

# Irony: Graded Salience and Indirect Negation

Rachel Giora and Ofer Fein

*Tel Aviv University*

Tamir Schwartz

*Haifa University*

According to the *graded salience hypothesis*, salient (e.g., conventional, frequent, familiar) meanings should be activated before less salient meanings are retrieved (R. Giora, 1997). The graded salience hypothesis thus predicts that irony should take longer to understand than nonironic language, because its comprehension should involve activation of the salient, literal meaning initially, before the less salient ironic meaning is derived. According to the *indirect negation view of irony* (R. Giora, 1995), irony comprehension should involve retention of the activated literal meaning so that the comprehender may compute the difference between the (usually desirable) state of affairs alluded to by the literal meaning and the less desirable, ironicized situation. Evidence from three experiments supports the graded salience hypothesis and the indirect negation view of irony. Experiment 1 showed that target sentences took longer to read in ironically than in literally biasing contexts. In Experiments 2 and 3, response times to test words that were either literally or ironically related to the target sentences were measured. Both literal and ironic target sentences facilitated only the literal meaning 150 and 1,000 msec after their offset. However, 2,000 msec after offset of the ironic target sentences, the ironic meaning became available and the literal meaning was still as active.

For about 2 decades, cognitive psychologists (e.g., Gibbs, 1994; Gibbs & Gerrig, 1989; Gildea & Glucksberg, 1983; Glucksberg 1989; Glucksberg, Gildea, & Bookin, 1982; Keysar, 1989, 1994; Ortony, Schallert, Reynolds, & Antos, 1978) and linguists (e.g., Sperber & Wilson, 1986/1995) have maintained that understanding literal and nonliteral language requires equivalent processes (henceforth the *processing-equivalence hypothesis*). They have argued against traditional theorists (e.g., Grice, 1975; Searle, 1979) who assumed that understanding nonliteral language requires a sequential, multiple-stage process, involving first a literal interpretation that is then revisited. In a previous study, Giora (1997) proposed an alternative approach to both views. According to Giora, it is meaning salience (rather than either literal or nonliteral meaning) that determines the type of processing invoked. The meaning of a linguistic expression is considered salient in case its interpretation can be directly computed from the lexical meanings automatically associated with entries before any extra inferences based on contextual assumptions have been derived. Though the salience of a word may be affected by, for example, familiarity, frequency, prior context, context has a limited role. Evidence from eye fixation (Rayner, Pacht, & Duffy, 1994), for instance, shows that even when prior context is heavily biased in favor of the less salient (e.g., less frequent) meaning of an ambiguous word, participants look at it longer than at its matched unambiguous control word. This finding suggests that the word's salient (e.g., more frequent) meaning has been activated as well, in spite of the bias of prior context in favor of the less salient meaning. (For a more extensive discussion of meaning salience, see Giora, 1997.) The *graded salience hypothesis* thus predicts that highly salient meanings should be processed initially (e.g., the idiomatic meaning of conventional idioms; see Gibbs, 1980). Alternative interpretations of similar salience should be processed in parallel (e.g., the literal and metaphorical meanings of conventional metaphors; see Blasko & Connine, 1993). Less salient meanings should be evoked after more salient meanings have been activated (e.g., the metaphorical meanings of unfamiliar metaphors; see Blasko & Connine, 1993).

One of the purposes of this study is to test the predictions of both the processing-equivalence hypothesis and the graded salience hypothesis regarding irony comprehension. According to the processing-equivalence hypothesis, given supportive context, ironic utterances should take no longer to read than their literal counterparts. Both the ironic meaning of an expression or utterance embedded in an ironically biasing context and the literal meaning of the same utterance embedded in a literally biasing context should be activated initially upon processing that utterance (e.g., Gibbs, 1986, 1994). In contrast, the graded salience hypothesis predicts that ironies should take longer to read than their literal counterparts. The assumption is that ironic interpretations are less salient than their literal meanings; hence, they should be derived after the literal meaning has been activated (and probably rejected as the intended meaning). Processing an utterance in a literally

biasing context, however, should activate the literal (i.e., salient) but not the ironic meaning. Along the lines suggested by Giora (1997), this article attempts to provide evidence in support of the graded salience hypothesis and against the processing-equivalence hypothesis. Thus, if target sentences take longer to read in ironically than in literally biasing contexts, this would be considered as evidence in support of the graded salience hypothesis and against the processing-equivalence hypothesis (Experiment 1). Furthermore, if ironies facilitate literally related concepts, this would also support the graded salience hypothesis (Experiment 2).

The second purpose of this study is to test the indirect negation view of irony (Giora, 1995). According to the indirect negation view, the literal meaning of an ironic utterance is functional in irony comprehension, because it enables the comprehender to compute the difference between the literal meaning (usually pertaining to a desirable state of affairs) and the ironized (less desirable) situation or object. Given that meanings relevant to comprehension are retained, while information that is no longer functional in comprehension is suppressed (e.g., Gernsbacher, 1990; Gernsbacher & Faust, 1990, 1991), the indirect negation view predicts that the activated literal meaning of an ironic utterance should be retained rather than suppressed. In this respect, the indirect negation view differs from the traditional account of irony (Grice, 1975; Searle, 1979), which assumes that the literal meaning should not be retained, but should be rejected and suppressed. Experiments 2 and 3 test this hypothesis.

## EXPERIMENT 1

The aim of Experiment 1 is to show that, as predicted by the graded salience hypothesis, but contrary to the processing-equivalence hypothesis, target sentences take longer to read in ironically than in literally biasing contexts. The assumption is that ironical interpretations are less salient than their literal meanings, hence only literal meanings should be evoked initially in either type of context. However, in the ironically biasing context, but not in the literally biasing context, the literal meaning must be rejected as the intended meaning before the ironic meaning is activated. This double take is expected to take longer than directly accessing the intended meaning only.

### Method

**Participants.** Twenty-four adults, ranging from 20 to 50 years old, working in a software company in Tel Aviv served as participants. They were all native speakers of Hebrew.

**Texts.** Twenty ironies (target sentences) were selected for the experiment and presented in Hebrew. For each target sentence (e.g., *You are just in time*, see [1] below), two contexts, three to four sentences long, were prepared. One biased the last clause—the target sentence—towards the ironic interpretation (e.g., 1a). The other biased it toward the literal interpretation (e.g., 1b). Some of the materials used here were taken from the classical studies of irony (e.g., Gibbs, 1986; Wilson & Sperber, 1992), and some were naturally occurring ironic texts for which we have contrived literal counterparts (see the Appendix). Each text was followed by a comprehension question. Seventy texts ending in either a metaphor or a literal target provided fillers. The following is an example of the two biasing contexts used:

- (1a) Anna is a great student, but she is very absent-minded. One day while I was well through my lecture, she suddenly showed up in the classroom. I said to her: "You are just in time."
- (1b) Anna is a great student and very responsible. One day she called to tell me she did not know when she would be able to show up for my lecture. However, just as I was starting, she entered the classroom. I said to her: "You are just in time."

**Apparatus.** Stimuli presentation and response collection were controlled by a PC with a 486 Intel processor, using a C++ program. It instructed the participants, presented them with the texts, and recorded their response times (RTs) and responses. The experiment began with the program randomly picking up one of the 55 pairs of segments. It then selected one of the pair and displayed it on the screen line by line. After the participant read the last line and pressed a key, a comprehension question was displayed. When the participant responded (by pressing either the "y" or the "n" key), the screen was blanked, and another text was displayed.

**Procedure.** The participants were tested individually. They were told to read the texts and answer a comprehension question that followed each text. They were told to press a key when they finished reading the text and to respond to a comprehension question by pressing either the "y" or the "n" key. Each participant read only one context version of a certain target sentence. They were given three training trials.

The texts were displayed line by line. Each line was presented for 75 msec  $\times$  the number of characters in the line. The last line—the target sentence—was marked by a page break. After participants had read this line, they pressed a key. The latency

between the onset of the target sentence and the pressing of the key was measured by the computer and served as reading time. After the response, a comprehension question was displayed. The participant responded "yes" or "no" by pressing either the "y" or the "n" key. The screen was then blanked, and another text was displayed.

## Results and Discussion

Two means were calculated—one for the reading times of the ironic targets and one for the reading times of the literal targets. As predicted by the graded salience hypothesis, participants took longer to read the ironic targets ( $M = 2.26$  sec,  $SD = 0.68$ ) than their literal counterparts ( $M = 1.93$  sec,  $SD = 0.42$ ),  $t(23) = 2.69$ ,  $p < .05$ . Contrary to current beliefs in equivalent processing of ironic and literal language (e.g., Gibbs, 1986, 1994; Gibbs, O'Brien, & Doolittle, 1995), but along the lines suggested by Giora (1995, 1997), our findings support the hypothesis that, rather than deriving the ironic meaning more or less directly (Gibbs, 1986), irony involves activating the salient literal meaning initially, before processing the less salient ironic meaning.<sup>1</sup>

## EXPERIMENT 2

According to the graded salience hypothesis, processing irony should usually activate the literal meaning initially, because most ironic meanings are context dependent and cannot be retrieved directly from the lexicon.<sup>2</sup> Similarly, processing the same utterance in a literally biasing context should activate the literal but not the ironic meaning, because only the literal meaning is salient (i.e., can be retrieved directly from the lexicon). The graded salience hypothesis thus predicts that processing an utterance in an ironically as well as in a literally biasing context

<sup>1</sup>Giora (1995) studied Gibbs (1986) at some length and reanalyzed the findings, showing that if one compares reading times of (the same) target sentences embedded in ironically versus literally biasing contexts (so that *sarcasm* is compared with *acknowledgment*—its most appropriate counterpart—rather than with its interpretation), ironies take longer to read. However, the same comparison did not reach significance in Gibbs et al. (1995).

<sup>2</sup>There are very few conventional ironies in Hebrew whose ironic meaning is salient (i.e., directly retrievable from the mental lexicon), such as, for example, *wise guy* that in Hebrew reads as *wise at night*; or *you discovered America*, which means that one is being uninformative; or *very funny*, which means aggravating. These, however, have not been included in our study.

should activate the literal meaning, whether it is compatible with the context (as in the case of a literally biasing context) or incompatible with it (as in the case of an ironically biasing context). One way to test this hypothesis is to measure RT to literally and ironically related concepts shortly (e.g., 150 msec) after the offset of the literally and ironically related utterances. Such a short delay (after the sentence was no longer on display) should enable us to see which concepts are facilitated by the target sentences. Further, to show that the activated concept is not suppressed, but remains active in both types of contexts, as predicted for the literal meaning by the indirect negation view of irony (Giora, 1995), RT to literally and ironically related concepts should be measured also (e.g., 1,000 msec after the offset of the literally and ironically related utterances). After 1,000 msec, it is contended, many meanings of lexical items that are irrelevant in a particular context are suppressed and, by some tests, are no longer available to comprehension (e.g., Gernsbacher, 1990; Swinney, 1979). Experiment 2 tests these hypotheses.

## Method

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

**Participants.** Forty-eight undergraduate and graduate students (39 women and 9 men) at Tel-Aviv University, ranging from 22 to 50 years old, served as participants. They participated in the experiment as part of their class assignments. They were all native speakers of Hebrew.

**Texts.** Sixteen ironies (target sentences), as in Experiment 1, were selected for the experiment and presented in Hebrew. For each target sentence (e.g., *You are just in time*; see [1]), two contexts, three to four sentences long, were prepared. One context biased the last clause—the target sentence—toward the ironic interpretation (e.g., see [1a] and the Appendix). The other biased it toward the literal interpretation (e.g., [1b]).

**Apparatus.** Stimuli presentation and response collection were controlled by an IBM-compatible 286 PC, using a Pascal program. Each of the 16 ironies (target

sentences), preceded by either an ironically or a literally biasing context, was followed by one of four (Hebrew) stimuli: (a) a word related to the literal meaning of the target sentence (e.g., *punctual*; see [1]), (b) a word related to the ironic meaning of the target sentence (e.g., *late*; see [1]), (c) a nonword created by a rearrangement of the letters of (a) (e.g., *cpnutlau*), or (d) a nonword created by a rearrangement of the letters of (b) (e.g., *etla*).

The combination of two context types and four stimulus types created eight conditions. Two ironies were assigned to each condition, in an  $8 \times 8$  Latin square design, in which each row was assigned to 6 participants. In addition to the 16 ironies, 32 metaphors made up the stimulus set with both metaphorically and literally biasing contexts, which served as fillers. 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 "I" 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 "I" 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 4 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 exposed for as long as it scored in the pretest, the screen went blank for an ISI of either 150 or 1,000 msec. For half of the participants (3 of each group of 6 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 word or nonword) was displayed in the center of the screen, and the participant had to respond by pressing one of two keys ("I" or "a"). The latency between the onset of the word-nonword and the pressing of the key was measured by the computer and served as RT. The response was followed by a 2-sec blank screen, and then the first line of the next text was displayed. The presentation of the 48 trials began after three training trials and was preceded by two buffer trials.

## Results and Discussion

The average RT of the two trials in each condition was the basic datum for the analysis. For each participant, a mean and a standard deviation for all eight conditions were calculated. RT of outliers above or below 2 SDs from the mean were excluded from the analysis. In addition, 2 participants who had means above 2 SDs from the overall mean were replaced. Only the responses to word stimuli were of interest to us, and only they were subjected to a  $2 \times 2 \times 2$  mixed analysis of variance (ANOVA), with one between-subjects factor, ISI (150 or 1,000 msec), and two within-subject factors, context type (ironically-literally biased) and word type (compatible-incompatible with context). Means and standard deviations for all conditions are presented in Table 1.

The ANOVA revealed no main effects (all  $F$ s < 1), but a significant Context Type  $\times$  Word Type interaction,  $F(1, 46) = 9.53, p < .005$ , as illustrated in Figure 1. 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, they

TABLE 1  
Mean Response Times (in Milliseconds) to Ironically and  
Literally Related Words in the Different Experimental Conditions

Response	<i>Interstimulus Interval</i>							
	<i>150 msec</i>				<i>1,000 msec</i>			
	<i>Ironiic Text</i>		<i>Literal Text</i>		<i>Ironiic 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	1,007	364	867	293	947	332	902	288
Incompatible	933	379	969	267	855	260	945	350







FIGURE 1 Mean response time (RT) of compatible and incompatible responses in ironically–literally biasing contexts.

responded faster to the compatible, literally related word than to the incompatible, ironically related word. In contrast, in the ironically biasing context, they responded faster to the incompatible, literally related word than to the compatible, ironically related word. That is, they always responded faster to the salient (e.g., literally related) word, irrespective of type of context. The same pattern holds for ISIs of both 150 and 1,000 msec. This is illustrated by the top and bottom panels of Figure 2 and is confirmed by the lack of significant Context Type  $\times$  Word Type  $\times$  ISI interaction,  $F(1, 46) < 1$ .

The only other significant interaction was Context Type  $\times$  ISI,  $F(1, 46) = 4.60$ ,  $p < .05$ . As illustrated in Figure 3, in the literally biasing context, participants' response times did not vary, whether they had an ISI of 150 or of 1,000 msec. However, the ironically biasing context elicited a different pattern of response. In the ironically biasing context, participants took much longer to respond when the ISI was 150 msec than when it was 1,000 msec.

Our findings support the graded salience hypothesis (Giora, 1997) and disconfirm the processing-equivalence hypothesis (e.g., Gibbs, 1986). They show that irony comprehension involves activation of the salient (literal) meaning initially, before activation of the less salient (ironic) meaning. They further demonstrate that the activated literal meaning of irony is not suppressed but remains active as predicted by the indirect negation view of irony (Giora, 1995).

### ISI = 150 msec



### ISI = 1000 msec




FIGURE 2 Mean response time (RT) of compatible and incompatible responses in ironically-literally biasing contexts, for interstimulus intervals (ISIs) of 159 msec (top panel) and 1,000 msec (bottom panel).




FIGURE 3 Mean response time (RT) of 150 and 1,000 msec responses in ironically–literally biasing contexts. ISI = interstimulus interval.

### EXPERIMENT 3

Experiment 3 was designed to further investigate the temporal aspects of irony comprehension. This time the test word was presented 2,000 msec after offset of the target sentence. According to the indirect negation view (Giora, 1995), results for the salient literally related word in the ironically biasing context should replicate Experiment 2. If comprehension of (the less salient ironic meaning of) irony requires additional time, as predicted by the graded salience hypothesis (Giora, 1997), ironically related concepts should show facilitation after such an interval.

#### Method

**Design.** A  $2 \times 2 \times 2$  factorial design was used with discourse type (ironically–literally biased), word type (ironically–literally related), and stimulus type (word–nonword) as within-subject factors.

**Participants.** Twenty-four undergraduate and graduate students (21 women and 3 men) at Tel-Aviv University, ranging from 22 to 40 years old, served as participants. They participated in the experiment as part of their class assignments. They were all native speakers of Hebrew.

*Texts and apparatus.* These were the same as in Experiment 2.

*Procedure.* This was the same as in Experiment 2. The only difference was the ISI, which was 2,000 msec for all participants.

## Results and Discussion

The average RT of the two trials in each condition was the basic datum for the analysis. For each participant, a mean and a standard deviation for all eight conditions were calculated. RT outliers above or below 2 SDs from the mean were excluded from the analysis. In addition, 1 participant who had a mean above 2 SDs from the overall mean was replaced. Only the responses to word stimuli were of interest to us. They were first subjected to a  $2 \times 2$  ANOVA, with two within-subject factors—discourse type (ironically–literally biased) and word type (compatible–incompatible with context). Means and standard deviations for all experimental conditions are presented in Table 2 and in Figure 4.

The ANOVA revealed a main effect of discourse type,  $F(1, 23) = 5.92, p < .05$ , but no main effect of word type,  $F(1, 23) < 1$ . In contrast to Experiment 2, there was no Discourse Type  $\times$  Word Type interaction,  $F(1, 23) < 1$ . To allow direct comparison between the two experiments, we performed a joined ANOVA for both. Because the patterns of results for ISIs of 150 and 1,000 msec in Experiment 2 were very similar, those conditions were pooled, and a  $2 \times 2 \times 2$  mixed ANOVA was performed, with one between-subjects factor, ISI (either 150 and 1,000 msec or 2,000 msec), and two within-subject factors, discourse type (ironically–literally biased) and word type (compatible–incompatible with context). The ANOVA revealed only one main effect, of discourse type,  $F(1, 70) = 4.20, p < .05$ , reflecting

TABLE 2  
Mean Response Times (in Milliseconds) to Words in the  
2000-msec Interstimulus Interval Experimental Conditions

Response	Irony Text		Literal Text	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Compatible	873	268	980	438
Incompatible	877	337	948	328




FIGURE 4 Mean response time (RT) of compatible and incompatible responses in ironically-literally biasing contexts for an interstimulus interval of 2,000 msec.

the fact that RTs were faster in the ironically biasing context. Most important, as predicted by the graded salience hypothesis, there was a significant three-way ISI  $\times$  Discourse Type  $\times$  Word Type interaction,  $F(1, 70) = 4.45, p < .05$ . This interaction can be understood when Figure 1 and Figure 4 are compared. When given an ISI of 150 and 1,000 msec (see Figure 1), participants were faster to respond to the salient (literally related) than to the less salient (ironically related) test word. They always responded faster to the salient test word, irrespective of type of context. However, when given an ISI of 2,000 msec, participants responded as fast to both the literally and the ironically related test word in the ironically biasing context. In the literally biasing context, this ISI elicited a slower response to the literally related test word (see Figure 4).

The only other significant interaction was Discourse Type  $\times$  ISI,  $F(1, 70) = 8.20, p < .01$ . As illustrated in Figure 5, in the literally biasing context, participants took longer to respond when the ISI was 2,000 msec than when it was 150 or 1,000 msec. In contrast, in the ironically biasing context, it took participants longer to respond when the ISI was 150 or 1,000 msec than when it was 2,000 msec.

Our findings support the graded salience hypothesis (Giora, 1997), which predicts that less salient meanings should be activated after more salient meanings have been evoked. They demonstrate that activation of the less salient ironic meaning of irony requires additional time compared to that consumed by the




FIGURE 5 Mean response time (RT) of 150 and 1,000 versus 2,000 msec responses in ironically–literally biasing contexts. ISI = interstimulus interval.

activation of the salient literal meaning of irony. In addition, they support the indirect negation view of irony (Giora, 1995), according to which the activated literal meaning should not be suppressed (as irrelevant, as predicted by the traditional account) but remain available for the computation of the difference between what is said and what is referred to.

These findings contest the processing-equivalence hypothesis (see Figure 5). Whereas irony comprehension takes longer to process than literal language (see also Experiment 1), derivation of its literal interpretation in the literally biasing context is instant (cf. Experiment 2) and might begin to fade, we speculate, after a delay of 2,000 msec. In contrast, in the ironically biasing context, that same (literal) interpretation remains active even after that length of delay, because, it is assumed, it is conducive to irony interpretation.

## GENERAL DISCUSSION

Overall, participants' response times in the literally biasing context did not vary, whether they had an ISI of 150 or 1,000 msec. In contrast, in the ironically biasing context, participants took much longer to respond when the ISI was 150 msec than when it was 1,000 msec. Further, whereas the salient literal meaning was available

immediately, the less salient ironic meaning became available rather late—2,000 msec after offset of the ironic target sentence. These findings support the graded salience hypothesis and contest the processing-equivalence hypothesis. They are inconsistent with the view that salient (e.g., literal) meanings can be bypassed given supportive context (as contended by, e.g., Gibbs, 1986, 1994; Glucksberg, Kreuz, & Rho, 1986; but for a different view, see Gibbs & Gerrig, 1989; Onifer & Swinney, 1981; Rayner et al., 1994).

Further, our findings show that the activated salient meaning of irony is retained even after a lengthy delay of 2,000 msec. These findings demonstrate that, contrary to the traditional account (e.g., Grice, 1975; Searle, 1979), the literal meaning of irony is not suppressed and replaced by the ironic meaning. Note that in this respect, irony comprehension differs from comprehension of literal ambiguities. Whereas disambiguation of literal ambiguities involves suppression of irrelevant meanings (e.g., Gernsbacher, 1990; Swinney, 1979), irony comprehension involves retention of the unintended literal meaning, as predicted by the indirect negation view (Giora, 1995).

The findings presented here are consistent with some previous findings (e.g., Dews & Winner, 1995, in press; Giora & Fein, in press) that demonstrate that some aspects of the literal meaning of irony are involved in irony comprehension. For instance, in Giora and Fein (in press), participants were presented with target sentences at the end of either ironically or literally biasing contexts. They were asked to complete one of two fragmented words (such as *t-b-e*) with the first word that came to mind. The target words were related to either the ironic or the literal meaning of the target sentence, so that activation of the different meanings could be assessed. Findings support the graded salience hypothesis. They show that in the ironically biasing context, there was no significant difference between ironically and literally related responses. However, there was a significant difference in the literal context. Processing the target sentences in the literally biasing context activated the literal meaning almost exclusively.

## ACKNOWLEDGMENTS

Support for this research was provided by grants from The Israel Science Foundation and Lion Foundation to the first author. We also thank Lior Noy for preparing the software for the second and third experiments and Shy Michaeli for his help in administering Experiments 2 and 3. Thanks are also due to Mira Ariel, John DuBois, Herb Colston, Ray Gibbs, Robert E. Haskell, Tanya Reinhart, and an anonymous reviewer for their 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. (1995). Muting the meaning: A social function of irony. *Metaphor and Symbolic Activity*, 10, 3–19.
- Dews, S., & Winner, E. (in press). Obligatory processing of the literal and nonliteral meanings of ironic utterances. *Journal of Pragmatics*.
- Gernsbacher, M. A. (1990). *Language comprehension as structure building*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Gernsbacher, M. A., & Faust, M. E. (1990). The role of suppression in sentence comprehension. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 97–128). Amsterdam: North-Holland.
- Gernsbacher, M. A., & Faust, M. E. (1991). The mechanism of suppression: A component of general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 245–262.
- 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. (1986). On the psycholinguistics of sarcasm. *Journal of Experimental Psychology: General*, 115, 3–15.
- Gibbs, R. W., Jr. (1994). *The poetics of mind*. Cambridge, England: Cambridge University Press.
- Gibbs, R. W., Jr., & Gerrig, R. J. (1989). How context makes metaphor comprehension seem "special." *Metaphor and Symbolic Activity*, 4, 145–158.
- 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.
- Gildea, P., & Glucksberg, S. (1983). On understanding metaphor: The role of context. *Journal of Verbal Learning and Verbal Behavior*, 22, 577–590.
- 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., & Fein, O. (in press). Irony comprehension: The graded salience hypothesis. *Humor*.
- Glucksberg, S. (1989). Metaphors in conversation: How are they understood? Why are they used? *Metaphor and Symbolic Activity*, 4, 125–143.
- Glucksberg, S., Gildea, P., & Bookin, H. G. (1982). On understanding nonliteral speech: Can people ignore metaphors? *Journal of Verbal Learning and Verbal Behavior*, 21, 85–98.
- 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.
- Grice, P. H. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Speech acts: Vol. 3. Syntax and semantics* (pp. 41–58). New York: Academic.
- Keysar, B. (1989). On the functional equivalence of literal and metaphorical interpretations in discourse. *Journal of Memory and Language*, 28, 375–385.
- Keysar, B. (1994). Discourse context effects: Metaphorical and literal interpretations. *Discourse Processes*, 18, 247–269.
- Onifer, W., & Swinney, D. A. (1981). Accessing lexical ambiguities during sentence comprehension: Effects of frequency of meaning and contextual bias. *Memory & Cognition*, 9, 225–236.



- 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.
- 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.
- Searle, J. (1979). *Expression and meaning*. Cambridge, England: Cambridge University Press.
- 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.
- Wilson, D., & Sperber, D. (1992). On verbal irony. *Lingua*, 87, 53-76.

## APPENDIX EXAMPLES OF TEST STIMULI

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

### Example 1

- (1a) Anna is a great student, but she is very absent-minded. One day while I was well through my lecture, she suddenly showed up in the classroom. I said to her: "You are just in time."
- (1b) Anna is a great student and very responsible. One day she called to tell me she did not know when she would be able to show up for my lecture. However, just as I was starting, she entered the classroom. I said to her: "You are just in time."

Ironically related test word: *late*

Literally related test word: *punctual*

### Example 2

- (2a) 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."

- (2b) 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 3

- (3a) Tom was building an addition to his house. He was working very 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."

- (3b) Tom was building an addition to his house. He was working very 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*

### 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: *little*

Literally related test word: *stop*

