

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Intentionality, speaker's attitude and the processing of verbal irony

Permalink

<https://escholarship.org/uc/item/624931c2>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 42(0)

Authors

Ronderos, Camilo Rodríguez

Noveck, Ira

Tomlinson, Jack

Publication Date

2020

Peer reviewed

Intentionality, speaker's attitude and the processing of verbal irony

Camilo Rodríguez Ronderos (camilo.ronderos@hu-berlin.de)

Humboldt-Universität zu Berlin, Berlin, Germany

Ira Noveck (ira.noveck@gmail.com)

Université Paris Diderot-Paris 7, Paris, France

Jack Tomlinson (otomlins@gmail.com)

Leibniz-zentrum Allgemeine Sprachwissenschaft, Berlin, Germany

Abstract

Does it take more or less time to read ironic sentences than to read literal equivalents? Though this question has been extensively discussed in the literature, the results are mixed (see eg. Filik & Moxey, 2010). The present work attempt to account for the differences in the literature by considering the variable effect of anticipating the intentions of a speaker during comprehension of ironic utterances used to answer yes/no questions, as well as the role of explicit cues regarding the attitude of a speaker. The results show that both of these factors interact and together modulate the interpretation of a sentence as ironic or literal as well as the utterance's reading times. We interpret the results are broadly in line with the predictions made by the echoic mention account.

Keywords: irony comprehension; self-paced reading; experimental pragmatics

Introduction

The most commonly addressed issue in the psycholinguistics of irony has been the compared speed of reading an ironic utterance and its literal equivalent:

- 1) *Juana: "Thanks a lot for your help today Maria!"*
[context: Maria promises to help her friend Juana move but shows up only when the move is done]

Will hearing (1) be easier to understand for Maria when it is meant literally (if she had actually helped Juana move, for example) than when it is understood ironically (as intended in 1)?

This question was originally asked in order to evaluate the psychological plausibility of Paul Grice's account of figurative language understanding (Grice, 1975) but then turned into a broader discussion regarding the specific contextual and linguistic properties of an utterance that modulate irony comprehension: When will an ironic *Thanks a lot!* be easy to understand? Does it depend on the properties of the stimulus, such as frequency or familiarity (Giora, 2003) or on contextual constraints, such as a speaker's occupation (Katz & Pexman, 1997)? Despite many years of research on this question, it is still unclear whether or not irony is per se faster or slower to process than its equivalent literal sentence, as evidenced by the mixed experimental results (Filik & Moxey, 2010; Gibbs, 1986; Ivanko & Pexman, 2003; Katz et al., 2004).

Spotorno and Noveck (2014) suggested that a way to overcome this impasse would be to consider the role that

intentionality recognition has throughout an experimental session, showing that, as an experiment progresses, participants get faster at grasping irony when they have reason to anticipate a speaker's upcoming ironic intention.

More generally, the ability to anticipate a speaker's intentions (known as mind reading, or theory of mind) has been shown to play a crucial role for irony understanding (Happé, 1993; McDonald, 2000), and is a critical component of a prominent theory of irony understanding, the echoic mention account (Wilson & Sperber, 1992, 2012). According to the echoic mention account, to understand irony a listener must understand that the speaker is attributing an utterance to someone else and simultaneously expressing a dissociative attitude towards it.

The current work intends to further explore the role of anticipating a speaker's intentions during irony comprehension by examining whether processing predictions derived from the echoic theory could account for differences in both the interpretation and speed of processing of literal and ironic utterances. Specifically, we hypothesized that participants' ability to anticipate the overall communicative intention of a speaker (in our experiments, the answer to a yes/no question), as well as explicit discourse cues regarding the speaker's attitude towards a proposition, could together modulate reading times of ironic utterances as well as determine whether or not an ironic interpretation is reached. Our aim is to continue bridging the gap between psycholinguistic experiments on irony processing and theoretical accounts of irony comprehension.

Reading irony: fast or slow?

Within his theoretical framework, Grice saw irony as a type of implicature triggered by a violation to his Maxim of Quality ('say only what you believe to be true'). As such, he believed ironic meaning to be a proposition that can be systematically derived from a literal utterance by substituting the utterance with its opposite meaning. To test the psychological validity of this approach, Gibbs (1986) translated this model into a sequential process with three stages: Computation of the literal meaning, recognition of the violation, and derivation of the implicature. Each of these stages was seen to require a specific amount of time, which resulted in the logical conclusion that deriving ironic meaning (3-stages process) would necessarily take longer than understanding an equivalent literal utterance (1-stage process). In his experiments, Gibbs (1986) went on to show

that a sentence such as *you're a fine friend* does not necessarily take longer when intended ironically compared to when it is meant literally, thus proving the standard pragmatic model wrong. He interpreted the finding as support for his Direct Access Model, according to which, given strong contextual constraints, ironic meaning can be grasped immediately without needing to compute a literal meaning of the same utterance first.

This finding triggered several other researchers to further explore the conditions under which the computation of the literal meaning can be bypassed, which resulted in the postulation of different psycholinguistic accounts of irony comprehension. Such is the case of the Graded Salience Hypothesis (Giora, 2003; Giora et al., 2007), which claims that specific features of individual words (such as familiarity or prototypicality) are responsible for biasing towards a literal or ironic interpretation of an utterance (with a literal interpretation being the default in the majority of the cases). The constraint satisfaction model (Pexman, 2008), on the other hand, claims that discourse cues, such as the occupation of the speaker (Katz & Pexman, 1997), provide evidence that can allow for a rapid identification of the ironic meaning of an utterance.

Some experimental tests of these models have shown that literal sentences are processed faster than their ironic equivalents in a majority of the cases, regardless of contextual constraints (Fein et al., 2015; Filik & Moxey, 2010). However, some other studies have suggested that under certain conditions irony is indeed just as easy to understand than their literal counterparts (Ivanko & Pexman, 2003; Katz et al., 2004), in line with the original findings of Gibbs (1986). The discrepancy in findings makes it hard to draw conclusions regarding the status of the standard pragmatic model, as well as to determine which factors play the biggest role during irony comprehension.

Intentions, Attitudes and Irony Comprehension

Instead of comparing reading times, an alternative line of research on irony comprehension has focused on the overarching cognitive mechanisms that allow for irony to be understood. Such is the case of the echoic mention account, developed within the framework of relevance theory (Sperber & Wilson, 1986). Relevance Theory differentiates between descriptive and interpretative use of language. Descriptive use of language takes place when language is used to represent an actual or possible state of events believed by the speaker to be true. Interpretative use of language, on the other hand, is used by speakers to represent another representation, i.e. an actual or possible utterance or thought.

Irony is seen as a type of interpretative language use in which a speaker 'points' the hearer towards a different representation by using an utterance that resembles said representation in terms of its content. The speaker does this to express a dissociative attitude towards said utterance, which gives irony its characteristic evaluative tone: In (1), Juana intends to 'echo' the social norm of thanking people

when they are helpful while simultaneously disassociating herself from the idea of using this utterance in this particular situation.

From the listener's perspective, echoic mention states that understanding irony is a product of understanding that a speaker holds a representation about another representation: Understanding *Thanks a lot* as ironic involves anticipating the communicative intent of the speaker (in 1, to express disappointment) and integrating this information during processing. Crucially, this is only possible if listeners monitor and anticipate a speaker's beliefs, desires and intentions, an ability known as Theory of Mind, or mindreading (Baron-Cohen et al., 1985).

Various studies have highlighted the importance of mindreading abilities for irony comprehension. Happé (1993) examined typically developing as well as autistic children's understanding of ironic utterances and found that those with underdeveloped mindreading abilities performed worse in an irony comprehension task than those with normally developing mindreading skills. McDonald (2000) reviewed several studies conducted with patients that suffered from damage to the right hemisphere and patients with traumatic brain injury and found a correlation between a patient's ability to understand irony and their ability to monitor and interpret a speaker's intentions. Spotorno et al. (2012) found that the neural networks typically associated with mindreading activity show increased activation during irony comprehension.

The relationship between reading times of irony comprehension and mindreading was first investigated by Spotorno & Noveck (2014). In experiment 1, they found that literal utterances are read significantly faster than their literal equivalents when the filler items of the experiment include so-called *decoys*: stories with a negative emotional valence (such as the ironic stories in the critical items) that are resolved with a banal final statement instead of an ironic sentence. However, when the same experiment was conducted without the decoys (experiment 2), they found that, towards the end of the experiment, participants were increasingly faster at reading the target ironic utterances, to the point where the reading times of ironic and literal sentences in the second half of the experiment were indistinguishable.

These results are important in as much as they show, for the first time, that manipulating the degree to which participants can use their metarepresentational abilities throughout an experiment will have a variable effect on reading times of ironic utterances. There are, however, some issues to be considered: First, though in experiment 1 there is a significant difference between literal and ironic sentences in the second half of the experiment (while this difference is not significant in experiment 2), both experiments show a similar early-late effect, as can be seen by the reported significant interaction between type of utterance (literal vs. ironic) and part of the experiment (first vs. second half) for both experiments. This means that it is possible that, had

experiment 1 been longer, reading times of irony and literals would have eventually been comparably short.

However, the more important issue with these results is the relationship between anticipation of intentions and reading times: In experiment 2, participants learnt to anticipate irony given the lack of decoys. Echoic mention, however, does not claim that listeners anticipate ‘irony’ as a trope, but instead anticipate the communicative intention behind the use of irony: In (1), Maria will be able to process Juana’s ironic statement faster if she can anticipate that Juana is going to express her disappointment (i.e. her communicative intent), not if she can anticipate that Juana is going to say something ironic, regardless of what it might be.

This is why we decided to improve on these issues in order to test the variable effect of mindreading on irony comprehension by (1) creating a design that does not rely on an early-late effect and (2) manipulating participants’ expectations of a speaker’s communicative intent, not their expectations of irony. We operationalized this by manipulating contextually-derived expectations of answers to yes/no questions. Crucially, we also intended to test the effect of explicit cues of the speaker’s attitude towards a proposition on interpretation and speed of processing.

In experiment 1, we set out to test whether participants could indeed generate expectations to a set of yes/no questions given context. Experiment 2 measured the reading times and interpretation of the answers to said yes/no questions.

Experiment 1: context norming

The first goal in our investigation was to create stories that could successfully elicit a specific expectation regarding the upcoming intentions of a speaker. We reasoned that this could be achieved by manipulating the contextual information leading up to a yes/no question: By relying on world knowledge, we created stories in which there is reason to develop either a strong expectation of a ‘NO’ answer or no particular expectation whatsoever. We used yes/no questions in order to limit the range of possible expectations regarding the possible answer to the question.

Experiment 1 was therefore conducted to ensure that participants were indeed able to generate expectations regarding a specific set of situations leading up to a yes/no question.

Participants

Using the Amazon Mechanical Turk platform, 90 participants were recruited to take part in the experiment. Participants were all native speakers of American English and received monetary compensation for their participation in the study.

Materials

We created ten different stories, each being 5-sentence long, representing a conversational exchange. Every story first provided background information on two interlocutors A and B, followed by a yes/no question asked to B by A. Participants had to rate on a scale from 1 to 10 how they

thought B would answer the question, with 1 being “definitely no”, 5 being “either yes or no” and 10 being “definitely yes”.

There were two versions to every story: One that intended to generate a strong bias towards expecting a ‘NO’ response (negative expectation condition, coded as -1), and one intended not to generate any expectation whatsoever (zero expectation condition, coded as 0). Each version was matched for total length (+ - 5 words between versions). The stories were distributed randomly in two lists and participants were assigned to either one of the lists, so that one participant only saw one of the two versions of each story.

Procedure

The experiment was web-based and conducted entirely on-line. Participants read the stories for as long as they wanted and selected their answers using the Amazon Mechanical Turk survey interface.

Predictions, Analysis and Results

We predicted that the stories in the negative expectation condition would be rated significantly lower than those in the zero-expectation condition. Furthermore, we expected all stories in the negative expectation condition to be rated on average below 3. To test this, A linear mixed-effects regression model was fitted to the data to measure the extent to which the contextual biases (negative vs zero expectation) could account for participants’ ratings (1-10). The model included random intercepts and slopes by participants and random intercepts by items. The model showed a significant difference between ratings in the negative and zero expectation conditions ($p < 0.0001$, $t = 19$). Results are summarized in the graph below.




Figure 1: Results of Experiment 1

Experiment 2: Self-paced reading

Experiment 2 meant to test whether the contextual biases established in experiment 1 could affect the processing of verbal irony. To do this, we added three additional lines to the stories in experiment 1: One that presented evidence as to the attitude of the speaker towards their upcoming utterance, one that presented an answer to the question asked by the interlocutor, and a final wrap-up sentence. The goal was to

investigate how explicit information about a speaker's attitude towards a proposition interacted with participants' contextually-derived expectations regarding the upcoming answer to the yes/no question. The conditions are summarized in table 1.

Table 1: Example of critical items in Experiments 1 and 2

negative bias condition (exp 1 & 2)	<i>Chris wants to buy his five-year-old daughter her first guitar. They go to a professional music shop together and she heads for the oldest and most valuable guitar in the store, which was behind a protective glass case. As she comes closer, one can see that the guitar is twice her size. The owner of the store, who really hates children, sees this and anxiously walks towards them. Chris sees him and says: "Sorry, could my daughter play this guitar?"</i>
Zero bias condition (exp 1 & 2)	<i>Chris wants to buy his 15-year-old daughter a new guitar, so they go to a music shop together. She is overwhelmed by all the different types of guitars they have, so she doesn't know which one to pick. They browse around for a while, and finally she finds one that she really likes, even though Chris doesn't understand why. He starts looking for the owner to ask him about it. Chris sees him and says: "Sorry, could my daughter play this guitar?"</i>
insincere attitude conditions (exp 2)	<i>The owner has a reputation for being a joker. He therefore replies:</i>
sincere attitude conditions (exp 2)	<i>The owner has a reputation for being frank. He therefore replies:</i>
target sentence (exp 2)	<i>„Yes, this guitar is here for everyone to play with. “</i>
spill-over region (exp 2)	<i>There were many other costumers in the store that day.</i>
Comprehension Question, multiple choice (exp 2)	<i>The owner will: (1) not let her play the guitar (2) let her play the guitar (3) Buy a guitar</i>

Participants

Using the Amazon Mechanical Turk platform, 68 participants were recruited to take part in the experiment. Participants were all native speakers of American English and received monetary compensation for their participation in the study.

Materials, design and procedure

The materials were identical to the ones used in experiment 1 with the exception of one item which was dropped from the experiment because of a technical problem. To each of the remaining 9 stories, 3 sentences were added: One that established whether the speaker had a sincere or an insincere attitude towards their upcoming answer to the question. A second one providing the answer to the question. This sentence was identical across versions of each story and was always a 'yes' response. Finally, an innocuous wrap-up sentence was added at the end and was also identical across versions.

This resulted in a 2X2 experiment with the factors Expectation bias (zero vs. negative expectation bias) and speaker's attitude (sincere vs. insincere attitude).

The study was programmed as a self-paced reading task using the *Ibex* scripting language created by Alex Drummond, and hosted on the *Ibex Farm* website. Participants read each story in a sentence-by-sentence fashion with no time constraints, hitting the space bar in order to advance to the next sentence. Crucially, participants had to answer a comprehension question after every critical trial. This question assessed both if participants were paying attention to the content of the stories and if they understood the answer to the yes/no question ironically or literally. For every participant, a new list was automatically created showing only one out of the 4 possible versions of each story using the built-in latin-square design function of the *Ibex* experimental software.

Participants also read 10 additional filler stories. These stories also involved a dialogue between two people, but were varied in the type of question asked and the answer given. There were comprehension questions in 7 out of the 10 filler trials as well. Filler and critical trials were pseudo-randomized, so that there would be at least one filler trial between every critical trial. Participants were told that the entire experiment would take about 20 minutes to complete.

As experiment 1, experiment 2 was web-based and conducted entirely online.

Predictions, Analysis and Results

We analyzed three aspects of the data: First, we looked at the reading times of the answer to the question (the seventh sentence in each critical story), which constituted our main dependent variable. We then looked at the reading times of the wrap-up sentence (the eighth sentence in each story), which were analyzed to check for spill-over effects. Finally, we checked the response patterns to the post-comprehension question, to assess whether participants understood the responses literally or ironically.

Prior to analysis, participants who answered less than 5 out of the 7 comprehension questions of the filler items correctly were excluded from the analysis, resulting in the exclusion of 18 participants.

Post-comprehension questions In order to establish whether participants constructed a literal or an ironic interpretation of the target sentences we asked them a post-comprehension

multiple-choice question regarding the outcome of the entire situation: the choice of answer would definitely indicate what type of interpretation they constructed, since this would necessarily result in two distinct situations. If the answer they selected could only come about via an ironic interpretation of the target sentence, we coded this as 1. If they selected the literal interpretation, we coded the response as 0.

We predicted that participants' interpretation of the sentences would be modulated by explicit information regarding the attitude of the speaker: According to the echoic mention account, understanding irony means understanding that the speaker is disassociating herself from the content of a proposition. Thus, we predicted that in the insincere conditions (where it is explicitly mentioned that the speaker is distancing herself from the truth of the utterance) participants would create an ironic interpretation of the sentences, whereas they would create a literal interpretation in the sincere conditions, in which speakers commit to the truth of the proposition.

We fitted a mixed-effects logistic regression model to the data including random intercepts by subjects and random intercepts and slopes (only main effects) by items, following the recommendations of Barr et al. (2013). The results showed a main effect of expectation bias ($p < 0.001$, $t = 8.5$) and a main effect of speaker's attitude ($p < 0.01$, $t = 3.2$) and no interaction effect ($t = 0.15$, $p = 0.8$). The results suggest that, overall, critical sentences in the insincere conditions were indeed understood as ironic, whereas sentences in the sincere conditions were perceived as literal. This effect was additionally modulated by expectation bias: When there was a strong negative expectation, participants perceived utterances as more ironic than when there was no bias whatsoever. Results are summarized in figure 2.

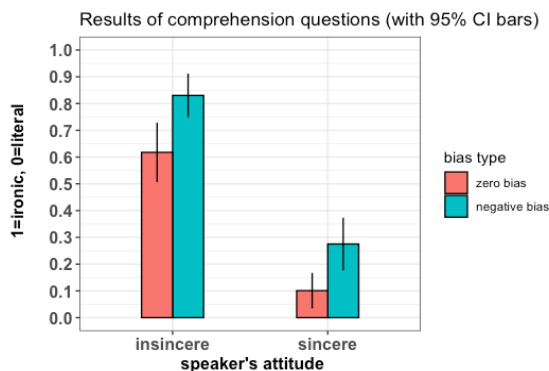


Figure 2: Results of Experiment 2, comprehension questions

Reading time data The reading-time data of the remaining 48 participants was log-transformed given the non-normality of the residuals of the model, following the results of a box-cox test. In order to compare reading times while keeping the interpretation of the conditions constant, we excluded all trials in which participants derived an ironic interpretation in the 'sincere' conditions and those in which they derived a literal interpretation in the 'insincere' conditions. This

resulted in an exclusion of 13% of the total data. This exclusion did not change the results pattern of the statistical analyses.

Our predictions were derived from the discussion presented earlier concerning the role of intentionality recognition during irony comprehension and its relation to reading times: If the strength of expectations about a speaker's upcoming intention determines ease of irony comprehension, we should see that participants take less time reading ironic utterances (i.e. utterances in the insincere conditions) when there is a negative expectation bias than when there is no bias whatsoever. Alternatively, when the utterances are understood literally (the sincere conditions), we should find the opposite pattern: Participants should be somewhat confused to read a 'Yes' response when they strongly expected a 'No', whereas they should have no problem reading the 'yes' response when they didn't expect any particular answer. This should lead to an interaction effect of our two factors. Additionally, if literal utterances are generally easier to understand than ironic utterances, we should find a main effect of speaker's attitude, with sincere conditions being overall shorter than the insincere conditions. We expected results to appear in the critical region, with possible spill-over effects present in the subsequent sentence. We fitted 2 linear mixed-effects regression models to the reading data to test our hypotheses, one for the critical region and one for the spill-over region. The models were fitted following the recommendations of Barr et al. (2013): We started off adding random intercepts and slopes for our main effects and their interaction by subjects and by items. We then subsequently reduced the random effects structure until we settled on the maximally converging model.

For the critical region, the model included random intercepts by subjects and random intercepts and slopes (both main effects and their interaction) by items. This model showed no main effects ($p = 0.6$, $t = 0.5$ & $p = 0.4$, $t = 0.8$) but a significant interaction of the two factors ($p < 0.001$, $t = 6.43$). These results are shown in figure 3 below. We followed-up on these results by re-fitting the model using a slider contrast coding scheme, which allowed us to directly compare each condition to one another. This model showed that the reading times of the insincere-negative bias condition were significantly shorter than in the insincere-zero bias condition ($t = 4.225$, $p < 0.001$) but not significantly different from the reading times in the sincere-zero bias condition ($t = 1.44$, $p = 0.15$). The sincere-zero bias condition had significantly shorter RTs than the insincere-zero bias condition ($t = 2.63$, $p < 0.01$) and the insincere-negative bias condition showed shorter RTs than the sincere-negative bias condition ($t = 5.029$, $p < 0.001$).

The model fitted to the spill-over data included random intercepts and slopes by items and by subjects. This model showed no significant effects.




Figure 3: boxplot of results of Experiment 2, target sentences

Conclusion

With the current work, we set out to investigate whether differences in the literature concerning reading times of ironic utterances could be accounted for by considering the variable way in which participants anticipate a speaker's communicative intentions. We built on the work of Spotorno & Noveck (2014), who had found preliminary evidence for a link between reading times of irony and participants' metarepresentational abilities via an early-late effect. We created a paradigm that allowed us to manipulate participants' expectations about a speaker's intentions on a trial-by-trial basis by using contexts that biased participants towards expecting a 'NO' answer (negative bias condition) to a yes/no question or to not expect an answer whatsoever (zero bias condition). In experiment 2, we measured participants' reading times to the answer of said questions as well as comprehension accuracy. We included a prior sentence that explicitly stated the speaker's attitude towards the upcoming answer (sincere vs. insincere conditions).

Our results showed that, overall, both cues about a speaker's attitude and expectational biases modulated comprehension of our target utterances: Utterances in the insincere conditions were understood mostly as ironic, with the insincere-negative bias triggering the most ironic interpretations. Utterances in the sincere conditions were mostly understood as literal, particularly those in the zero-bias condition. This is in line with the echoic mention account, in as much as it posits that understanding irony crucially involves both a recognition of intentions and of a speaker's dissociative attitude: When there is a strong expectation which is violated together with an explicit cue about a speaker's dissociative attitude, participants understood the sentences as ironic.

The results of the reading times provide support for the idea that biasing expectations about a speaker's intentions can modulate reading times of ironic utterances: When participants understood a sentence as ironic they read it faster when they had generated strong expectations about the answer to a yes/no question than when they had no such expectations.

This finding in particular is crucial because not only does it provide support to the idea that ToM abilities are involved during irony comprehension (in line with the preliminary results of Spotorno and Noveck, 2014), but it also shows that engaging in mindreading abilities occurs on a trial-by-trial basis and not only throughout an entire experimental session. This suggests that listeners can rapidly re-calibrate their expectations and integrate an ironic response with context differently as a function of how certain they are of a speaker's upcoming communicative intention.

The results can also help account for differences in the literature on irony processing: We found that if participants had strong expectations and understood an utterance as ironic (insincere-negative bias condition) they were faster at reading it than when they understood it literally and also had strong negative expectations (sincere-negative bias condition). However, they were faster at understanding the utterance as literal when they had no expectations (sincere-zero bias) compared to when they had no expectations and understood it as ironic (insincere-negative bias). Finally, they were similarly fast at reading an ironic utterance when they had strong expectations (insincere-negative bias condition) than when they read a literal utterance without any expectations (sincere-zero bias conditions). This results in a case in which ironic takes less time than literal, one in which literal takes less time than ironic, and one in which they take roughly the same time: The crucial mediating factors being the variable way in which they engage in their metarepresentational abilities, in line with the echoic mention account.

Overall, the results of this work are a contribution towards bridging the gap between experiments and theory in the literature on irony comprehension by explicating the way in which listeners engage in ToM abilities differentially depending on the type of contextual information that shapes their understanding of a speaker's intention prior to encountering the actual ironic utterance. Models of irony comprehension should therefore consider the variable effect that ToM biases can have on processing when discussing differences between literal and ironic language understanding.

References

- Baron-Cohen, S., Leslie, A. M., & Frith, U. (1985). Does the autistic child have a "theory of mind"? *Cognition*, 21(1), 37–46.
- Fein, O., Yeari, M., & Giora, R. (2015). On the priority of salience-based interpretations: The case of sarcastic irony. *Intercultural Pragmatics*, 12(1), 1–32.
- Filik, R., & Moxey, L. M. (2010). The on-line processing of written irony. *Cognition*, 116(3), 421–436.
- Gibbs, R. W. (1986). On the psycholinguistics of sarcasm. *Journal of Experimental Psychology: General*, 115(1), 3–15.
- Giora, R. (2003). *On Our Mind: Salience, Context, and Figurative Language*. Oxford University Press.

- Giora, R., Fein, O., Laadan, D., Wolfson, J., Zeituny, M., Kidron, R., Kaufman, R., & Shaham, R. (2007). Expecting Irony: Context Versus Salience-Based Effects. *Metaphor and Symbol, 22*(2), 119–146.
- Grice, H. P. (1975). Logic and Conversation. *Speech Acts, 41–58*.
- Happé, F. G. E. (1993). Communicative competence and theory of mind in autism: A test of relevance theory. *Cognition, 48*(2), 101–119.
- Ivanko, S. L., & Pexman, P. M. (2003). Context Incongruity and Irony Processing. *Discourse Processes, 35*(3), 241–279.
- Katz, A. N., Blasko, D. G., & Kazmerski, V. A. (2004). Saying What You Don't Mean: Social Influences on Sarcastic Language Processing. *Current Directions in Psychological Science, 13*(5), 186–189.
- Katz, A. N., & Pexman, P. M. (1997). Interpreting Figurative Statements: Speaker Occupation Can Change Metaphor to Irony. *Metaphor and Symbol, 12*(1), 19–41.
- McDonald, S. (2000). Neuropsychological Studies of Sarcasm. *Metaphor and Symbol, 15*(1–2), 85–98.
- Pexman, P. M. (2008). It's Fascinating Research: The Cognition of Verbal Irony. *Current Directions in Psychological Science, 17*(4), 286–290.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Spotorno, N., Koun, E., Prado, J., Van Der Henst, J.-B., & Noveck, I. A. (2012). Neural evidence that utterance-processing entails mentalizing: The case of irony. *NeuroImage, 63*(1), 25–39.
- Spotorno, N., & Noveck, I. A. (2014). When is irony effortful? *Journal of Experimental Psychology: General, 143*(4), 1649–1665.
- Wilson, D., & Sperber, D. (1992). On verbal irony. *Lingua, 87*(1), 53–76.
- Wilson, D., & Sperber, D. (2012). *Meaning and Relevance*. Cambridge University Press.