The meaning of negated adjectives

TAMAR FRAENKEL and YAACOV SCHUL

Abstract

The study investigates the way people use and make sense of negated adjectives. Past research showed that by using a negated adjective, instead of an available antonym, one is able to communicate a mitigated sense of that antonym. To illustrate, by saying 'not hot'*one can communicate* 'neither hot, but not quite cold'*. This effect has been termed the* mitigation hypothesis*. Our theoretical analysis suggests that the extent of mitigation should vary as a function of two factors. First, mitigation should be more pronounced for contraries (adjectives that lie on a continuum) than for contradictories (adjectives that form a dichotomy); Second, the extent of meaning mitigation of marked adjectives should be stronger than that of unmarked adjectives. Finally, we hypothesized that these two factors interact, so that the markedness effect should be stronger for contraries than for contradictories. We report results from three experiments that tested these hypotheses with native speakers of Hebrew and discuss alternative mechanisms that might lead to mitigation of negated terms. We also address the practical importance of our findings for questionnaire design and communication.*

1. The meaning of negated adjectives

Negations are prevalent in communication. Mehl and Pennebaker (2003) analyzed natural conversations of students and found that negations were slightly more prevalent than words connoting positive emotions, two times more frequent than words connoting negative emotions, and almost three times more prevalent than words used to denote causality. This occurs in spite of findings which show that processing negations is often harder, slower, and more error-prone than processing affirmations (e.g., Carpenter & Just 1975; Clark & Chase 1972; Just & Clark 1973; Lea &

Intercultural Pragmatics 5-4 (2008), 517–540 1612-295X/08/0005–000517
DOI 10.1515/IPRG.2008.025

Mulligan 2002; MacDonald & Just 1989; Wason 1965). This raises the question of why. Why use a negation when the message can be conveyed by using the corresponding opposite term affirmatively?

Past research has suggested several reasons, which are not mutually exclusive. Negations might be used as a means of being polite (Colston 1999; Giora et al. 2005; Horn 1989). That is to say, communicators can avoid being perceived as blatant by negating an adjective instead of using an affirmation which has a strong derogatory connotation (e.g., saying "John is not smart" is more polite than saying "John is stupid"). Negations might also be employed when one wants to contradict a common expectation or belief held by the receiver (Allowd 1977; Clark & Clark 1977; Givon 1978; Jordan 1988; Leech 1983; Wason 1965). Thirdly, negations are sometimes used to convey understatement or irony (Giora, Balaban, Fein, & Alkabetz 2005; Giora et al. 2005). The current work focuses on a fourth reason: According to the *mitigation hypothesis*, which was formalized close to 90 years ago (e.g., Jespersen 1976 [1924]), when one member in a pair of antonymic adjectives (e.g., hot/cold) is negated (e.g., "not hot") it conveys a weakened sense of the antonym (i.e., "cold"). Therefore, when people want to convey a mitigated sense of an adjective, they negate its opposite.

In what follows, we first describe past research that bears on the hypothesis. Then we discuss two factors that may moderate the level of mitigation produced by negation: the type of the antonymic pair and the markedness of the adjective that is negated.

1.1. *The mitigation hypothesis*

According to the hypothesis, negations are used to convey weakened messages. For example, someone may say that the coffee is "not hot" to indicate that the coffee is lukewarm, being neither cold nor hot. Although the hypothesis was discussed by Jespersen (1976 [1924]) and even by Plato, empirical demonstrations of this sense of mitigation could be numbered (see Giora 2006 for a review). Indeed, the mitigation hypothesis is often taken for granted, as when negations are used in Horn scales (Fogelin 1967; Horn 1978; Merin 2003).

Research suggests that the mitigation hypothesis can be attributed to two fundamentally different mechanisms that are involved in (i) encoding of a negated adjective and (ii) interpreting it in an interpersonal interaction. Schematically, encoding a negated statement could be described as developing in two phases. During the initial phase of comprehension, recipients activate the associations of each linguistic element in the message. Accordingly, while encoding the statement "the coffee is not hot",

recipients activate inferences that are associated with "coffee", "hot", and the "not" operator (e.g., Hasson & Glucksberg 2006; Giora 2006; Kaup et al. 2006). As they continue processing, however, recipients integrate the different words in the statement and in doing so they can activate inferences that are congruent with "not hot" or "cold" (Fischler et al. 1983; Gannon & Ostrom 1996; Mayo, Schul, & Burnstein 2004; Whitmeyer 1997). The simultaneous presence of associations to "hot" and "cold" may induce a moderated meaning of the negated concept, with "not hot" being understood as neither very hot nor very cold.

The mitigation of the meaning of a negated message may also stem from the pragmatic inferences recipients make during the interaction. According to Grice (1975) and Horn (1984), communicators prefer expressions that incorporate all of the information they want to convey and at the same time are parsimonious. As a result, upon hearing a non-parsimonious communicative element such as negation, the listener often assumes that it has been generated for a purpose. Such conversational implicatures allow one to go beyond the literal meaning of the statement and provide a rich field of shades of meanings (Horn 2004). For example upon hearing the statement "the coffee is not hot" the listener may wonder why the affirmative statement, which has a simpler form, was not used. The listener may conclude that negation was used purposely to convey a weakened sense. Accordingly, "not hot" is understood to mean neither hot, nor cold but something in between.

We are not concerned here with demonstrating one of these mechanisms or the other. In fact, we believe that both work in tandem, giving rise to the mitigated meaning accorded to negated statements. The current research seeks to refine the mitigation hypothesis by considering two characteristics of negated adjectives that can serve as moderators for the extent of weakening. To this end, we study how the nature of the negated adjective—whether it is contrary or contradictory—and its degree of markedness influence the extent of mitigation.

1.2. *Mitigation of meaning in negation of contraries and contradictories*

A pair of antonymic adjectives are considered contraries if they can be used in two sentences that cannot simultaneously be true but may simultaneously be false (e.g., the coffee is cold or hot, or neither hot nor cold). Contradictories, in contrast, are antonyms which follow the *law of missing middle*, meaning that they can be used in two sentences, exactly one of which must be true (e.g., the door is open or closed, but the door must be either open or closed). This distinction is quite old, dating back to Aristotle, but also appearing in more recent essays (for example, Ladusaw

1996). We hypothesize that the extent of the mitigation of negated adjectives is moderated by whether the adjectives are members of a contradictory or a contrary pair. In particular, mitigation should be greater when a negated adjective anchors a continuum—belongs to a contrary pair—than when the adjective is a part of a dichotomy—belongs to a contradictory pair. The hypothesis is consistent with the intuition behind the *mitigation hypothesis* because, in a pure contradictory pair, the negation of one adjective must imply its antonym, as there is no middle point. Paradis and Willners (2006) recently reported findings consistent with this assumption. Unfortunately, the interpretation of their findings is problematic because contradictions and contraries were investigated in different experiments, thereby compromising random assignment.

1.3. *Mitigation of meaning in the negation of marked and unmarked adjectives*

In this section we consider the effect of the markedness of an adjective on how its negation is interpreted. To anticipate the conclusion of this discussion, we hypothesize that the negation of marked adjectives leads to a greater mitigation of meaning than the negation of unmarked adjectives (see Figure 1).

Mitigation is often operationalized through similarity of meaning, since mitigation is inversely related to similarity. In the extreme case, the assertion that the meaning of a negated concept (e.g., "not dead") is not mitigated at all is equivalent to the assertion that it has the same meaning as its antonym (i.e., "alive"). Because similarity of meaning is easier to follow, we shall often relate to the existence of mitigation in terms of "dissimilarity", and lack of mitigation in terms of "similarity" of the negated adjective and the adjective's antonym. Thus, assuming that "good" is the unmarked member of the good/bad pair, our hypothesis implies that "not good" should resemble "bad" more than "not bad" resembles "good". This is equivalent to asserting that the meaning of "not bad" is more strongly mitigated than the meaning of "not good".

In order to explain why the hypothesis makes sense, let us briefly discuss the notion of markedness. Linguists use the term markedness in many ways: phonologically, syntactically, and semantically. Hartmann and Stork (1972) defined the *marked* member of a binary pair as the one which carries a distinctive feature that distinguishes it from the other member. The *unmarked* member is typically the usual, the normal, the positive, the common, and the neutral or less specific, compared to the *marked* member (Battistella 1996; Boucher & Osgood 1969; Levinson 1983). The term, though controversial, has been found useful not

only as a means to understanding linguistic structures (Jacobson & Pomorska 1983), but also in explaining errors in second-language learning (Santos 1987), or even the organization of kinship universals (Hage 2001).

Lehrer (1985) discussed markedness in relation to gradable antonyms. She examined different criteria for defining the marked member in pairs of antonymic adjectives, and concluded that markedness is not one structural characteristic but rather consists of several independent but correlated characteristics. Using her suggestions, we operationalized the *unmarked* member of a pair of antonyms as the one that is used neutrally in asking a question, is more common, is used to name the entire scale, and is associated with the positive meaning of the scale.

We hypothesize that the extent of meaning mitigation induced by negation depends on the markedness of the adjectives. Specifically, we propose that the negation of the *unmarked* adjective weakens the message *less* than the negation of the *marked* adjective. Returning to our previous example, given that "good" is the unmarked member in the good/bad pair, we hypothesize that "not good" resembles "bad" (small discrepancy—slight mitigation) more than "not bad" resembles "good" (high discrepancy—strong mitigation).

The intuition behind this proposal involves the way the marked adjective is understood. By definition, the marked member is a modification of the unmarked member, thus having a unique "mark". The unmarked member, in contrast, is defined independently of that mark. Accordingly, when an unmarked adjective is modified by negation, it acquires the mark of being "not unmarked" which makes it similar to the marked member. When the marked adjective is negated, it loses its mark. However, losing the mark need not bring meaning of the negated adjective closely to that of the unmarked antonym since the unmarked adjective is defined independently of the mark.

From a more psychological point of view, this prediction could be justified by considering the range of meanings associated with the unmarked and marked adjectives. A marked adjective has a narrower and more specific range of meanings than its unmarked counterpart. This can lead to two related outcomes. First, because of the greater richness of potential meanings that are associated with the unmarked adjective, when it is negated its meaning is susceptible to a greater shift than the marked adjective. Second, since the marked adjective is mainly the opposite of its [unmarked] antonym, it should more closely resemble the negation of the unmarked adjective. The unmarked adjective has a broader meaning (being not merely the negation of the marked antonym), and therefore it should resemble the negated marked antonym less closely.

Support for the significant role markedness plays in the interpretation of negation is suggested by several lines of work. Clark (1974) refers to the asymmetry within antonymic pairs, where one member has a contrast meaning (as the opposite of its antonym) and the other—a neutral meaning. He defines the contrast member as the negation of the neutral one, but not vice versa. Although not defined in terms of markedness, the suggestion is consistent with our hypothesis. Dease (1964) found that the associative connection between the marked adjective and the corresponding unmarked antonym is stronger than that between the unmarked adjective and the corresponding marked antonym. Mann (1968) and Huttenlocher and Higgins (1971) found that people often use negations in defining marked rather than unmarked adjectives (e.g., ''sad'' means ''not happy'').

Marked adjectives are more likely to be negative than their unmarked counterparts (Boucher & Osgood 1969; Horn 1989; Lehrer 1985). With negativity as an indicator of markedness, our hypothesis is indirectly supported by results showing that negating an adjective with a positive meaning is not symmetric to negating its antonym, with a negative meaning (Blutner 2004; Blutner & Solsdat 2001; Horn 1989). Specifically, the negation of a positive adjective implies the negative one, whereas the opposite need not be true.

So far we have discussed two potential factors assumed to moderate the extent of the mitigation of the meanings of negated adjectives. The next question, then, is whether the two moderators are independent of each other, or whether they interact. We hypothesize that the type of adjective (being a member of a contrary or a contradictory pair) and its markedness interact, so that the effect of markedness is smaller in contradictory adjectives than in contraries. The intuition behind this hypothesis has to do with the fact that the two adjectives in a contradictory pair are logically defined as the negations of each other. Consequently, the meaning of the negation of one of the adjectives is likely to gravitate toward the other, whether it is marked or unmarked. In contrast, with contraries, intermediate meanings are expected. In this case, therefore, the effect of markedness should be more pronounced.

Figure 1 presents our hypotheses in a graphic format. According to the first hypothesis, negations of contraries (top figure) weaken the antonym to a greater extent than negations of contradictories (bottom figure). The second hypothesis suggests that negations of marked adjectives (right side of top figure) weaken their meaning to a greater extent than negations of unmarked adjectives (left side of top figure). According to the third hypothesis, the latter effect should be more pronounced with contraries than with contradictory adjectives (compare with bottom figure, where

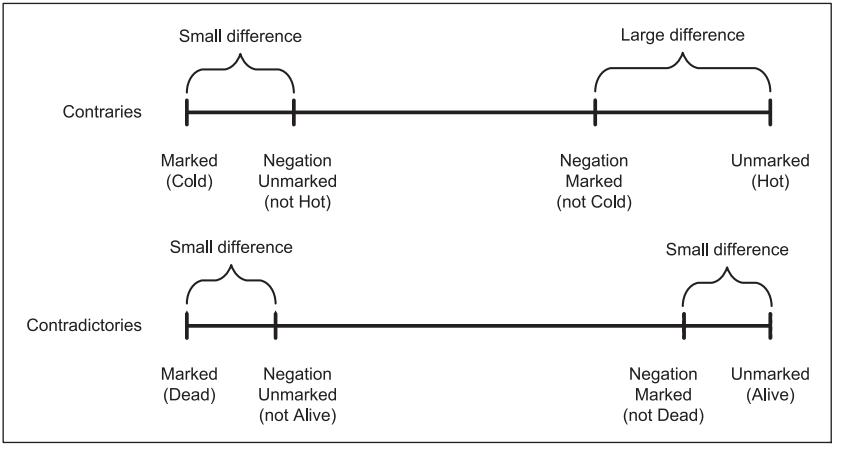

Figure 1. *Summary of the hypotheses*

there is no difference). Three experiments were conducted to test these hypotheses.

2. Experiment 1

This experiment investigated participants' assessments of the meaning resemblance between a negated adjective and its antonymic counterpart. Pairs of sentences were presented on a computer screen, one above the other. One sentence contained a negated adjective (e.g., the coffee is *not hot*) and the other contained the adjective's affirmative antonym (e.g., the coffee is *cold*). Participants rated the resemblance in meaning between the sentences. We intentionally provided a minimal context in which each adjective or negation was embedded. This was done to balance the need to provide a pragmatic context for interpreting each adjective (or its negation) and the need not to constrain the meaning of the adjectives (or negations) by an elaborated context (but see Experiment 3).

Specifically, although markedness is typically discussed as a feature of an adjective pair, it might be sensitive to context variation. Since the unmarked adjective is the normal, common, or expected attribute in a particular setting, it is quite possible that what is normal in one setting can be abnormal in another setting. Consider, for instance the hot/cold pair. *Hot* is expected when someone refers to hot drinks like coffee, but unexpected when one refers to the weather in the Arctic Zone. *Happy* might be the unmarked member of the happy/sad pair in normal situations

where happiness prevails (Matlin & Stang 1978), but *sad* may become the unmarked adjective in situations where sadness prevails. To overcome such context-sensitivity we provided the respondents in the current experiment with minimal contextual information and counterbalanced the contextual information across participants (see below). We further elaborate on this issue in the General Discussion.

2.1. *Method*

2.1.1. *Participants.* Thirty-two students from the Hebrew University participated in the experiment, either for partial course credit, or for a small monetary compensation (the equivalent of US$2). They were all native speakers of Hebrew.

2.1.2. *Stimuli and procedure.* Participants saw pairs of sentences that were displayed on a computer screen one above the other. One sentence contained a negated adjective and the other its affirmative antonym. Participants had to rate the meaning resemblance of the sentences on a 21-point scale (0 = entirely different; 20 = identical). The scale appeared below the sentences and responses were made by clicking on a particular point on the scale with the computer mouse. Table 1 contains an example for the sentences and the scale.

2.1.3. *Stimuli construction.* The adjectives were selected through a multi-stage selection process. The first phase involved the creation of a list of antonym pairs based on a Hebrew thesaurus. Next, we filtered out repetitions or words which are not standard Hebrew, resulting in 572 pairs of antonyms. This list was further refined in 4 stages of pre-tests which verified that the antonyms (i) are perceived as antonyms; (ii) are not rare; (iii) are perceived unambiguously as either contrary or

Table 1. *An example of the stimulus material and the scale used in Experiment 1*

First block:	The coffee is hot.	The coffee is not cold.
Second block:	The soup is cold.	The soup is not hot.
Completely Identical		Completely Different
20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0		

Note: Stimulus sentences and scales were presented in Hebrew, which is written in a right-to-left direction.

contradictory; and (iv) there was a consensus among four expert judges (doctoral students in linguistic) as to which adjective in the pair was marked. This process left us with a list of 45 antonymic pairs—24 contrary antonyms and 21 contradictory antonyms. The antonyms for all the experiments were selected from this list. The adjectives used in Experiment 1 appear in the Appendix.

2.1.4. *Counter-balancing.* For each pair of adjectives (e.g., hot-cold), we constructed two pairs of sentences, one containing the unmarked adjective and the negation of the marked adjective (e.g., "The soup was *hot*"\"The soup was *not cold*"), and the other containing the marked adjective and the negation of the unmarked adjective (e.g., "The coffee was *cold*"\"The coffee was *not hot*"). This resulted in 60 pairs of sentences that included the experimental adjectives. Because the interpretation of negations might be influenced by the context (e.g., Colston 1999), half of the participants received the unmarked adjectives and the negation of the marked adjectives embedded in one pair of sentences, and the other half saw them embedded in the other pair of sentences (e.g., "The coffee was *hot*"\"The coffee was *not cold*" and "The soup was *cold*"\"The soup was *not hot*"). In addition, we created 12 filler pairs, each of which contained an adjective and the negation of an unrelated adjective. The fillers were used to anchor the lower end of the resemblance scale.

The 72 pairs of sentences that each participant saw were divided into two blocks. The sentences in each block contained 15 pairs of contraries, 15 pairs of contradictories, and 6 fillers. Half of the sentence pairs contained the negation of the marked adjective and half contained the negation of the unmarked adjective. The order of presentation was controlled.

2.1.5. *Design.* The experiment manipulated two within-participant factors: Adjective type (contrary \ contradictory), and the Markedness of the negated adjective (marked \ unmarked). The dependent measure was the rated resemblance between a negated adjective and its antonym.

2.2. *Results*

One participant was eliminated from the analysis because he failed to follow the instructions. For each of the remaining 31 participants we calculated four resemblance scores by averaging each participant's resemblance ratings within each of the four types of sentence pairs (contrary/contradiction $\times$ marked/unmarked). Thus, each mean was based on 15 pairs of sentences. The means computed across participants are presented in Table 2.

Table 2. *Means (SD) of resemblance judgments as a function of Adjective Type and Markedness (Experiment 1)*

	Contraries	Contradictories
Marked is negated	11.26 (2.78)	17.03 (1.80)
Unmarked is negated	10.26 (2.45)	16.94 (2.02)

Note: Greater resemblance indicates that the negated adjective and the corresponding antonym were seen as more similar to each other, implying a lesser mitigation of meaning.

A two-way repeated-measures ANOVA revealed significant main effects for adjective type ($F(1, 30) = 180.3$, $p < 0.01$. $d = 4.89$) and markedness ($F(1, 30) = 9.93$, $p < 0.01$, $d = 1.15$), and an adjective type $\times$ markedness interaction ($F(1, 30) = 7.53$, $p < 0.01$, $d = 1.01$). In line with the first hypothesis, participants rated the similarity in meaning of a negated adjective to its antonym as greater when the two adjectives were contradictory (e.g., "*not alive*" and "*dead*", $M_{similarity} = 16.96$) than when they were contraries (e.g., "*not cold*" and "*hot*", $M_{similarity} = 10.76$). In line with our second hypothesis, negated unmarked adjectives were rated as more similar to the corresponding antonym (e.g., "*not hot*" and "*cold*", $M = 14.15$) than negated marked adjectives (e.g., "*not cold*" and "*hot*", $M = 13.60$). Finally, in line with the third hypothesis, the markedness effect was qualified by the interaction. Whereas the differences between negating a marked adjective ($M = 10.26$) and negating an unmarked adjective ($M = 11.26$) were significant for contraries, they failed to reach significance for contradictory adjectives.

3. Experiment 2

Experiment 2 investigated the mitigation hypothesis with a different measurement procedure. We were inspired by a debate in the questionnaire-survey literature about the use of bipolar and unipolar scales (Gannon & Ostrom 1996; Yorke 2001). Briefly, bipolar scales have the scale endpoints labeled with two antonymic adjectives. In contrast, unipolar scales have the endpoints labeled with an adjective and its negation. Based on the mitigation hypothesis, we reasoned that the range of values covered by a unipolar scale should be smaller than the one covered by a bipolar one. Moreover, in line with our hypotheses about the adjective type and markedness, the mapping from a unipolar scale to the corresponding bipolar scale should depend on the type and markedness of the adjectives being used as the scale's endpoints. Respondents in Experiment 2 were

asked to mark the range covered by the anchors of a unipolar scale (for example, hot—not hot) on the corresponding bipolar scale (hot—cold). We manipulated the type and markedness of the adjective used for labeling the unipolar scale.

3.1. *Method*

3.1.1. *Participants.* Twenty students from the Hebrew University participated in the experiment, for either course credit or a small monetary fee, the equivalent of US$2. The participants were all native speakers of Hebrew.

3.1.2. *Stimulus construction.* Adjectives for the second experiment were selected from the same pool of antonyms whose construction was discussed in Experiment 1 (see Appendix).

3.1.3. *Counter-balancing.* For each of the bipolar scales (e.g., rich-poor) we constructed two corresponding unipolar scales, one in which the endpoints were labeled by the marked adjective and its negated version (e.g., not poor / poor) and another in which the endpoints were labeled by the unmarked adjective and its negation (rich / not rich). Overall, the 30 antonymic pairs gave rise to 30 bipolar scales and 60 unipolar scales. These were divided into two blocks. Each bipolar scale appeared in both blocks—in one block with its corresponding unipolar unmarked scale, and in the other block with its corresponding unipolar marked scale. In each block, 15 scales used contraries and 15 scales used contradictories, 15 unipolar scales were marked and 15 were unmarked. The order and distribution of the scales between the two blocks were varied among participants to control for presentation order effects.

3.1.4. *Design.* The experiment manipulated two within-participant factors: adjective type (contrary \ contradictory), and the markedness of the adjective used in the unipolar scale (marked \ unmarked). The resemblance between a negated adjective and its antonym was indicated by the length of the bipolar scale that was marked by participants as corresponding to the unipolar scale.

3.1.5. *Procedure.* Each participant was given a 60-page booklet. On each page the participant saw a bipolar scale (e.g., a scale anchored by "*rich*" and "*poor*") and the labels anchoring a unipolar scale underneath. These labels consisted of one of the adjectives from the bipolar scale and

Table 3. *Mean (SD) portion of bipolar scale covered by the unipolar scale, calculated separately by adjective type and markedness of the adjective used in the unipolar scale (Experiment 2)*

	Contraries	Contradictories
Unmarked is unipolar	54.42% (12.58)	78.99% (9.72)
Marked is unipolar	51.29% (12.29)	80.37% (9.26)

Note: Larger numbers indicate that the negated adjective and the corresponding antonym were seen as more similar to each other. Greater similarity indicates a lesser mitigation of meaning.

its negated version (e.g., "*rich*" and "*not rich*"). Participants were asked to color (with a marker) a range on the bipolar scale which corresponded to the unipolar scale. Participants who thought that the negated adjective (e.g., "*not rich*") was very similar to the corresponding antonym (e.g., "*poor*") marked a larger portion of the bipolar scale than participants who thought that the negated adjective was a weakened version of the antonym.

3.2. *Results*

For each participant we averaged the proportions of the bipolar scale that was seen as corresponding to the unipolar scale. This was done separately for each of the four categories of scales defined by the adjective type (contrary\contradictory) and the markedness of the negated adjective. The means of these proportions are presented in Table 3.

We were interested in the extent of the resemblance between the unipolar and bipolar scales, because it would be revealing about the similarity of meaning between the negated adjective and the corresponding antonym. A two-way repeated-measures ANOVA revealed main effects for Adjective Type ($F(1, 19) = 106.6,\ p < 0.01,\ d = 4.72$) and Markedness ($F(1, 19) = 4.75,\ p < 0.05,\ d = 1.00$), and an Adjective-Type by Markedness interaction ($F(1, 19) = 40.9,\ p < 0.01,\ d = 2.93$). In line with the first hypothesis, unipolar scales whose endpoints were labeled by a contradictory adjective and its negation were seen as more similar to the bipolar scales than unipolar scales whose endpoints were labeled by a contrary adjective and its negation (79.68% vs. 53.33%). Put differently, a negated adjective was rated as more similar in meaning to its antonym when the adjectives were contradictory.

In accordance with the second hypotheses, when the unipolar scale was labeled by the unmarked adjective and its negation, it was seen as more similar to the bipolar scale than when it was labeled by the marked

adjective and its negation (67.23% vs. 65.78%). That is, overall, negated unmarked adjectives resembled their antonyms to a greater extent than negated marked adjectives did. Finally, the two-way interaction indicates that the effect of markedness differed for contraries and contradictions. For contraries, unipolar scales anchored by unmarked adjectives and their negations were seen as more similar to the bipolar scales than unipolar scales anchored by the marked adjectives and their negations, $t(19) = 5.14$, $p < 0.05$, $d = 2.34$. For contradictions, there was a non-significant trend in the opposite direction, $t(19) = -0.92$ $p = 0.37$, $d = .42$).

4. Experiment 3

Experiment 1 demonstrated that participants understood negated adjectives as mitigated versions of their corresponding antonyms and that the magnitude of the mitigation depended on the type and markedness of the adjective being negated. Experiment 2 provided converging support by showing that the resemblance between unipolar and bipolar scales was also influenced by whether the anchors of the bipolar scales were contradictory or contrary, and whether the unipolar scale was anchored by a negation of the marked or the unmarked adjective. Thus, when participants were queried directly about the meaning of negation, their responses showed that they understood negations in line with our three hypotheses. However, do people apply this understanding when they use negations to make inferences? Experiment 3 attempts to generalize our earlier findings by investigating the impact of negated (vs. non-negated) information on decisions.

Experiment 3 also differs from Experiments 1 and 2 in the richness of the contextual information. In the first two experiments we minimized the amount of contextual information. In contrast, in Experiment 3, the items were embedded in a richer context which allowed for a more ecologically valid investigation of the way negations are used and understood in communication.

Respondents in Experiment 3 were asked to make judgments on the basis of short verbal descriptions that included an adjective that was either contrary or contradictory, either marked or unmarked, and, critically, either negated or not. For example, respondents had to decide whether to hire a gardener whose description included one of the following: being *thorough*, *not thorough*, *careless*, or *not careless* (a between-participants factor). We investigated the differences in the decision to hire the person in question and, in particular, whether the impact of

negated adjectives is weaker than the impact of their corresponding antonyms. According to the mitigation hypothesis, decisions based on the former should be less extreme than decisions based on the latter. Moreover, as in Experiments 1 and 2, we tested whether the mitigation of meaning is sensitive to the type of adjective (contrary/contradictory) and/or its markedness.

4.1. *Method*

4.1.1. *Participants.* Eighty students at the Hebrew University participated in the experiment. Some of them received partial course credit for their participation; others received a small monetary compensation. They were all native speakers of Hebrew.

4.1.2. *Procedure.* Participants were presented a booklet containing 16 scenarios, each on a separate page. For each scenario participants were requested to imagine themselves in a particular role and make a decision. For example, in one scenario participants were told, "Imagine that you are considering whether to go on a date with Yael. Your friend tells you 'Yael is a second year law student at the Hebrew University. She is renting an apartment in an upper-middle class neighborhood in Jerusalem with a friend. Yael is (pretty/not pretty/ugly/not ugly, depending on the experimental condition). She is intelligent and has a sense of humor. She likes to dance and she is interested in music and movies.'" Each scenario was followed by a question (e.g., "Would you like to meet Yael?") and a 21-point response scale (0 = not at all; 20 = very much).

4.1.3. *Stimulus construction.* Adjectives from 16 antonym pairs, 8 contraries and 8 contradictories were used as targets in the 16 scenarios (see Table 4). For each pair of antonyms we constructed 4 identical descriptions which differed only in the target adjective. Each description, which was 3–4 sentences long, contained a target adjective which was either marked or unmarked, and either negated or not. For example, in the "dating Yael" scenario, participants were asked to consider a description of Yael which included one of the attributes pretty, not pretty, ugly, or not ugly.

4.1.4. *Design.* The Experiment involved a $2 \times 2 \times 2$ repeated measures design. Each participant read two scenarios in each combination of (1) adjective type (contrary\contradictory); (2) markedness (marked\unmarked); and (3) negation (negated/not-negated). The allocation

Table 4. *Mean judgment response (Experiment 3)*

	Adjectives	Marked		Negated Unmarked	Negated Marked		Un-marked
Contrary	Industrious-Lazy	7.30	−	5.85	14.55	+	16.35
	Thorough-Careless	11.45	−	9.26	14.00	+	16.15
	Happy-Sad	9.20	+	11.15	16.65	+	17.05
	Curious-Apathetic	9.40	+	9.55	14.89	+	15.94
	Clear-Vague	9.30	+	9.95	14.15	+	14.60
	Pretty-Ugly	8.65	+	10.10	14.15	+	17.95
	Deep-Superficial	10.85	+	11.65	16.60	+	17.00
	Sensitive-Impassive	7.35	+	7.45	14.65	−	11.80
Mean contrary adj.		9.19		9.37	14.96		15.86
Contradictory	Together-Apart	12.80	+	12.90	16.40	+	16.80
	Successful-Failing	11.90	−	11.40	15.40	+	17.90
	Right-Wrong	10.95	+	13.40	16.55	+	18.55
	Visible-Hidden	12.63	−	9.40	12.50	+	15.20
	Awake-Asleep	3.00	+	4.30	12.32	+	12.80
	Open-Jammed	11.55	−	9.05	16.40	−	14.05
	Hit-Miss	7.85	+	9.20	17.00	−	16.10
	Vacant-Occupied	5.40	−	5.30	11.65	+	12.95
Mean contradictory adj.		9.51		9.37	14.78		15.54

Notes: Translated from Hebrew (see footnote 1). Marked adjectives are underlined. "+" indicates that judgments based on a negated adjective were less extreme on the average than judgments based on the antonym. Responses were given on a 0–20 scale, with endpoint labels indicating that higher numbers were aligned with the unmarked adjectives.

of the 16 scenarios to the $2 \times 2 \times 2$ design was done systematically. Using a Latin Square we created eight different combinations in which the 16 scenarios were joined. Each combination was presented to 10 participants.

4.2. *Results*

Table 4 presents the mean decision response for each of the scenarios, separately for each adjective and its negation. Inspection of the averages of the mean responses reveals that decisions based on contrary adjectives were in line with the mitigation hypothesis, that is, responses based on negated adjectives were less extreme than those based on the corresponding

antonyms, and this mitigation was greater when the negated adjective was marked. At the same time, the pattern of decisions for contradictory adjectives only partially supported the hypothesized pattern, which is illustrated in Figure 1. Specifically, decisions based on negated unmarked adjectives were, on the average, lower in magnitude (9.37) than those based on the corresponding (marked) antonym (9.51). Thus, in this case, the negation did not mitigate the interpretation, but instead slightly but insignificantly amplified it.

To test our hypotheses statistically we performed both participant (F_1) and item (F_2) ANOVAs. In the participant analysis, we computed eight decision scores for each participant, each based on two scenarios in the combination of type of adjective pair (contrary vs. contradictory) $\times$ markedness (marked vs. unmarked) $\times$ negation/affirmation. In the item analysis, we averaged decision responses of all participants in each combination of adjective-type, markedness, and negation/affirmation, and computed a mitigation score for each adjective as the difference between decisions based on the adjective's negation (e.g., *not lazy*) and its antonym (e.g., *industrious*). The item analysis seemed particularly useful in the present design because, unlike participants in Experiments 1 and 2, each participant in the Experiment 3 read descriptions containing *either* a negation or its antonym, so that the negation/antonym difference could not be computed within participant.

Both analyses showed a strong mitigation effect, $F_1(1, 78) = 311.47$, $d = 3.99$; $F_2(1, 14) = 182.98$, $d = 7.23$. Still, even though the means were ordered as expected (see Table 4), we could not reject the null hypothesis of no difference for either the adjective type, $F_1(1, 78) = 0.29$, $d = 0.12$; $F_2(1, 14) = 0.14$, $d = 0.20$, or markedness, $F_1(1, 78) = 1.45$, $p = .24$, $d = 0.27$; $F_2(1, 14) = 1.67$, $p = 0.22$, $d = 0.69$. In part, the failure to reject the null hypotheses reflects the small number of adjective pairs and the large variability due to context effect on the interpretation of the adjectives and/or their negations.

In order to partially overcome this problem, we examined the pattern of the mitigation effects, ignoring the magnitude of the mitigation. Overall, decisions based on negations were less extreme than those based on the corresponding antonym in 23 out of the 32 pairs (see "+" in Table 4), permitting rejection of the overall null hypothesis of no mitigation, $\chi^2(1) = 6.25$, $p = .01$. Separating this effect according to the markedness of the negated adjective, we find that in 13 of the 16 cases, decisions based on negated marked adjectives were less extreme than those based on the corresponding antonym, $\chi^2(1) = 6.25$, $p = .01$. However, only 10 of the 16 decisions based on the unmarked adjectives were more extreme, $\chi^2(1) = 1.00$, $p = .32$. Thus, the results of the more qualitative analysis

are consistent with our previous findings, showing that negating a marked adjective does not shift its meaning toward its corresponding antonym to the same extent that negating an unmarked adjectives does.

Similar analysis, examining contraries and contradictions, revealed equivalent results. As the entries in Table 4 show, in 13 of the 16 contrary pairs (averaging over markedness), the decision based on the negation of an adjective was less extreme than the decision based on its antonym, $\chi^2(1) = 6.25$, $p = .01$. However, for contradictory pairs, this occurred in only 10 of the 16 cases, $\chi^2(1) = 1.00$, $p = .32$. This pattern is consistent with our earlier findings which suggest that negating an adjective of a contrary pair is more likely to lead to mitigation of meaning than negating an adjective of a contradictory pair. In fact, as we noted earlier, the decisions based on negated adjectives in contradictory pairs were, on the average, slightly more extreme than those based on the antonyms.

5. General discussion

Communicators have a rich array of linguistic devices that can be combined to express shades of meaning (Giora 2006). The current paper investigates a particular linguistic device—negation. Results from three experiments show that messages containing a negated adjective are understood as a weakened version of a message containing the antonym of that adjective, and that this meaning mitigation depends on whether the adjective and its corresponding antonym form a contradictory or a contrary pair and whether the negation is applied to the marked or the unmarked adjective.

The influence of the type of antonym pair—contrary versus contradictory—is not surprising. If a specific situation can be *strictly* viewed as either X or Y, then, logically, if it is not-X, it must be Y, with little room for mitigation. From this perspective, the mitigation hypothesis does not make sense with truly contradictory pairs. However, it is quite clear that contradictory pairs are not quite that dichotomous; just anyone can come up with possible realistic interpretations in which an apparent logical dichotomy has meaningful mid-values (Paradis & Willners 2006). Terms like ''half dead'' or ''barely alive'', ''almost right'' or ''not quite wrong'', are easily understood and often used in conversations. Still, our findings reveal that, although contradictory pairs are not truly dichotomous, they are seen as significantly more complementary than contrary pairs, in the sense that meaning mitigation occurs more strongly in contrary pairs.

The other factor involved in moderation of meaning mitigation—the markedness of the adjective—is worthy of further consideration in this Special Issue, because markedness might be sensitive to pragmatic and social inferences made in specific contexts. For instance, we already noted that in many cultures marked adjectives tend to be evaluatively unfavorable and unmarked adjectives evaluatively favorable. Addressees may reason that a communicator is using a negation of a favorable (unmarked) adjective to describe others because asserting directly that a particular person has an unfavorable quality is not pleasant or polite (Colston 1999). Such reasoning cannot be applied, however, to statements about favorable qualities, since describing another person positively is both polite and pleasant.

An opposite communication bias may occur when communicators describe themselves. In making self-descriptions communicators may wish to appear modest, and this desire may vary as a function of the culture and the person's societal role (e.g., Cialdini et al. 1998). Modesty can be achieved by using negations of unfavorable attributes (e.g., "my performance was not bad") instead of the corresponding favorable attributes ("my performance was good").

In order to infer the communicator's "true" sentiment, addressees must remove the "contamination" due to politeness or modesty. When they suspect politeness, addressees should adjust the interpretation of a negation of a favorable (unmarked) attribute toward the antonym ("she is not pretty" means "she is ugly"), to a greater extent than they should adjust the interpretation of a negated unfavorable (marked) adjective. Accordingly, predictions from politeness coincide with predictions from level of markedness. However, if addressees suspect modesty, for example, when communicators are describing themselves, the opposite effect should occur, and thus, the effect of the adjective's markedness could be dissociated from the adjective's unfavorability. To the best of our knowledge, the impact of the self/other factor on the asymmetry between negating favorable and unfavorable attributes has not been tested.

The present research considered markedness as a linguistic phenomenon. Can we generalize from this effect to non-linguistic domains? The most promising direction, in our view, involves comparing norms and exceptions (Kahneman & Miller 1986). Speculatively, let us assume that normal events are analogous to unmarked concepts, while exceptions are analogous to marked ones. From this perspective it follows that negating a normal event makes it abnormal to a greater extent than negating an abnormal (or exceptional) event makes it normal. In part, this has to do with the tendency to draw inferences from some abnormal characteristics to abnormality in other aspects—termed the ubiquitous halo effect (Cooper 1981). Accordingly, when people try to negate a specific

characteristic, they fail to undo all the inferences that are associated with that characteristic (Schul & Burnstein 1985; Schul & Mayo 1999). Given this perspective, the effect of markedness on the interpretation of negated concepts is a specific case of people's difficulty in discounting information, implying that mitigation should disappear under the conditions specified for successful discounting (Schul & Burnstein 1998).

The analogy between unmarkedness and normality brings us back to the relevance of contextual information in determining markedness. Because context affects what is normal or expected, it can influence which member of an adjective pair is the unmarked member, and as a result, how negations are understood. Consider, for example, the following two situational contexts:

1: Kim told her father that she had wrecked his car. Her boyfriend asked her how he took it and she said: (a la Colston 1999)
2: Kim told her father that she had won a scholarship for college. Her boyfriend asked her how he took it and she said:
   a: He was mad
   b: He was not mad
   c: He was happy
   d: He was not happy

"Mad" (in a) is expected or normal in the context of (1); "happy" (in c) is normal in the context of (2). As the impact of negation depends on the normality or expectedness of the negated concept, the context should affect how "not-mad" (in b) is interpreted. Specifically, in context (1), where "mad" is normal, it should resemble "happy" more than in context 2, where it is abnormal. This reasoning complements Colston's (1999) theoretical analyses and results, at least where negative adjectives are concerned.

Our findings about the mitigation of the meaning of negated terms have implications for survey researchers. In a recent review, Schaeffer and Presser (2003: 77) compared the unipolar scale (interesting/not-interesting) to its bipolar counterpart (interesting/boring). They noted, "One might assume that the category 'not at all interesting' in the unipolar version includes all the positions between 'extremely boring' and 'neither boring nor interesting' in the bipolar version, but little is known about how respondents actually perceive the difference between the two versions." Our findings provide relevant data to resolve their uncertainty. Briefly, it depends on the type of adjective. Specifically, participants in Experiment 2 of our study indicated that the range of the unipolar scale for contrary adjectives covers roughly 50% of the bipolar scale. This was not the case when the unipolar and the bipolar scales contained contradictory adjectives. In this

case, the unipolar scale included about 80% of the range of the bipolar scale. Thus, pollsters who use unipolar scales involving contraries might be missing a significant chunk of the corresponding bipolar scales.

Our view of mitigation differs from that of Giora (e.g., Giora et al. 2005) in a subtle yet psychologically significant aspect. Consider, for example, the statement "the coffee is not cold". Our analysis, as well as Giora's, proposes that this statement might be interpreted as "the coffee is neither cold nor warm". However, whereas our analysis considers the mitigation as a deviation from the assertion "the coffee is warm," Giora and Fein consider it as a deviation from "the coffee is cold".

There is a strong argument in favor of Giora's view. Processing explicit negations must start from the negated core—the adjective which is being negated (e.g., Deutsch, Gawronski, & Strack 2006; Giora et al. 2005; Hasson & Glucksberg 2006; Kaup, Ludtke, & Zwaan 2006). Accordingly, negation might be considered an adjustment of that adjective, which turns out to be insufficient, in the sense that the two alternatives, X and not-X, are not complementary (Bonini, Orsherson, Viale, & Williamson 1999; Yaniv, Schul, Raphaelli-Hirsch, & Maoz 2002).

Why, then, do we view mitigation as an attenuation of the antonym? Basically, our view is motivated by two arguments. First, consider how negations are understood according to Figure 1. Negations are interpreted as closer to the meaning of the adjective's antonym than to the meaning of the adjective itself. In the Figure, as well as in our findings, this is seen by the crossover of the two negated terms—each approaches the adjective's antonym. For example, such crossover occurred in all the scenarios in Table 4, including both negations of contrary and contradictory adjectives. To describe this phenomenon as mitigation of the meaning of the original adjective might be confusing.

Second, the negation operator undoubtedly signals a need for meaning adjustment. However, the magnitude of the adjustment may depend not on the adjective that is being negated but rather on the extreme alternative (e.g., the antonym). Upon hearing the statement "The coffee is not cold," listeners may ask, "So what might it be? Is it freezing? Is it hot?" It makes (pragmatic) sense to opt for one of these alternatives in terms of how they fit the context rather than how far they deviate from the original adjective ("cold"). Clearly, this is *not* done in the initial phases of processing (where one focuses on the negated core), but rather at a later stage of processing. Indeed, considerations of politeness and even irony require that listeners consider the extreme alternative and draw inferences from it.

Our study compared direct and indirect measures of mitigation. The direct effects (Experiments 1 and 2) were strong and robust, indicating that when participants were asked directly about the meaning of negated

adjectives, the negated adjectives were interpreted as mitigated versions of the corresponding antonyms. The two experiments showed very similar patterns of findings, in which the extent of the mitigation was chiefly influenced by whether the negated term belonged to a contrary or a contradictory pair. The effect of markedness was more minor, occurring only for contraries. Experiment 3 investigated mitigation in a more complex fashion, as it was assessed indirectly, via the impact of negations on judgments. Unlike Experiments 1 and 2, where context effects were kept minimal, in Experiment 3 the negated statements were embedded in specific contexts which could have influenced how the negated concept was interpreted. Perhaps this was the reason that only the relatively crude analyses, assessing whether or not mitigation occurred, demonstrated the hypothesized effects. We interpret this, speculatively, as indicating that in a few cases the context led to augmentation rather than mitigation of the negated adjectives, perhaps due to contrast effects.

Taken together, the statistically significant effects, as well as the effects which failed the statistical criterion, point to the importance of context for understanding how negations are interpreted. Our data suggest that negations are generally interpreted as inviting a mitigated sense of the alternative, namely, connoting a quality that is not as extreme as the antonym. Nevertheless, the extent of mitigation varies as a function of the properties of what is being negated (marked vs. unmarked), the relationship between the negated information and the alternative (contrary or contradictory pairs), and the context in which the negated concept is embedded. The latter might well be the most potent factor, with the least amount of research done to uncover its effect.

Acknowledgment

This article is based on a doctoral dissertation submitted to the Hebrew University by the first author under the supervision of the second author. The research was supported by grants from the Israeli Science Foundation (ISF) and the K-Mart Center for Retailing and International Marketing to YS. We would like to thank Ruth Mayo, Rachel Giora, and two anonymous reviewers for comments on an earlier version of this manuscript.

Appendix

Adjectives used in Experiments 1 and 2
(translated from Hebrew[1])

Contrary pairs (Experiment 1): clear-vague, pretty-ugly, friendly-solitary, rich-poor, delicate-coarse, sensitive-impassive, deep-shallow, joyous-sad, industrious-lazy, fair-dark, hairy-bald, thorough-superficial, happy-depressed, warm-cool, smooth-rough
Contradictory pairs (Experiment 1): alive-dead, connected-disconnected, right-wrong, real-fabricated, mixed-separated, awake-asleep, fixed-variable, intact-broken, opened-closed, vacant-occupied, whole-torn, free-imprisoned, successful-failing, open-jammed, hit (target)-miss
Contrary pairs (Experiment 2): fair-dark, warm-cool, industrious-lazy, thorough-superficial, mature-childish, happy-depressed, smooth-rough, pleasant-annoying, curious-apathetic, rich-poor, hairy-bald, joyous-sad, pretty-ugly, friendly-solitary, sensitive-impassive
Contradictory pairs (Experiment 2): finite-infinite, real-fabricated, together-apart, conclusive-ambiguous, alive-dead, free-imprisoned, successful-failing, connected-disconnected, right-wrong, visible-hidden, direct-indirect, awake-asleep, opened-closed, intact-broken, vacant-occupied

Note

1. The sense of antonymity might be weakened due to translation. Hebrew does not have negation modifiers such as "un" or "in" and all pairs of antonyms used in the experiments consisted of two completely different words.

References

Allwood, Jens. 1977. Negation and the strength of presuppositions or there is more to speaking than words. In Osten Dahl (ed.), *Logic, pragmatics and grammar*. Lund: Studentliteratur. 11–52.

Battistella, Edwin L. 1996. *The logic of markedness*. Oxford: Oxford University Press.

Blutner, Reinhard and Torgrim Solstad. 2001. Two case studies in lexical pragmatics. In Karoly Bibok, and Eniko E. Nemeth (eds.), *Pragmatics and the flexibility of word meaning*. Amsterdam: Elsevier. 11–28.

Blutner, Reinhard. 2004. Pragmatics and the lexicon. In Laurence R. Horn and Gergory Ward (eds.), *Handbook of pragmatics*. Oxford: Blackwell. 488–514.

Bonini, Nicolao, Daniel Osherson, Riccardo Viale, and Timothy Williamson. 1999. On the psychology of vague predicates. *Mind & Language* 14: 377–393.

Boucher, Jerry and Charles E. Osgood. 1969. The Pollyanna hypothesis. *Journal of Verbal Learning and Verbal Behavior* 8: 1–8.

Carpenter, Patricia and Marcel A. Just. 1975. Sentence comprehension: A psycholinguistic processing model of verification. *Psychological Review* 82: 45–73.

Cialdini, Robert B., Wilhelmina Wosinska, Amy J. Dabul, Robin Whetstone-Dion and Irena Heszen. 1998. When social role salience leads to social role rejection: Modest self-

presentation among women and men in two cultures. *Personality and Social Psychology Bulletin* 24: 473–481.
Clark, Herbert H. 1974. Semantics and comprehension. In Thomas A. Sebeok (ed.), *Current Trends in Linguistics. Volume 12: Linguistics and Adjacent Arts and Sciences*. The Hague: Mouton. 1291–1498.
Clark, Herbert H. and William G. Chase. 1972. On the process of comparing sentences against pictures. *Cognitive Psychology* 3: 472–517.
Clark, Herbert H. and Eve V. Clark. 1977. *Psychology and Language: An Introduction to Psycholinguistics*. San Diego: Harcourt Brace Jovanovich.
Colston, Herbert L. 1999. "Not Good" is "Bad", but "Not bad" is not "Good": An analysis of three accounts of negation asymmetry. *Discourse Processes* 28: 237–256.
Cooper, William H. 1981. Ubiquitous halo. *Psychological Bulletin* 90: 218–244.
Deutsch, Roland, Bertram Gawronski and Fritz Strack. 2006. At the boundaries of automaticity: Negation as reflective operation. *Journal of Personality and Social Psychology* 91: 385–405.
Fischler, Ira, Paul A. Bloom, Donald G. Childers, Salim E. Roucos and Nathan W. Perry. 1983. Brain potentials related to stages of sentence verification. *Psychophysiology* 20: 400–408.
Fogelin, Robert J. 1967. *Evidence and meaning*. New York: Humanities Press.
Gannon, Katherine M. and Thomas M. Ostrom. 1996. How meaning is given to rating scales: The effects of response language on category activation. *Journal of Experimental Social psychology* 32: 337–360.
Giora, Rachel. 2006. Anything negatives can do affirmatives can do just as well, except for some metaphors. *Journal of Pragmatics* 38: 981–1014.
Giora, Rachel, Noga Balaban, Ofer Fein and Inbar Alkabetz. 2005. Negation as positivity in disguise. In Herbert L. Colston and Albert N. Katz (eds.), *Figurative language comprehension: Social and cultural influences*. Hillsdale, NJ: Erlbaum. 233–258.
Giora, Rachel, Ofer Fein, Jonathan Ganzi, Natali Alkeslassy Levi and Hadas Sabah. 2005. On negation as mitigation: The case of negative irony. *Discourse Processes* 39: 81–100.
Givon, Talmy. 1978. Negation in language: Pragmatics, function, ontology. In Peter Cole (ed.), *Syntax and semantics 9: Pragmatics*. New York: Academic Press. 69–112.
Grice, H. Paul. 1975. Logic and conversation. In Peter Cole and Jerry L. Morgan (eds.), *Syntax and Semantics: Vol 3, Speech Acts*. New York: Academic Press. 43–58.
Hage, Per. 2001. Marking theory and kinship analysis: Cross-cultural and historical applications. *Anthropological Theory* 1: 197–211.
Hartmann, Reinhard R. K. and Francis C. Stork. 1972. *Dictionary of language and linguistics*. London: Applied Science Publishers.
Horn, Laurence R. 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicatures. In Deborah Schiffrin (ed.), *Meaning, form, and use in context*. Washington: Georgetown University Press. 11–42.
Horn, Laurence R. 1978. Some aspects of negation. In Joseph H. Greenberg (ed.), *Universals of Human Language, Vol 4: Syntax*. Stanford: Stanford University Press. 127–210.
Horn, Laurence R. 1989. *A natural history of negation*. Chicago: Chicago University Press.
Horn, Laurence R. 2004. Implicature. In Gergory Ward and Laurence R. Horn (eds.) *The handbook of pragmatics*. Oxford: Blackwell Publishing. 3–28.
Huttenlocher, Janellen and E. Tory Higgins. 1971. Adjectives, comparatives, and syllogisms. *Psychological Review* 78: 487–504.
Jespersen, Otto. 1976 (1924). *The philosophy of grammar*. New York: Norton.
Jordan, Michael P. 1998. The power of negation in English: Text, context and relevance. *Journal of Pragmatics* 29: 705–752.

Just, Marcel A. and Herbert H. Clark. 1973. Drawing inferences from the presuppositions and implications of affirmative and negative sentences. *Journal of Verbal Learning and Verbal Behavior* 12: 21–31.

Kaup, Barbara, Jana Ludtke and Rolf A. Zwaan. 2006. Processing negated sentences with contradictory predicates: Is a door that is not open mentally closed? *Journal of Pragmatics* 38: 1033–1050.

Ladusaw, William A. 1996. Negation and polarity items. In Shalom Lappin (ed.), *The handbook of contemporary semantic theory*. Oxford: Blackwell. 321–341.

Lea, R. Brooke and Elizabeth J. Mulligan. 2002. The effect of negation on deductive inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 28: 303–17.

Leech, Geoffrey N. 1983. *Principles of pragmatics*. London: Longman.

Lehrer, Adrienne. 1985. Markedness and antonymy. *Journal of Linguistics* 21: 397–429.

Levinson, Stephen C. 1983. *Pragmatics*. Cambridge: Cambridge University Press.

MacDonald, Maryellen C. and Marcel A. Just. 1989. Changes in activation levels with negation. *Journal of Experimental Psychology* 15: 633–642.

Mann, John W. 1968. Defining the unfavorable by denial. *Journal of Verbal Learning and Verbal Behavior* 7: 760–766.

Matlin, Margaret W. and David J. Stang. 1978. *The Pollyanna principle: Selectivity in language, memory, and thought*. Cambridge, Mass.: Schenkman.

Mayo, Ruth, Yaacov Schul, and Eugene Burnstein. 2004. "I am guilty" vs. "I am innocent": Successful negation may depend on the schema used for its encoding. *Journal of Experimental Social Psychology* 40: 443–449.

Mehl, Matthias R. and James W. Pennebaker. 2003. The sounds of social life: A psychometric analysis of students' daily social environments and natural conversations. *Journal of Personality and Social Psychology* 84: 857–870.

Merin, Arthur. 2003. Replacing 'Horn scales' by act-based relevance orderings to keep negation and numerals meaningful. Available online in http://semanticsarchive.net/Archive/2JmMzM2N/rhsabro.pdf

Paradis, Carita and Caroline Willners. 2006. Antonymy and negation: The boundedness hypothesis. *Journal of Pragmatics* 38: 1051–1080.

Schaeffer, Nora C. and Stanley Presser. 2003. The science of asking questions. *Annual Review of Sociology* 29: 65–88.

Schul, Yaacov and Eugene Burnstein. 1985. When discounting fails: Conditions under which individuals use discredited information in making a judgment. *Journal of Personality and Social Psychology* 49: 894–903.

Schul, Yaacov and Eugene Burnstein. 1998. Suspicion and discounting: Ignoring invalid information in an uncertain environment. In Jonathan M. Golding and Colin M. MacLeod (eds.), *Intentional forgetting: Interdisciplinary approaches.* Hillsdale, NJ: Erlbaum. 321–348.

Schul, Yaacov and Ruth Mayo. 1999. Two sources are better than one: The effects of ignoring one message on using a different message from the same source. *Journal of Experimental Social Psychology* 35: 327–345.

Wason, Peter C. 1965. The contexts of plausible denial. *Journal of Verbal Learning and Verbal Behavior* 4: 7–11.

Whitmeyer, Joseph M. 1997. Contrasting trait adjectives as distinct psychological entities. *Psychological reports* 80: 739–743.

Yaniv, Ilan, Yaacov Schul, Ronna Raphaelli-Hirsch and Ifat Maoz. 2002. Inclusive and exclusive modes of thinking: Studies of prediction, preference, and social perception during parliamentary elections. *Journal of Experimental Social Psychology* 38: 352–367.

Yorke, Mantz. 2001. Bipolarity ... or not? Some conceptual problems relating to bipolar rating scales. *British Educational Research Journal* 27: 171–186.