Expecting Irony: Context Versus Salience-Based Effects

Rachel Giora

Department of Linguistics Tel Aviv University

Ofer Fein

Department of Behavioral Sciences The Academic College of Tel Aviv Yaffo

Dafna Laadan, Joe Wolfson, Michal Zeituny, Ran Kidron, Ronie Kaufman, and Ronit Shaham Department of Linguistics Tel Aviv University

Results from 4 experiments support the view that, regardless of contextual information, when an end-product interpretation of an utterance does not rely on the salient (lexicalized and prominent) meanings of its components, it will not be faster than nor as fast to derive as when it does. To test this view, we looked into interpretations of salience-based (here, literal) interpretations and expectation-based (here, ironic) interpretations in contexts inducing an expectation for irony. In Experiment 1, expectancy was manipulated by introducing an ironic speaker in vivo who also uttered the target utterance. Findings show that ironic targets were slower to read than literal counterparts. Experiment 2 shows that ironies took longer to read than literals and that response times to ironically related probes were longer than to literally related probes, regardless of context. Experiments 3 and 4 show that, even when participants were given extra processing time and were exclusively presented ironically biasing contexts, the expectancy for irony acquired throughout such exposure did not facilitate expectancy-based compared to salience-based interpretations.

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

Destroy them until no inhabitant is left

Turn them into a dried-up desert. Mounds of rubble ...

... make bombs

and shower them on villages and on cities and on houses until they crumble down.

Kill them, draw-out their blood, terrorize their lives ...(Ilan Sheinfeld, July 21–28, 2006, translated by Rachel Giora)

Sheinfeld's Hebrew poem, entitled *Let Them Be Strong* (alluding to a well known poem of the nation's ancestral poet, Bialik), was made public during the Israeli reoccupation of Lebanon, termed the 2nd Lebanon War. On July 30, 2006, Sheinfeld e-mailed his poem to his list with a preface explaining its inception (http://www.gogay.co.il/today/files/5072.asp). Some Israeli readers, who had little knowledge of the poet's current political inclinations and who read the poem but not the preface clarifying its nature, interpreted it ironically. They assumed that such an attitude, if taken at face value, is inconceivable, especially among "the likes of us," poets included. Little did they know that the poem was intended literally. Richer contextual information might have helped those readers come up with the appropriate reading of what is said. It's not just nonliteral language, then, that might rely heavily on context. Literal interpretations might just as well depend on contextual information for their derivation.

It is indeed widely acknowledged that contextual information affects endproduct interpretations, and that when it is strong enough, it results in appropriate readings quite early on. In the literature, such effects are labeled "context effects." They are visible in case a target is facilitated compared to a control. Although there is a consensus that rich contextual information facilitates appropriate interpretations, there is a disagreement as to whether (a) such context facilitates interpretations immediately by governing initial lexical processes so that only contextually appropriate meanings are accessed; or whether (b) it is ineffective in interfering with initial lexical processes and therefore cannot sieve out inappropriate meanings initially.

On the first view, highly specific contexts will always result in swift appropriate interpretations, regardless of inappropriate but salient meanings. Thus, having read *The gardener dug a hole. She inserted the bulb*, comprehenders will access only the contextually appropriate "flower" meaning of the ambiguous (*bulb*) target (Vu, Kellas, & Paul, 1998). On the second view, even in highly supportive environments, contextually appropriate interpretations might temporarily be hampered in case lexical processes result in inappropriate but salient outputs that will have to be redressed. According to this view, *The gardener dug a hole. She inserted the bulb* will initially involve the contextually inappropriate "light" meaning of the ambiguous target (*bulb*), which will have to be suppressed and replaced by the appropriate ("flower") meaning (see Peleg, Giora, & Fein, 2005; for a review, see Giora, 2003).

The different predictions regarding the temporal aspects of context facilitative effects have been extensively tested with regard to irony. With the exception of conventional irony, ironic interpretations are not salient (Giora, 1997, 1999, 2003). They are not listed in the mental lexicon and are not foremost on our mind due to, for example, conventionality, familiarity, frequency, or prototypicality. Rather, they are divorced from the salient *meanings* of the utterances' components and rely heavily on contextual information for their derivation. For example, when a coach, nervous about an upcoming game, finds out that three of his team's lead players will not be able to show up for the game and mutters: "This is terrific news!" he communicates an interpretation of *terrific* that is not coded but, instead, relies on contextual information for its decoding. It is this difference between what is said (*terrific*) and what is referred to by what is said (bad news about prospective losses for the team) that irony hinges on (Giora, 1995). This is why irony has often been used to explore the effects of a supportive context on deriving contextually appropriate interpretations. Given that the kind of irony interpretation studied here is induced by a contextual expectation for an ironic utterance, irony is treated here as "expectation-based interpretation."

Note that the *interpretation* of what is said is also context dependent. However, because it is closely related to and often includes the salient *meanings* of the utterance's components, it is termed here "salience-based interpretation." Thus, the salient meaning of (the key word of) what is said here is associated with "great"—the salient (although nonliteral) sense of *terrific*. In the context of a game, the salience-based interpretation of what is said relates to "good news about winning." However, in this ironically biased context, this interpretation is inappropriate. A crucial question to be addressed here is whether, although inappropriate, such interpretations will nonetheless be derived, or whether they will not, thus allowing the contextually motivated (in our case, "expectation-based" ironic) interpretations to be derived exclusively.

To test context facilitative effects on irony interpretation, various studies have manipulated context strength by involving various factors. For instance, in Gibbs (1986), contexts contained some inherent contrast typically inviting ironic remarks (Gibbs, 2002, p. 462). Katz and his colleagues (Katz, Blasko, & Kazmerski, 2004; Katz & Pexman, 1997; Pexman, Ferretti & Katz, 2000) introduced discourse elements that evoked expectations of either a literal or a nonliteral interpretation as a function of speakers' occupation. However, despite controlling for context strength in various ways, findings were not always supportive of the view that context can facilitate irony interpretation early enough to allow for only appropriate outputs to become available. For instance, in Gibbs (1986), reading times of ironic and literal interpretations indeed did not differ significantly, and, at times, ironies were even faster to understand than literals. At the same time, however, ironic compliments took longer to read than their equivalent literals. Similarly, in Katz and colleagues' studies, ironies always took longer to read than literal controls. Such findings are more consistent with the rival view, which holds that even a strong context involving an expectation for irony cannot neutralize initial lexical processes and, as a result, when inappropriate, such meanings interfere with the interpretation process.

In a recent study, Ivanko and Pexman (2003) tested the expectation factor more closely. According to Ivanko and Pexman, expectation for a specific articulation determines a statement's ease of processing (henceforth "the expectation hypothesis"). They thus predicted that even when a context is highly supportive of a negative stance but nonetheless invites a literal statement, processing of an ironic statement will be hampered. Literal statements, however, will be facilitated. In contrast, weak contexts, which do not raise any specific expectation with regard to either articulation, would result in no processing difference between ironies and literals. Findings indeed showed that strongly negative contexts facilitated literal interpretations, which were faster to read than ironies both initially and downstream. However, weak contexts facilitated ironic interpretations, which were only initially faster to read than literals. Although findings demonstrating no initial facilitative effects of strong contexts on irony can also be accounted for by the view that even strong contexts cannot sieve out inappropriate lexical outputs [see (b) earlier], the finding that weak contexts facilitated irony or alternatively slowed down literals requires an explanation on both views (Ivanko & Pexman's included).

Regardless, what is crucial for testing Ivanko and Pexman's proposal is an examination of a complementary prediction that an expectation for an ironic statement will facilitate interpretation of ironies compared to literals, which, in turn, will be slowed down by that same expectation. The purpose of this study is, therefore, to further examine the expectation hypothesis, according to which a strong context inducing an expectation for an ironic target will facilitate irony comprehension initially.

But what kind of context might induce an expectation for an ironic utterance? In line with studies showing that expectancy may be built up by the preceding stimulus sequences (Jentzsch & Sommer, 2002; Kirby, 1976; Laming, 1968, 1969; Soetens, Boer, & Hueting, 1985), we assume that, among other things, ironic utterances on their own induce an expectation for additional ironic utterances. Therefore, with the exception of Experiment 2, which forms the basis for Experiments 3 and 4, the contexts and experimental design we use in our studies proliferate ironic uses within (Experiment 1) and between items (Experiments 3 and 4).

EXPERIMENT 1

Experiment 1 examines the expectation hypothesis proposed by Ivanko and Pexman (2003), which suggests that it is the type of utterance anticipated that might account for the differences found earlier (e.g., Giora & Fein, 1999; Giora, Fein, & Schwartz, 1998; Schwoebel, Dews, Winner, & Srinivas, 2000) or lack of them (Gibbs, 1986). To test this hypothesis, we introduced anticipation for ironic utterances into conversational contexts. Our materials comprised two types of context—dialogues that bi-

ased an identical penultimate target sentence either toward its ironic or toward its literal interpretation. These dialogues always contained an ironic utterance in dialogue midposition. Introducing an ironic speaker in vivo should boost an expectation for another ironic turn on the part of that speaker. This expectation is satisfied in the ironic context, in which that speaker also utters the ironic target sentence, but is frustrated in the literal context, in which that speaker's turn is literal.

If indeed it is the expectation for a literal target that slowed down processing of unexpected ironic targets in Ivanko and Pexman's (2003) studies, then enhancing an expectation for ironic targets should speed up processing of such ironies and slow down processing of literal targets. Alternatively, if such an expectation does not eliminate earlier differences between literals and ironies (cf. Giora, 2002, 2003; Giora & Fein, 1999a; Giora, Fein, & Schwartz, 1998; Schwoebel et al., 2000), but rather maintains them, this would argue against the expectation hypothesis and in favor of the graded salience hypothesis (Giora, 1997, 1999, 2003). According to the graded salience hypothesis, even a strong context, predictive of ironic utterances, cannot sieve out salient meanings, which emerge on account of their salience. Constructing an appropriate utterance interpretation that is divorced from these meanings will, therefore, result in complex interpretive processes.

Based on the graded salience hypothesis, we predict that, despite an expectation for ironic utterances, induced here by a speaker exercising irony, ironic utterances will not be facilitated immediately and will take longer to read than equivalent literals. Although the literal interpretations of the items tested here are closely related to the salient meanings of their components and should therefore be rather accessible, their ironic interpretations are removed from them and should therefore be less accessible. Given that the task we employ here is self-paced reading of whole utterances, difficulties in deriving the ironic interpretation of the target utterance might, but need not, spill over to the next utterance.

Method

Participants. Participants were 24 (mostly) students (13 women and 11 men) of Tel-Aviv University, between the ages of 19 and 32, who volunteered to participate in the experiment.

Materials. Materials included 16 pairs of Hebrew dialogues (1–2), which were 11 utterances long, and which took place between friends. The dialogues' 6th utterance was always ironic (bolded here for convenience) and biased the 10th utterance—the target sentence—(bolded here for convenience), spoken by the same ironic speaker, either toward its ironic (1) or toward its literal (2) interpretation. In addition, there were 10 fillers, which included no ironies:

(1) Barak: I finish work early today. Sagit: So, do you want to go to the movies?

Barak: I don't really feel like seeing a movie. Sagit: So maybe we could go dancing? Barak: No, at the end of the night my feet will hurt and I'll be tired. Sagit: You're a really active guy ... Barak: Sorry, but I had a rough week. Sagit: So what are you going to do tonight? Barak: I think I'll stay home, read a magazine, and go to bed early. Sagit: Sounds like you are going to have a really interesting evening. Barak: So we'll talk sometime this week. (2) Barak: I was invited to a film and a lecture by Amos Gitai. Sagit: That's fun. He is my favorite director. Barak: I know, I thought we'll go together. Sagit: Great. When is it on? Barak: Tomorrow. We will have to be in Metulla¹ in the afternoon. Sagit: I see they found a place that is really close to the center. Barak: I want to leave early in the morning. Do you want to come? Sagit: I can't, I'm studying in the morning. Barak: Well, I'm going anyway. Sagit: Sounds like you are going to have a really interesting evening.

Barak: So we'll talk sometime this week.

To ascertain that these texts induced an expectation for an ironic statement, we presented 24 participants (15 women and 9 men, aged 23–38) with the targetless contexts and asked them to select the ending which they thought was the most suitable completion (either the ironic or the literal one). Averaging across items, we found that the ironic ending was chosen 19.6 times (out of 24), which is significantly higher than the expected 12 (indicating no clear preference), $Z_{\text{score}} = 12.46$, p < .0001.

To further make sure that our dialogues induced an expectation for oncoming ironic utterances, we ran an additional pretest. This time we presented two groups of 16 readers—in all, 32 participants (19 women and 13 men, aged 23–38)—with two types of booklets, each containing 16 experimental contexts short of the two final utterances and 10 filler contexts. One type of booklet included experimental contexts with an ironic utterance in midposition (e.g., 1); the other type included the same contexts with an equivalent literal utterance in midposition. In both, the 10 filler items were identical, half included an ironic utterance in midposition and half did not. Each context was followed by two endings presented randomly: One was the ironic target (see 1) and the other was its literal equivalent. The fillers were followed by either a literal or a metaphoric ending. Participants were asked which of the two they would expect to end the dialogue (see Ivanko & Pexman, 2003 for the same task). Findings showed that, following a context featuring an ironic utter-

¹Metulla is the most northern town of Israel and far away from the center.

ance, subjects selected an ironic ending 12.2 (SD = 2.8) times out of the 16 cases. However, following the contexts not featuring an ironic speaker, subjects selected an ironic ending only 7.3 (SD = 4.4) times out of the 16 cases. Comparing the mean times readers selected an ironic ending following the two types of context demonstrates that the context featuring an ironic speaker induced a significantly stronger expectation for an oncoming ironic utterance, $t_1(30) = 3.81$, p < .0005; $t_2(15) =$ 7.33, p < .0001. These two pretests, then, establish that all the dialogues used in our experiment induced an expectation for an ironic utterance.

Procedure. Participants self-paced their reading, advancing the context sentence by sentence. They were told the investigators were interested in reading times for comprehension. Reading times of the target sentences were measured by the computer.

Results and Discussion

Although the ironic targets were the preferred and expected endings, they took longer to read (1,703 msec; SD = 590) than the literal counterparts (1,570 msec; SD = 436) by subject, $t_1(23) = 2.42$, p < 0.05, though not by item analysis—probably due to the small number of items. Reading times of the last sentence following the target did not disclose spillover effects, all ts < 1. Such results argue against the expectation hypothesis, which predicts speedier reading times for ironies and slower reading times for literals in contexts inducing an expectation for irony. Instead, replicating earlier results (Giora & Fein, 1999a; Giora et al., 1998), these results argue in favor of the view that context cannot bypass salient but inappropriate meanings. As a result, salience-based (here literal) interpretations are speedier than context-based (here ironic) interpretations, even when context prompts an expectation for an ironic utterance.

That irony requires intensive processing even in the presence of a strong context is not an unprecedented finding. As mentioned earlier, a number of studies have already indicated that comprehension of nonsalient ironies could be a complex process. For instance, in Giora et al. (1998) and Schwoebel et al. (2000), reading times of ironic utterances were longer than their (salience-based) literal interpretations. In other studies, ironies took longer to process than (salience-based) conventional metaphors (Colston & Gibbs, 2002; Pexman et al., 2000), or (salience-based) literal controls (Pexman et al., 2000). Moreover, when lexical decision tasks were administered, ironically biasing contexts did not affect initial processing. In fact, literally and ironically biased targets were initially (150 msec after offset of the target sentences) processed along the same lines: Both invoked only salient (literal) meanings, regardless of context. However, only in the case of the ironically biasing contexts, these meanings were contextually incompatible. It is no wonder, then, that when ERPs were measured, ironies elicited larger negativity brainwaves (N400) than literal counterparts (Katz et al., 2004). Elevated N400 brainwaves are associated with the difficulty to integrate activated information with an ongoing mental representation, which, in the case of irony, might reflect the contextual incompatibility of the salient meanings evoked at the initial stage of comprehension.

Later on, however, (1,000–2,000 msec after offset of the target sentences), long after salience-based (literal) interpretations had been computed, nonsalient, contextually appropriate (ironic) interpretations were derived, though not at the cost of suppressing these salient meanings (Giora & Fein, 1999a; Giora et al., 1998).

Such findings are highly revealing, because they are informative about both the time-course of irony interpretation and the kind of meanings involved at the various temporal stages of the process. If a strong context could indeed block salient (here mostly literal) meanings and facilitate ironic interpretations very early on, lexical decisions to irony-related probes should have been as fast as to those invoked by literally related probes in literally biasing contexts. The fact that in many of the studies of irony, initial processing of irony, preceded by a heavily supportive context, is slow, involving activation of salient but inappropriate meanings, demonstrates that even a strong context does not have initial facilitative effects: It does not sieve out salient but inappropriate meanings that conflict with contextual information, resulting, therefore, in lengthy latencies.

Note that our earlier results were criticized for employing weak contexts, which, as a result, did not have speedy facilitative effects (Ivanko & Pexman, 2003; Utsumi, 2000). Such contexts, it was contended, might account for our longer response times to ironically than to literally related probes and for the longer reading times of ironies compared to literals. However, results in Ivanko & Pexman themselves already challenge the validity of this contention. In Ivanko & Pexman, it was the weakly negative rather than the strongly negative context that eliminated such differences. In contrast, strong contexts resulted in longer latencies for ironies than for literals.

One could argue, though, that the response times to the ironically related probes found in Giora et al.'s studies were relatively long because these probes conveyed negative meanings and in addition were presented following negative contexts, which could also slow down response times (Gibbs, 2002; Ivanko & Pexman, 2003). However, our studies controlled for such possibilities (Giora et al., 1998). In these studies, we also obtained response times to related probes in both the literally and ironically biasing contexts presented without their final target sentences. Results showed that when context alone was at stake, no context effects were visible: Both kinds of probes took similarly long to respond to, regardless of context bias and negativity. Such findings guaranteed that the response times found for these probes in the contexts containing the target sentences were not affected by the positivity or negativity of the contexts themselves; neither were they affected by the negativity of the probes. This line of criticism is in fact also foiled by Ivanko & Pexman's (2003) own findings in which it was the strongly negative contexts that speeded up reading times of literal items, whose critical word is negative (e.g., rotten in Sam is a rotten friend).

It is nonetheless plausible that, in our earlier studies, response times to ironically and literally related probes were different, because these probes tapped different types of meaning. Whereas the literally related probes were more closely associated with *lexical entries* and could be accessed directly from the mental lexicon, the ironically related probes were associated with utterance level *interpretations*, which should take longer to process than *lexical meanings* (Gibbs, 2002, p. 468, 2005, p. 225).

To examine this line of criticism before further testing the expectation hypothesis, Experiment 2 was designed. It aimed to explore the time-course of utterance-level interpretation, using probes related to this level of interpretation, whether literal or ironic. Thus, following a target such as "*This is terrific news!*" the probe following a literally biasing context would be *winning*; the probe following an ironically biasing context would be *losses*. In this way we aimed to guarantee that response times to such probes would tap context-based inferences, both literal and ironic. In the case of our targets, we assumed that, because the literal interpretations of the statements closely relied on the salient meanings of their components, they should be easier to derive than the ironic ones that did not. On the assumption that irony stands in some contrast to (the salient meanings of the components of) what is said (Carston, 2002; Colston, 2002; Colston & O'Brien, 2000; Giora, 1995), it would require complex processes of adjustment to contextual information even when the latter is highly predictive.

EXPERIMENT 2

Experiment 2 attempts to investigate the time-course of utterance-level interpretation. Its aim is to rectify for the possibility that earlier studies, rather than tapping irony's late interpretive processes, explored early lexical processes. To do that, we measured reading times of critical target words and the word that followed, and response times to probes related to higher-level interpretations of both literals and ironies. Assuming that interpretation is a late process, we expected processing difficulties to be visible for ironies only downstream, following the critical word, thereby disclosing spillover effects. We further predicted that facilitating lexical decisions to probes would occur late (at 1,400 msec interstimulus interval) rather than immediately (at 250 msec interstimulus interval) and would be visible only for the salience-based probes. Early facilitation of probes was not expected because no lexical priming was anticipated.

Method

Design. A $2 \times 2 \times 3$ factorial design was used with interstimulus interval (ISI) of either 250 or 1,400 msec as a between-subjects factor, and context type

(ironically/literally biasing) and word type (either ironically or literally related, or unrelated) as within-subjects factors.

Participants. Participants were 48 students of Tel-Aviv University, between the ages of 18 and 35 (28 women and 20 men), who were paid 15 NIS (about \$3) for their participation in the experiment.

Materials. Materials included 30 Hebrew texts (see Appendix), about 7–9 sentences long (3a–b), 15 of which biased the last target sentence (bolded here for convenience) toward its literal interpretation (3a) whereas the remaining 15 biased it toward its ironic interpretation (3b). The critical target word (*noble*; italicized here for convenience) was always placed in penultimate sentential position (in Hebrew). Materials also included 30 equally long filler texts with either a literal or a metaphoric ending and 45 yes/no comprehension questions, which appeared following all the experimental texts and half of the filler texts. For each target sentence, two types of probes were prepared, either related (generous; stingy) or unrelated (sleepy) to its interpretation. Nonwords were presented at offset of filler texts:

(3a) Sarit worked as a waitress in a small restaurant in central Naharia. The evening was slow, and even the few customers she did wait on left negligible tips. She didn't think that the elderly man who walked in alone and ordered just a couple of small sandwiches would be any different. But when he had left, and she collected his pay for the meal from off his table, she found no less than 60 NIS tip! When she showed her friends how much she got, Orna commented: "**That was real** *noble* of him!"

(3b) Sarit worked as a waitress in a small restaurant in central Naharia. The evening was slow, and even the few customers she did wait on left negligible tips. She didn't think that the elderly man who walked in alone and ordered just a couple of small sandwiches would be any different. Indeed, after making her run back and forth throughout the meal, he left, and she collected his pay for the meal from off his table and found 2.5 NIS tip! When she showed her friends how much she got, Orna commented: **"That was real** *noble* **of him!"**

Probes: Literally related—generous; ironically related—stingy; unrelated—sleepy.

The materials were selected following a pretest, which involved 24 students of Tel-Aviv University, between the ages of 18 and 35 (14 women and 10 men), who rated targets following biasing contexts on a 7-point ironiness scale. Only ironic

items scoring above 5 and only literal items scoring below 3 were selected for the experiment.

To test the related and unrelated probes for their relative salience, 30 students of Tel-Aviv University, between the ages of 18 and 35 (13 women and 17 men) volunteered to participate in an online pretest. They were seated in front of a computer monitor (the one used for the main experiment) and were asked to make lexical decisions as to whether a letter string, presented at a short ISI (250 msec) following offset of a neutral sentence context, made up a Hebrew word or not. Materials were displayed centrally and contained 60 neutral sentence contexts and 4 probes: literally related, ironically related, unrelated, and nonwords. Of the sentence contexts, 30 were followed by the experimental words and 30 were followed by the nonwords. Participants self-paced their reading of the neutral context and then made a lexical decision by pressing a "yes" or a "no" key. They were given 10 practice trials. Results exhibited no salience difference between the literally related $(654 \operatorname{msec}, SD = 113)$, ironically related (668 $\operatorname{msec}, SD = 131$), and unrelated (674, SD = 123) probes. This was confirmed by both subject and item analysis of variance (ANOVAs), $F_1(2, 58) < 1$, $F_2(2, 58) < 1$. We thus guaranteed that the probes to be used in the experiment would be of comparable salience.

To make sure that the literally and ironically related probes were similarly related to the interpretation of their relevant target sentences in their respective contexts, and the unrelated probes were indeed unrelated, we ran another pretest. In this pretest, 30 volunteer participants, between the ages of 20 and 60 (11 women and 19 men) were presented the experimental texts each followed by three test words—the related (whether literal or ironic), the unrelated, and an additional semirelated word, which was used only to allow a choice of midposition on the scale. Participants were asked to rate each word on a 7-point relatedness scale. Specifically, they were asked to rate the extent to which the probe was related to the interpretation of the last sentence in its given context. Results showed that, in the ironically biasing contexts, the ironically related word was highly related (6.00, SD = 0.68). The semirelated word was less related (2.59, SD = 0.78). The unrelated word was unrelated (1.45, SD = 0.43). Similarly, in the literally biasing contexts, the literally related word was highly related (5.91, SD = 0.62). The semirelated word was less related (3.18, SD = 0.91). The unrelated word was unrelated (1.38, SD = 0.91). SD = 0.36). A 2 × 2 ANOVA with context type (ironic, literal) and relatedness (related, unrelated) as within-subject factors was run. It revealed no effect of context type and no interaction. The subject and item analyses revealed only a relatedness effect, $F_1(1, 29) = 1,501.8$, p < .0001; $F_2(1, 29) = 618.1$, p < .0001. These results thus guaranteed that the literally and ironically related probes were similarly related to the interpretation of the experimental items and targets they were designed to test.

To ascertain that processing of the probes will not be influenced only by the context itself but also by the target sentence in its context, 30 volunteer participants between the ages of 24 and 65 (15 women and 15 men) were presented the experi-

mental texts. They were asked to mark the point at which they felt quite certain that the protagonist thought that *X* (where *X* was the probe word to be used in the experiment). Results showed that, on the whole, participants indicated that the probes were more closely related to the interpretation of the target sentence than to the context alone, marking that point in 60% of the times within the ironic target sentences and in 66% of the times within the literal target sentences. Both are significantly higher, $Z_{\text{score}} = 3.68$, p < .0005; $Z_{\text{score}} = 6.32$, p < .0001, respectively, than the expected 50% (which would have been the case had there been no clear preference for marking that point either within the target sentence or within the context).

Procedure. The experimental texts were displayed centrally. Readers advanced the context sentences in three-word segments. The last target sentence, however, was advanced word by word so that measuring reading times of single words was allowed (Just, Carpenter, & Woolley, 1982). We were interested in the reading times of the critical words (*terrific*), whose final interpretation was affected by contextual information. Reading times of these words should basically reflect early lexical processes, which, according to the graded salience hypothesis, should not be sensitive to contextual information and, therefore, initially, should not distinguish literally from ironically biased words. However, reading times of the word that followed that critical word might reflect spillover effects for this word ensuing from later, context-sensitive interpretive processes, which should therefore be lengthier for ironically than for literally biased targets.

Following the reading phase, participants were asked to make a lexical decision to letter strings displayed centrally either at a 250 or a 1,400 msec ISI. They had to decide whether the letter strings made up words or not. The latency between the display of the probes and the pressing of the "yes"/"no" key was measured by the computer and served as response time. The pressing of the key further displayed a yes/no comprehension question, which appeared in 75% of the cases.

Results

Reading times. Mean reading times of the critical word and the word that followed are presented in Table 1. As predicted, reading times of the critical words did not differ significantly whether they were ironically or literally biased, t(29) < 1. In contrast, reading time of the word that followed the critical word was longer in the ironically biasing contexts than in the literally biasing contexts, by item analysis, $t_2(29) = 1.67$, p = 0.053, revealing spillover effects in this condition.

Response times. Mean response times at the short ISI are presented in Table 2 and illustrated in Figure 1. Mean response times at the long ISI are presented in Table 3 and illustrated in Figure 2. A three-way ANOVA was performed for both subject (F_1) and item (F_2) analyses, with ISI, context type, and probe type as factors. The ANOVA by subjects showed two significant effects: A significant effect

	Word Position		
Context Type	Penultimate	Ultimate	
Literal context	387	549	
	(27)	(101)	
Ironic context	380	605	
	(25)	(153)	

TABLE 1 Means of Reading Times of Penultimate and Ultimate Words in the Target Sentence in Experiment 1 (*SD* in Parentheses)

TABLE 2
Mean Response Times at 250 msec ISI—Experiment 2
(SD in parentheses)

Context Type	Probe Type		
	Literally Related	Unrelated	Ironically Related
Literal context	886	865	939
	(290)	(206)	(267)
Ironic context	856	831	985
	(269)	(281)	(314)

of probe type, $F_1(2, 92) = 8.15$, p < 0.001, and a significant probe type × ISI interaction effect, $F_1(2, 92) = 3.45$, p < 0.05. The ANOVA by items showed two significant effects: A significant effect of ISI, $F_2(1, 29) = 25.31$, p < 0.0001 and probe type, $F_2(2, 58) = 15.22$, p < 0.0001.

Because there was no three-way interaction between the variables, we performed an ANOVA analysis within each ISI, with probe type and context type as factors. In the shorter ISI (250 msec), the only significant effect was of probe type, $F_1(2, 46) = 6.66, p < 0.005, F_2(2, 58) = 4.41, p < 0.05$. In the longer ISI (1,400 msec), there was both a probe-type effect, $F_1(2, 46) = 5.35, p < 0.01, F_2(2, 58) =$ 6.14, p < 0.005, and a context-type effect in the subject analysis, $F_1(1, 23) = 5.26, p$ < 0.05. No interaction effects were found.

To study more closely the differences apparent in the results, we ran a contrast analyses for probe type (Literal, Unrelated, Ironic) within each ISI. In the short ISI, both subject and item analyses with the coefficients [-.5, -.5, 1], were significant, $F_1(1, 23) = 11.04, p < 0.005, F_2(1, 29) = 6.79, p < 0.05$. In the longer ISI, the analyses with the coefficients [-1, .5, .5], were also significant, $F_1(1, 23) = 7.05, p < 0.05, F_2(1, 29) = 7.32, p < 0.05$.



FIGURE 1 Mean response times at 250 msec ISI-Experiment 2

Discussion

Earlier studies of irony comprehension and interpretation (Giora & Fein, 1999a; Giora et al., 1998) might have compared different processes when exploring interpretations of literals and ironies. Although lexical decisions to literally related probes could be associated with lexical entries only, irony-related probes must have been associated with higher-level interpretations (Gibbs, 2002, p. 468). In this experiment, we aimed to remedy this possible flaw by using probes that were related only to higher-level interpretations of both literals and ironies and could not be primed by lexical entries. Lexical decisions to probes related to utterance interpretations made at short (250 msec) and long (1,400 msec) ISIs could reflect early and late interpretation processes. Given the nature of the probes, which could not benefit from lexical priming, we did not anticipate early context facilitative effects.

Early and late processes were also tapped through measuring reading times. Reading times of critical words—those the literal and ironic interpretations rest on—reflect initial lexical processes; reading times of the word that follows that

(<i>SD</i> in Parentheses)			
	Probe Type		
Context Type	Literally Related	Unrelated	Ironically Related
Literal context	820	977	1001
	(274)	(485)	(528)
Ironic context	934	1023	1029
	(337)	(412)	(510)

TABLE 3 Mean Response Times at 1400 msec ISI—Experiment 2 (*SD* in Parentheses)



FIGURE 2 Mean response times at 1,400 msec ISI—Experiment 2

critical word might disclose spillover effects which are suggestive of later interpretive processes.

Pretests guaranteed that the probes shared similar salience and were similarly accessible in a neutral context. They also controlled for the literality and ironiness of the targets embedded in their respective contexts and for the similar relatedness of the probes to their target utterances. They further controlled for the relatedness of the probes to the interpretation of the target sentences rather than to the context alone.

Results from reading times confirm that, as shown earlier (Giora & Fein, 1999a; Giora et al., 1998; Katz et al., 2004), context had no effect on the early stages of comprehension: In both the ironically and literally biasing contexts, reading times of critical words were the same, reflecting similar initial processes. Later processes, however, revealed a difference. Reading times of the word following the critical word disclosed spillover effects only in the ironically biasing context, suggesting that, as anticipated, ironic interpretations were more difficult to derive than their literal equivalents, because, unlike their literal counterparts, their accessible, salience-based interpretations conflicted with contextual information and had to be adjusted.

Results from response times corroborated findings from reading times. As for earlier processes, they showed that, as anticipated, 250 msec after offset of the target sentence, literal and ironic interpretations were not available. This demonstrates the absence of lexical priming between targets and probes. Indeed, this should come as no surprise, given that the probes were designed to tap utterance-level interpretations, which should take longer to derive. What is unclear, though, is why they were less available than the unrelated probes. It is possible that, at this stage, processes were slowed down by the interference of the salient lexical meanings in the computation of the appropriate interpretation, which is more complex for the ironic than for the literal interpretation. Importantly, however, response times occurring later in the interpretation process, 1,400 msec after offset of the target sentence, revealed that, at this stage, literal interpretations were already computed and ironic interpretations lagged behind. As anticipated, the literal interpretations were faster to derive than the ironic ones, regardless of context (as attested by the lack of probe type × context type interaction). That is, in both, the literally and ironically biased context, literal interpretations were derived earlier in the interpretation process, because, we contend, they rely rather closely on the salient meanings of the utterance components. In contrast, ironies require redressing salience-based interpretations. Their derivation was probably slowed down by the interference of these interpretations with their appropriate interpretations, which induced a kind of Stroop effect (see Dews & Winner, 1997; Giora & Fein, 1999a; Giora et al., 1998).

Results further argue against possible criticism, controlled for by a pretest (see earlier), that the related probes were related mostly to their contexts rather than tapping their targets' interpretations. If this were true, then, the literally related probes could not be more accessible than the ironically related probes in the context biasing the target toward its ironic interpretation. The fact that, in these contexts, the ironically related probes took longer to respond to suggests that, although they were related to contextual information, their retrieval involved processing the salience-based literal interpretation of the target, which then had to be redressed.

Now that we have established that salience-based (here literal) interpretations are faster to derive than context-based (here ironic) interpretations, it is necessary to test whether strengthening an expectation for a context-based (ironic) interpretation would facilitate that interpretation so that it is derived as fast as or faster than a salience-based (literal) counterpart. To do that, we ran Experiment 3, in which we presented participants with two subsets of the materials used in Experiment 2. One subset comprised half literally and half ironically biasing contexts; another comprised ironically biasing contexts only. Although in one, the mixed nature of the items should not favor any expectation, neither for an ironic nor for a literal statement, in the other, an expectation for an ironic ending should be acquired in the course of the experiment. Given that all the items, including the fillers, ended in an ironic statement, this should raise the probability and thus the expectation for ironic endings (Jentzsch & Sommer, 2002). These different subsets should allow comparing responses to items not designed to raise an expectation for either an ironic or a literal target to items that implicitly train readers to expect an ironic statement.

EXPERIMENT 3

The aim of Experiment 3 was to test the expectation hypothesis with regard to salience-based (literal) versus expectation-based (ironic) message level interpretations. The procedure and items were the same as in Experiment 2, only this time the items were divided differently to create an expectation in one subset (+Expectation condition) but no expectation in another (-Expectation condition). In addition, this time, the ISI was 750 msec, which can reflect initial interpretive processes following lexical processes of both literal and ironic interpretations. On the basis of the graded salience hypothesis (Giora, 1997, 1999, 2003), we anticipated no interaction between expectancy and probe type. Specifically, we expected salience-based probes to be activated faster than expectation-based ones, regardless of expectancy.

Method

Design. A 2×3 factorial design was used with expectancy (+/– Expectation) as a between-subjects factor and word type (ironically/literally related and unrelated) as a within-subjects factor.

Participants. Participants were 72 volunteer students (40 women and 32 men) of Tel-Aviv University and their friends, between the ages of 20 and 42.

Materials. Materials were the same as in Experiment 2, only here they were divided into two subsets, each comprising two blocks of items—one inducing a specific expectation for an ironic target and treated as "+ Expectation" condition and one not inducing a specific expectation and treated as "-Expectation" condition. Thus one block—the set inducing an expectation for irony—comprised 36 contexts, all biased toward their ironic interpretation. This block included 3 ironic buffers, 15 experimental ironic contexts, and another 18 ironic filler contexts. The other block-the set not inducing a specific expectation-also comprised 36 contexts. They included the same 18 contexts appearing in the other block—15 experimental ironic contexts and 3 buffers, and another 18 contexts biasing their final statements toward their literal interpretation and which functioned as filler items. The items were arranged in four versions, and, with the exception of the buffers, their order of presentation was random. In this way, half of the subjects received half of the ironies as experimental contexts and the rest of the subjects received the other half of the ironies as experimental contexts. Each time, the ironies that were not used as experimental items were used as expectancy generators to form the "+Expectation" conditions, whereas the literal contexts were used to create the "-Expectation" condition.

Because a small number of the probes were replaced, we reran the salience test to make sure salience levels were kept constant. Thirty-three students (18 women and 15 men) of Tel-Aviv University, between the ages of 19 and 40 volunteered to participate in an online pretest, performed in the same way as in Experiment 2. Results exhibited salience differences between the literally related (726 msec, SD = 207), the ironically related (681 msec, SD = 157), and the unrelated probes (696 msec, SD = 178). This was confirmed by both subject and item

ANOVAs, $F_1(2, 64) = 3.27$, p < .05, $F_2(2, 58) = 3.44$, p < .05. Given that the ironically related probes were faster than the literally related ones and that using them could bias our experimental results in a direction opposite to our prediction and cannot therefore be counted as an artifact, we decided to use the probes, despite of the salience differences.

Procedure. As in Experiment 2, the experimental texts were displayed centrally. Readers advanced the contexts sentence by sentence. The last target sentence, however, was advanced word by word. Following the reading phase, participants were asked to make a lexical decision to letter strings displayed centrally after an ISI of 750 msec. The pressing of the key further displayed a yes/no comprehension question, 12 in all, which appeared in 33% of the cases.

Results and Discussion

Response times larger than 3 SDs above the mean of each participant were considered outliers and were excluded from the analyses. Overall, only 12 response times, out of 1,080, were excluded (1.1%). Mean response times were used as the basic datum for the analyses. Results are presented in Table 4.

A 2 × 2 ANOVA was performed for both subject (F_1) and item (F_2) analyses, with expectancy and probe type as factors. The ANOVA by subjects showed only a significant effect of probe type, $F_1(2, 140) = 4.94$, p < .01 and, equally important, no expectancy × probe-type interaction, $F_1(2, 140) < 1$. To see where the differences in the probe-type factor are, we ran a contrast analysis for probe type (Literal, Unrelated, Ironic) with the coefficients [-.5, -.5, 1], which turned out to be significant, $F_1(1, 70) = 8.51$, p << .005.

The ANOVA by items did not reveal any significant effects. However, in this analysis we could take into account the baseline for each item found in our pretest. Indeed, after subtracting the baselines, there was only a probe-type effect, $F_2(2, 58) = 4.19$, p < .05, as in the subject analysis, and the same contrast was likewise

TABLE 4 Mean Response Times (Before Subtraction of Baseline Means) at 750 msec ISI—Experiment 3 (*SD* in Parentheses)

Expectancy	Probe Type		
	Literally Related	Unrelated	Ironically Related
+Expectancy	857	860	913
	(278)	(202)	(276)
-Expectancy	831	868	918
	(81)	(141)	(160)



FIGURE 3 Mean response times (after subtraction of baseline means) differences at 750 msec ISI-Experiment 3

significant, $F_2(1, 29) = 5.55$, p < .05, showing that the ironic probe took longer to process than both the literal and the unrelated probes (see Figure 3).

Such findings argue against the expectation hypothesis, which predicts initial facilitation of targets invited by an expectation. They show that, even when repeated experience evoked an expectation for a context-based (ironic) interpretation, its derivation was not speeded up to the extent that it overrode or even matched that of a salience-based (literal) interpretation, although the latter was both contextually inappropriate and unexpected. Because the interpretation of literals is (mostly) salience-based whereas that of (nonconventional) irony is not, irony derivation is more complex. Even the presence of a strong environment invoking an expectation for irony did not facilitate it initially.

To test whether additional processing time might result in expectancy effects, we ran Experiment 4, which resembled Experiment 3 in every respect except for the ISI, which, this time, was longer. Note that even if, at this late processing stage, ironic interpretation eventually becomes as fast as or faster than literal interpretation, this will not argue against the temporal priority of salience-based interpretations, which applies only initially.

EXPERIMENT 4

The aim of Experiment 4 was to test the expectation hypothesis with regard to message-level interpretations of salience versus expectation-based interpretations when longer (1,000 msec) processing time is allowed. As in Experiment 3, here too, we predicted no expectancy \times probe-type interaction.

Method

Design. As in Experiment 3.

Participants. Participants were 72 (mostly) students (37 women and 35 men) of Tel-Aviv University, between the ages of 16 and 42 who were paid 15 NIS (about \$3) for their participation in the experiment.

Materials. As in Experiment 3.

Procedure. As in Experiment 3, only this time an ISI of 1,000 msec was used.

Results and Discussion

Response times larger than 3 SDs above the mean of each participant were considered outliers and were excluded from the analyses. Overall, only 14 response times, out of 1,080, were excluded (1.3%). Mean response times were used as the basic datum for the analyses. Results are presented in Table 5.

A two-way ANOVA was performed for both subject (F_1) and item (F_2) analyses, with expectancy and probe type as factors. The ANOVA by subjects did not reveal any significant results. In the ANOVA by items, there was only an expectancy effect, $F_2(1, 29) = 9.66$, p < .005. When, as in Experiment 3, we subtracted the baselines (obtained in the pretest) from each item's mean (see Figure 4), there was, in addition to the expectancy effect, a marginally significant probe-type effect, $F_2(2, 58) = 2.76$, p = .07.

Importantly, as in Experiment 3, there was no expectancy \times probe-type interaction, neither in the subject, nor in the item analyses, showing that expectancy did not affect ironic, literal, and unrelated responses differently.

TABLE 5 Mean Response Times (Before Subtraction of Baseline Means) at 1,000 msec ISI—Experiment 4 (*SD* in Parentheses)

	Probe Type		
Expectancy	Literally Related	Unrelated	Ironically Related
+Expectancy	881	886	884
	(167)	(145)	(209)
-Expectancy	937	1007	943
	(183)	(228)	(150)



FIGURE 4 Mean response times (after subtraction of baseline means) differences at 1,000 msec ISI—Experiment 4

GENERAL DISCUSSION

In this study we tested the hypothesis that, regardless of contextual expectation, inferences closely associated with the salient meanings of an utterance's components would be easier to derive than those that are not. Specifically, we examined facilitation effects of a strong context inducing an expectation for an ironic utterance on the processing of such an utterance. We compared ironic utterances, whose interpretations are nonsalient but contextually compatible, with their literal interpretations, which are salience-based but contextually incompatible.

The debate about context effects revolves around the extent to which context can facilitate appropriate meanings of words and expressions very early on so that their integration with contextual information becomes seamless (for a review, see Giora, 2003, Chapter 3). In the field of nonliteral language, this translates into whether strong contexts allow comprehenders to circumvent an inappropriate but salience-based interpretation of an utterance-mostly, but not always, its literal interpretation-so that the end-product involves only contextually appropriate (mostly nonliteral) interpretations (Gibbs, 1986, 1994; Ortony, Schallert, Reynolds, & Antos, 1978). Although locally salient but inappropriate lexical meanings might pop up, the crux of the matter is whether derivation of the appropriate utterance-level interpretation is derived exclusively without involving such inappropriate interpretations (Gibbs, 1983, 1984). Empirically, findings demonstrating equal reading times of salience-based (often literal) and context-based (often nonliteral) utterances embedded in highly supportive contexts support the view that rich contextual information may have early effects, facilitating contextually appropriate interpretations immediately without having to go through contextually inappropriate ones (Gibbs, 1986; Ortony et al., 1978). In contrast, findings showing speedier reading times for literal than for nonliteral utterances under these conditions attest to lexical effects rather than to context effects. They disclose context's failure to bypass salient but inappropriate meanings which, consequently, slow down derivation of contextually appropriate interpretations (Giora & Fein, 1999a, 1999b; Giora et al. 1998).

Would context also fail to block inappropriate but accessible utterance-level interpretations? If context effects do not interfere with lexical processes, inferences based on these meanings might be difficult to block on account of their accessibility even when inappropriate. In our studies, the salience-based inferences were related to literal items whereas inferences removed from the salient meanings of the utterance's components were related to ironic targets. Our studies adduced evidence demonstrating that even a rich and supportive context, predictive of an ironic interpretation, did not allow circumventing contextually inappropriate but accessible (literal) interpretations.

In Experiment 1, we looked into the time-course of contextual effects on nonsalient, context-based (ironic) interpretations and salience-based (literal) interpretations as a consequence of a context inducing an expectation for ironic turns. We showed that raising an expectation for an ironic turn by embedding an ironic speaker in vivo did not facilitate ironic interpretations to the extent that they were read faster than or as fast as literal interpretations. Despite of an involvement of an ironic speaker, turns expressed by this speaker took longer to read in ironically than in literally biasing contexts. Such results argue against the claim that a rich context inducing an expectation for a certain articulation of information should facilitate processing that information initially in case this expectation, it is the salience-based interpretation that is derived initially.

In Experiment 2, we further investigated the temporal aspects of processing context-based (ironic) and salience-based (literal) interpretations:

(a) We tapped *initial lexical processes* by applying a reading measure of the critical words the literal and ironic interpretations rest on;

(b) we tapped *early interpretation derivation* by searching for spillover effects which should be visible when reading the word that follows the critical word is accomplished;

(c) we tapped *a bit later interpretation processes* by administering lexical decisions to inference-related probes at a short (250 msec) ISI;

(d) we tapped *late interpretation processes* by administering lexical decisions to inference-related probes at a long (1,400 msec) ISI. At this rather late stage, we expected the products of interpretation to be observable (Gibbs, 1993).

The first stage of processing (a) was disclosed by using a moving windows procedure (Just, Carpenter, & Woolley, 1982). This methodology records reading times of each word in target sentences, disclosing immediate online processes. Indeed, at this very early stage, critical words were read equally fast, regardless of whether context was literally or ironically biased. This finding ties in well with earlier results showing that the early stages of literal and ironic interpretations were identical, involving access of same salient meanings, regardless of context (Giora et al., 1998; Giora & Fein, 1999a).

The second stage of processing (b, c) was tapped by measuring both reading times of words that followed the critical words and latencies to lexical decisions made to probes related to literal and ironic interpretations at a short delay—250 msec after offset of targets. Both measures defied the claim that a supportive context could facilitate nonsalient but appropriate interpretations by blocking accessible but inappropriate ones. Findings from reading times demonstrate that, at this stage, ironically biased targets took longer to process than literally biased targets, attesting to spillover effects. And although, at this stage, both literal and ironic interpretations were still unavailable (attesting to their inference-related nature), the ironic-related probes took longer to respond to than both the unrelated and the literally related probes.

The third stage of processing (d) was explored by administering lexical decisions to interpretation-related probes made at a long delay of 1,400 msec. Indeed, at this late stage, salience-based (literal) interpretations were already made available, regardless of contextual bias. In contrast, context-based (ironic) interpretations were not yet available. Thus, regardless of contextual information, responses to literally related probes were faster than to ironically related probes, suggesting that literal interpretations of whole utterances were faster to derive than ironic ones. They emerged earlier even in ironically biasing contexts. Such results attest to the temporal priority of salience-based over nonsalient context-based inferences even in a context strongly biased in favor of the nonsalient interpretation. Such findings then suggest that context-based (ironic) interpretation involves derivation of contextually inappropriate but salience-based utterance-level inferences. It is not the case, then, that a supportive context allows bypassing inappropriate but accessible information, not even at the level of utterance interpretation.

Having established that initially a supportive context did not facilitate context-based (ironic) interpretations to the extent that they overrode salience-based (literal) interpretations, we turned to test the expectation hypothesis under circumstances that trained comprehenders to anticipate an ironic target as opposed to circumstances that did not. In Experiment 3, participants were presented with the same items used in Experiment 2, only this time we compared between (a) a set equally divided between contexts ending either in an ironic or in a literal interpretation and (b) a set whose contexts always ended in an ironic target. Findings showed that the two sets were indistinguishable even though comprehenders were allowed 750 msec processing time after offset of the probes. Similar to results obtained in Experiment 2, response times to ironically related probes were always longer than to literally related ones. An attempt to replicate these findings by allowing comprehenders a longer processing time of 1,000 msec did not change the overall picture (Experiment 4). Even after such a long delay, responses to contextually appropriate but nonsalient (ironic) interpretations were not faster than responses to inappropriate but salience-based (literal) interpretations. In fact, they were slower. Based on these findings, it is quite safe to contend that context-based (irony) interpretation involves an earlier stage of computing the salience-based (here literal) interpretation of the utterance, despite its contextual misfit.

Contrary to appearances, however, advocating the temporal priority of salience-based interpretations differs from the assumptions of the Standard Pragmatic Model (Grice, 1975) in various respects. First, although the Standard Pragmatic Model posits the priority of literal interpretations, salience-based interpretations are not necessarily literal. They can just as well be nonliteral. For instance, the salience-based ironic interpretation of Einstein! is nonliteral. In addition, the Standard Pragmatic Model posits automatic suppression of contextually inappropriate literal outputs and their replacement by appropriate nonliteral interpretations. In contrast, the graded salience hypothesis, supplemented by the suppression/retention hypothesis (Giora, 2003; Giora & Fein, 1999a, 1999b) does not assume unconditional suppression of initial outputs when contextually inappropriate. Instead, it takes a functional view of suppression, proposing that contextually inappropriate meanings and interpretations, which have been activated on account of their salience or salience-based grounds, will be retained in case they are deemed instrumental to, for example, the appropriate interpretation. They will be suppressed in case they are deemed disruptive (see also Giora, 2006). In the case of irony, salience-based interpretations play a role in constructing the ironic interpretations. They constitute a reference point, which allows the computation of the difference between what is said and the situation described; they help flesh out a speaker's wishful thinking and the extent to which it has fallen short of expectation (Giora, 1995). Therefore, salience-based interpretations were not only derived here but also retained for as long as 1,000 msec following offset of targets (see also Giora, 2002, 2003; Giora & Fein, 1999a, 1999b).

In sum, our findings argue in favor of the temporal priority of salience-based interpretations, regardless of context. They show that contextually appropriate (ironic) interpretations did not sieve out accessible but inappropriate (literal) interpretations. Rather, such inappropriate interpretations were derived quite early on and effected longer latencies for expectation-based (ironic) interpretations than for salience-based (literal) interpretations, regardless of contextual expectation for an ironic utterance. These findings argue against the hypothesis that it is an expectation for a certain articulation of information that affects initial speed and manner of processing (Ivanko & Pexman, 2003). Our findings thus suggest that it is not only the case that inappropriate lexical meanings are not blocked by contextual processes (Giora & Fein, 1999; Giora et al., 1998) but that, in fact, even contextually inappropriate utterance-level interpretations are not inhibited. They show that interpretations relying on the salient meanings of their components are initially resistant to context effects.

ACKNOWLEDGMENTS

The studies reported here were supported by a grant to the first author by Adams Super Center for Brain Tel Aviv University, and by Tel Aviv University Basic Research Fund. We thank Tal Eyal for very helpful comments, Veronica Yudkevich for help with the second experiment, Adi Ezra for help with the third experiment, and Robyn Carston, Ray Gibbs, and two anonymous reviewers for their very helpful criticisms.

REFERENCES

- Carston, R. (2002). Thoughts and utterances. Oxford: Blackwell.
- Colston, H. L. (2002). Contrast and assimilation in verbal irony. Journal of Pragmatics, 34, 111-142.
- Colston, H. L., & Gibbs, R. W., Jr. (2002). Are irony and metaphor understood differently? *Metaphor and Symbol*, 17, 57–80.
- Colston, H. L., & O'Brien, J. (2000). Contrast of kind versus contrast of magnitude: The pragmatic accomplishments of irony and hyperbole. *Discourse Processes*, 30, 179–199.
- Dews, S., & Winner, E. (1997). Attributing meaning to deliberately false utterances: The case of irony. In C. Mandell, & A. McCabe (Eds.), *The problem of meaning: Behavioral and cognitive perspectives* (pp. 377–414). Amsterdam: Elsevier.
- Gibbs, R. W., Jr. (1983). Do people always process the literal meanings of indirect requests? Journal of Experimental Psychology: Learning, Memory, and Cognition, 9, 524–533.

Gibbs, R. W., Jr. (1984). Literal meaning and psychological theory. Cognitive Science, 8, 575-304.

- Gibbs, R. W., Jr. (1986). On the psycholinguistics of sarcasm. *Journal of Experimental Psychology: General*, 115, 3–15.
- Gibbs, R. W., Jr. (1993). Processes and products in making sense of tropes. In A. Ortony (Ed.), Metaphor and thought (2nd ed., pp. 252–276). Cambridge: Cambridge University Press.
- Gibbs, R. W., Jr. (1994). The poetics of mind. Cambridge: Cambridge University Press.
- Gibbs, R. W., Jr. (2002). A new look at literal meaning in understanding what is said and implicated. *Journal of Pragmatics*, *34*, 457–486.
- Gibbs, R. W., Jr. (2005). Literal and nonliteral are corrupt ideas. In S. Coulson & B. Lewandowska-Tomaszczyk (Eds.), *The literal and nonliteral in language and thought* (pp. 221–238). Frankfurt am Main: Peter Lang.
- Giora, R. (1995). On irony and negation. Discourse Processes, 19, 239-264.
- Giora, R. (1997). Understanding figurative and literal language: The graded salience hypothesis. Cognitive Linguistics, 7, 183–206.
- Giora, R. (1999). On the priority of salient meanings: Studies of literal and figurative language. *Journal of Pragmatics*, *31*, 919–929.
- Giora, R. (2002). Literal vs. figurative language: Different or equal? *Journal of Pragmatics*, 34, 487–506.
- 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. http://dx.doi.org/10.1016/j.pragma.2005.12.006
- Giora, R., & Fein, O. (1999a). Irony: Context and salience. Metaphor and Symbol, 14, 241-257.
- Giora, R., & Fein, O. (1999b). On understanding familiar and less-familiar figurative language. *Journal of Pragmatics*, 31, 1601–1618.
- Giora, R., Fein, O., & Schwartz, T. (1998). Irony: Graded salience and indirect negation. *Metaphor and Symbol*, 13, 83–101.
- Grice, H. P. (1975). Logic and Conversation. In P. Cole & J. Morgan (Eds.), Speech acts: Syntax and semantics (Vol. 3, pp. 41–58). New York: Academic Press.
- Inhoff, A. W., Lima, S. D., & Carroll, P. J. (1984). Contextual effects on metaphor comprehension in reading. *Memory & Cognition*, 12, 558–567.
- Ivanko, L. S., & Pexman, M. P. (2003). Context incongruity and irony processing. *Discourse Processes*, 35, 241–279.
- Jentzsch, I., & Sommer, W. (2002). The effect of intentional expectancy on mental processing a chronopsychological investigation. Acta Psychologica, 111, 265–282.
- Just, M. A., Carpenter, P. A., & Woolley, J. D. (1982). Paradigms and processes in reading comprehension. Journal of Experimental Psychology: General, 111, 228–238.
- Katz, A. N., Blasko, G. D., & Kazmerski, A. V. (2004). Saying what you don't mean: Social influences on sarcastic language processing. *Current Directions in Psychological Science*, 13, 186–189.
- Katz, A. N., & Pexman, P. (1997). Interpreting figurative statements: Speaker occupation can change metaphor into irony. *Metaphor and Symbolic Activity*, 12, 19–41.
- Kirby, N. H. (1976). Sequential effects in two-choice reaction time: Automatic facilitation or subjective expectancy? *Journal of Experimental Psychology: Human Perception and Performance*, 2, 567–577.
- Laming, D. R. J. (1968). Information theory of choice-reaction times. London: Academic Press.
- Laming, D. R. J. (1969). Subjective probability in choice-reaction experiments. *Journal of Mathemati*cal Psychology, 6, 81–120.
- Ortony, A., Schallert, D. L., Reynolds, R. E., & Antos, S. J. (1978). Interpreting metaphors and idioms: Some effects of context on comprehension. *Journal of Verbal Learning and Verbal Behavior*, 17, 465–477.
- Peleg, O., Giora, R., & Fein, O. (2004). Contextual strength: The whens and hows of context effects. In: I. Noveck & D. Sperber (Eds.), *Experimental pragmatics* (pp. 172–186). Basingstoke: Palgrave.
- Pexman, P. M., Ferretti, T., & Katz, A. (2000). Discourse factors that influence irony detection during on-line reading. *Discourse Processes*, 29, 201–222.
- Schwoebel, J., Dews, S., Winner, E., & Srinivas, K. (2000). Obligatory processing of the literal meaning of ironic utterances: Further evidence. *Metaphor and Symbol*, 15, 47–61.
- Sperber, D., & Wilson, D. (1986/1995). Relevance: Communication and Cognition. Oxford: Blackwell.
- Soetens, E., Boer, L. C., & Hueting, J. E. (1985). Expectancy or automatic facilitation? Separating sequential effects in two-choice reaction time. *Journal of Experimental Psychology: Human Perception and Performance*, 11, 598–616.
- Utsumi, A. (2000). Verbal irony as implicit display of ironic environment: Distinguishing ironic utterances from nonirony. *Journal of Pragmatics*, 32, 1777–1806.
- Vu, H., Kellas, G., & Paul, S. T. (1998). Sources of sentence constraint in lexical ambiguity resolution. *Memory & Cognition*, 26, 979–1001.

APPENDIX

Sample items Experiments 2–4 (translated from Hebrew); (a) versions are literal; (b) versions are ironic.

(1a) John was a basketball coach. For the past week he was feeling restless, worrying about the upcoming game. It was yet unclear how the two teams matched up, and he was anxious even on the day of the game. When he got a call telling him that the three lead players on the opposing team will not be able to play that night, John wiped the sweat off of his forehead and said to his friend, "This is *terrific* news!"

(1b) John was a basketball coach. For the past week he was feeling restless, worrying about the upcoming game. It was yet unclear how the two teams matched up, and he was anxious even on the day of the game. When he got a call telling him that the three lead players on his team will not be able to play that night, John wiped the sweat off of his forehead and said to his friend: "This is *terrific* news!"

Probes: Literally related—winning; ironically related—losses; unrelated—friends.

(2a) Yuval and Omry went out for their lunch break after a morning of work. They went to the cafeteria in their office building and each filled a platter with food. They stood in line for a long while and were eager to start the meal. When they had sat down to eat, Yuval saw that his colleague filled his platter with salad, tofu, and sprouts and chose natural carrot juice for a drink. Then Yuval said: "I see that you picked the *ideal* meal today!"

(2b) Yuval and Omry went out for their lunch break after a morning of work. They went to the cafeteria in their office building and each filled a platter with food. They stood in line for a long while and were eager to start the meal. When they had sat down, Yuval saw that his colleague chose fried sausage, chips, a glass of coke for a drink, and a sugar-glazed doughnut for desert. Then Yuval said: "I see that you picked the *ideal* meal today!"

Probes: Literally related—healthy; ironically related—harmful; unrelated lighter.

(3a) Becky's anxiety over her last test for the semester turned out to be justified, and she fared miserably despite more than a week of studying for the exam. Her friends had planned a trip to celebrate the end of the school year but they were afraid that thoughts of the test would cloud her mood the whole way. To their surprise, she was as cheerful as ever, and Tami just had to say: "I see your mood is really uplifted today!"

(3b) Becky's anxiety over her last test for the semester turned out to be justified, and she fared miserably despite more than a week of studying for the

exam. Her friends had planned a trip to celebrate the end of the school year but they were afraid that thoughts of the test would cloud her mood the whole way. They were sad to see that she was gloomy, melancholy, and uncommunicative, and Tami just had to say: "I see your mood is really uplifted today!"

Probes: Literally related—happy; ironically related—depressed; unrelated—costume.