was to test context effects on negation in a design that does not compare coherent and incoherent discourse segments but involves contexts that resemble natural discourses more closely.

EXPERIMENT 2

Findings in Experiment 1 demonstrate that suppression is sensitive to global rather than local discourse cues. They show that anticipating forward coherence—coherence of currently processed information with an oncoming segment—affects retention of information within the scope of negation for at least as long as 1,000 ms following the negated concept. These findings thus argue against the view that negation induces suppression unconditionally.

In Experiment 2, we further tested the effects of contextual information on the retention of concepts within the scope of negation. This time, however, we aimed to demonstrate the effect of prior rather than late supportive context on the retention of information within the scope of negation. Although it might be the case that out of a specific context, negated information may lose accessibility between 500 to 1,000 ms following its offset (as shown by Hasson & Glucksberg, 2006; Kaup et al., 2006), following a supportive context, we argue that such information might be retained even for as long as 750 ms following its offset. To this end we presented participants with items, including negated concepts, that were relevant to information recently being processed (see 10). We predicted that, although at a very early stage of comprehension (100 ms following offset of the negated concept), facilitation of such concepts might not be allowed, yet at a longer delay of 750 ms, such information will show facilitation, in spite of the negation marker.

Method

Participants. Forty-eight undergraduates of the Academic College of Tel Aviv Yaffo participated in the experiment for course credit; 32 were women and 16 were men, and their ages ranged from 19 to 30. They were all native speakers of Hebrew and had a normal or corrected-to-normal vision.

Materials. Materials were 72 strings each containing a target sentence and a prior context. Half of the items were fillers followed by a nonword probe and the other half, made up of the experimental items, were followed by one of two probe words, related (*rich*) or unrelated (*quick*) to the negated target concept (*wealthy*) relevant to a prior supportive context (*millionaires*):

(10) I live in the neighborhood of millionaires who like only their own kind. Nonetheless on Saturday night, I also invited to the party at my place a woman who is not *wealthy.* (Probes: *related—rich; unrelated—quick*)

In addition, a yes–no comprehension question followed each item. The experimental items were made up of Giora et al.'s (2005) items, which here were furnished with a coherent prior context.

To control for the priming effects expected to be obtained in the experiment and block the possibility that they might originate either in the prior context or in the probes' relative salience, we ran a pretest. Thirty-six undergraduates of the Academic College of Tel Aviv Yaffo, 27 women and 9 men ages 20 to 35, participated in the pretest for a course credit. They were all native speakers of Hebrew and had a normal or corrected-to-normal vision. They were presented a set of both fillers and experimental items, which were identical to those to be used in the experiment, with the exception of the experimental items, which were slightly revised. Although the target sentences retained the original prior context, they ended in a word that was neutral (*religious*) in terms of its relatedness to the related probe (*rich*):

(11) I live in the neighborhood of millionaires who like only their own kind and my sister lives in Haifa in a neighborhood that is *religious*. (Probes: *related—rich; unrelated—quick*)

Participants were seated in front of a computer screen. They were tested individually in a well-lit and quiet room. The viewing distance was approximately 40 to 50 cm (between the monitor and the participant). The session took approximately 20 min. They were first informed about the task by the experimenter and then read the instructions. Participants were asked to read the text strings and make a lexical decision as to whether a letter string (displayed either 100 or 750 ms after offset of the target sentence) made up a word or not. After they had made a decision by pressing a yes or no key, they were presented a comprehension question. Six hundred milliseconds following the reply to the comprehension question, the next text strings were displayed. The experiment included five item trials and eight buffer trials, which preceded the experimental items.

Results showed no facilitation for any of the probe words in either of the ISI conditions. Mean response times to the to-be related and unrelated probes were very similar—1,237 ms (SD = 309) for the to-be-related and 1,261 ms (SD = 313)

for the to-be-unrelated probes in the 100 ms ISI; 1,245 ms (*SD* = 347) for the to-be-related and 1,231 ms (*SD* = 406) for the to-be-unrelated in the 750 ms ISI (all *t*s < 1). Such results guaranteed that any effects to be obtained in the experiment would not be explainable by the salience of the probes or by the prior context.

Procedure. As in the pretest, participants were seated in front of a computer screen and were informed about the task by the experimenter. They were asked to read the text strings and make a lexical decision as to whether a letter string displayed either 100 or 750 ms following offset of the target sentence made up a word or not. After they had indicated their decision by pressing a yes or no key, they were presented a comprehension question. Six hundred milliseconds following the reply to the comprehension question, the next text strings were displayed. The experiment included five item trials and eight buffer trials, which preceded the experimental items.

Results and Discussion

Three participants whose errors exceeded 25% were replaced. Response times of incorrect responses (23 out of 1,728, 1.3%), and response times larger than 3 standard deviations above the mean of each participant (37 out of 1,728, 2.1%) were excluded from the analyses. Mean response times were used as the basic datum for the analyses. As shown in Table 2, at an early ISI of 100 ms, there was no significant difference between response times to related and unrelated probes, $t_1(47) < 1$, $t_2(35) < 1$. Such results suggest that the supportive context did not speed up initial facilitative effects shown earlier to be visible at a 100 ms delay, when these target sentences were presented in isolation (Giora et al., 2005). However, at a longer ISI of 750 ms, the difference was significant by participant, $t_1(47) = 1.94$, $p < .05$, and (marginally so) by item analysis, $t_2(35) = 1.34$, $p = .09$, indicating that responses to related probes were faster than responses to unrelated ones. A 2 × 2 analysis of variance confirmed this pattern of results, showing a significant ISI × Probe Type

TABLE 2
Mean and Standard Deviation Response Times
at 100 and 750 ms ISIs in Experiment 2

	Probe Type			
	Related		*Unrelated*	
ISI	*M*	*SD*	*M*	*SD*
100 ms	1,134	307	1,128	241
750 ms	1,089	280	1,130	270

Note. ISI = interstimulus intervals.

interaction, only in the item analysis, $F_1(1, 47) = 1.90$, $p = .17$, $F_2(1, 35) = 4.03$, $p = .052$.

Such results demonstrate that, unlike negated concepts presented in isolation, which, at an ISI of 750 ms, were susceptible to suppression effects (Kaup et al., 2006), here they were not. When embedded in a relevant prior context, such items were retained rather than deactivated.

EXPERIMENT 3

Would complying with backward coherence, which involves mapping currently processed information onto a recently processed substructure, affect retention of negated information to the same extent that anticipation of forward coherence does? To test the duration of the effects of prior relevant context on the retention of information within the scope of negation, we ran Experiment 3. Experiment 3 was identical to Experiment 2 in every respect apart from the length of the ISI, which this time was 1,000 ms.

Method

Participants. Twenty-four undergraduates of Tel Aviv University, 10 women and 14 men ages 19 to 32, served as paid participants. They were all native speakers of Hebrew and had a normal or corrected-to-normal vision.

Materials. Materials were the same as in Experiment 2.

Procedure. Procedure was the same as in Experiment 2, only this time the letter string to which participants had to respond was displayed 1,000 ms following offset of the target sentence.

Results and Discussion

As in Experiment 2, response times of incorrect responses (13 out of 864, 1.5%), and response times larger than 3 standard deviations above the mean of each participant (21 out of 864, 2.4%) were excluded from the analyses. Mean response times were used as the basic datum for the analyses. This time, responses to related probes (925, $SD = 212$) were not faster than responses to unrelated probes (890, $SD = 165$), $t_1(23) < 1$, $t_2(35) < 1$.

Results of this experiment show that, 1,000 ms after offset of a negated concept supported by a prior context, retention of this concept was no longer operative. At this late stage, related probes were no longer more accessible than unrelated ones. Such findings suggest that, given a preceding supportive context, negated concepts

begin to lose accessibility between 750 to 1,000 ms after their offset. It is possible that once they were mapped onto recently processed information, surface information was dampened and made way for the next piece of information, which, in this case, was irrelevant to what had been recently processed. That is, given that the target (negated) concept appeared at the end of a given segment, the next experimental segment the processor encountered was entirely irrelevant to the one recently being processed.

GENERAL DISCUSSION

How is negation processed in context? The moral of these studies seems to suggest that how negation is processed depends on global discourse considerations. Interpreting negation is a function-oriented procedure and does not operate unconditionally. On the basis of Gernsbacher's (1990) Structure Building Framework, we tested the hypothesis that negation would not induce suppression once this local cue conflicted with a global cue. Indeed, when an oncoming discourse segment seemed relevant to and about to utilize information encoded in the substructure recently constructed (*The train to Boston was **no rocket**. The trip to the city was **fast** though*), no suppression of concepts (*fast*) within the scope of negation was triggered even as late as 1,000 ms following offset of these concepts. Rather, contrary to findings in Hasson and Glucksberg (2006), these concepts were retained and primed the next discourse segment. However, when the next discourse segment seemed irrelevant to the substructure recently being built (*The train to Boston was no rocket. The old man in the film spoke fast*), no facilitation of related concepts (*fast*) was visible, because comprehenders discarded previously mentioned information as unusable (Experiment 1).

Similarly, when a discourse segment currently being processed seemed relevant to and mappable onto a recently processed substructure (*I live in the neighborhood of millionaires who like only their own kind. Nonetheless on Saturday night, I also invited to the party at my place a woman who is **not wealthy***), suppression of information within the scope of negation (*wealthy*) was not triggered even as long as 750 ms following its offset. This finding allows us to narrow down the range of the temporal stage at which suppression becomes effective when prior relevant context is available. We can now say more accurately that, based on our findings, the assumption that negated information loses accessibility between 500 to 1,000 ms following its offset (Hasson & Glucksberg, 2006) is correct. In fact, what we show is that, given a relevant prior context, information within the scope of negation is actually retained at least as long as 750 ms after its offset, following which it might be suppressed (Experiment 2).

Indeed, Experiment 3 showed that negated information relevant to prior context may lose accessibility between 750 to 1,000 ms following its offset. Our attempt,

then, to replicate the findings of Experiment 2 at a longer ISI of 1,000 ms resulted in null effects. This finding corroborates Hasson and Glucksberg's (2006) finding at such a long delay, demonstrating deactivation of initially accessible information to baseline levels even out of a specific context. Thus, at such a long delay, a relevant prior context no longer has any effect on the retainability of negative concepts appearing in the next discourse segment, despite their relevance to it.

Should we deduce that backward coherence is established faster than forward coherence? Although our studies have not been specifically designed to test this hypothesis, they seem to suggest that backward coherence is indeed established faster than forward coherence. They are consistent with earlier findings that backward-coherence inferences are as helpful as explicitly stated information when it comes to verifying information. In contrast, verifying information assumed through predictive or elaborative inferences has been shown to be slow (Gernsbacher, 1990).

Examined from a different angle, our findings showing that forward coherence is maintained for a long time is reminiscent of the "Zeigarnik effect," according to which people remember unfinished tasks better than completed ones (Zeigarnik, 1927, 1967). In our task, people retained information within the scope of negation when the task of comprehending a discourse segment was ongoing and incomplete and when they suspected that this information might become instrumental in completing the task.

Do global considerations shown to account for our findings also account for the findings in the literature attesting to some suppressive effects of negation? Support for the global considerations hypothesis comes from Kaup and Zwaan (2003). Their findings showed that when a negated concept was present in the situation model, it was accessible after a long delay, regardless of negation. When it was not, its availability was reduced, regardless of negation. Mayo et al.'s (2004) findings also demonstrated sensitivity to global considerations. Recall that comprehenders had to decide whether a discourse segment was congruent or incongruent with prior context. This task by itself imposes global constraints. Findings indeed showed that when following negation (*not tidy*), if the next discourse segment was compatible with an available alternative schema ('messy'), this schema was activated and facilitated congruence judgments. However, when no alterative was available (as in *not adventurous*), no facilitation of congruent judgments was demonstrated. It is possible to view these results as attesting to comprehenders' abidance by a global requirement to consider congruence, which triggered the activation of an antonymic schema when this was available.

Our findings demonstrate effects of late context on retention and suppression of concepts mentioned earlier. They show that comprehenders retain information within the scope of negation until alerted to the contrary. They further show that this is true even when concepts are bipolar (cf. Mayo et al., 2004) or scalar (cf. Fillenbaum, 1966) such as *fast*. Our findings also attest to effects of prior context

on information currently being processed. They reveal that, when mappable onto a recently processed substructure, information within the scope of negation is retained as long as 750 ms following its offset. In the absence of such coherence relation, this information is no more available than unrelated concepts (as shown in the pretest of Experiment 2). In all, such findings demonstrate the superiority of global over local considerations in shaping text representation following negation.

ACKNOWLEDGMENTS

The studies reported here were supported by grants to Rachel Giora by Adams Super Center for Brain, Tel Aviv University and by Tel Aviv University Basic Research Fund. We thank Mira Ariel, Sam Glucksberg, Barbara Kaup, Orna Peleg, Michael Schober, and two anonymous reviewers for very helpful comments on an earlier version of the article.

REFERENCES

Ariel, M. (1990). *Accessing noun phrase antecedents.* London: Routledge.

Carston, R. (2002). *Thoughts and utterances.* Oxford, England: Blackwell.

Duffy, S. A. (1986). Role of expectations in sentence integration. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8,* 195–209.

Fillenbaum, S. (1966). Memory for gist: some relevant variables. *Language and Speech, 9,* 217–227.

Frisson, S., & Pickering, M. (2001). Obtaining a figurative interpretation of a word: Support for underspecification. *Metaphor and Symbol, 16,* 149–172.

Garnham, A. (1992) Minimalism versus constructionism: A false dichotomy in theories of inference during reading. *Psycoloquy, 3,* 63.

Gernsbacher, M. A. (1990). *Language comprehension as structure building.* Hillsdale NJ: Lawrence Erlbaum Associates, Inc.

Giora, R. (1985). A text-based analysis of nonnarrative texts. *Theoretical Linguistics, 12,* 115–135.

Giora, R. (2003). *On our mind: Salience, context and figurative language.* New York: Oxford University Press.

Giora, R. (2006). Anything negatives can do affirmatives can do just as well, except for some metaphors. *Journal of Pragmatics, 38,* 981–1014.

Giora, R., Balaban, N., Fein, O., & Alkabets, I. (2005). Negation as positivity in disguise. In H. L. Colston & A. Katz (Eds.), *Figurative language comprehension: Social and cultural influences* (pp. 233–258). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Giora, R., & Fein, O. (1999). On understanding familiar and less-familiar figurative language. *Journal of Pragmatics, 31,* 1601–1618.

Glenberg, M. A., Robertson, D. A., Jansen, J, L., & Johnson-Glenberg, M. C. (1999). Not propositions. *Journal of Cognitive Systems Research, 1,* 19–33.

Hasson, U., & Glucksberg, S. (2006). Does understanding negation entail affirmation? *Journal of Pragmatics, 38,* 1015–1032.

Just, M. A., Carpenter, P. A., & Woolley, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General, 111,* 228–238.

Kaup, B. (2001). Negation and its impact on the accessibility of text information. *Memory & Cognition, 29,* 960–967.

Kaup, B., Lüdtke, J., Zwaan, A. R. (2006). Processing negated sentences with contradictory predicates: Is a door that is not open mentally closed? *Journal of Pragmatics, 38,* 1033–1050.

Kaup, B., Yaxley, R. H., Madden, C. J., Zwaan, R. A., & Lüdtke, J. (in press). Experiential simulations of negated text information. *Quarterly Journal of Experimental Psychology.*

Kaup, B., & Zwaan, R. A. (2003). Effects of negation and situational presence on the accessibility of text information. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 439–446.

Keysar, B. (1994). Discourse context effects: Metaphorical and literal interpretations. *Discourse Processes, 18,* 247–269.

Levy, G. (2003, June 27). The rest was up to God. *Ha'aretz Magazine,* 6.

Lüdtke, J., & Kaup, B. (2006). Context effects when reading negative and affirmative sentences. In R. Sun & N. Miyake (Eds.), *Proceedings of the 28th Annual Conference of the Cognitive Science Society* (pp. 1735–1740). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

MacDonald, M. C., & Just M. A. (1989). Changes in activation levels with negation. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15,* 633–642.

Mayo, R., Schul, Y., & Burnstein, E. (2004). "I am not guilty" versus "I am innocent": The associative structure activated in processing negations. *Journal of Experimental Social Psychology, 40,* 433–449.

McKoon, G., & Ratcliff, R. (1986). Priming in item recognition: The organization of propositions in memory for text. *Journal of Verbal Learning and Verbal Behavior, 19,* 369–386.

O'Brien, E. J., Shank, D. M., Myers, J. L., & Rayner, K. (1988). Elaborative inferences during reading: Do they occur online? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 410–420.

Paradis, C., & Willners, C. (2006). Antonymy and negation: The boundedness hypothesis. *Journal of Pragmatics, 38,* 1051–1080.

Potts, G. R., Keenan, J. M., & Golding, J. M. (1988). Assessing the occurrence of elaborative inferences: Lexical decision versus naming. *Journal of Memory and Language, 27,* 399–415.

Resistance in Germany: A call to all Germans! (1943). [Leaflets of the *White Rose*]. Retrieved from http://209.85.135.104/search?q=cache:ydGmYQ2d2NUJ:www.whiterosesociety.org/WRS_pamphlets_fifth.html+%22a+call+to+all+Germans%22&hl=en&ct=clnk&cd=2

Sinai, R. (2006). *No disabled included in Israeli delegation to United Nations debate on their rights* [In Hebrew]. Retrieved July 9, 2006, from http://www.haaretz.co.il/hasite/spages/736219.html

Singer, M. (1979). Processes of inference in sentence encoding. *Memory & Cognition, 7,* 192–200.

Singer, M. (1980). The role of case-filling inferences in the coherence of brief passage. *Discourse Processes, 3,* 185–200.

Singer, M., & Ferreira, F. (1983). Inferring consequences in story comprehension. *Journal of Verbal Learning and Verbal Behavior, 22,* 437–448.

Szymborska, W. (1996). A "Thank You" note (J. M. Trzeciak, Trans.). In *Miracle Fair* (p. 24). New York: Norton.

Zandberg, E. (2006). *British architects plan barrier boycott*. Retrieved February 13, 2006, from http://www.Haaretz.com/hasen/objects/pages/PrintArticleEn.jhtml?itemNo=681935

Zeigarnik, B. (1927). Über das Behalten von erledigten und unerledigten Handlungen [On finished and unfinished tasks]. *Psychologische Forschung, 9,* 1–85.

Zeigarnik, B. (1967). On finished and unfinished tasks. In W. D. Ellis (Ed.), *A sourcebook of Gestalt psychology* (pp. 300–314). New York: Humanities.