

Irony: Context and Salience

Rachel Giora

*Department of Linguistics
Tel Aviv University*

Ofer Fein

*Department of Psychology
Princeton University*

Two experiments test a graded salience account of irony processing (Giora, Fein, & Schwartz, 1998). Experiment 1 shows that, as predicted, less familiar targets embedded in ironically biasing contexts facilitate only the salient literal meaning initially: 150 msec after their offset. However, 1,000 msec after their offset, the less salient ironic meaning becomes available and the literal meaning is still as active. In contrast, familiar ironies facilitate both their salient literal and ironic meanings initially: 150 msec after their offset. Results do not change significantly after a 1,000-msec delay. In the literally biasing contexts, less familiar ironies facilitate only the salient literal meaning. In contrast, familiar ironies facilitate both their salient literal and ironic meanings under both interstimulus interval conditions, as predicted. Experiment 2 confirms that these findings were affected by the target sentences rather than by the contexts themselves. In Experiment 2, the contexts were presented without the targets, inducing no difference in response patterns.

Irony comprehension is believed to rely heavily on context (cf. Katz & Lee, 1993). Researchers, however, fail to agree on the temporal stage at which context affects irony comprehension. Some studies report results consistent with the interactive, direct access view (e.g., Gibbs, 1986a, 1986b, 1994; Gibbs, O'Brien, & Doolittle, 1995; see Sperber & Wilson, 1986/1995, p. 239, for similar assumptions) that assumes that contextual information affects comprehension very early on. They question the traditional assumptions often known as the Standard Pragmatic Model (Grice, 1975; Searle, 1979; see Temple & Honeck, 1999, for recent findings regard-

ing proverb comprehension) that posit the precedence of the compositional interpretation of the sentence over its nonliteral interpretation. The claim is that in a rich and supportive context, irony is comprehended more or less directly, bypassing the contextually incompatible literal interpretation of the ironic utterance. According to the direct access view, then, appropriate contexts should enhance activation of the contextually appropriate meaning, so that only that meaning becomes available for comprehension. Embedded in appropriate contexts, then, ironic and literal interpretations should involve equivalent processes (see also Kumon-Nakamura, Glucksberg, & Brown, 1995). Evidence supporting the direct access view of irony comes from equal reading times of ironic and nonironic utterances (Gibbs, 1986a, 1986b; Gibbs et al., 1995; see Giora, 1995, and Dews & Winner, 1997, for a critique of some of the findings).

In contrast, some studies are consistent with some aspects of the modular view (Dews & Winner, 1997, 1999; Giora, 1995; Grice, 1975; Searle, 1979) that assumes that the lexical processes involved at the initial stage of comprehension should not be affected by nonlexical information. Rather, they are automatic and autonomous, and impervious to context effects (cf. Fodor, 1983). The initial stage of irony comprehension must, therefore, be literal. Contextual information should affect irony comprehension only at a later stage and should trigger revisitation of the contextually incompatible literal meaning. On this view, irony comprehension should involve more complex inferential processes than literal interpretation. Findings consistent with a modular-based view of irony comprehension show that utterances took longer to read in ironically than in literally biased contexts (Giora, in press; Giora, Fein, & Schwartz, 1998; Schwoebel, Dews, Winner, & Srinivas, 1999) and longer to be judged as positive or negative relative to their literal counterparts (Dews & Winner, 1997). They further show that irony comprehension involved longer response times to ironically than to literally related probes (Giora et al., 1998).

The different views and findings may be reconciled by a more general principle of salience. According to the graded salience hypothesis (Giora, 1997, 1999, in press; Giora et al., 1998; Giora & Fein, in press; see also Goldvarg & Glucksberg, 1998), the factor determining initial activation is neither literality nor compatibility with context, but rather the salience of the verbal stimulus: Salient meanings of words and expressions should always be accessed and always first. A meaning of a word or an expression is salient if it is coded in the mental lexicon. Salience, however, admits degrees. Factors affecting degree of salience are, for example, conventionality, frequency, familiarity, or prototypicality. For instance, both meanings of *bank* (i.e., the “financial institution” and the “river edge” meanings) are listed in the mental lexicon. However, for those of us from urban communities, in which rivers are less common than financial institutions, the commercial sense of *bank* is foremost (i.e., salient). By the same token, the riverside sense is less salient. In contrast, inferences computed on the fly are nonsalient because they are not coded in the mental lexicon.

Although prior context may enhance a word's meaning, it is relatively ineffective in inhibiting activation of salient meanings. For example, in *I needed money, so I went to the bank*, the prior occurrence of the word *money* may speed up activation of the financial institution meaning of *bank*. In *Standing on the riverbank I saw some fish*, the word *river* may facilitate activation of the riverside meaning of *bank*. However, although "river" in *riverbank* may enhance the less salient riverside meaning of *bank*, it may not prevent activation of its more salient, financial institution meaning on its encounter: The salient, financial institution meaning would pop up in spite of contextual misfit. The major claim of the graded salience hypothesis, then, is that salient meanings of words and collocations are always accessed initially (although not necessarily solely), irrespective of contextual information or bias (see also Rayner, Pacht, & Duffy, 1994; for a more detailed discussion, see Giora, 1997, in press; but see Vu, Kellas, & Paul, 1998, for a different view).

According to the graded salience hypothesis, then, the salient meaning of a word or an expression is accessed directly. When it is contextually compatible, no more processes are required. However, when a less salient meaning has to be activated to make sense of an utterance (as in the case of the literal meaning of conventional idioms or the ironic interpretation of familiar metaphors), comprehension should involve an ordered access: The more salient, albeit inappropriate meaning should be processed initially, before the less salient, appropriate meaning can be retrieved. Indeed, in Pexman, Ferretti, and Katz (1999), participants took longer to read ironic than metaphoric interpretations of familiar (but also less familiar) metaphors. Similarly, in Gibbs (1980, 1986c), participants took longer to read idioms in a literally than in an idiomatically biased context, and their literal paraphrases took longer to be judged as meaningful than their nonliteral interpretations, regardless of contextual bias (Gibbs, 1986c). In this connection, we should introduce a cautionary note about possible confounds. Appropriate paraphrases for literal utterances are always problematic. Although idioms may have a nonliteral interpretation at their disposal, which is different from their literal compositional meaning, literal utterances do not. Whereas "*Kick the bucket*" has a coded interpretation that can be computed directly from the mental lexicon ("die"), its literal interpretation (kick the bucket, literally) does not have a ready-made literal paraphrase that can be similarly computed (e.g., "tip the pail"). Consequently, literal paraphrases tend to be made up of a lot less salient or frequent words. Consider the commonality of the interpretation of the nonliteral meaning of the idiom "*He kept it under his hat*" ("He did not tell anyone") as opposed to the oddity and scarcity of its literal interpretation and words ("It is beneath his cap"). This could be one reason why literal targets took longer to read in Gibbs (1986c).

In the same vein, when two or more meanings are salient, they should be accessed in parallel. Thus, conventional metaphors whose figurative and literal meanings are similarly salient should be processed directly both literally and metaphorically (as shown by Blasko & Connine, 1993; see Giora, in press, for a cri-

tique; see also Williams, 1992, regarding metaphorically based polysemies). Consequently, they should take equally long to read in literally and metaphorically biasing contexts (Giora & Fein, 1999; see Turner & Katz, 1997, for similar findings regarding familiar proverbs).

The graded salience hypothesis (Giora, 1997, 1999, in press; see also Récanati, 1995, and Turner & Katz, 1997, for somewhat similar views) is thus consistent with an ordered access account (see Gorfein, 1989; Hogaboam & Perfetti, 1975; Kawamoto, 1993; Rayner et al., 1994; Simpson & Burgess, 1985), which maintains that lexical processes are autonomous but sensitive to frequency. It holds that the factor relevant to comprehension is not literality (or nonliteral), but the degree of salience of the utterance processed. It predicts that less familiar ironies would be initially processed only literally, in both literally and ironically biasing contexts, because less familiar ironies have only one salient meaning—the literal meaning (made up, among other things, of the salient, literal meanings of the lexical components). Ironically biasing contexts should affect their processing only at a later stage, in which the ironic meaning may be inferred. In contrast, familiar ironies, whose ironic (utterance) meaning and literal (lexical) meanings are coded in the lexicon, would be processed in parallel in both types of contexts.

EXPERIMENT 1

In Experiment 1, 48 students were shown both familiar and less familiar ironies, embedded in contexts biasing their interpretation either toward a literal or an ironic meaning. Having read the texts, the participants had to make a lexical decision as to whether a letter string was a word or a nonword. The critical manipulations were the type of word—which was either related to the ironic or the literal interpretation of the irony—and the interval between the display of the irony and the display of the letter string. This enabled us to assess the activation of both the literal and ironic meanings in both types of ironies after two different intervals (150 msec and 1,000 msec).

Method

Design. A $2 \times 2 \times 2 \times 2 \times 2$ factorial design was used with interstimulus interval (ISI; 150 msec and 1,000 msec) as a between-subjects factor, and irony type (less familiar and familiar), context type (ironically and literally biased), word type (ironically and literally related), and stimulus type (word and nonword) as within-subjects factor.

Participants. Forty-eight undergraduate students (23 women and 25 men) of Tel Aviv University, ranging from 22 to 30 years old, served as paid participants. They were all native speakers of Hebrew.

Texts. Thirty-two ironies (“target sentences”) were selected for Experiment 1 and presented in Hebrew. Sixteen were familiar ironies, and 16 were less familiar ironies. The ironies were classified as “familiar” and “less familiar” on the basis of a familiarity pretest. In the familiarity pretest, 24 undergraduates, all of whom are native speakers of Hebrew, were presented 40 contextless sentences. They participated in the test as part of their class assignments. They were asked to write down the coded meaning or meanings of the sentences. A sentence that received an ironic interpretation from more than half of the tested population was classified as “a familiar irony.” Sentences not reaching that threshold were classified as “less familiar ironies.” Sixty-four contexts, three to four sentences long, were created, two for each target sentence. One biased the last clause—the target sentence—toward the ironic interpretation (e.g., 1a and 2a translated from Hebrew; see also the Appendix), and the other biased it toward the literal interpretation (e.g., 1b and 2b translated from Hebrew; see also the Appendix):

Familiar irony:

(1a) Iris was walking on her own in the dark alley, when all of a sudden a hand was laid on her back. Startled, she turned around to find out that the hand was her young brother’s who sneaked behind her to frighten her. She said to him: “Very funny.”

(1b) Tal and Ortal, the twins, wanted to go to the movies. Their mother recommended a movie she had seen shortly before. When they came home, she was eager to know how they found the movie. They both agreed: “Very funny.”

Ironically related test word: *annoying*

Literally related test word: *amusing*

Less familiar irony:

(2a) After he had finished eating pizza, falafel, ice cream, wafers and half of the cream cake his mother had baked for his brother Benjamin’s birthday party, Moshe started eating coated peanuts. His mother said to him: “Moshe, I think you should eat something.”

(2b) At two o’clock in the afternoon, Moshe started doing his homework and getting prepared for his Bible test. When his mother came home from work at eight p.m., Moshe was still seated at his desk, looking pale. His mother said to him: “Moshe, I think you should eat something.”

Ironically related test word: *stop*

Literally related test word: *little*

Apparatus. Stimuli presentation and response collection were controlled by an IBM-compatible 386 PC, using a Pascal program. Each of the 32 target sentences was followed by one of four (Hebrew) stimuli:

- a. A word related to the literal meaning of the target sentence (e.g., *amusing*; see 1b).
- b. A word related to the ironic meaning of the target sentence (e.g., *annoying*; see 1a).
- c. A nonword created by a rearrangement of the letters of 1a (e.g., *uamsing*).
- d. A nonword created by a rearrangement of the letters of 1b (e.g., *oyignnan*).

The combination of two context types and four stimulus types created eight conditions. Two familiar and two less familiar ironies were assigned to each condition, in an 8×8 Latin square design, in which each row was assigned to six participants. All texts were arranged randomly and presented in a different order for each participant.

Procedure. Participants were tested individually. They were each seated in front of a computer and were instructed as follows:

During the experiment, you will have to assess whether a letter string that will be displayed on the screen is a word (e.g., *table*) or a nonword (e.g., *latbe*). You will be presented with short stories that will be displayed sentence by sentence, which you will have to read. After the last sentence of each story, the letter string will be displayed. You will have to press the “l” key if the string makes up a word, and the “a” key if it is a nonword. You have to press the key as fast as possible, but make sure that you do not make mistakes. Now you will be presented with 3 trial texts for training. Please put your right finger on the “l” key and the left finger on the “a” key and press one of them to start the training.

The texts were presented line by line. Each line, mostly corresponding to a sentence, appeared in the center of the screen for 3 sec. It then disappeared, and the next line was displayed. The last line—the target sentence—was displayed for a length of time that was determined by its score in a pretest. This pretest, which included 10 participants, measured the average reading time of each sentence out of context. After the target sentence was displayed for as long as it had scored in the pretest, the screen went blank for an ISI of either 150 msec or 1,000 msec. For half of the participants (three of each group of six participants assigned to each Latin square row), the ISI was 150 msec, and for half it was 1,000 msec. After the ISI, the test word (either a word or a nonword) was displayed in the center of the screen, and the participant had to respond by pressing one of two (“l” or “a”) keys. The latency between the onset of the word–nonword and the pressing of the key was measured by the computer and served as response time (RT). A 2-sec blank screen followed the response, and then the first line of the next text was displayed. The

presentation of the 32 trials began after three training trials and was preceded by two buffer trials.

Results

Both participant and item analyses were conducted. For the participant analyses, we averaged the RT of the two trials in each condition. For the item analyses, we averaged the RT of the three trials in each condition. RT outliers above or below 2 standard deviations from either the participants' or the items' mean (about 7%) were excluded from the analysis. In addition, five participants who had means 2 standard deviations above the overall mean were replaced. In both the participant and item analyses, only the responses to word stimuli were of interest, and only they were subjected to four-way analyses of variance (ANOVAs). Means and standard deviations for all conditions are presented in Table 1. The participant ANOVA included one between-subjects variable (ISI: 150 msec or 1000 msec) and three within-subjects factor: irony type (less familiar and familiar), context type (ironically and literally biased), and word type (compatible and incompatible with context). This ANOVA showed only two significant effects. First, there was a significant Context Type \times Word Type interaction, $F(1, 44) = 4.28, p < .05$. However, this effect results primarily from the three-way Irony Type \times Context Type \times Word Type interaction, $F(1, 44) = 6.31, p < .05$. The equivalent item ANOVA showed only a significant effect of ISI, $F(1, 30) = 11.45, p < .005$. To test more precisely the different predictions about Context Type \times Word Type interactions, four separate ANOVAs were performed, two for less familiar ironies (one for an ISI of 150 msec and another for an ISI of 1,000 msec) and two for familiar ironies (for both ISIs).

TABLE 1
Mean Response Times (in Milliseconds) to Words Compatible or Incompatible
With Context—Experiment 1

Word	150 msec				1,000 msec			
	Ironic Text		Literal Text		Ironic Text		Literal Text	
	M	SD	M	SD	M	SD	M	SD
Less familiar ironies								
Compatible	1,099	253	993	232	937	313	920	261
Incompatible	1,055	259	1,099	373	934	278	1,078	321
Familiar ironies								
Compatible	1,053	252	1,063	320	877	239	964	281
Incompatible	1,031	327	1,094	293	932	255	963	346

Each 2×2 ANOVA included two within-subjects factors: context type (ironically and literally biased) and word type (compatible and incompatible with context).

Less familiar ironies. For less familiar ironies (“*Moshe, I think you should eat something*”), the participant ANOVA for the 150 msec ISI revealed no main effects (all F s < 1), but a significant Context Type \times Word Type interaction, $F(1, 22) = 5.07, p < .05$, as illustrated in Figure 1 (top panel). The same interaction in the cor-

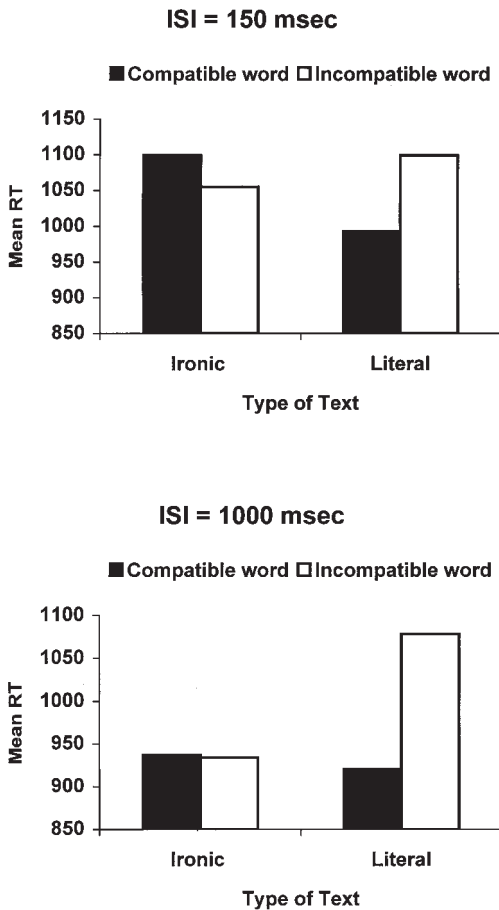


FIGURE 1 Mean response time (in milliseconds) to compatible and incompatible words related to less familiar ironies embedded in ironically and literally biasing contexts, for ISIs of 150 msec (top panel) and 1,000 msec (bottom panel).

responding item analysis was almost significant, $F(1, 15) = 4.38, p = .054$. As predicted by the graded salience hypothesis, participants were faster to respond to the salient than to the less salient test word. In the literally biasing context (2b), it was the contextually compatible (i.e., literally related) test word (*little*) rather than the contextually incompatible (i.e., ironically related) test word (*stop*) that was responded to faster. In the ironically biasing context (2a), however, it was the contextually incompatible (i.e., literally related) test word (*little*) rather than the contextually compatible (i.e., ironically related) test word (*stop*) that was responded to faster. At this early stage of processing, participants always responded faster to the salient (literally related) test word, irrespective of contextual compatibility.

The participant ANOVA for the 1,000 msec ISI produced different results. This time there was a significant context type effect, $F(1, 22) = 5.65, p < .05$, which can be explained by the significant Context Type \times Word Type interaction, $F(1, 22) = 8.30, p < .01$. The pattern of this interaction is quite different from the pattern of the previous (150 msec ISI) interaction, as the bottom panel of Figure 1 illustrates. As predicted by the graded salience hypothesis, after a 1,000 msec delay, contextual information affected comprehension of less familiar irony, and participants no longer responded faster to the salient, contextually incompatible, literally related (*little*) test word in these contexts. This was confirmed by the lack of significant effect when comparing the compatible and the incompatible responses in the ironically biasing context alone, $F(1, 22) < 1$. In contrast, the same comparison performed on results obtained from the literally biasing context reveals a significant word type effect, $F(1, 22) = 15.90, p < .001$, reflecting the fact that in this context, literally related test words were still processed faster than ironically related test words. The item analysis, however, failed to show any significant effect (all $ps > .15$). These results show that after a long delay, less salient but contextually compatible meanings no longer lag behind salient meanings.

Familiar ironies. For familiar ironies, all ANOVAs (for the 150 msec and 1,000 msec ISIs, both across participants and items) indicate no significant effect (all $ps > .20$), as predicted. According to the graded salience hypothesis, given the coded, salient status of both the literal and the ironic meanings of familiar ironies, they should both be activated initially, regardless of context.

Discussion

Our findings show that salient meanings are always processed initially, regardless of contextual information. The salient literal meaning of both familiar and less familiar ironies was swiftly available in the ironically biased context, even though it was incompatible with contextual information. Furthermore, familiar ironies facil-

itated their salient albeit contextually incompatible ironic meaning in the literally biasing contexts. These findings are consistent with the view that contextual information does not affect initial access (cf. Fodor, 1983; Swinney, 1979): At the early stage of comprehension, context neither availed the meaning compatible with it, nor did it block the meaning incompatible with it. Salient information was accessed directly and automatically. When it did not reach contextual fit, it was adjusted to contextual information. This adjustment stage occurred at a later moment of comprehension—1,000 msec after the offset of the target sentence. When it did, as in the case of familiar ironies, search for the appropriate meaning was stopped. Direct access, then, is not necessarily a function of context monitoring access of appropriate meanings (see Gibbs, 1994; Glucksberg, Kreuz, & Rho, 1986). Rather, direct access may be a function of meaning salience (for a similar view of comprehension, see Gibbs, 1980, 1982, 1983, 1986c; for recent research, see Horton & Keysar, 1996; Keysar, Barr, Balin, & Paek, 1998; Keysar, Barr, & Horton, 1998).

EXPERIMENT 2

A number of alternative explanations for the results obtained in Experiment 1 may come to mind. It is plausible that word frequency of the probes affected our results. It is also possible that the context itself, rather than the target sentences, induced the pattern of results achieved under the short ISI condition in Experiment 1. To control for these confounds, we repeated the 150 msec condition, where we found differences, with one exception. We simply presented participants with the contexts without the target sentences. We assumed that if the results of Experiment 1 were replicated in the absence of the target sentences, this would suggest that these results were not determined by the degree of salience of the targets, but rather by either word frequency or contextual information. However, if no differences ensue, this would support our claim that the early moments of comprehension are impervious to context effects, yet sensitive to degree of salience.

Method

Design. A $2 \times 2 \times 2 \times 2$ factorial design was used with context type (ironically and literally biased), word type (ironically and literally related), irony type (familiar and less familiar), and stimulus type (word and nonword) as within-subjects factors.

Participants. Twenty-four undergraduate students (14 women, 10 men) of Tel Aviv University, ranging in age from 21 to 24 years old, served as paid participants. They were all native speakers of Hebrew. They did not participate in Experiment 1.

Texts. As in Experiment 1, except for the target sentences, which were left out.

Procedure. As in Experiment 1, except for the ISI, which this time was only 150 msec.

Results and Discussion

The data were analyzed in the same way as in Experiment 1: 5.5% of the RTs that were outliers were excluded from the analyses. Means and standard deviations for all conditions are presented in Table 2. Overall, three-way participant and item ANOVAs did not produce any significant result (all $ps > .15$). Even when the data were analyzed separately for familiar and less familiar ironies, both ANOVAs indicated no significant effect (all $ps > .15$) either across participants or across items.

Experiment 2 served as a control. The results show that when no target sentences (i.e., ironies) were presented, test words were responded to with similar RTs (although with higher latencies than in Experiment 1). Neither test words' saliency and frequency nor contextual information (which, if anything, was slightly, though insignificantly, biased toward the ironic interpretation) can account for the results of Experiment 1. Such results, then, have to be attributed to the effect of the target sentences themselves (albeit in context).

GENERAL DISCUSSION

Our findings support a salient-first comprehension model (Giora, 1997, 1999, in press; Giora et al., 1998). Salient meanings were processed initially, regardless of contextual information. They were always responded to faster at the initial comprehension stage, irrespective of contextual information. For instance, in the ironically

TABLE 2
Mean Response Times (in Milliseconds) to Words Compatible or Incompatible With Context, in the Absence of Targets—Experiment 2

Word	<i>Less Familiar Ironies (150 msec)</i>				<i>Familiar Ironies (150 msec)</i>			
	<i>Ironic Text</i>		<i>Literal Text</i>		<i>Ironic Text</i>		<i>Literal Text</i>	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Compatible	1,213	304	1,183	289	1,207	363	1,237	294
Incompatible	1,200	301	1,200	330	1,286	440	1,265	340

biasing contexts, familiar and less familiar ironies were processed literally initially, in spite of contextual incompatibility. Similarly, in the literally biasing contexts, familiar ironies were processed ironically initially, regardless of contextual bias (see Experiment 1).

Experiment 2 further demonstrates that the various patterns of response exhibited in Experiment 1 were not a result of context effects or (test) word frequency. The targetless contexts in Experiment 2 did not result in different response pattern to the test words, as did the same contexts when followed by the target sentences (see Experiment 1). Taken together, these findings suggest that salient meanings, both literal and nonliteral, are initially processed alike. In contrast, utterances diverging in salience (e.g., less salient ironic interpretation and more salient literal interpretations of less familiar ironies) behave differently.

Because the intended ironic meaning of less familiar irony is not salient, it requires extra processing time to be retrieved, as shown. The literal interpretation of the same utterances, however, did not benefit from extra processing time, because in the literally biasing context, salient (literal) meaning and contextual information accidentally matched. These findings contest a direct access view of irony comprehension that assumes that context constrains comprehension even at the initial stage of comprehension so that only contextually compatible meanings reach sufficient levels of activation. They also contest a literal-first model of comprehension (cf. Grice, 1975; Searle, 1979). They show, instead, that it is salient rather than either literal or contextually compatible meaning that is activated initially.

However, it could be argued that, although there is no doubt that it is the target sentences rather than the contexts that primed the test words, our probes were not, in fact, tailored to distinguish word from message level meanings. Thus, literally related test words could be primed by words, whereas ironically related test words could only be primed by the sentence (ironic) interpretation. For instance, whereas *funny* in “Very funny” (see Example 1) could prime *amusing* (the literally related test word), interpreting the utterance “Very funny” was required before *annoying* (the ironically related test word) could be primed, because *annoying* taps the message level meaning. If this is the case, then, the version of the direct access view that argues against the priority of the utterance’s literal interpretation over the utterance’s ironic interpretation (Gibbs, 1986b, 1994; Kumon-Nakamura et al., 1995) cannot be rejected.

Note, however, that it was not always the case that the literally related test words were responded to faster. Rather, it was salience-related test words that were responded to fastest, regardless of literality and nonliterality. These findings do not stand in isolation, but are corroborated by previous findings, which include reading times of whole sentences. In Giora et al. (1998), we measured the reading times of the set of unfamiliar stimuli used here. We found that these targets took

longer to read in the ironically than in the literally biasing contexts, suggesting that nonsalient (ironic) utterances took longer to comprehend than their more salient (literal) interpretations. Recall that the contexts were not found to be more heavily biased in favor of the literal interpretations (see Experiment 2). Such findings, then, suggest that participants analyzed the salient literal meanings of sentences before deriving their nonsalient ironic interpretation.

Longer reading times for less salient (ironic) interpretations than for more salient (literal) interpretations were also found by Schwoebel et al. (1999), in which participants advanced the text phrase by phrase across the screen by pressing a key. Pexman et al. (1999) also presented evidence in favor of a salient-first model of comprehension. They used self-paced moving windows in which participants advanced a text word by word, reflecting the moment by moment processing that occurs naturally. Their findings showed that utterances took longer to read when embedded in contexts biasing their interpretation toward the nonsalient ironic meaning than when embedded in contexts biasing their interpretation toward the salient (conventional metaphoric) meaning. Such findings support the graded salience hypothesis. They go beyond lexical decision tasks and help tease apart word-level, sentence-level, and message-level effects.

Our findings replicate previous findings. In Giora et al. (1998), lexical decision tasks induced similar patterns to those found in this article, only emerging more slowly. In Giora et al., the salient literal meaning was facilitated immediately in both types of context, as found in this article, but it was also the only one available after an ISI of 1,000 msec. Less salient ironic meanings were facilitated later than a 1,000-msec delay and were available 2,000 msec after the offset of the target sentence. We speculate that this change in speed of response is a matter of fatigue. In Giora et al., participants were presented with three times as many texts as in this article, a load that may have affected the participants' alertness.

In sum, our findings show that the processes involved in irony comprehension are a function of their salience. Less familiar ironies, whose literal but not ironic meaning is coded in the mental lexicon, were processed literally first, and consequently took longer to comprehend (ironically) than their literally intended counterparts. Processing more familiar ironies, however, involved accessing both their literal and ironic interpretations initially, because both these meanings are coded in the mental lexicon. Such ironies did not take longer to process than their literal counterparts, because their intended ironic meaning was accessed directly, in parallel with the unintended literal meaning. Direct access, then, may be a function of meaning salience, rather than of context effects. According to the graded salience hypothesis, more salient ironies than the Hebrew set tested in this article may, like idioms, be processed ironically first and literally second. Future research will show whether this is true of conventional English ironies, which are more common and frequent than Hebrew conventional ironies.

ACKNOWLEDGMENTS

Support for this research was provided by grants from The Israel Science Foundation and Lion Foundation to Rachel Giora.

We thank Lior Noy and Olga Zeigelman for preparing the software for the experiment, Shay Michaely for preparing the materials, Noga Balaban and Naama Har'el for their help in administering the experiments, and Ray Gibbs, Albert Katz, Ken McRae, and two anonymous referees for very helpful comments.

REFERENCES

- Blasko, G. D., & Connine, C. (1993). Effects of familiarity and aptness on meta processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 295–308.
- 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.
- Dews, S., & Winner, E. (1999). Obligatory processing of the literal and nonliteral meanings of ironic utterances. *Journal of Pragmatics*, *31*.
- Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Gibbs, R. W., Jr. (1980). Spilling the bean on understanding and memory for idioms in conversation. *Memory & Cognition*, *8*, 449–456.
- Gibbs, R. W., Jr. (1982). A critical examination of the contribution of literal meaning to understanding nonliteral discourse. *Text*, *2*, 9–27.
- 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. (1986a). Comprehension and memory for nonliteral utterances: The problem of sarcastic indirect requests. *Acta Psychologica*, *62*, 41–57.
- Gibbs, R. W., Jr. (1986b). On the psycholinguistics of sarcasm. *Journal of Experimental Psychology: General*, *115*, 3–15.
- Gibbs, R. W., Jr. (1986c). Skating on thin ice: Literal meaning and understanding idioms in conversation. *Discourse Processes*, *9*, 17–30.
- Gibbs, R. W., Jr. (1994). *The poetics of mind*. Cambridge, England: Cambridge University Press.
- Gibbs, R. W., Jr., O'Brien, J. E., & Doolittle, S. (1995). Inferring meanings that are not intended: Speakers' intentions and irony comprehension. *Discourse Processes*, *20*, 187–203.
- 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. (in press). *On our mind: Salience, context, and figurative language*. New York: Oxford University Press.
- Giora, R., & Fein, O. (1999). On understanding familiar and less-familiar figurative language. *Journal of Pragmatics*, *31*.
- Giora, R., & Fein, O. (in press). Irony comprehension: The graded salience hypothesis. *Humor*.
- Giora, R., Fein, O., & Schwartz, T. (1998). Irony: Graded salience and indirect negation. *Metaphor and Symbol*, *13*, 83–101.
- Glucksberg, S., Kreuz, R., & Rho, S. H. (1986). Context can constrain lexical access: Implications for models of language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*, 323–335.

- Goldvarg, Y., & Glucksberg, S. (1998). Conceptual combinations: The role of similarity. *Metaphor and Symbol, 13*, 243–255.
- Gorfein, S. D. (Ed.). (1989). *Resolving semantic ambiguity*. New York: Springer-Verlag.
- Grice, P. H. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Speech acts: Syntax and semantics* (Vol. 3, pp. 41–58). New York: Academic.
- Hogaboam, T. W., & Perfetti, C. A. (1975). Lexical ambiguity and sentence comprehension. *Journal of Verbal Learning and Verbal Behavior, 14*, 265–274.
- Horton, W. S., & Keysar, B. (1996). When do speakers take into account common ground? *Cognition, 59*, 91–117.
- Katz, A., & Lee, C. J. (1993). The role of authorial intent in determining verbal irony and metaphor. *Metaphor and Symbolic Activity, 8*, 257–279.
- Kawamoto, A. H. (1993). Nonlinear dynamics in the resolution of lexical ambiguity: A parallel distributed processing account. *Journal of Memory and Language, 32*, 474–516.
- Keysar, B., Barr, D. J., Balin, J. A., & Paek, T. S. (1998). Definite reference and mutual knowledge: Process models of common ground in comprehension. *Journal of Memory and Language, 39*, 1–20.
- Keysar, B., Barr, D. J., & Horton, W. S. (1998). The egocentric basis of language use: Insights from a processing approach. *Current Directions in Psychological Sciences, 7*, 46–50.
- Kumon-Nakamura, S., Glucksberg, S., & Brown, M. (1995). How about another piece of the pie: The allusionary pretense theory of discourse irony. *Journal of Experimental Psychology: General, 124*, 3–21.
- Pexman, P., Ferretti, T., & Katz, A. (1999). *Discourse factors that influence irony detection during on-line reading*. Manuscript submitted for publication.
- Rayner, K., Pacht, J. M., & Duffy, S. A. (1994). Effects of prior encounter and global discourse bias on the processing of lexically ambiguous words: Evidence from eye fixations. *Journal of Memory and Language, 33*, 527–544.
- Récanati, F. (1995). The alleged priority of literal meaning. *Cognitive Science, 19*, 207–232.
- Schwoebel, J., Dews, S., Winner, E., & Srinivas, K. (1999). *Obligatory processing of the literal meaning of ironic utterances: Further evidence*. Manuscript submitted for publication.
- Searle, J. (1979). *Expression and meaning*. Cambridge, England: Cambridge University Press.
- Simpson, G. B., & Burgess, C. (1985). Activation and selection processes in the recognition of ambiguous words. *Journal of Experimental Psychology: Human Perception and Performance, 11*, 28–39.
- Sperber, D., & Wilson, D. (1995). *Relevance: Communication and cognition*. Oxford, England: Blackwell. (Original work published 1986)
- Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior, 18*, 645–659.
- Temple, J. G., & Honeck, R. P. (1999). Proverb comprehension: The primacy of literal meaning. *Journal of Psycholinguistic Research, 28*, 41–70.
- Turner, N. E., & Katz, A. (1997). Evidence for the availability of conventional and of literal meaning during the comprehension of proverbs. *Pragmatics and Cognition, 5*, 203–237.
- Vu, H., Kellas, G., & Paul, S. T. (1998). Sources of sentence constraint on lexical ambiguity resolution. *Memory & Cognition, 26*, 979–1001.
- Williams, J. N. (1992). Processing polysemous words in context. Evidence from interrelated meanings. *Journal of Psycholinguistic Research, 21*, 193–218.

APPENDIX

Examples of Test Stimuli

Translated sample items: (a) versions are ironically biasing contexts; (b) versions are literally biasing contexts.

Familiar Ironies

Example 1

(1a) Iris was walking on her own in the dark alley, when all of a sudden a hand was laid on her back. Startled, she turned around to find out that the hand was her young brother's who sneaked behind her to frighten her. She said to him: "Very funny."

(1b) Tal and Ortal, the twins, wanted to go to the movies. Their mother recommended a movie she had seen shortly before. When they came home, she was eager to know how they found the movie. They both agreed: "Very funny."

Ironically related test word: *annoying*

Literally related test word: *amusing*

Example 2

(2a) Ziv visited his friend, Ran, in New York. Ran advised him to use the subway, but Ziv insisted on renting a car. Three days later, Ziv gave up and told Ran: I have had enough. The traffic jam here is incredible. Ran said: "Tell me about it."

(2b) In the middle of the night Royi woke up and started crying. His mother heard him and went up to his room. "What happened?" she asked. Royi said that he had had a nightmarish dream. His mother said: "Tell me about it."

Ironically related test word: *known*

Literally related test word: *disclosing*

Example 3

(3a) At dinner Erez was talking to his family, bragging about how he beat all the neighborhood kids at chess. After his parents expressed their admiration, he admitted that actually all the competitors were at least four years younger than him. His sister retorted: "Very smart."

(3b) Two year old Galit played with the new Lego pieces she had just received. Trying hard, she finally managed to build a nice and big Lego house. Her mother said: "Very smart."

Ironically related test word: *simple*

Literally related test word: *cleverness*

Less Familiar Ironies

Example 4

(4a) After he had finished eating pizza, falafel, ice cream, wafers and half of the cream cake his mother had baked for his brother Benjamin's birthday

party, Moshe started eating coated peanuts. His mother said to him: “Moshe, I think you should eat something.”

(4b) At two o’clock in the afternoon, Moshe started doing his homework and getting prepared for his Bible test. When his mother came home from work at eight p.m., Moshe was still seated at his desk, looking pale. His mother said to him: “Moshe, I think you should eat something.”

Ironically related test word: *stop*

Literally related test word: *little*

Example 5

(5a) Just how far have women risen in the film community?

According to M. P., who was at Woman in Film luncheon recently in Los Angeles, it has actually been a very good year for women. Demi Moore was sold to Robert Redford for \$1 million in the movie *Indecent Proposal* ... Uma Thurman went for \$40,000 to Robert De Niro in the recent movie, *Mad Dog and Glory*. “Just three years ago, in *Pretty Woman*, Richard Gere bought Julia Roberts for—what was it? \$3,000?”

“I’d say women have had real progress.”

(5b) Just how far have women risen in the film community?

According to M. P., who was at Woman in Film luncheon recently in Los Angeles, it has actually been a very good year for women: Demi Moore earned \$10 million in the movie *Indecent Proposal* ... Uma Thurman made \$400,000 in the recent movie, *Mad Dog and Glory*. “Just three years ago, in *Pretty Woman*, Julia Roberts earned—what was it? \$130,000?”

“I’d say women have had real progress.”

Ironically related test word: *regress*

Literally related test word: *success*

Example 6

(6a) Tom was building an addition to his house. He was working real hard putting in the foundation. His younger brother was supposed to help. But he never showed up. At the end of a long day, when Tom’s brother finally appeared, Tom said to his brother:

“Thanks for your help.”

(6b) Tom was building an addition to his house. He was working real hard putting in the foundation. Suddenly his younger brother showed up and started to work too. At first Tom was afraid his brother would just be a nuisance. But at the end of a long day, Tom said to his brother:

“Thanks for your help.”

Ironically related test word: *angry*

Literally related test word: *useful*