Language and Cognition (2026), 18, e64, 1–22
doi:10.1017/langcog.2026.10116

ARTICLE

When irony is (mis)understood: Vigilance cues affect irony comprehension in written discourse

Ana Milosavljevic[1], Nausicaa Pouscoulous[2] and Diana Mazzarella[1]

[1]Cognitive Science Centre, University of Neuchâtel, Switzerland and [2]Linguistics, University College London, UK
Corresponding author: Diana Mazzarella; Email: diana.mazzarella@unine.ch

(Received 08 October 2025 UTC; Revised 26 August 2026 UTC; Accepted 28 August 2026 UTC)

Abstract
Irony comprehension requires going beyond literal meaning to infer a speaker's dissociative stance. Pragmatic accounts propose that this entails evaluating a speaker's vigilance towards the falsity or irrelevance of the literal proposition, a cognitively demanding process even for adults. Yet, previous research often presupposes successful irony comprehension, focusing on processing measures rather than directly testing interpretive accuracy. This study adopts a theory-driven approach by manipulating speaker cues as pragmatic markers of epistemic vigilance and assessing their impact on irony processing and comprehension. Participants read vignettes with ironical or literal utterances produced by speakers described as epistemically vigilant or non-vigilant. Online processing was assessed through reading times, while irony understanding was measured with polar questions probing interpretations of the speaker's belief and communicative intent. Results showed that participants often misinterpreted ironical statements but were significantly more accurate when speakers were described as vigilant, indicating that speaker-based epistemic cues guide pragmatic interpretation. This effect emerged in comprehension but not in online processing, suggesting that evaluations of speaker intent occur later during more reflective stages of interpretation. These findings underscore the interpretive complexity of irony, challenge assumptions about its automatic understanding in adults, and highlight the socially situated nature of pragmatic inference, with implications for theories of communication and language processing.

Keywords: attitude; epistemic vigilance; irony comprehension; language processing

1. Introduction

Verbal irony is a complex and elusive pragmatic phenomenon in which speakers intentionally convey meanings that diverge from, and often directly contradict, the literal content of their utterances. This interpretive disjunction places considerable demands on the addressee, who must not only recognise that the speaker's literal statement is false but also infer that it is transparently and intentionally false, thereby

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

signalling a dissociative or critical stance. According to the Relevance Theory framework, this attitude is directed towards a thought echoed by the ironical utterance – one attributed to another source and judged to be false, absurd, or irrelevant in the given context (Wilson & Sperber, 2012). To illustrate this, consider the statement 'What an interesting lecture', said after a dull and monotonous talk. While this could be taken literally as a white lie or polite remark, the ironic reading of this statement qualifies as an 'echoic use' of language, in which a proposition is put forth to show that the speaker is thinking about it and wants to share their stance towards it (Wilson, 2009). Unlike descriptive uses, which aim to report facts, irony often involves rejecting a socially expected evaluation by echoing it with critical distance. What defines irony, therefore, is not the proposition itself, but the speaker's stance towards it. This stance is typically described as encompassing both a propositional attitude (towards the thought made manifest by the statement) and a non-propositional, affective disposition (towards that thought and its presumed source) (Yus, 2016). The centrality of such a dissociative attitude is acknowledged beyond relevance-theoretic accounts of irony. For instance, from a Cognitive Linguistics perspective, Lozano-Palacio and Ruiz de Mendoza Ibáñez (2022) similarly identify the dissociative attitude as a key feature of ironical interpretations. In their cognitive-pragmatic account, this attitude arises inferentially from a clash between an expected and an attested state of affairs and may subsequently be parametrised according to contextual factors, for instance in terms of humour or sarcasm.

Recognising that the speaker is being ironical is not trivial, and a number of self-paced reading and eye-tracking studies have shown that irony processing incurs significant cognitive costs. These are often manifested in longer total reading times (e.g., Au-Yeung et al., 2015; Filik et al., 2018; Filik & Moxey, 2010; Olkoniemi et al., 2019), along with more regressions and rereading, as readers often revisit earlier text to reconcile ironical meaning with context, which is a pattern rarely observed with literal language (Kaakinen et al., 2014; Olkoniemi et al., 2016, 2019). Insights from electrophysiological and neuroimaging studies support the existence of increased processing demands for verbal irony compared to literal language. For instance, ERP research reveals increased N400 activity, reflecting enhanced semantic and emotional processing, and elevated P600 responses, which suggest increased effort in reanalysing and integrating ironical statements with contextual and world knowledge (e.g., Caillies et al., 2019; Filik et al., 2014; Regel et al., 2010; Spotorno et al., 2013). Furthermore, findings from fMRI research indicate that the comprehension of irony involves more widespread and differentiated patterns of brain activation, both in terms of neural regions engaged and the magnitude of activation (see, e.g., Akimoto et al., 2014; Shibata et al., 2010; Spotorno et al., 2012). These results underscore the increased cognitive complexity associated with processing ironic, as opposed to literal, language.

While the study of irony processing is flourishing, the question of whether this increased cognitive effort actually leads to successful comprehension remains insufficiently explored. This highlights two key methodological gaps in adult literature. First, although many processing studies assume near-ceiling levels of irony comprehension, they seldom incorporate independent explicit measures to verify that irony was truly understood (e.g., Fein et al., 2015; Filik et al., 2017; Filik & Moxey, 2010; Spotorno & Noveck, 2014). In the absence of direct assessments, shorter reading or reaction times are often interpreted as indicators of processing ease or familiarity. However, without confirming whether the irony was actually understood, such

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

inferences remain speculative. Crucially, faster responses may not reflect fluent comprehension but instead signal heuristic processing or strategic disengagement in the face of interpretive complexity – potentially obscuring underlying difficulties rather than revealing them (Ferreira & Yang, 2019). Second, the few studies that have directly measured irony comprehension suggest that readers frequently fail to accurately grasp irony, with performance levels well below ceiling even in carefully controlled experimental conditions (see e.g., Kaakinen et al., 2014; Kyriacou & Köder, 2024; Olkoniemi et al., 2016). Furthermore, in their meta-analytic and systematic review of eye-tracking studies that included comprehension measures, Olkoniemi and Kaakinen (2021) found a substantial negative overall effect (Cohen's $d < 0$), indicating greater difficulty in understanding irony. Another important point is that even when irony comprehension measures are employed, they are often theoretically limited or insufficiently sensitive to exclude alternative interpretations or to adequately capture the inferential processes involved. For instance, some studies have employed comprehension questions that probe the understanding of the speaker's belief (e.g., 'Does X think *p*?'), requiring participants to infer whether the speaker genuinely endorses the literal proposition. While such measures tend to yield high accuracy rates, they may lack the diagnostic precision needed to distinguish between irony and deception, both of which involve a discrepancy between the literal statement and the speaker's actual belief (e.g., Au-Yeung et al., 2015; Barich et al., 2025; Olkoniemi et al., 2023). Similarly, normative rating scales used to judge how ironical or natural a statement appears may tap into superficial impressions rather than the deeper mental state reasoning involved in irony (e.g., Cao et al., 2024; Filik et al., 2019; Spotorno & Noveck, 2014). Other studies have assessed the speaker's emotional stance (e.g., 'Was X happy or glad when he/she said *p*?'), thereby tapping into the attitudinal dimension that is central to ironic interpretation, but these measures often yield only moderate accuracy levels (e.g., around 75%) (e.g., Kaakinen et al., 2014; Kyriacou et al., 2025; Kyriacou & Köder, 2024). Taken together, these findings highlight the value of a more comprehensive approach to assessing irony comprehension. In particular, jointly assessing beliefs and intentions can help rule out plausible alternative interpretations in which the ironical speaker is understood as either mistaken or deceptive. This, in turn, would support more theoretically grounded interpretations of processing outcomes.

Nevertheless, a growing body of research suggests that the challenges associated with irony comprehension and processing are not uniform but instead hinge on a constellation of modulating factors. Broadly, the literature has evolved along two largely distinct trajectories. The first adopts an experimental-pragmatic approach, manipulating contextual cues – such as the presence of discourse antecedents, explicit mention of the speaker's expectations, or information about the speaker's traits – to examine their role in facilitating access to the speaker's intended meaning, often circumventing the need for prior engagement with the literal interpretation (for a review, see Olkoniemi & Kaakinen, 2021). For instance, they show that irony may be more efficiently understood when it directly echoes prior discourse (e.g., Filik et al., 2017; Ţurcan & Filik, 2017; but see Spotorno & Noveck, 2014), while it becomes more challenging as these echoes become less explicit, an effect particularly pronounced in older adults (Mazzaggio et al., 2023). Also, given that irony is a deeply social construct, requiring context-sensitive and socially grounded inferencing, information about the speaker's social identity, such as status, communicative style, personality, or professional role, can significantly bias interpretation towards irony and

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

heighten the perceived intent behind it (e.g., Pexman & Olineck, 2002; Regel et al., 2010; for a review, see Mazzarella & Pouscoulous, 2023). For instance, Katz and Pexman (1997) demonstrated that listeners draw on occupational stereotypes to infer a speaker's propensity for ironical expression. Professions such as actors were associated with a higher likelihood of irony use, whereas roles perceived as more formal or sincere, such as university professors or priests, were judged less likely to employ irony. These socially grounded expectations shape comprehension, facilitating or impeding irony detection depending on the congruence between the speaker's perceived identity and the pragmatic plausibility of ironical intent. While these studies elucidate how social information can modulate interpretive outcomes, they tend to conceptualise such cues as external heuristics, seldom addressing the extent to which they invoke deeper inferential processes, such as attributing beliefs, intentions, or stances to the speaker.

The second trajectory takes a more cognitive-individual differences perspective, correlating irony comprehension with broader socio-cognitive capacities such as working memory capacity (WMC), fluid intelligence, emotional sensitivity, and Theory of Mind (ToM) (for a review, see also Kałowski et al., 2023). For instance, individuals with stronger working memory often demonstrate faster first-pass processing and fewer rereading instances (Kaakinen et al., 2014; Olkoniemi et al., 2016, 2019). However, some studies have found that this may matter only in larger discourse (Olkoniemi et al., 2019; Parola & Bosco, 2022) and highlight fluid intelligence, or the ability to reason and adapt flexibly, as a more reliable predictor of irony comprehension and processing (e.g., Kyriacou & Köder, 2024). Research on both typically and atypically developing individuals has shown that individual differences in Theory of Mind (ToM) predict irony comprehension (see, e.g., Bosco et al., 2018; Spotorno et al., 2012; Zhu & Wang, 2020). This likely reflects irony's echoic nature, which requires representing the speaker's attitude about the thoughts echoed by the literal statement. Neuroimaging research further supports this link, showing that processing ironical utterances activates brain regions traditionally associated with ToM, such as the medial prefrontal cortex, temporo-parietal junction, and posterior cingulate cortex (e.g., Shibata et al., 2010; Spotorno et al., 2012). However, these correlational findings, while informative, offer limited insight into the mechanisms by which such cognitive resources are recruited or activated in real-time comprehension, particularly through discourse-level cues.

Therefore, what remains underexplored is how specific contextual cues, especially those that signal the speaker's attitude, may serve as pragmatic prompts that activate or reveal underlying cognitive mechanisms. Bridging this gap may offer valuable insights into the social-cognitive mechanisms engaged during irony processing and help clarify how readers infer speaker intent and attitude through context-sensitive processing. To our knowledge, the only study directly addressing this is Ronderos et al. (2024), who underscore the central role of mentalising, and particularly cues signalling the speaker's attitude, in irony comprehension. By framing speakers as either sincere or insincere through brief trait attributions (e.g., 'a reputation for being a jokester'), the study provided contextual anchors that shaped how participants inferred speaker stance. Ironic interpretations were more likely when speakers were described as insincere, and reading time data showed that irony was processed more efficiently than literal statements when the speaker's attitude aligned with contextually supported expectations. This interaction suggests that irony comprehension is facilitated not simply by contextual coherence, but by access to cues that engage

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

mentalising processes, specifically, reasoning about the speaker's affective and epistemic stance.

The manipulation of speaker cues offers a valuable yet underexplored opportunity to investigate the role of epistemic vigilance mechanisms in irony (for a discussion, see Mazzarella & Pouscoulous, 2023). These mechanisms underpin the assessment of the speaker's reliability as a source of information (its competence and benevolence) and the credibility of the incoming information (Sperber et al., 2010). By drawing on these mechanisms, readers can avoid misconstruing the ironical speaker as either misinformed or intentionally misleading (Matsui, 2019; Mazzarella & Pouscoulous, 2021; Wilson, 2009). According to Mazzarella and Pouscoulous (2021, 2023), to fully grasp irony and attribute a dissociative stance, the addressee must exercise not only vigilance towards others' competence and benevolence, but also what they call 'second-order' vigilance (vigilance towards others' vigilance), that is, alertness to the speaker's ability to recognise implausible information or unreliable sources. This, in turn, allows readers to evaluate whether the speaker's expression of a dissociative attitude is intentional and warranted. It is only by assuming that the speaker has undergone such an assessment that we can conclude that they have detected the falsity or irrelevance of the proposition and the unreliability of its source and used it as a justification to dissociate from it. Therefore, this account predicts that the audience may be more likely to interpret the speaker as expressing a critical, ironical, attitude when this is perceived as a vigilant individual rather than a gullible one. Although previous research has examined the influence of information about the speaker's vigilance on irony comprehension in children and in audiovisual contexts (e.g., Falkum & Köder, 2026; Milosavljevic et al., 2026a, 2026b), to the best of our knowledge, this question has not been directly investigated in adults using a reading-based paradigm. The present study addresses this gap by directly manipulating information providing access to the speaker's vigilance to examine whether the participants exploit this information when interpreting ironical utterances.

1.1. The present study

This study aims to investigate how contextual cues signalling speaker vigilance influence both the comprehension and the processing of verbal irony. We manipulated contextual information in such a way as to experimentally test the hypothesis of Mazzarella and Pouscoulous (2021) that the evaluation of the speaker's vigilance supports the attribution of a dissociative attitude to the speaker, thus making the ironical interpretation available.

To investigate this, we present participants with short narrative stories adapted from Spotorno and Noveck (2014), ending in either a positive (literal-supportive context condition) or a negative outcome (irony-supportive context condition). These stories were modified to include contextual information about the speaker's traits, specifically their general level of vigilance. The speaker was introduced as a vigilant individual, a gullible one, or via a neutral characterisation, allowing us to examine whether these features influenced irony comprehension and processing. The study had a 2×3 within-subjects design with the factors Condition (Literal vs. Irony) and Speaker Vigilance (Vigilant, Gullible, Neutral). All stories incorporated theoretically grounded features while ensuring narrative consistency and naturalness. For instance, the target utterance, identical across conditions (literal vs. ironic), echoed a

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

proposition introduced earlier in the story, aligning with the relevance-theoretic emphasis on the role of echo in irony comprehension (Wilson & Sperber, 2012). Each narrative involved two closely related characters, a factor known to facilitate irony interpretation through shared context (e.g., Pexman & Zvaigzne, 2004; Weng et al., 2023), and of the same gender to avoid any potential gender bias. We measured irony understanding using two Yes/No questions, targeting understanding of speaker belief and speaker intent (e.g., Filippova & Astington, 2008, 2010). This dual-question approach was designed to provide a more discriminating assessment of irony comprehension by helping to exclude plausible alternative interpretations. Specifically, the belief question distinguished irony from a sincere-but-mistaken interpretation, whereas the intent question distinguished it from deception. Together, these measures also allowed us to assess whether the manipulation of speaker's epistemic information, providing cues to the speaker's dissociative attitude, influenced irony understanding beyond mere sensitivity to contextual mismatch. Therefore, the study addresses two methodological gaps identified in the Introduction. First, it investigates the role of speaker cues as pragmatic markers of epistemic vigilance in ironical contexts. Second, it incorporates both comprehension and processing measures within a multidimensional, theory-driven approach that can reveal potential patterns of misinterpretation.

We hypothesised that contextual information about the speaker's vigilance/gullibility should be a cue to the speaker's dissociative attitude and modulate both irony understanding and processing. More specifically, we predicted that in irony-supportive contexts, it would be easier for participants to understand irony in the Vigilant than in the Gullible condition. As for irony processing, we expected participants to be faster in the Vigilant condition than in the Gullible one. Moreover, the Neutral condition was expected to provide additional insight. Specifically, if reading times for ironical utterances were faster in the Vigilant condition than in both the Neutral and Gullible conditions, this would suggest that vigilance facilitates irony processing. By contrast, if reading times were slower in the Gullible condition than in both the Vigilant and Neutral conditions, this would indicate that gullibility impedes processing. The predictions, along with all the study materials, datasets, and analysis, were pre-registered and can be found on the project's OSF page link: (https://osf.io/ak9dq/).

2. Methods

2.1. Participants and power analysis

To determine sample size requirements, we conducted simulation-based power analyses using the *simr* package (Green & MacLeod, 2016). For reaction times, we modelled conservative variance parameters and a wide range of plausible interaction effects (20–250 ms), simulating designs with 200–400 participants and 12 items each. For comprehension analysis, we used a logistic mixed-effects model with baseline accuracy set at 70% and explored interaction effects corresponding to ±0–20 percentage points. Across 1,000 iterations per condition, both analyses indicated that samples of ~300 participants provide at least 80% power to detect the anticipated interaction effects. Increasing the sample to 350 participants would further enable the detection of more subtle but theoretically meaningful interactions should they exist. Full simulation details and code are available on the project's OSF (Supplementary material, https://osf.io/gkn79/).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

To account for potential exclusions due to technical issues or predefined criteria, we initially recruited 381 participants online via Prolific, slightly exceeding the suggested sample size based on our power analysis. Of these, 19 participants (4.99%) were excluded for failing to meet the a priori exclusion criterion of answering correctly at least 4 out of 6 attention check questions. Additionally, 12 participants (3.15%) were excluded due to anomalously rapid reading patterns (i.e., average overall response times below 300 milliseconds), which likely reflect non-human or inattentive responding and are therefore considered to constitute noise. Following these exclusions, the final sample consisted of 350 participants who were native speakers of American English, aged between 25 and 41 years (mean age = 32.00 years, *SD* = 4.28, female = 177). Participants gave their informed consent and received monetary compensation of £2.88 for their participation. Ethical approval for the study was granted by the Ethics Committee of the University of Neuchâtel.

2.1.1. Materials

The experiment comprised three types of stimuli: familiarisation stories, critical stories and control stories. The familiarisation trials introduced participants to the reading task and response format before the main experiment. The critical items varied based on experimental manipulations, while the control items remained consistent across participants.

Critical items. Each participant was presented with 12 written stories featuring an interaction between two characters who knew each other well. In each story, one of the two characters uttered a target utterance that could be interpreted as ironical or literal as a function of a minor modification made to the prior context (adapted from Spotorno & Noveck, 2014). All the stories were 7 lines long and followed the same pattern (see Table 1 for an example of a critical story). The first two lines introduced the story context, followed by a third line presenting an explicit antecedent for the ironical statement. This antecedent was not repeated verbatim: the key adjective in the target utterance was replaced with a synonym, and the key noun was repositioned from the sentence-initial to the sentence-final position (e.g., 'Tonight, the performance will be fantastic' vs. 'The singers gave a superb performance,' see Table 1). The following two lines were the only ones that would change depending on the condition. More precisely, the fourth line would present the outcome of the situation and could be either positive (in the literal-supportive context) or negative (in the irony-supportive context). The fifth line introduced the manipulation of the speaker's vigilance (see Table 1 for exact descriptions). The sixth line was the target utterance that was the same in all conditions and always positive in valence. It was always 6 words long (between 9 and 10 syllables), and its focus was on the evaluation of a third party or an external event, not directly related to the two main characters. The seventh line was an ordinary wrap-up conclusion of the story that was meaningful in both the Irony- and Literal-supportive context conditions and was always 12 words long (between 15 and 16 syllables). Each line in the critical stories did not exceed 95 characters to ensure it fit on a single line on the screen, thereby reducing additional cognitive effort. After reading the story, the participants were asked two Yes/No test questions assessing the interpretation of the target utterance. The belief question was always presented first (see Belief question, Table 1) to test participants' understanding of the speaker's belief and to exclude the possibility that (in the irony-supportive context) the speaker has been interpreted as sincerely wrong or mistaken. The intent

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

Table 1. Example of a critical story in English in both Irony-supportive and Literal-supportive context conditions and across all three Vigilance levels (vigilant, gullible, and neutral), followed by two comprehension questions

Line 1	Sarah and Emma go together to the opera.		
Line 2	On the night of the premiere, they meet at the main entrance of the theatre.		
Line 3	Before the show starts, Emma says to Sarah: 'Tonight, the performance will be fantastic'.		
	Context		
	Irony-supportive context	*Literal-supportive context*	
Line 4	During the performance, some singers get easily distracted and often sing off-key.	During the performance, all the singers sing beautifully and get standing ovations.	
	Speaker Vigilance		
	Vigilant	*Gullible*	*Neutral*
Line 5	After the show, Sarah, who has a reputation for being a sceptical and critical person, says:	After the show, Sarah, who has a reputation for being a naïve and uncritical person, says:	After the show, Sarah, who has been patiently waiting for the taxi outside the opera, says:
Line 6	'The singers gave a superb performance'.		
Line 7	Line 7 Sarah and Emma meet the next day for lunch at the restaurant.		
	Belief question		
	Does Sarah believe that the singers gave a superb performance?		
	Intent question		
	Does Sarah want Emma to believe that she thinks that the singers gave a superb performance?		

Each participant saw each story in one of the two conditions and with one of the three manipulations of the speaker's vigilance.

question (see Intent question, Table 1) was administered second to check how participants judged the speaker's intent behind the target utterance and to assess the possibility (in the irony-supportive context) that the participants may interpret the target utterance as a lie. Their choice in these two questions cumulatively indicated whether they constructed an ironical or a literal interpretation.

Control items. We created 6 negative control items (decoys), in which a negative outcome was followed by a banal literal target utterance (adapted from Spotorno & Noveck, 2014) and 6 positive control items (fillers) (adapted from Rubio-Fernández et al., 2019). These control stories followed the same structure as the critical ironical stories except that the target utterance did not have an ironical reading and there was no manipulation of speaker vigilance in the fifth line. The length of the target utterance in the control stories matched that of critical items. The control stories were followed by two comprehension questions that matched the structure of the belief and intent comprehension questions used in critical trials, but that required literal interpretation. A control (memory) question was presented after six randomly selected control stories (fillers and/or decoys) and served as a pre-registered exclusion criterion. These questions focused on an element of the story that did not reference the target utterance.

All the experimental stories and materials (in English) used in the experiment are available on the OSF platform (https://osf.io/gkn79/).

2.1.2. Procedure

The experiment was conducted online using the Ibex platform (Drummond, 2013) in combination with PennController (Zehr & Schwarz, 2018). Participants performed a self-paced reading task, advancing through the text by pressing the spacebar.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

They were instructed to keep their index fingers on the F and J keys and their thumb on the spacebar throughout the task. After providing informed consent, participants completed two familiarisation trials featuring non-ironical content, during which they received feedback if they responded incorrectly to the comprehension questions. They were then informed that the main task involved reading short, everyday conversational stories at a natural pace, followed by some comprehension questions. Participants were encouraged to complete the experiment in a single uninterrupted session to ensure consistency in attention and performance. Reading times were recorded for each sentence, with particular focus on the target (Line 6) and wrap-up (Line 7) sentences for the analysis. After each story, participants answered two comprehension questions. The second question appeared only after the first was answered. A control question was administered after randomly selected control stories. For all questions, participants responded using the F (left option) and J (right option) keys. No feedback was provided during the experimental trials. The average duration of the experiment was 24 minutes, followed by a debriefing screen.

Each participant read a total of 24 stories: 12 critical and 12 control stories. For critical items, each participant was assigned a unique list generated via a manually implemented Latin square design, selecting one version per item. Participants saw half the critical items in the Literal-supportive context condition and half in the irony-supportive context condition, with two instances of each vigilance level. We settled on this number to keep the web-based task concise and reduce potential trial effects. Story versions were counterbalanced so that each participant saw only one version of a given story (e.g., ironical in the vigilant condition). The story order was pseudo-randomized to prevent more than two consecutive critical trials. All participants saw the same six decoy and six filler control stories.

2.1.3. Coding

In critical trials, accurate understanding required correctly answering both comprehension questions. The first question ensured that participants could interpret the speaker's belief correctly, and in the irony-supportive context condition, it allowed us to exclude the possibility that the speaker was judged as sincerely mistaken. The second question targeted the speaker's intent and ruled out a deceptive interpretation of the speaker's ironical utterance. The correct response to both questions was 'Yes' in the Literal-supportive condition and 'No' in the Irony-supportive context condition. Therefore, responses were coded as 0 for 'Yes' in Irony and 1 for 'Yes' in Literal for both questions. For comprehension questions following control stories, correct answers were predetermined for each story. Specifically, in three decoy stories, the correct response was 'Yes' to the first question and 'No' to the second question, while in the other three, it was 'No' for both questions. Similarly, for three filler stories, the correct answer was 'Yes' for both questions and in the other three stories, it was 'No' for the first question and 'Yes' for the second. This design prevented a systematic expectation of positive answers for positive filler stories and negative answers for negative decoy stories, thereby reducing predictability.

2.2. Data treatment and analysis

Data processing, analyses, and plotting were conducted using R software (version 4.2.3). The dataset and the R scripts for generalised linear mixed-effects models (GLMMs) and

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

linear mixed-effects models (LMMs), descriptive statistics, and supplementary materials are available on the project's Open Science Framework (OSF) web platform at the following link: https://osf.io/gkn79/. In line with our pre-registration, we fitted theoretically driven models for our main hypotheses, with further simplification of the random part in case they failed to converge. We checked model assumptions using the *DHARMa* package for diagnostic inspection (Hartig, 2018). Analysis of Deviance (with type III Wald Chi-square test) was used to check the significance of each term, and only those with a significant Wald Chi-square test were considered significant. To further investigate and interpret specific contrasts, we used Tukey Corrected least square means post-hoc, using the *emmeans* package (Lenth, 2016). All reading time measures were log-transformed and analysed with linear mixed-effects models to normalise the distribution.

All responses were first screened for completeness and data quality. For the control items, descriptive statistics (mean accuracy, standard deviation, and 95% confidence intervals) were computed to verify that participants correctly understood and followed the task. For the critical items, mean accuracy scores were computed per participant for each Condition (Irony, Literal) and each Vigilance level (Gullible, Neutral, Vigilant). Descriptive statistics were calculated for each group, including mean accuracy, standard deviation, standard error, and 95% confidence intervals. To assess differences in performance between Vigilance levels, pairwise *t*-tests (with Bonferroni correction) were conducted separately for the Irony-supportive context and Literal-supportive context conditions. These descriptive summaries and inferential comparisons are reported alongside the main mixed-effects analyses to provide a comprehensive overview of participant performance.

To investigate whether Vigilance facilitated understanding of ironical compared to literal utterances, we fitted a theoretically driven generalised linear mixed-effects model (GLMM, binomial family, logit link) using the *lme4* package on a dataset involving only critical items (Irony and Literal) (Bates et al., 2015). We started by fitting a model that included Accuracy as the dependent variable, with Vigilance, Condition, and Trial, along with the resulting interaction between Vigilance and Condition, as fixed effects. The random-effects structure included random intercepts for Subjects and Items, and a by-subject random slope for Vigilance as recommended by Barr (2013). This model failed to converge with the default optimiser (*Nelder–Mead*), so it was refit using *bobyqa*, and it converged. However, the model returned a singular fit, indicating that one or more random-effect variance components were estimated near zero. To assess the robustness of this model, we applied the function allFit() to refit the model using multiple optimisers. Although several optimisers converged, all returned singular fits, and some fixed effects estimates varied moderately across optimisers. This suggests that the full model was overparametrised relative to the data and potentially unstable. Given this variance and near-zero variance in random slopes, we simplified the model by removing the by-subject random slopes for Vigilance. Likelihood ratio tests confirmed that excluding the slopes did not significantly reduce model fit, and the resulting model, with random intercepts for subjects and items, converged reliably, showed no singularity, and produced stable fixed-effect estimates consistent with theoretical expectations. We note that while removing random slopes may limit generalisability (Barr, 2013), it was necessary to ensure model interpretability and convergence.

We also performed an explorative analysis of different interpretations in the Irony-supportive context condition. Based on participants' answers to the two comprehension questions (Belief and Intent question), we coded their interpretations

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

Table 2. Behavioural interpretation patterns based on participants' responses to the Belief and Intent comprehension questions

Behavioural pattern	ComQ1 (belief question)	ComQ2 (intent question)
A. Irony	No	No
B. Sincere Literal	Yes	Yes
C. Lie Literal	No	Yes
D. Other	Yes	No

as 'Irony', Sincere Literal' (corresponding to a literal interpretation of the target utterance, taken as sincere) and 'Lie Literal' (corresponding to a literal interpretation of the target utterance, taken as intentionally deceptive), see Table 2.

To explore whether the distribution of interpretive response types varied as a function of Vigilance, we fit a multinomial logistic regression using the multinom() function from the nnet package in R (Venables & Ripley, 2002). The nominal outcome was response type (Behaviour), while Vigilance (Gullible, Neutral, Vigilant) and Trial were predictors. The Sincere Literal response served as the reference category, and the Neutral condition was used as the baseline for comparisons across Vigilance levels. Coefficients are reported as log relative risks (and exponentiated as relative risk ratios) with Wald 95% CIs and two-sided *p*-values.

For reading times, in accordance with our preregistration, we excluded log-transformed values that exceeded ±3 standard deviations from the mean. This was done to remove outliers and reduce noise from implausible responses prior to the main analyses. This procedure resulted in the removal of 1.31% of data for the Target utterance analysis and 0.95% for the Wrap-up sentence analysis. As raw RTs were positively skewed and heteroscedastic, all statistical models were fit using log-transformed RTs, which yielded a more normal distribution. Estimated marginal means were back-transformed for interpretability.

As specified in the study preregistration, accurate irony understanding was a prerequisite for inclusion in the reading time analyses. Therefore, to examine whether Speaker Vigilance modulated the processing of ironical utterances, we focused on reading times for the target utterance and the wrap-up sentence within the Irony-supportive context condition, analysing only the subset of trials where participants demonstrated accurate comprehension, thereby ensuring that reading time analyses reflected processing of genuinely understood ironical content.

To examine if Speaker Vigilance influenced reading times for the target and the wrap-up sentence, we fitted separate linear mixed-effects regression models that included Vigilance and Trial and their resulting interaction in the fixed part, and random intercepts for subject and item and a by-subject random slope for Vigilance, but they failed to converge. We then refitted the models using an alternative optimiser, 'bobyqa,' but convergence issues persisted. Consequently, we simplified the random-effects structure by removing the random slope for Vigilance and retaining only random intercepts for subjects and items. The simplified models converged successfully and were used for subsequent analyses.

As an exploratory analysis, we also wanted to examine whether reading times for the target and the wrap-up sentence differed across Condition and Vigilance. We fitted two linear mixed-effects models predicting log-transformed RTs on accurate trials from both conditions. For both analyses, the fixed part of the model included

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

Condition, Vigilance, Trial, and their interaction. The random-effects structure was initially specified with random intercepts for Subjects and Items, and a by-subject random slope for Vigilance. However, both models showed convergence issues (singular fits for the target utterance; non-convergence for the wrap-up sentence), even after testing different optimisers. To ensure stability and interpretability, we simplified the random part by removing the slopes, retaining only random intercepts for Subjects and Items. The simplified models converged successfully and were used for inference.

3. Results

3.1. Irony comprehension

As a preliminary step, we checked whether Trial, Age, and Gender were associated with Accuracy. Separate logistic mixed-effects models were fit for each predictor. Trial significantly improved model fit, $\chi^2(1) = 4.12$, $p = 0.043$, with accuracy increasing from Trial 1 to Trial 6. Age ($p = 0.126$) and Gender ($p = 0.159$) were not significant predictors. Based on these results, only Trial was included as a covariate in the main analyses (see the R script for details, https://osf.io/gkn79/).

Descriptive analyses of the control items were conducted to assess baseline performance. As expected, participants demonstrated high accuracy on these items. For the decoy items, mean accuracy was 90% ($SD = 13.6$) for ComQ1and 84% ($SD = 20.0$) for ComQ2(overall 76%). For the filler items, mean accuracy was 89% ($SD = 14.9$) for ComQ1 and 94% ($SD = 10.7$) for ComQ2 (overall 84%). One-sample t-tests confirmed that performance was significantly above chance for both decoys, $t(2099) = 28.50$, $p < 0.001$, and fillers, $t(2099) = 42.90$, $p < 0.001$. These results indicate that participants understood the task and were engaged throughout.

The Analysis of Deviance for the accuracy model revealed significant main effects of Condition (χ^2 (1) =1114.24, $p < 0.001$), and Vigilance (χ^2 (2) =6.42, $p = 0.040$) and Trial, $\chi^2(1) = 9.47$, $p = 0.002$, as well as a significant interaction between Vigilance and Condition (χ^2 (2) =28.61, $p < 0.001$). Estimated marginal means revealed that Accuracy increased steadily across trials. The predicted probability of a correct response rose from 0.61 ($SE = 0.033$, 95% CI [0.549, 0.677]) on Trial 1 to 0.71 ($SE = 0.029$, 95% CI [0.654, 0.766]) by Trial 6, suggesting a modest learning or adaptation effect as the task progressed. Tukey Corrected least square means showed that Accuracy was significantly lower in the Irony-supportive context condition compared to the Literal-supportive context condition ($B = -4.68$; $z = -33.38$; $p < 0.001$). Post hoc pairwise comparisons conducted to examine differences in Accuracy between levels of Vigilance showed that none of the comparisons reached statistical significance after Tukey adjustment. However, pairwise comparisons revealed several significant differences in accuracy across levels of Condition and Vigilance. In the Irony-supportive context condition, Accuracy differed significantly across Vigilance groups. Specifically, participants in the Gullible condition were significantly less accurate than those in the Vigilant condition ($B = -0.87$; $z = -5.81$; $p < 0.001$), and those in the Neutral condition were less accurate than those in the Vigilant condition ($B = -0.52$; $z = -3.70$; $p = 0.003$). In contrast, within the Literal condition, no significant difference in accuracy was found between Vigilance groups (see Table 3 for all the comparisons).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

Table 3. Pairwise comparisons of estimated marginal means for Vigilance × Condition combinations

Contrast	*Estimate*	*SE*	*z*	*p*
Irony Gullible – Literal Gullible	−5.100	0.216	−23.64	<0.001
Irony Gullible – Irony Neutral	−0.350	0.155	−2.25	0.215
Irony Gullible – Literal Neutral	−5.384	0.229	−23.51	<0.001
Irony Gullible – Irony Vigilant	−0.870	0.150	−5.81	<0.001
Irony Gullible – Literal Vigilant	−4.765	0.201	−23.71	<0.001
Literal Gullible – Irony Neutral	4.751	0.206	23.04	<0.001
Literal Gullible – Literal Neutral	−0.283	0.231	−1.22	0.825
Literal Gullible – Irony Vigilant	4.231	0.197	21.43	<0.001
Literal Gullible – Literal Vigilant	0.335	0.207	1.62	0.588
Irony Neutral – Literal Neutral	−5.034	0.223	−22.58	<0.001
Irony Neutral – Irony Vigilant	−0.520	0.141	−3.70	0.003
Irony Neutral – Literal Vigilant	−4.416	0.194	−22.78	<0.001
Literal Neutral – Irony Vigilant	4.514	0.214	21.09	<0.001
Literal Neutral – Literal Vigilant	0.618	0.220	2.81	0.056
Irony Vigilant – Literal Vigilant	−3.896	0.185	−21.05	<0.001

Note: Pairwise comparisons are based on estimated marginal means from the generalised linear mixed-effects model (GLMM) predicting accuracy. Estimates represent differences in log odds; positive estimates indicate higher accuracy in the first group of each contrast. *p-Values* are adjusted for multiple comparisons using the Tukey method.

Descriptive analysis supports this, showing that, as expected, the participants were very accurate (*M:* 92.7%, *SD* = 14.8) in the Literal-supportive context condition. Surprisingly, though, the rate of irony interpretations in the Irony-supportive context condition was much lower than expected (*M* = 21.2%, *SD* = 24.5). Pairwise *t*-tests with Bonferroni correction were used to compare Accuracy across Vigilance levels within each condition. In the Irony condition, the Vigilant group performed significantly better than both the Gullible group ($p < 0.001$) and the Neutral group ($p = 0.003$), whereas no difference emerged between the Gullible and Neutral groups ($p = 0.25$). In the Literal condition, no significant differences were detected between different Vigilance levels (see Figure 1).

Given the unexpectedly low rate of ironical interpretations (even in the Vigilant condition, 28%), we decided to perform an exploratory analysis to further investigate the rate of different interpretations in the Irony-supportive context condition. A Chi-squared goodness-of-fit test indicated that the observed distribution of behavioural responses differed significantly from uniform expectations, $\chi^2(3) = 1406.00$, $p < 0.001$). Distribution of participants' responses across these four categories reveals that Sincere Literal was the most frequent interpretation (58%), Irony and Lie Literal were equally likely (21% and 19%), while the Other was extremely rare (2%) (see Figure 2 for distribution of interpretive behaviours).

The multinomial logistic regression model showed that compared to the Neutral baseline condition, participants in the Gullible condition were significantly less likely to select Irony responses (OR = 0.68, $p = 0.008$) and Lie Literal responses (OR = 0.67, $p = 0.006$) relative to Sincere Literal. In the Vigilant condition, participants were significantly more likely to choose Irony responses (OR = 1.57, $p < 0.001$) than Sincere Literal. However, their likelihood of selecting Lie Literal responses ($p = 0.569$) or Other responses ($p = 0.394$) did not differ significantly from the Neutral baseline condition. Trial number significantly predicted increased selection of both Irony (OR = 1.11, $p = 0.002$) and Lie Literal responses (OR = 1.08, $p = 0.032$), indicating a gradual shift towards these interpretations over time. No significant effect of Trial was observed for the Other response category ($p = 0.837$; OR = 1.02). To facilitate

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

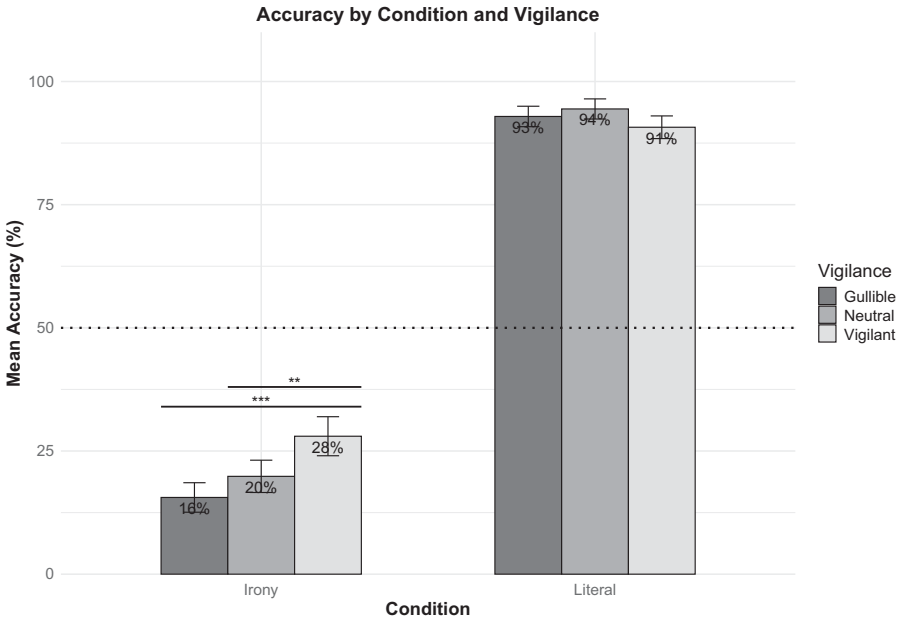


Figure 1. Mean accuracy (%) as a function of Condition (*Irony*, *Literal*) and Vigilance level (*Gullible*, *Neutral*, *Vigilant*).
Note: The dotted line indicates chance-level accuracy. Bars represent model-estimated marginal means with 95% confidence intervals. Asterisks denote results of Bonferroni-corrected pairwise *t*-tests between Vigilance groups within each condition (* is assigned when $p < 0.05$; ** is assigned when $p < 0.01$; *** is assigned when $p < 0.001$).

transparent interpretation via odds ratios and *p*-values, we used multinomial logistic regression without random effects to model the influence of Vigilance on response behaviour. For comparison, we also fitted a hierarchical Bayesian model using the brms package in R (Bürkner, 2017). We started by fitting a full Bayesian multinomial logistic regression model with Vigilance and Trial as fixed effects. The random-effects structure included intercepts for Subjects and Items, and random slopes for Vigilance by Subject. Due to convergence difficulties (unstable posteriors), we removed the by-Subject slope. The simplified model converged well and produced stable estimates, reinforcing our conclusions without overfitting (the R script can be found in Supplementary material, OSF link: https://osf.io/gkn79/.

3.2. Irony processing

The Analysis of Deviance for the linear mixed-effects model on reading times for the target utterance revealed an effect of Trial on reading times of the target utterances as reaction times decreased significantly across trials, $\chi^2(1) = 27.81$, $p < 0.001$, while no effect of Vigilance ($\chi^2(2) = 0.277$, $p = 0.871$) nor of Vigilance and Trial interaction ($\chi^2(2) = 0.072$, $p = 0.964$) was observed. Estimated marginal means (EMMs) for log-transformed RT showed a consistent decrease across trials, from 7.81 ($SE = 0.06$) at Trial 1 to 7.46 ($SE = 0.06$) at Trial 6. This pattern suggests a gradual reduction in reading times over the course of the task.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

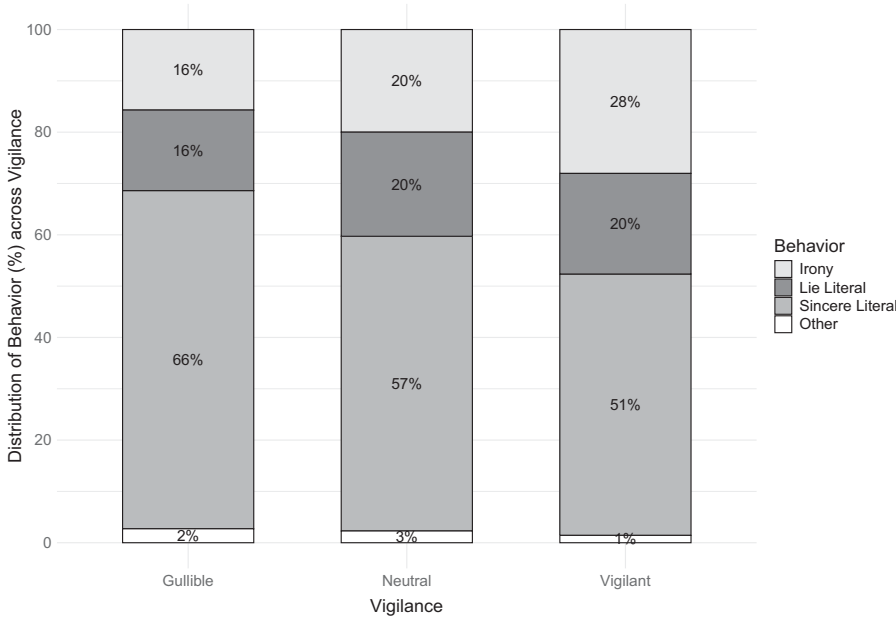


Figure 2. Distribution of interpretive behaviours across Vigilance conditions in irony-supportive contexts. *Note*: Bars represent the proportion of each interpretive behaviour across the three speaker conditions: Vigilant, Neutral, and Gullible. Interpretations were based on responses to two Yes/No comprehension questions assessing belief and intent. Irony responses (No/No) indicate correct interpretation of irony. Sincere Literal (Yes/Yes) reflects a misunderstanding of irony as sincere. Lie Literal (No/Yes) suggests that the speaker's belief was correctly understood, but their intent was misread as a white lie. Other (Yes/No) reflects cases where the speaker's intent was inferred, but their belief was misunderstood.

As for the analysis of reading times of the wrap up sentence, the Analysis of Deviance revealed an effect of Trial, $\chi^2(1) = 6.72$ $p = 0.009$), no significant effect of Vigilance ($\chi^2(2) = 1.75$, $p = 0.417$) nor Vigilance and Trial interaction ($\chi^2(2) = 2.16$, $p = 0.340$) was observed (see Table 4 for estimated marginal means of log-transformed RT6 (target utterance) and RT7 (wrap up sentence) by Vigilance level, back-transformed to milliseconds (ms) with 95% confidence intervals).

Regarding the exploratory analysis assessing the effect of Vigilance on reading times of the target utterance across both Conditions, the analysis of Deviance revealed a significant main effect of Condition, $\chi^2(1) = 14.89$, $p < 0.001$, and a significant effect of Trial, $\chi^2(1) = 90.72$, $p < 0.001$. No main effect of Vigilance was observed, $\chi^2(2) = 0.75$, $p = 0.688$. In addition, none of the two-way interactions (Vigilance × Condition, Vigilance × Trial, or Condition × Trial) nor the three-way interaction (Vigilance × Condition × Trial) reached significance (all ps > 0.09). A pairwise comparison revealed that reading times were significantly longer for Irony compared to Literal items, $t(2056) = 5.64$, $p < 0.001$, with an estimated difference of 0.13 ($SE = 0.02$) on the log-transformed scale. Estimated marginal means (EMMs) indicated a steady decline in log-transformed reading times across trials, suggesting improved processing efficiency over time. Specifically, mean log RTs decreased from 7.74 ($SE = 0.043$, 95% CI [7.65, 7.83]) on Trial 1 to 7.43 ($SE = 0.042$, 95% CI [7.35, 7.52]) on Trial 6.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

Table 4. Estimated marginal means of RT6 (target sentence) and RT7 (wrap-up sentence) across vigilance (log-transformed and back-transformed to ms)

	LogRT6 (target)	RT6 (ms)	LogRT7 (wrap-up)	RT7 (ms)
Vigilance	M ± SE	[95% CI]	M ± SE	[95% CI]
Gullible	7.59 ± 0.06	1985 [1761, 2238]	7.74 ± 0.07	2315 [2004, 2673]
Neutral	7.64 ± 0.06	2084 [1862, 2332]	7.82 ± 0.07	2484 [2168, 2846]
Vigilant	7.65 ± 0.05	2093 [1883, 2326]	7.81 ± 0.06	2479 [2179, 2820]

Note: Values represent estimated marginal means from linear mixed-effects models predicting log-transformed response times for the target sentence (log_RT6) and wrap-up sentence (log_RT7). Back-transformed values (in milliseconds) are shown for interpretability. CIs represent 95% confidence intervals based on fixed effects.

As for the wrap-up sentence analysis, the analysis of Deviance revealed only a significant effect of Trial, $\chi^2(1) = 10.42$, $p = 0.001$, indicating a reduction in reading time across trials. Estimated marginal means (EMMs) for log-transformed RTs showed a consistent decline across trials, with predicted values decreasing from 7.92 ($SE = 0.048$, 95% CI [7.83, 8.02]) on Trial 1 to 7.62 ($SE = 0.046$, 95% CI [7.53, 7.72]) on Trial 6, suggesting a potential learning or adaptation effect over time.

4. Discussion

The goal of the current study was twofold. First, we aimed to investigate whether providing contextual cues that pointed to the speaker's vigilance could facilitate comprehension of ironical relative to literal utterances. Second, we wanted to examine whether the presence of such cues would also affect irony processing, with the expectation that the participants would process ironical utterances more quickly when the speaker was portrayed as vigilant rather than gullible.

Our results provide partial support for these predictions. Although participants' overall performance in the irony comprehension task was lower than expected, our accuracy analysis revealed a significant interaction between speaker vigilance and context: participants were significantly more likely to reach ironical interpretations in irony-supportive contexts when the speaker was framed as vigilant, compared to when the speaker was gullible or when no cues were provided. In contrast, the manipulation of speaker vigilance did not influence reading times of the target or wrap-up sentence. These findings support the hypothesis that vigilance cues modulate the detection of the speaker's dissociative attitude and the recognition of the ironical intent. We now turn to a more detailed examination of these results in the remainder of the section.

While we did observe an overall effect of speaker vigilance cues on irony comprehension, performance rates were unexpectedly low across all conditions. This pattern warrants careful attention and methodological reflection. To begin with, the format of the task – devoid of prosodic, visual, or interactive cues that support irony detection in face-to-face conversations – may have likely contributed to some extent to the difficulty participants experienced in accessing the intended meanings (Bryant, 2010). Furthermore, our rigorous comprehension measure, designed to assess the possibility of alternative interpretations of the ironical statements, may have come with certain trade-offs. While appropriate to probe core aspects of irony understanding, it may have placed substantial metarepresentational and inferential demands on participants. Even if our comprehension measure draws on a format successfully used

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

in developmental research (e.g., Demorest et al., 1984; Filippova & Astington, 2008), its use with an adult population, particularly in embedded discourse, may raise some methodological issues. It could be that this recursive belief-intent questioning may be effective for children, as it scaffolds their developing ability of Theory of Mind by encouraging deliberate reflection on mental states – engaging slow effortful reasoning to aid in understanding the speaker's dissociative stance. However, for adults, irony is typically processed automatically and intuitively through fast, heuristic-based processing, especially when contextual cues make the speakers ironic intent salient. Imposing reasoning in this case may disrupt the natural flow of comprehension, forcing adults into an analytical mode they would not ordinarily use in everyday discourse. As such, while recursive tasks can enhance learning in children, they risk reducing ecological validity when used with adults and may lead to overanalysing. Moreover, explicitly embedding the literal proposition in questions about the speaker's beliefs and intentions may have increased its accessibility as an interpretive option, contributing to the relatively high rate of literal responses observed here.

Importantly, the relatively low rates of explicit irony comprehension cannot be attributed solely to the strict coding criteria, as performance remained low even under a more lenient classification (combining ironical and lie-literal responses). An additional explanation may lie in the broader experimental context created by the epistemic-vigilance manipulation. Although vigilance conditions were balanced across trials, repeated exposure to speakers with different epistemic profiles may have heightened the salience of speakers' epistemic states as a relevant dimension of interpretation. In particular, encountering speakers presented as gullible may have increased the accessibility of a sincere-but-mistaken interpretation, making it more plausible that an utterance-context mismatch reflected a genuine mistake speaker rather than an ironical intention. Such heightened attention to speakers' epistemic states may also have extended to neutral trials, potentially contributing to the predominance of literal responses in that condition and to the lower-than-expected rates of explicit irony comprehension more generally.

Nevertheless, the present results suggest that experiments which assumed irony comprehension without directly verifying it may have overestimated participants' pragmatic understanding. Our findings underscore the need for irony research to move beyond assumptions of successful comprehension and towards a more integrated approach – one that combines both processing and outcome-based measures and reduces concerns about the interpretability of processing data in relation to readers' actual mental representations (Ferreira & Yang, 2019). Doing so will yield a more realistic assessment of adult pragmatic competence and support the development of methodologies capable of capturing irony with greater accuracy and less cognitive strain.

While speaker vigilance did not significantly modulate reading times of the target and wrap-up sentences in irony-supportive contexts, our exploratory analysis revealed a consistent difference in processing effort when comparing literal and ironical utterances: on trials interpreted accurately, participants took longer to read ironical target utterances than literal ones. This pattern is consistent with prior research suggesting that irony comprehension involves additional inferential complexity (e.g., Barich et al., 2025; Spotorno & Noveck, 2014). These findings reinforce the view that irony processing is cognitively effortful and context-dependent, particularly in written formats that lack supportive cues like tone or facial expression (Bryant, 2010).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

Crucially, our findings revealed a significant relationship between speaker vigilance and context, suggesting that epistemic cues meaningfully shape participants' interpretive outcomes and support pragmatic inferencing. Specifically, when cues to speaker vigilance were present – indicating a critical or dissociative disposition – participants were more likely to interpret the utterance as ironical in irony-supportive contexts. In line with Mazzarella and Pouscoulous' s (2021) framework, information about the speaker's vigilance appears to facilitate the recognition of the speaker's ironical attitude, an attitude which would well instantiate this general trait attributed to the speaker. This appraisal may have made ironical interpretations more plausible, particularly in contexts where a critical stance is expected. This pattern also aligns well with developmental findings showing that epistemic vigilance facilitates irony comprehension, and that vigilant children are more likely to understand irony than less vigilant ones (Milosavljevic et al., 2026a). The convergence of findings from the additional multinomial analysis of interpretative patterns in irony-supportive context supports the conclusion that participants in the vigilant speaker condition were particularly attuned to detecting irony compared to those exposed to gullible speaker cues or no speaker cues at all. Additionally, the gradual increase in Irony responses and, to a lesser extent, Lie Literal over time across all vigilance levels indicates a broader shift towards interpretations reflecting growing interpretive engagement or increased awareness of sophisticated underlying motives as the task progressed. Notably, this effect was evident in comprehension rather than processing speed, but this shift in locus is itself theoretically informative. It suggests that speaker-related contextual cues, particularly those tied to epistemic vigilance, may exert their influence more at the level of interpretive resolution than early-stage processing.

Our findings more broadly highlight the socially situated nature of irony, whose understanding depends not only on recognising contextual incongruity, but also on contextual cues about the speaker's mindset and communicative goals (see also Ronderos et al., 2024). The sensitivity to speaker traits and intentions may be particularly essential for irony, distinguishing it from other forms of nonliteral language, such as metaphor, metonymy, or hyperbole, whose nature is descriptive rather than echoic. Irony typically requires a more explicit appraisal of the speaker's affective stance and interpersonal positioning, i.e., not only what the speaker really means but also who the speaker is and why they would choose to express a dissociative attitude. Thus, we should bear in mind that understanding irony is not just a matter of language processing, but requires grasping its socially situated and relational architecture, where the information about the speaker-hearer dynamics, as well as the identity of the speaker and their communicative motives, may play a central role in irony interpretation (Scianna, 2024).

Taken together, these observations underscore the need to refine both experimental design and measurement strategies in future research. Accurately capturing irony comprehension in written paradigms remains a methodological challenge, one that demands careful attention to narrative clarity, trait framing, and the specificity of comprehension assessments to disentangle the subtle inferential processes involved. Crucially, future approaches should deepen our understanding of the socially situated nature of irony. Irony is inherently relational, shaped by shared knowledge, speaker-hearer dynamics, and contextual cues that often extend beyond the text itself. Neglecting these dimensions may obscure how irony is processed and understood in real-world communication.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

The present study focused specifically on the role of epistemic vigilance in verbal irony comprehension. Future research should broaden this scope by examining its contribution to a wider range of pragmatic and communicative phenomena, thereby clarifying which vigilance mechanisms are specific to irony and which support pragmatic interpretation more generally. Humour represents a particularly relevant domain in this respect, given both its frequent intersection with irony and its independent occurrence. Comparing humorous and non-humorous irony with non-ironical humour may help disentangle the role of epistemic vigilance in ironical and humorous interpretation (see Milosavljevic & Mazzarella, forthcoming), while also clarifying whether and how epistemic vigilance contributes to the fine-tuning or contextual parametrisation of the specific attitude expressed by an ironical language use within the broader dissociative range.

5. Conclusion

The present study provides a compelling demonstration that speaker trait information, specifically speaker vigilance, has a significant and interpretable impact on irony comprehension. The fact that participants exposed to vigilant speaker cues were able to more reliably identify ironical intent reinforces the importance of speaker modelling in pragmatic inference. Beyond this, the findings support the role of epistemic vigilance and, critically, the capacity to monitor others' vigilance, that is, to assess whether a speaker is likely to express a dissociative or critical stance. Our findings challenge the assumption that irony comprehension can be taken for granted in experimental settings and underscore the need for more direct, validated measures of irony processing. While our focus was on the influence of speaker vigilance as a pragmatic cue, irony comprehension is inherently multifaceted. It can be influenced by a range of additional factors, including individual differences in Theory of Mind, executive functions, language proficiency, cultural norms, and even mood or interpersonal context (for a review, see Kałowski et al., 2023). Future research should consider how epistemic vigilance interacts with other elements of this complex socio-cognitive repertoire to shape both the interpretation and processing of irony in real-world communication.

Data availability statement. All materials related to this study, including datasets, analysis scripts, and experimental stimuli, are openly available on the Open Science Framework (OSF) platform. These can be accessed at https://osf.io/gkn79/.

Funding statement. This research was supported by a Swiss National Science Foundation Eccellenza Grant (186931) awarded to Diana Mazzarella.

Competing interests. We have no conflicts of interest to disclose.

References

Akimoto, Y., Sugiura, M., Yomogida, Y., Miyauchi, C. M., Miyazawa, S., & Kawashima, R. (2014). Irony comprehension: Social conceptual knowledge and emotional response. *Human Brain Mapping*, 35(4), 1167–1178. https://doi.org/10.1002/hbm.22242

Au-Yeung, S. K., Kaakinen, J. K., Liversedge, S. P., & Benson, V. (2015). Processing of written irony in Autism spectrum Disorder: An eye-movement study. *Autism Research*, 8(6), 749–760. https://doi.org/10.1002/aur.1490

Barich, T., Kyriaki, L., Forndran, A., Williamson, P., & Arciuli, J. (2025). Children's processing of irony during reading in English: An eye-tracking study. *Journal of Experimental Child Psychology*, 254, 106202. https://doi.org/10.1016/j.jecp.2025.106202

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

Barr, D. J. (2013). Random effects structure for testing interactions in linear mixed-effects models. *Frontiers in Psychology*, 4. https://doi.org/10.3389/fpsyg.2013.00328

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. https://doi.org/10.18637/jss.v067.i01

Bosco, F. M., Gabbatore, I., Angeleri, R., Zettin, M., & Parola, A. (2018). Do executive function and theory of mind predict pragmatic abilities following traumatic brain injury? An analysis of sincere, deceitful and ironic communicative acts. *Journal of Communication Disorders*, 75, 102–117. https://doi.org/10.1016/j.jcomdis.2018.05.002

Bryant, G. A. (2010). Prosodic contrasts in ironic speech. *Discourse Processes*, 47(7), 545–566. https://doi.org/10.1080/01638530903531972

Bürkner, P.-C. (2017). brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, 80(1), 1–28. https://doi.org/10.18637/jss.v080.i01

Caillies, S., Gobin, P., Obert, A., Terrien, S., Coutté, A., Iakimova, G., & Besche-Richard, C. (2019). Asymmetry of affect in verbal irony understanding: What about the N400 and P600 components? *Journal of Neurolinguistics*, 51, 268–277. https://doi.org/10.1016/j.jneuroling.2019.04.004

Cao, N., Zhou, L., & Zhang, S. (2024). Uncovering echoic mechanisms in verbal irony comprehension. *Language and Cognition*, 16(4), 2148–2168. https://doi.org/10.1017/langcog.2024.53

Demorest, A., Meyer, C., Phelps, E., Gardner, H., & Winner, E. (1984). Words speak louder than actions: Understanding deliberately false remarks. *Child Development*, 55(4), 1527. https://doi.org/10.2307/1130022

Drummond, A. (2013). *Ibex farm*. http://spellout.net/ibexfarm

Falkum, I. L., & Köder, F. (2026). The development of irony and epistemic vigilance: Content inappropriateness and speaker trustworthiness. *Pragmatics & Cognition*, 33(1), 75–98. https://doi.org/10.1075/pc.24045.fal

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. https://doi.org/10.1515/ip-2015-0001

Ferreira, F., & Yang, Z. (2019). The problem of comprehension in psycholinguistics. *Discourse Processes*, 56(7), 485–495. https://doi.org/10.1080/0163853X.2019.1591885

Filik, R., Brightman, E., Gathercole, C., & Leuthold, H. (2017). The emotional impact of verbal irony: Eye-tracking evidence for a two-stage process. *Journal of Memory and Language*, 93, 193–202. https://doi.org/10.1016/j.jml.2016.09.006

Filik, R., Howman, H., Ralph-Nearman, C., & Giora, R. (2018). The role of defaultness and personality factors in sarcasm interpretation: Evidence from eye-tracking during reading. *Metaphor and Symbol*, 33(3), 148–162. https://doi.org/10.1080/10926488.2018.1481258

Filik, R., Leuthold, H., Wallington, K., & Page, J. (2014). Testing theories of irony processing using eye-tracking and ERPs. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(3), 811–828. https://doi.org/10.1037/a0035658

Filik, R., & Moxey, L. M. (2010). The on-line processing of written irony. *Cognition*, 116(3), 421–436. https://doi.org/10.1016/j.cognition.2010.06.005

Filik, R., Ţurcan, A., Ralph-Nearman, C., & Pitiot, A. (2019). What is the difference between irony and sarcasm? An fMRI study. *Cortex*, 115, 112–122. https://doi.org/10.1016/j.cortex.2019.01.025

Filippova, E., & Astington, J. W. (2008). Further development in social reasoning revealed in discourse irony understanding. *Child Development*, 79(1), 126–138. https://doi.org/10.1111/j.1467-8624.2007.01115.x

Filippova, E., & Astington, J. W. (2010). Children's understanding of social-cognitive and social-communicative aspects of discourse irony. *Child Development*, 81(3), 913–928. https://doi.org/10.1111/j.1467-8624.2010.01442.x

Green, P., & MacLeod, C. J. (2016). SIMR: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution*, 7(4), 493–498. https://doi.org/10.1111/2041-210X.12504

Hartig, F. (2018). *DHARMa: Residual diagnostics for hierarchical (multi-level/mixed) regression models* (R package version 0.2.0). https://cran.r-project.org/package=DHARMa

Kaakinen, J. K., Olkoniemi, H., Kinnari, T., & Hyönä, J. (2014). Processing of written irony: An eye movement study. *Discourse Processes*, 51(4), 287–311. https://doi.org/10.1080/0163853X.2013.870024

Kałowski, P., Zajączkowska, M., Branowska, K., Olechowska, A., Siemieniuk, A., Dryll, E., & Banasik-Jemielniak, N. (2023). Individual differences in verbal irony use: A systematic review of quantitative

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

psycholinguistic studies. *Metaphor and Symbol*, 38(1), 81–111. https://doi.org/10.1080/10926488.2022.2160253

Katz, A. N., & Pexman, P. M. (1997). Interpreting figurative statements: Speaker occupation can change metaphor to irony. *Metaphor and Symbol*, 12(1), 19–41. https://doi.org/10.1207/s15327868ms1201_3

Kyriacou, M., & Köder, F. (2024). The cognitive underpinnings of irony comprehension: Fluid intelligence but not working memory modulates processing. *Applied PsychoLinguistics*, 45(6), 1219–1250. https://doi.org/10.1017/S0142716424000444

Kyriacou, M., Rummelhoff, C., & Köder, F. (2025). Irony processing in adults with ADHD: Evidence from eye-tracking and executive attention tasks. *Journal of Attention Disorders*, 29(9), 724–744. https://doi.org/10.1177/10870547251333819

Lenth, R. V. (2016). Least-squares means: The R package LS means. *Journal of Statistical Software*, 69(1), 1–33. https://doi.org/10.18637/jss.v069.i01

Lozano-Palacio, I., & Ruiz de Mendoza Ibáñez, F. J. (2022). *Modeling irony: A cognitive pragmatic account*. John Benjamins Publishing Company. https://doi.org/10.1075/ftl.12

Matsui, T. (2019). Component processes of irony comprehension in children: Epistemic vigilance, mind-Reading and the search for relevance. In K. Scott, B. Clark, & R. Carston (Eds.), *Relevance, pragmatics and interpretation* (1st ed., pp. 231–239). Cambridge University Press. https://doi.org/10.1017/9781108290593.020

Mazzaggio, G., Zappoli, A., & Mazzarella, D. (2023). Verbal irony and the implicitness of the echo: The processing of young and older adults. *Pragmatics & Cognition*, 30(2), 412–443. https://doi.org/10.1075/pc.00036.maz

Mazzarella, D., & Pouscoulous, N. (2021). Pragmatics and epistemic vigilance: A developmental perspective. *Mind & Language*, 36(3), 355–376. https://doi.org/10.1111/mila.12287

Mazzarella, D., & Pouscoulous, N. (2023). Ironic speakers, vigilant hearers. *Intercultural Pragmatics*, 20(2), 111–132. https://doi.org/10.1515/ip-2023-2001

Milosavljevic, A., Castelain, T., Pouscoulous, N., & Mazzarella, D. (2026a). The developmental puzzle of irony understanding: Is epistemic vigilance the missing piece? *Journal of Child Language*, 53(3), 684–705. https://doi.org/10.1017/S0305000925000091

Milosavljevic, A., Pouscoulous, N., & Mazzarella, D. (2026b). Irony comprehension: Pragmatics meets epistemic vigilance. *First Language*, 46(3), 393–413. https://doi.org/10.1177/01427237251396783

Milosavljevic, M., & Mazzarella, D. (forthcoming). Playing with words, watching for meaning: Epistemic vigilance in humour and irony. In S. Attardo & A. Piata (Eds.), *Humor, cognition, and pragmatics*. John Benjamins Publishing Company.

Olkoniemi, H., Halonen, S., Pexman, P. M., & Häikiö, T. (2023). Children's processing of written irony: An eye-tracking study. *Cognition*, 238, 105508. https://doi.org/10.1016/j.cognition.2023.105508

Olkoniemi, H., Johander, E., & Kaakinen, J. K. (2019). The role of look-backs in the processing of written sarcasm. *Memory & Cognition*, 47(1), 87–105. https://doi.org/10.3758/s13421-018-0852-2

Olkoniemi, H., & Kaakinen, J. K. (2021). Processing of irony in text: A systematic review of eye-tracking studies. *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale*, 75(2), 99–106. https://doi.org/10.1037/cep0000216

Olkoniemi, H., Ranta, H., & Kaakinen, J. K. (2016). Individual differences in the processing of written sarcasm and metaphor: Evidence from eye movements. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(3), 433–450. https://doi.org/10.1037/xlm0000176

Olkoniemi, H., Strömberg, V., & Kaakinen, J. K. (2019). The ability to recognise emotions predicts the time-course of sarcasm processing: Evidence from eye movements. *Quarterly Journal of Experimental Psychology*, 72(5), 1212–1223. https://doi.org/10.1177/1747021818807864

Parola, A., & Bosco, F. M. (2022). An eye-tracking investigation of the cognitive processes involved in the comprehension of simple and complex communicative acts. *Quarterly Journal of Experimental Psychology*, 75(10), 1976–1995. https://doi.org/10.1177/17470218221079629

Pexman, P. M., & Olineck, K. M. (2002). Understanding irony: How do stereotypes cue speaker intent? *Journal of Language and Social Psychology*, 21(3), 245–274. https://doi.org/10.1177/0261927X02021003003

Pexman, P. M., & Zvaigzne, M. T. (2004). Does irony go better with friends? *Metaphor and Symbol*, 19(2), 143–163. https://doi.org/10.1207/s15327868ms1902_3

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116

Regel, S., Coulson, S., & Gunter, T. C. (2010). The communicative style of a speaker can affect language comprehension? ERP evidence from the comprehension of irony. *Brain Research*, 1311, 121–135. https://doi.org/10.1016/j.brainres.2009.10.077

Ronderos, C. R., Tomlinson, J. M., & Noveck, I. (2024). When irony is faster than its literal control: The role of mindreading during irony comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 50(3), 509–522. https://doi.org/10.1037/xlm0001219

Rubio-Fernández, P., Mollica, F., Oraa Ali, M., & Gibson, E. (2019). How do you know that? Automatic belief inferences in passing conversation. *Cognition*, 193, 104011. https://doi.org/10.1016/j.cognition.2019.104011

Scianna, C. (2024). Irony as a complex social phenomenon. In A. Capone, P. Perconti, & R. Graci (Eds.), *Philosophy, cognition and pragmatics* (Vol. 34, pp. 231–249). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-50109-8_13

Shibata, M., Toyomura, A., Itoh, H., & Abe, J. (2010). Neural substrates of irony comprehension: A functional MRI study. *Brain Research*, 1308, 114–123. https://doi.org/10.1016/j.brainres.2009.10.030

Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, 25(4), 359–393. https://doi.org/10.1111/j.1468-0017.2010.01394.x

Spotorno, N., Cheylus, A., Van Der Henst, J.-B., & Noveck, I. A. (2013). What's behind a P600? Integration operations during irony processing. *PLoS One*, 8(6), e66839. https://doi.org/10.1371/journal.pone.0066839

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. https://doi.org/10.1016/j.neuroimage.2012.06.046

Spotorno, N., & Noveck, I. A. (2014). When is irony effortful? *Journal of Experimental Psychology: General*, 143(4), 1649–1665. https://doi.org/10.1037/a0036630

Țurcan, A., & Filik, R. (2017). Investigating sarcasm comprehension using eye-tracking during reading: What are the roles of literality, familiarity, and echoic mention? In A. Athanasiadou & H. L. Colston (Eds.), *Figurative thought and language* (Vol. 1, pp. 255–276). John Benjamins Publishing Company. https://doi.org/10.1075/ftl.1.13tuc

Venables, W. N., & Ripley, B. D. (2002). *Modern applied statistics with S* (4th ed). Springer. https://doi.org/10.1007/978-0-387-21706-2

Weng, X., Jiang, X., & Liao, Q. (2023). Using information of relationship closeness in the comprehension of Chinese ironic criticism: Evidence from behavioral experiments. *Journal of Pragmatics*, 215, 55–69. https://doi.org/10.1016/j.pragma.2023.07.005

Wilson, D. (2009). Irony and metarepresentation. *UCL Working Papers in Linguistics*, 21, 183–226.

Wilson, D., & Sperber, D. (2012). *Meaning and relevance* (1st ed.). Cambridge University Press. https://doi.org/10.1017/CBO9781139028370

Yus, F. (2016). Propositional attitude, affective attitude and irony comprehension. *Pragmatics & Cognition*, 23(1), 92–116. https://doi.org/10.1075/pc.23.1.05yus

Zehr, J., & Schwarz, F. (2018). *PennController for Internet Based Experiments (IBEX).* https://doi.org/10.17605/OSF.IO/MD832

Zhu, N., & Wang, Z. (2020). The paradox of sarcasm: Theory of mind and sarcasm use in adults. *Personality and Individual Differences*, 163, 110035. https://doi.org/10.1016/j.paid.2020.110035

Cite this article: Milosavljevic, A., Pouscoulous, N., & Mazzarella, D. (2026). When irony is (mis) understood: Vigilance cues affect irony comprehension in written discourse, *Language and Cognition*, 18, e64, 1–22. https://doi.org/10.1017/langcog.2026.10116

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 17 Nov 2026 at 07:35:16, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/langcog.2026.10116