

ARTICLE

The Culprit-Hypothesis: how blame attribution influences the interpretation of expressive adjectives

Kalle Glauch 

Germanistische Linguistik, Ruhr-Universität Bochum, Bochum, Germany
Email: kalle.glauch@rub.de

(Received 22 March 2024; Revised 24 August 2024; Accepted 04 September 2024)

Abstract

Frazier and colleagues, in 2015, proposed the speech-act hypothesis as an inferentially rich pragmatic account for the interpretational flexibility of expressive adjectives (EAs) (e.g., *damn*, *frigging*). One pragmatic cue in EA interpretation proposed by Frazier and colleagues is the Culprit-Hypothesis, which predicts that the likelihood of EAs targeting the subject-referent of an utterance increases with the degree of its perceived culpability or blameworthiness in negative events. This article aims to refine the Culprit-Hypothesis by embedding it in a robust theoretical framework based on the psychological models of blame attribution and providing reliable empirical validation. Focusing on the role of intentionality, one of the major components of blame attribution, this article reports a forced-choice study investigating the influence of blameworthiness on EA interpretation. The study followed a 2×3 within-subject repeated measures design, with sentences manipulated by the factors intentionality (intentional versus unintentional versus under-specified) and EA placement (subject-internal versus object-internal) (*The [damn] NOUN1 [intentionally | unintentionally | ϕ] verbs the [damn] NOUN2*). Participants ($n = 100$) read the sentences and selected their preferred interpretation of the EA among the subject-referent, the object-referent and the event-referent. A generalized linear mixed effects model fitted to the data reveals that intentional actions are significantly more likely to result in subject-readings compared to unintentional actions, thus corroborating the Culprit-Hypothesis.




Keywords: blame attribution; blameworthiness; Culprit-Hypothesis; expressive adjectives; intentionality; speech-act hypothesis

1. Introduction

Expressive adjectives (EAs) (e.g., *damn*, *frigging*) express a negative speaker attitude. Unlike descriptive attributive adjectives (e.g., *blue*), the interpretation of EAs






is flexible and not strictly tied to their syntactic position, allowing them to target (be interpreted relative to) discourse referents other than their syntactic sister. Drawing on Frazier et al. (2015), Potts (2004, 2007) propose the ‘speech-act hypothesis’ to account for the interpretational flexibility of EAs. The speech-act hypothesis is a pragmatic account of EA interpretation that treats EAs as quasi-independent non-at-issue speech-acts. As independent speech-acts, EAs can freely target any referent within an utterance, irrespective of their syntactic position.¹ For instance, in sentence (1), the EA *damn* can target the subject-referent, the object-referent (local reading) or the event-referent, despite being realized object-internally.²

- (1) The waiter shattered the damn glass.
: (The waiter) subject-reading
: (The glass) object-reading
: (The waiter shattering the glass) event-reading

Frazier et al. (2015) postulate that addressees derive the speaker-intended EA target via pragmatic inference based on two linguistic cues.

First, they argue that EAs tend to target the local referent denoted by their syntactic sister. Addressees possess the implicit knowledge that EAs can target any discourse referent, regardless of their syntactic position. From this, the addressees infer that the speaker’s choice of a specific placement for the EA is deliberate, aiming to indicate the intended target through syntactic association. For example, compared to sentence (1), which contains an object-internal EA, subject-readings are more likely in sentences like (2) where the EA is realized subject-internally.

- (2) The damn waiter shattered the glass.
: (The waiter) subject-reading
: (The glass) object-reading
: (The waiter shattering the glass) event-reading

Second, they introduce the ‘Culprit-Hypothesis’, which posits that the likelihood of EAs targeting the subject-referent depends on the extent of blame attributed to it. The Culprit-Hypothesis is grounded in the assumption that addressees infer that speakers are inclined to express negative attitudes toward referents perceived as causally responsible, and thus blameworthy or culpable, for undesirable outcomes or events. For instance, in sentence (3), subject-readings are less likely than in sentence (2), as only animate referents can be held causally responsible for negative events.

¹Gutzmann (2019a, 2019b) argues against the pragmatic approach by Frazier et al. (2015) and instead advocates a syntactic approach to EA interpretation based on an upward-looking version of agree (Zeijlstra, 2012). Bross (2021) experimentally tests the predictions derivable from both approaches, concluding that the syntactic approach cannot adequately account for the data as it undergenerate possible readings. A theoretical argument against Gutzmann (2019a, 2019b) syntactic account is its inability to explain EAs targeting linguistic content, such as implicatures or presuppositions, that are not overtly realized in a sentence and, therefore, not potential targets of syntactic movement (Guercio & Orlando, 2022).

²The sad smiley notation to illustrate a negative speaker attitude toward a discourse referent is taken from Gutzmann (2019b).

(3) The damn wind shattered the glass.

☹️: (The waiter)	subject–reading
☹️: (The glass)	object–reading
☹️: (The waiter shattering the glass)	event–reading

While the primary aim of Frazier et al. (2015) was to argue for the speech-act hypothesis to account for the interpretational flexibility of EAs (which they successfully did by demonstrating that EAs have nonlocal readings), the tentatively proposed Culprit-Hypothesis has been taken for granted by subsequent research on EA interpretation (e.g., Bross, 2021; Guercio & Orlando, 2022; Gutzmann, 2019b; Padilla Cruz, 2021, 2022).³ Despite its wide acceptance and the fact that it seems to accurately identify a key component in EA interpretation, there are areas in the theoretical underpinning and methodology of the empirical investigation into the Culprit-Hypothesis, as presented in Frazier et al. (2015), that could benefit from further refinement to develop a more precise theory of EA interpretation.

This article addresses this desideratum, concretely aiming to elaborate on the Culprit-Hypothesis by:

1. Embedding it within the broader context of blame attribution and demonstrating how incorporating components of blame attribution can enhance its explanatory power.
2. Corroborating the theoretical relevance of blameworthiness as a pragmatic cue in EA interpretation by empirically testing its effects through a conceptual replication of Frazier et al. (2015) in sentences manipulated for intentionality.

The remainder of this article is structured as follows. Section 2 discusses how incorporating psychological research on blame attribution into the Culprit-Hypothesis can strengthen its theoretical foundation and increase its predictive power for a wider range of pragmatic cues potentially relevant to EA interpretation. Section 3 explains the rationale for the conceptual replication of Frazier et al. (2015) and describes how the current study addresses potential confounds identified in the original research that prevent the results from being unambiguously ascribed to culpability. Section 4 empirically examines the effect of intentionality, a major component of blame attribution, on EA interpretation in a forced-choice task. Section 5 interprets the findings and Section 6 concludes this article.

2. Blame attribution and the Culprit-Hypothesis

The Culprit-Hypothesis posits that the interpretation of EAs is affected by the perceived degree of blameworthiness of a subject-referent in negative/norm-violating events because addressees pragmatically infer that speakers are more likely to negatively evaluate the causal agents of such events. The authors' intuition that some form of blame attribution influences the preferred interpretation of EAs via pragmatic inference seems highly plausible. The Culprit-Hypothesis highlights a natural connection between EAs expressing negative speaker attitudes, and the negativity

³To clarify, these articles do not further investigate the Culprit-Hypothesis or discuss it in greater detail. However, they refer to it as an established cue in EA interpretation.

associated with blameworthiness for norm-violating events. Additionally, the empirical findings of their study also support their claims.

Frazier et al. (2015) experimentally tested the influence of the Culprit-Hypothesis by presenting participants with short sentences manipulated by the animacy of the subject-referent in the two levels animate (4) versus inanimate (5). The underlying assumption was that only animate referents can be blamed (be considered culpable) for negative events. Participants read the sentences and were prompted to select the preferred EA target among the subject-referent, the object-referent and the event-referent.

(4) The secretary handed in her damn resignation.

- ☹️: (The secretary) subject-reading
 - ☹️: (The resignation) object-reading
 - ☹️: (The secretary handing in her resignation) event-reading
- Frazier et al. (2015, p. 303)

(5) The glasses are covered with damn mineral residue.

- ☹️: (The glasses) subject-reading
 - ☹️: (The mineral residue) object-reading
 - ☹️: (The glasses being covered with mineral residue) event-reading
- Frazier et al. (2015, p. 302)

The study found a significant influence of animacy, with a higher proportion of subject-readings for animate referents compared to inanimate referents.

Based on the results of their study, Frazier et al. (2015) argue for the Culprit-Hypothesis as follows:

Based on our intuitions, examples with ‘causal’ relations seemed easy to interpret in terms of laying blame on the person or entity responsible (henceforth, the ‘culprit’ hypothesis). In other words, the reader may infer that the author has a negative attitude toward the entity responsible for bringing about the unwanted situation. In this case, the negative attitude toward the situation may transfer to the person or entity responsible for the situation – in a sentence with a truly causal structure, the agent (Frazier et al., 2015, p. 294).

As stated before, the intuition that perceived culpability influences EA interpretation seems very plausible given that EAs express a negative speaker attitude and the tendency for norm-violating actions to elicit negative evaluations of the causally responsible agent, which in turn, form the basis of blame attribution (e.g., Guglielmo, 2015; Martin & Cushman, 2016).

However, while the intuitions of Frazier et al. (2015) are based on perceived culpability, components and processes involved in blame attribution identified by prior literature in social and moral psychology are not taken into account. Numerous models addressing blame attribution have been proposed, including more recent ones such as Greene’s (2007) dual process model of permissibility, Cushman’s (2008) model of wrongness and blame and Malle et al.’s (2014) path model of blame. Discussing the Culprit-Hypothesis within the framework of established models of blame attribution provides a more structured and empirically grounded basis for the

observation that referents perceived as culpable for norm-violating events are preferred EA targets because they are perceived more negatively.

Key components of blame attribution identified by previous models and supported by extensive empirical evidence as influencing the perceived degree of blameworthiness of causal agents in negative actions include:

1. Intentionality (whether the action and its outcome are intended or unintended),
2. Foreseeability (the cognitive capacity to prevent an event from happening),
3. Obligation (the social expectation to prevent an event due to social role or relationship),
4. Agency (the capacity of an entity to conduct an action in a self-propelled manner),
5. Reasons (the motivation behind an intentional action).

By integrating these factors (for an extensive overview, see Malle, 2021; Malle et al., 2014) influencing the assessment of a referent's blameworthiness into the Culprit-Hypothesis, more nuanced and precise predictions regarding EA interpretation for a broader spectrum of potential pragmatic cues can be derived.

For instance, the Culprit-Hypothesis predicts a higher proportion of subject-readings stemming from increased blameworthiness for sentence (6) compared to (7) due to the presence of the frequency adverb *again*, which indicates that the waiter had previously slipped and, therefore, should have been able to foresee and prevent it from happening again (foreseeability component of blame attribution); a higher proportion of subject-readings for sentence (8) compared to (9) because it is expected that a teacher, rather than a classmate, is socially obliged to prevent a student from shattering the glass (obligation component of blame attribution); a higher proportion of subject-readings for sentence (10) than for (11), as the intentionally acting waiter is more blameworthy than one who shatters the glass unintentionally (intentionality component of blame attribution); and a higher proportion of subject-readings for (12) than for (13), as purely self-serving reasons aggravate blameworthiness, whereas acting for a friend mitigates it (reason component of blame attribution).

(6) The waiter shattered the damn glass because he slipped on the wet floor again.








(7) The waiter shattered the damn glass because he slipped on the wet floor.

- | | |
|-------------------------------------|-----------------|
| 😞:(The waiter) | subject-reading |
| 😞:(The glass) | object-reading |
| 😞:(The waiter shattering the glass) | event-reading |

(8) The teacher stood by as the bully shattered the damn glass.

(9) The classmate stood by as the bully shattered the damn glass.

- | | |
|--|-----------------------------|
| 😞:(The teacher classmate) | matrix subject-reading |
| 😞:(The teacher classmate standing by...) | matrix event-reading |
| 😞:(The bully) | subordinate subject-reading |
| 😞:(The glass) | object-reading |
| 😞:(The bully shattering a glass) | subordinate event-reading |

- (10) The waiter intentionally shattered the damn glass.
- (11) The waiter unintentionally shattered the damn glass.
: (The waiter) subject–reading
: (The glass) object–reading
: (The waiter shattering a glass) event–reading
- (12) The waiter shattered the damn glass to avenge himself.
- (13) The waiter shattered the damn glass to avenge his friend.
: (The waiter) subject–reading
: (The glass) object–reading
: (The waiter shattering the glass) matrix event–reading
: (The waiter avenging himself | his friend) subordinate event–reading

Some components of blame attribution that can be subsumed under the Culprit-Hypothesis are extremely fine-grained, allowing for predictions across very specific events. For instance, the Culpable Control Model by Alicke (2000) specifies general (e.g., blindness); task-specific (e.g., inability to swim); cognitive (e.g., moral distinctions) and emotional (e.g., impulse control) capacity constraints that influence the perceived personal control of an agent over a negative event, with reduced control mitigating blameworthiness. For all of these components, the Culprit-Hypothesis predicts varying proportions of subject-readings in EA interpretation due to different blame attribution processes in appropriate contexts.

In addition to the broader range of specific predictions derivable from the Culprit-Hypothesis when discussed within the context of blame attribution, another advantage of considering a wider array of factors influencing blameworthiness is the possibility to identify components that can be straightforwardly manipulated linguistically.

While the animacy manipulation employed in Frazier et al. (2015) is a valid operationalization of blameworthiness, targeting the agency component of blame attribution (e.g., Guglielmo & Malle, 2017; Lagnado & Channon, 2008), and the authors conducted a norming study showing that animate referents are perceived as more causally responsible than inanimate referents, it does not convincingly demonstrate on its own that the observed effect is due to changes in blameworthiness resulting from different levels of animacy. This is partly due to a well-established cross-linguistic processing bias favoring animate referents across various linguistic phenomena, such as relative clause attachment (e.g., Desmet et al., 2006; Mak et al., 2002), anaphora resolution and production (e.g., Fukumura & Van Gompel, 2011; Vogels et al., 2014) and differential object marking (e.g., Aissen, 2003; Næss, 2007) due to increased bottom-up salience⁴ (stimulus-driven salience that draws attention to a referent).

As EA interpretation has never been investigated in terms of its interaction with the salience of discourse referents, it cannot be definitively stated that animacy (as a major salience-lending cue) should influence EA interpretation due to the heightened salience of the associated discourse referent. However, prior psychological research indicates that more salient individuals are more likely to be viewed as causally responsible for actions (e.g., McArthur & Ginsberg, 1981; Robinson & McArthur, 1982; Taylor & Fiske, 1975). While these studies manipulate salience via audio-visual cues (e.g., sound intensity, color of shirt), their findings show that causality

⁴For discussions of the term *salience* in linguistics, see Boswijk (2022) and Zarcone et al. (2016).

attribution – an integral process in blame attribution – is susceptible to similar bottom-up salience cues relevant for pronoun resolution, such as prosodic or structural prominence (e.g., Falk, 2014; Järvikivi et al., 2005; Rohde & Kehler, 2014). Therefore, it cannot be ruled out that EA interpretation, especially from the perspective of blame attribution, may be biased toward animacy as a salience-lending cue in a manner similar to pronoun resolution, which goes beyond the intended manipulation of blameworthiness in Frazier et al. (2015).

To convincingly argue that variations in blameworthiness lead to a higher proportion of EAs targeting the subject-referent in negative events, it is essential to demonstrate that manipulating other components of blame attribution proposed in social and moral psychology literature, which are not associated with a potential processing bias, yields similar effects. If manipulating these factors induces differential effects on EA interpretation across different levels, the findings would have greater internal validity and make a more compelling case for blameworthiness as the underlying cognitive cause.

The study presented in Section 4 addresses this desideratum by taking the ‘intentionality’ component of blame attribution and testing its influence on the interpretation of EAs. Intentionality is selected for its central role in essentially all models of blame attribution (e.g., Alicke, 2000; Cushman, 2008; Malle et al., 2014; Monroe & Malle, 2017; Ohtsubo, 2007) and the ease with which it can be manipulated linguistically.

Social perceivers distinguish between intentional and unintentional actions in a binary fashion. Intentionality aggravates blame attribution for norm-violating actions while unintentionality mitigates blame attribution (e.g., Malle et al., 2014; Patil & Cushman, 2017; Young & Saxe, 2009). Given that intentionality is positively correlated with blameworthiness, and blameworthiness, in turn, is positively correlated with negative speaker attitudes, the Culprit-Hypothesis predicts that negative actions perceived as intentional will yield a higher rate of EAs targeting the subject-referent than actions perceived as unintentional.

3. Motivation for Conceptual Replication

As the study presented in Section 4 is essentially a conceptual replication of the study presented in Frazier et al. (2015) with the aim of corroborating the Culprit-Hypothesis, we would like to further motivate the necessity of this replication and show how the design, particularly the manipulation of culpability via intentionality, in the current study avoids certain methodological issues of Frazier et al. (2015).

Apart from the potential processing bias for animacy in EA interpretation already discussed in Section 2, we believe that two additional factors in the study of Frazier et al. (2015) hinder the straightforward association of their findings with the culpability manipulation. Both factors are linked to the animacy manipulation, which strongly restricts the formulation of minimal pairs:

1. The mostly uncontrolled experimental items and
2. The difference in valence⁵ of the subject-referents between conditions.

⁵Valence refers to the emotional dimension reflecting the degree to which a stimulus (here a discourse referent) is perceived as negative or positive. Valence is a crucial element of affective space, integral to nearly all emotional models, regardless of their complexity (e.g., Gillioz et al., 2016; Stanley & Meyer, 2009; Verma & Tiwary, 2017).

Table 1. Sample experimental item

	Subject	Object
Intentional	The damn waiter intentionally shattered the glass.	The waiter intentionally shattered the damn glass.
Unintentional	The damn waiter unintentionally shattered the glass.	The waiter unintentionally shattered the damn glass.
Underspecified	The damn waiter shattered the glass.	The waiter shattered the damn glass.

The contrast between the sentences (4) and (5) shown at the beginning of Section 2, represent the animate and the inanimate conditions of the same experimental item. The pair is representative for most items used in Frazier et al. (2015) and demonstrates that the study diverges from the standard practice of constructing highly controlled sentences as experimental items, where only the relevant manipulation varies across conditions. The sentences used to assess animacy's impact on EA interpretation differ in many potentially relevant aspects. The experimental items feature entirely different lexical items, denote disparate events, and often lack structural similarity. This lack of control introduces the risk of various potential confounders. For instance, differences in how situation models are conceptualized (e.g., Radvansky & Zacks, 2014; Zwaan & Radvansky, 1998; Zwaan et al., 1998) for different event types denoted by action verbs (e.g., *hand in*) or stative verbs (e.g., *covered with*), could potentially influence EA interpretation.

Another possibly relevant aspect accompanying the animacy manipulation is the necessity to use different subject-referents within experimental items. In Frazier et al. (2015), this resulted in differently valenced subject-referents across conditions, with subjects in the animate condition being considerably more positive than those in the inanimate condition across experimental items. Based on the sentiment dictionary of Warriner et al. (2013), containing the valence ratings of roughly 14,000 words on a 1 (negative) to 9 (positive) scale, the animate subject-referents used in Frazier et al. (2015) have a median valence of 6.63, whereas the inanimate subject-referents are notably more negative, with a median valence of 5.18.

For our sample experimental item shown in Table 1, this essentially amounts to using the positive word *friend* in the intentional condition while retaining the neutral word *waiter* for the unintentional condition, as illustrated in sentences (14) and (15).⁶ At least intuitively, the positive word *friend* is the less likely EA target, irrespective of the blameworthiness manipulation.

⁶As briefly mentioned by Frazier et al. (2015), p. 298, a similar point was raised by a reviewer who noted that some negative object-referents in their study (e.g., *mildew*, *stinkbugs*) might have contributed to the relatively high percentage of object-readings in non-causal sentences with subject-internal EAs, as in sentence (a).

(a) *The damn roses are covered in that yellow mildew.* (Frazier et al., 2015, p. 302)

However, taking all object-referents across both conditions, they are similarly valenced compared to the subject-referents that differ more strongly.

Another brief remark regarding example (a) is that demonstrative pronouns (e.g., *that*, *those*), which appear more frequently in the inanimate condition (though not systematically), are known to convey an evaluative component (e.g., Davis & Potts, 2010; Potts & Schwarz, 2010) and may have served as an additional confound by negativizing the object-referent in certain sentences.

- (14) The waiter intentionally shattered the damn glass.
- (15) The friend unintentionally shattered the damn glass.
- | | |
|--|-----------------|
| ☹️: (The waiter friend) | subject–reading |
| ☹️: (The glass) | object–reading |
| ☹️: (The waiter friend shattering a glass) | event–reading |

Considering the difference in subject-referent valence between animate and inanimate sentences in Frazier et al. (2015), the observed preference for subject-readings in animate sentences might be at least partly due to the subject-referents being more negative than those in inanimate sentences.

Taken together, these points highlight the importance of using a within-subject design with tightly controlled experimental items that minimally differ between conditions in empirical pragmatic investigations of speaker attitudes and their targets. While we do not claim that all of the named factors strongly influence EA interpretation, their impact cannot be dismissed outright. In addition to the previously addressed issues linking animacy as a salience-lending cue to a general processing bias, which may also impact EA interpretation, this discussion demonstrates that the Culprit-Hypothesis benefits from further evidence obtained through a better-controlled conceptual replication. This is not because we expect results contrary to Frazier et al. (2015), but to obtain data of higher validity on which to base future research on EA interpretation in this area.

To avoid the confounds outlined in this section, blameworthiness is operationalized via ‘intentionality’ in the replication. The main advantage of intentionality is that it can be manipulated linguistically in a straightforward manner using mental-attitude adverbials (e.g. Maienborn & Schäfer, 2011) such as *intentionally* or *accidentally* which explicitly specify the intentionality component of blame attribution. Using mental-attitude adverbs, the desired differences in perceived blameworthiness can be achieved by varying a single concrete element that semantically conveys the intended meaning component within otherwise unchanged sentences. This approach allows the experiment to manage potential cognitive and emotional cues within an utterance that may influence a speaker’s perceived attitude and its target. Due to their variety, such cues become increasingly difficult to control as the variability between conditions increases.

4. Study: Intentionality and EA interpretation

4.1. Method

The study examines the effect of intentionality on EA interpretation through a three-answer forced-choice task, conducted using the Magpie experimental framework (<https://magpie-experiments.org/>). The study employs a 2 × 3 within-subject repeated measures design, resulting in six conditions derived from fully crossing the two factors ‘intentionality’ (intentional versus unintentional versus underspecified) and ‘EA placement’ (subject-internal versus object-internal). Participants were presented with short sentences following the structure ‘The [EA] subject-noun [intentional-AdvP | unintentional-AdvP | ϕ] verbed the [EA] object-noun’ and prompted to select the most likely target of the speaker’s negative attitude. Inferring the target of a negative speaker attitude within an utterance is the typical

The waiter intentionally shattered the damn glass.

Which is the speaker most likely to have a negative attitude toward?



Figure 1. Example experimental item from participants' perspective.

operationalization of EA interpretation in experimental research (e.g., Bross, 2021; Frazier et al., 2015; Gutzmann, 2019b; Ronderos & Domaneschi, 2023). Participants were instructed to respond quickly based on their initial intuitions. Following Frazier et al. (2015), the response options in the forced-choice task included the subject-referent, the object-referent and the event-referent. Participants provided responses by pressing the keys 1, 2 and 3 to select the corresponding interpretation.

Figure 1 illustrates a sample experimental item from the participants' perspective.

4.2. Items

In total, 24 experimental items were constructed in all six conditions. The items were evenly distributed across six lists using a Latin-Square Design, ensuring that each participant encountered every item only once and that each condition was processed with equal frequency.

In contrast to Frazier et al. (2015), who exclusively used the EA *damn*, the current study varied among the following EAs:

- *blasted, darn, darned, damn, fucking, goddamn*

Intentionality was operationalized by mental-attitude adverbials. For intentional actions, the adverbials used were:

- *intentionally, on purpose, knowingly, purposefully, deliberately, with intent, calculatedly, selectively, methodically.*

For unintentional actions, the adverbials used were:

- *unintentionally, by accident, accidentally, inadvertently, by mistake, unknowingly, absentmindedly, mistakenly.*

The purpose of the variation, both for EAs and for mental-attitude adverbs, was to conceal the experimental manipulation. For underspecified sentences, no adverbials were used.

All experimental sentences denote slightly negative or norm violating events (e.g., *shattering the glass, running a red light*) that are in principle compatible with both intentionality values. The subject-nouns denote social roles or professions (e.g., *designer, roommate*) that according to large-scale sentiment dictionaries (e.g.,

Mohammad, 2018; Warriner et al., 2013) are neutrally to (slightly) positively valenced.⁷

Table 1 shows an experimental item in all conditions.⁸

The experimental items were not normed for intentionality prior to the study. The choice of verbs was informed by their compatibility with the mental-attitude adverbs *intentionally* and *unintentionally* in the sentence frame ‘*The subject-noun verbs the object-noun [intentionally|unintentionally]*’, indicating they do not imply an intentionality value. The selection of mental state adverbs was based on 1) dictionary entries and 2) the strong intuition that using both an intentional and an unintentional adverb from the specified sets results in contradictions (e.g., \perp *The waiter shattered the glass intentionally and unintentionally*), while using two adverbs from the same set leads to infelicity due to redundancy (e.g., $\#$ *The waiter shattered the glass on purpose and intentionally*). We consider it a key advantage of the intentionality operationalization of blameworthiness that, especially for intentional and unintentional actions, it does not require norming and can be implemented in any sentence frame simply by using the appropriate adverbial.⁹

In addition to the experimental items, each list contained 51 filler items and 3 trial items. The filler items followed the same sentence structure as the experimental items, with half featuring a subject-internal adjective and the other half an object-internal adjective (e.g., *The rude guest loudly complained about the dinner menu.*). Two thirds of the filler items contained a manner adverb modifying the action (e.g. *quickly, cautiously*). The trial items were presented at the start of each list to familiarize the participants with the task. Afterward, the experimental and filler items were sequentially displayed in an order uniquely randomized for each participant.

4.3. Hypotheses

As already outlined in Section 2, the basic hypothesis derived from applying the Culprit-Hypothesis to intentionality in EA interpretation is that negative actions perceived as intentional are more likely to result in subject-readings compared to unintentional actions. The underlying premises are that:

1. EAs express a negative speaker attitude toward a discourse referent in the given utterance.

⁷One exception to this is the noun *delinquent* used in experimental item 6.

⁸All other items can be found in the Supplementary Material.

⁹As suggested by Reviewer 2, we conducted a post hoc norming study for the intentionality adverbs *selectively* and *calculatedly*. In the norming study, we asked participants ($n = 30$) if actions denoted by sentences containing these adverbs were performed intentionally or unintentionally. While we agree with the reviewer that—depending on the context—their intentionality component might not be as transparent as in *intentionally* or *on purpose*, the norming data show that for the experimental items in the current study containing these adverbs, the manipulation was successful, with *selectively* inducing an intentional reading in all cases and *calculatedly* with a proportion of 0.96. Furthermore, we normed the experimental items without EAs and without intentionality adverbs (e.g., *The waiter shattered the glass*). The results indicate that the sentences differ considerably regarding perceived intentionality, with a mean proportion of 0.51 ($SD = 0.5$). In general, we think that the use of intentionality norming for the current study is questionable, as it says little about perceived intentionality in sentences with EAs, given the interaction between the factors EA-placement and intentionality found in the study (see Figure 2) and the completely different task demands between the norming and the main study.

2. A higher degree of blameworthiness of causal agents is associated with a higher degree of negativity.
3. Intentionally acting agents are perceived as more blameworthy than unintentional agents.

For the levels ‘intentional’ and ‘unintentional’ in the current study, this straightforwardly predicts a higher proportion of subject-readings for sentences containing a mental-attitude adverbial like *on purpose*, setting the value of intentionality to ‘intentional’, than for sentences containing an adverbial like *by accident*, which sets the value to ‘unintentional’.

Sentences lacking such adverbials are semantically ‘underspecified’ in terms of intentionality. However, given that blame attribution has been shown to operate (also) automatically (e.g., Reeder, 2009; Van Bavel et al., 2012) and its potential relevance to the inference of the speaker-intended EA target, it seems likely that an intentionality value is nonetheless assigned to the subject-referents. Notably, increased blameworthiness in underspecified sentences occurs exclusively in the subset of trials in which the action is interpreted as intentional, leading to the assumption of an intermediate frequency of subject-readings overall.

(H1) Main effect of intentionality

EAs are more likely to target subject-referents that intentionally perform a norm-violating action (negative > neutral > positive).

Regarding EA placement, Frazier et al. (2015) postulate that the position of the EA indicates the local referent as the preferred target of the negative speaker attitude via defeasible syntactic association. This hypothesis was confirmed both in their study and in subsequent research (e.g., Bross, 2021; Gutzmann, 2019b). Therefore, an increased proportion of subject-readings is predicted when the EA is realized subject-internally, and an increased proportion of object-readings is predicted when the EA is realized object-internally.

(H2) Main effect of EA placement

EAs are more likely to target the referent denoted by their syntactic sister.

No interaction effect is predicted between the two factors.

4.4. Participants

Participants ($n = 100$; M age = 43.27, $SD = 15.5$; 52 female, 43 male, 5 nonbinary) were recruited through Prolific. Participants were prescreened for English as a native language, currently residing in the United States, a submission rejection rate of $\leq 1\%$, and prior participation in at least 200 experiments on Prolific. Participants were randomly distributed to one of the six experimental lists. Each list was completed by at least 13 participants. The median completion time was 10 minutes and 12 seconds. Participants were remunerated equal to £10 per hour. Participants were recruited in two sequential phases: an initial group of 80, followed by an additional 20. The extra 20 participants were included to more accurately assess the significance of the observed UNINTENTIONAL \times EA

Table 2. Subject replies by condition

Position	Intentionality	<i>n</i>	Abs. Freq.	Rel. Freq.	SE	CI low	CI high
object	underspecified	376	112	0.30	0.02	0.25	0.35
object	intentional	379	150	0.40	0.03	0.35	0.45
object	unintentional	371	92	0.25	0.02	0.20	0.30
subject	underspecified	361	237	0.66	0.03	0.61	0.71
subject	intentional	368	229	0.62	0.03	0.57	0.67
subject	unintentional	373	182	0.49	0.03	0.44	0.54

PLACEMENT interaction, which minimally varied around the .05 threshold depending on the theoretical approach to modeling maximal random effect structures in data analysis (for discussions, see Barr et al. (2013) and Matuschek et al. (2017)).

4.5. Data exclusion

Observations were excluded if the response time was ≤ 1500 ms. This lower limit was set based on a trial run with three reliable participants, indicating reaction times below this threshold to be implausible. Based on this criterion, 172 observations were removed.¹⁰

4.6. Results

Table 2 displays the absolute and relative frequencies of subject interpretations, along with their 95% confidence intervals for each condition.

Figure 2 presents an interaction plot visualizing the data. The plot depicts the proportion of subject interpretations on the y-axis against EA placement on the x-axis. Different levels of intentionality are distinguished by varying line types. The red points represent the proportion of subject interpretations, and the error bars indicate the standard error for each distinct factor combination.

A generalized linear mixed effects model using a binomial distribution with a logit link function was fitted to the data using the R package lme4 (Bates et al., 2015). Dummy coding was used to obtain binary response values, by treating the object-reading and the event-reading uniformly as the nonsubject response category. The model predicts EA interpretation as a function of the fixed effects EA placement ('position.n') and Intentionality ('intentional.n' and 'unintentional.n'). Given that Figure 2 shows nonparallel distributions, indicating a potential interaction effect, interaction terms for the fixed effects are included in the model.

The model's random effect structure consists of random intercepts for items ('itemname') and random intercepts and slopes for participants ('submission_id'). The random effects account for individual response tendencies for participants and experimental items, independent of the fixed effects. To determine the random effect structure best fitting the data, we followed the backward-selection heuristic suggested by Matuschek et al. (2017), starting with the most complex model and progressively

¹⁰Removing the observations had no impact on the results.

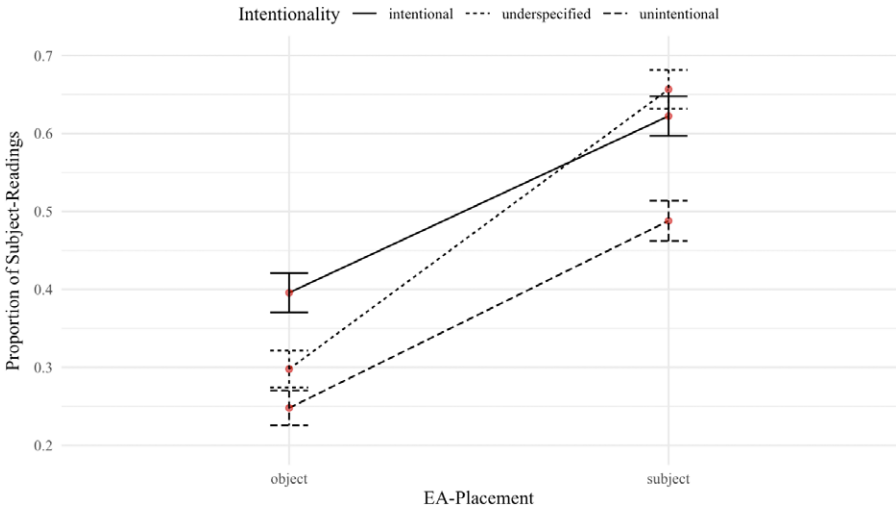


Figure 2. Proportion of subject interpretations by EA placement and intentionality with SE.

reducing complexity until further reduction would result in a significant loss in goodness-of-fit.¹¹

The whole model translates to:

```
glmer(interpretation_category ~ intentional.n + unintentional.n + position.n +
unintentional.n:position.n + intentional.n:position.n +
(1 + intentional.n + unintentional.n + position.n || submission_id) +
(1 | itemname),
data = data,
family = binomial,
control = glmerControl(optimizer = 'bobyqa', optCtrl = list(maxfun = 1e7))).
```

Table 3 presents the model output with underspecified object-internal sentences as the reference category (Intercept).¹²

For the factor ‘intentionality’, the model reveals a main effect of intentional sentences with a significantly higher proportion of subject-readings compared to underspecified sentences (*Estimate* = 0.666, *Std. Error* = 0.212, *z* = 3.140, *p* = 0.002**).

¹¹The model presented here was derived from a model suggested by Reviewer 2, who pointed out some issues in the random effect structure of our initial model.

¹²The results reported in Table 3 and discussed below are based on uncorrected p-values, despite the sequential data collection. However, even when applying very conservative corrections, such as the Bonferroni correction, the significance of the reported effects remains consistent. Detailed p-values adjusted for multiple testing, using various correction methods, are provided in the accompanying R script.

Table 3. Summary of fixed effects analysis

Term	Estimate	Std. error	z Value	Pr(> z)
(Intercept)	-1.336	0.258	-5.176	2.27e-7***
intentional.n	0.666	0.212	3.140	0.00169**
unintentional.n	-0.422	0.252	-1.673	0.09424
position.n	2.515	0.285	8.836	< 2e-16***
unintentional.n:position.n	-0.834	0.313	-2.661	0.00780**
intentional.n:position.n	-0.910	0.300	-3.033	0.00242**

Note: Signif. codes: 0 '****' 0.001 '***' 0.01 '**' 0.05 '.' 0.1 ' ' 1.

No overall significant difference is found between underspecified and unintentional sentences (*Estimate* = -0.422, *Std. Error* = 0.252, $z = -1.673$, $p = 0.095$).

For the factor 'EA placement', the model shows a main effect with a significantly higher probability of subject-readings for subject-internal EAs compared to object-internal EAs (*Estimate* = 2.515, *Std. Error* = 0.285, $z = 8.836$, $p < 2e-16$ ***).

Additionally, the model indicates a significant interaction effect between EA placement and intentionality, with a stronger increase in subject-readings for underspecified sentences compared to unintentional (*Estimate* = -0.834, *Std. Error* = 0.313, $z = -2.661$, $p = 0.008$ ***) and intentional sentences (*Estimate* = -0.910, *Std. Error* = 0.300, $z = -3.033$, $p = 0.002$ **) when EA placement changes from object-internal to subject-internal.

5. Discussion

The present study investigated the influence of intentionality as an operationalization of blameworthiness on EA interpretation.

(H2) predicted a significantly higher proportion of subject-readings for subject-internal EA placement because EA placement functions as a pragmatic cue to the target of the negative speaker attitude via syntactic association. The findings support the hypothesis, showing a significantly higher proportion of subject-readings when the EA is realized subject-internally compared to object-internally.

(H1) predicted a significantly higher proportion of subject-readings for intentional sentences compared to unintentional sentences, with an intermediate probability for sentences underspecified for intentionality, because intentionally conducted negative actions lead to a higher degree of blameworthiness for the causally responsible agent.

The predictions of (H1) are only partly confirmed. While intentional sentences show a significantly higher proportion of subject-readings compared to unintentional sentences, underspecified sentences do not follow the expected pattern. This deviation is due to a significant interaction effect between intentionality and EA placement, with underspecified sentences behaving like intentional sentences for subject-internal EAs and like unintentional sentences for object-internal EAs.

For intentional and unintentional sentences, the difference in subject-readings remains consistent across both levels of EA placement, with a significantly higher proportion of subject-readings for intentional sentences. In contrast, a differential influence of EA placement on EA interpretation is observed when intentionality is underspecified and must be inferred by the addressee. While for intentional and unintentional sentences, EA-placement results in a 22–24% increase in subject

interpretations when changed from object-internal to subject-internal, the increase for underspecified sentences is significantly stronger with almost 36%.

The interaction effect can be readily explained by considering the flexibility of the intentionality value across the different levels of intentionality. While for intentional and unintentional sentences, the intentionality value is semantically specified by a mental-attitude adverb; however, in the case of underspecified sentences, the value must be inferred based on prior knowledge about specific actions and other contextual cues. Given that the subject-referents used in the current study are mostly positively valenced and most verb-object combinations (e.g., *oversalt the soup*) readily allow for the interpretation as accidental mishaps, the intentionality value is preferentially set to unintentional in the absence of concrete linguistic evidence indicating otherwise.

In the case of object-internal EA placement, no linguistic evidence biasing toward an intentional interpretation is given. As a result, both unintentional and underspecified sentences lead to comparable interpretations concerning the intentionality of the subject-referent, resulting in no significant differences in the proportion of subject-readings between these two conditions. In contrast to intentional sentences, unintentional and underspecified sentences consequently feature a lower probability of subject interpretations in case of object-internal EA placement.¹³

However, when the EA is realized subject-internally, it functions as a concrete pragmatic cue pointing to the subject-referent as the target of the negative speaker attitude via syntactic association. When intentionality is underspecified, addressees infer that the speaker's negative evaluation of the subject-referent, as signaled by EA placement, is more probable when the norm-violating action is performed intentionally, thereby biasing the interpretation of the intentionality value accordingly. Subject-internal EA placement thus plays a dual role in EA interpretation in underspecified sentences. First, and as is the case across all levels of intentionality, it affects the proportion of local interpretations by acting as a cue to the EA target by syntactic association. In the second step, this inference then influences the interpretation of the intentionality variable, setting its value to 'intentional', thereby leading to increased blameworthiness of the subject-referent. As a result, the subject-internal EA placement causes both pragmatic cues to converge on the subject-referent as the EA target in underspecified sentences. Consequently, underspecified sentences behave the same as intentional sentences and have a significantly higher proportion of subject-readings compared to unintentional sentences.

When intentionality is unambiguous, on the other hand, EA placement enhances the likelihood of subject-readings solely through the pragmatic inference of syntactic association inducing local readings. With intentionality having a fixed value, the position of the EA cannot exert a varying influence on intentionality across different levels of EA placement. This leads to a uniform pattern with identical slopes in the probability distributions for both intentional and unintentional sentences shown in Figure 2.

Overall, the effect of intentionality on EA interpretation observed in the current study successfully replicates the findings of Frazier et al. (2015), confirming that

¹³Given that, at least descriptively, underspecified sentences result in a higher proportion of subject-readings than unintentional sentences for object-internal EAs, it might be argued that the intentionality value is not set to 'unintentional' but simply not specifically set to 'intentional'. While this conclusion cannot be derived from the current data, as underspecified sentences do not differ significantly from unintentional sentences for object-internal EAs, it does not seem unlikely and may be tested in further studies.

blameworthiness is a relevant pragmatic cue used by addressees to infer the speaker-intended EA target.

The interaction effect between EA placement and intentionality, furthermore points to intricate ways in which EAs influence underspecified variables relevant for the interpretation of speaker attitudes that are not derivable from sentences not containing the EA.

6. Conclusion

This article aimed to elaborate on the Culprit-Hypothesis of Frazier et al. (2015) by illustrating how its theoretical foundation can be refined and its explanatory power in EA interpretation increased when discussed in the context of psychological models of blame attribution.

Improving the design of Frazier et al. (2015), the experiment reported in the current study employed 1) highly controlled experimental items and 2) a straightforward operationalization of blameworthiness via intentionality that aligns with models of blame attribution and is free from inherent linguistic processing biases. Consequently, the findings serve as robust evidence supporting the Culprit-Hypothesis proposed by Frazier et al. (2015).

While the present study focused solely on intentionality as a measure of blameworthiness, the Culprit-Hypothesis predicts that blameworthiness influences EA interpretation via a wide range of pragmatic cues related to blame attribution. Concrete examples of such factors and their predicted influence on EA interpretation were provided in sentences (6)–(13). Given the internal complexity of the blame attribution process, many more cues can likely be added to derive experimentally testable hypotheses.

The interaction effect between EA placement and intentionality observed in the current study, although unexpected, indicates that the influence of other culpability cues on EA interpretation may also be modulated by EA placement and not uniform across all levels. This suggests that their influence as components of blame attribution on EA interpretation is not derivable from their influence on perceived blameworthiness in sentences that do not contain EAs.

Altogether, this article not only substantiates the validity of the Culprit-Hypothesis but also points to new directions for exploring the interplay between EA interpretation and blameworthiness.

Data availability statement. The following OSF link provides access to (1) the CSV files with the data, (2) the CSV file listing all experimental items for the main study and (3) the Rmd files containing the data analyses in R: <https://osf.io/eq7wc/>.

Competing interest. The author declares none.

Ethical standard. No ethics committee approval was necessary for this study since the data collected were purely behavioral, and all participants were adults who provided informed consent for the anonymized use of their data prior to participation.

References

- Aissen, J. (2003). Differential object marking: Iconicity vs. economy. *Natural Language & Linguistic Theory*, 21, 435–483.
- Alicke, M. D. (2000). Culpable control and the psychology of blame. *Psychological Bulletin*, 126(4), 556.

- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278.
- 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>
- Boswijk, V. H. (2022). *The salient elephant in the room: Exploring the concept of linguistic salience*.
- Bross, F. (2021). On the interpretation of expressive adjectives: Pragmatics or syntax? *Glossa: A Journal of General Linguistics*, 6(1), 1–11.
- Cushman, F. (2008). Crime and punishment: Distinguishing the roles of causal and intentional analyses in moral judgment. *Cognition*, 108(2), 353–380.
- Davis, C., & Potts, C. (2010). Affective demonstratives and the division of pragmatic labor. In *Logic, language and meaning: 17th Amsterdam Colloquium, Amsterdam, The Netherlands, December 16–18, 2009, Revised Selected Papers* (pp. 42–52). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Desmet, T., De Baecke, C., Drieghe, D., Brysbaert, M., & Vonk, W. (2006). Relative clause attachment in dutch: On-line comprehension corresponds to corpus frequencies when lexical variables are taken into account. *Language and Cognitive Processes*, 21(4), 453–485.
- Falk, S. (2014). On the notion of salience in spoken discourse-prominence cues shaping discourse structure and comprehension. *TIPA. Travaux interdisciplinaires sur la parole et le langage*, 30, 1–23.
- Frazier, L., Dillon, B., & Clifton, C. (2015). A note on interpreting damn expressives: Transferring the blame. *Language and Cognition*, 7(2), 291–304.
- Fukumura, K., & Van Gompel, R. P. G. (2011). The effect of animacy on the choice of referring expression. *Language and Cognitive Processes*, 26(10), 1472–1504. <https://doi.org/10.1080/01690965.2010.506444>
- Gillioz, C., Fontaine, J. R., Soriano, C., & Scherer, K. R. (2016). Mapping emotion terms into affective space. *Swiss Journal of Psychology*, 75(3), 141–148. <https://doi.org/10.1024/1421-0185/a000180>
- Greene, J. D. (2007). Why are vmPFC patients more utilitarian? A dual-process theory of moral judgment explains. *Trends in Cognitive Sciences*, 11, 322–323. <https://doi.org/10.1016/j.tics.2007.06.004>
- Guercio, N. L., & Orlando, E. (2022). Expressives and argument extension. In *Proceedings of SALT 32: Semantics and Linguistic Theory* (pp. 44–62).
- Guglielmo, S. (2015). Moral judgment as information processing: An integrative review. *Frontiers in Psychology* 6, 166023.
- Guglielmo, S., & Malle, B. F. (2017). Information-acquisition processes in moral judgments of blame. *Personality and Social Psychology Bulletin*, 43(7), 957–971.
- Gutzmann, D. (2019a). The dog ate the damn cake! The syntax of expressive adjectives. In *New Frontiers in Artificial Intelligence: JSAI-isAI 2018 Workshops, JURISIN, AI-Biz, SKL, LENLS, IDAA, Yokohama, Japan, November 12–14, 2018, Revised Selected Papers* (pp. 240–255). Springer International Publishing.
- Gutzmann, D. (2019b). *The grammar of expressivity* (Vol. 72). Oxford University Press.
- Järvikivi, J., Van Gompel, R. P., Hyönä, J., & Bertram, R. (2005). Ambiguous pronoun resolution: Contrasting the first-mention and subject-preference accounts. *Psychological Science*, 16(4), 260–264.
- Lagnado, D. A., & Channon, S. (2008). Judgments of cause and blame: The effects of intentionality and foreseeability. *Cognition*, 108, 754–770. <https://doi.org/10.1016/j.cognition.2008.06.009>
- Maienborn, C., & Schäfer, M. (2011). Adverbs and adverbials. In C. Maienborn, K. Heusinger, & P. Portner (Eds.), *Semantics: An international handbook of natural language meaning* (pp. 1390–1420). de Gruyter.
- Mak, W. M., Vonk, W., & Schriefers, H. (2002). The influence of animacy on relative clause processing. *Journal of Memory and Language*, 47, 50–68. <https://doi.org/10.1006/jmla.2001.2837>
- Malle, B. F. (2021). Moral judgments. *Annual Review of Psychology*, 72, 293–318.
- Malle, B. F., Guglielmo, S., & Monroe, A. E. (2014). A theory of blame. *Psychological Inquiry*, 25(2), 147–186.
- Martin, J. W., & Cushman, F. (2016). Why we forgive what can't be controlled. *Cognition*, 147, 133–143.
- Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., & Bates, D. (2017). Balancing type I error and power in linear mixed models. *Journal of Memory and Language*, 94, 305–315.
- McArthur, L. Z., & Ginsberg, E. (1981). Causal attribution to salient stimuli: An investigation of visual fixation mediators. *Personality and Social Psychology Bulletin*, 7(4), 547–553.
- Mohammad, S. M. (2018). Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 English words. In *Proceedings of the 56th annual meeting of the association for computational linguistics* (pp. 174–184). Melbourne, Australia: Association for Computational Linguistics.
- Monroe, A. E., & Malle, B. F. (2017). Two paths to blame: Intentionality directs moral information processing along two distinct tracks. *Journal of Experimental Psychology: General*, 146(1), 123.

- Næss, Å. (2007). *Prototypical transitivity* (Vol. 72). John Benjamins.
- Ohtsubo, Y. (2007). Perceived intentionality intensifies blameworthiness of negative behaviors: Blame-praise asymmetry in intensification effect 1. *Japanese Psychological Research*, 49(2), 100–110.
- Padilla Cruz, M. (2021). On the interpretation of utterances with expressive expletives. *Pragmatics & Cognition*, 28(2), 252–276.
- Padilla Cruz, M. (2022). Is free enrichment always free? Revisiting ad hoc-concept construction. *Journal of Pragmatics*, 187, 130–143.
- Patil, I., & Cushman, F. (2017). The behavioral and neural basis of empathic blame. *Scientific Reports*, 7(1), 5200. <https://doi.org/10.1038/s41598-017-05299-9>
- Potts, C. (2004). *The logic of conventional implicatures* (Vol. 7). OUP Oxford.
- Potts, C. (2007). The expressive dimension. *Theoretical Linguistics*, 33(2), 165–198.
- Potts, C., & Schwarz, F. (2010). Affective ‘this’. *Linguistic Issues in Language Technology*, 3, 1–35.
- Radvansky, G. A., & Zacks, J. M. (2014). *Event cognition*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199898138.001.0001>
- Reeder, G. D. (2009). Mindreading: Judgments about intentionality and motives in dispositional inference. *Psychological Inquiry: An International Journal for the Advancement of Psychological Theory*, 20(1), 1–18. <https://doi.org/10.1080/10478400802615744>
- Robinson, J., & McArthur, L. Z. (1982). Impact of salient vocal qualities on causal attribution for a speaker’s behavior. *Journal of Personality and Social Psychology*, 43(2), 236.
- Rohde, H., & Kehler, A. (2014). Grammatical and information-structural influences on pronoun production. *Language, Cognition and Neuroscience*, 29(8), 912–927.
- Ronderos, C. R., & Domaneschi, F. (2023). Anticipating the damn referent: How comprehenders rapidly retrieve the speaker’s attitude when processing negative expressive adjectives. *Cognitive Science*, 47(5), e13295.
- Stanley, D. J., & Meyer, J. P. (2009). Two-dimensional affective space: A new approach to orienting the axes. *Emotion*, 9(2), 214.
- Taylor, S. E., & Fiske, S. T. (1975). Point of view and perceptions of causality. *Journal of Personality and Social Psychology*, 32(3), 439.
- Van Bavel, J. J., Xiao, Y. J., & Cunningham, W. A. (2012). Evaluation is a dynamic process: Moving beyond dual system models. *Social and Personality Psychology Compass*, 6(6), 438–454. <https://doi.org/10.1111/j.1751-9004.2012.00438.x>
- Verma, G. K., & Tiwary, U. S. (2017). Affect representation and recognition in 3d continuous valence–arousal–dominance space. *Multimedia Tools and Applications*, 76, 2159–2183.
- Vogels, J., Maes, A., & Krahmer, E. (2014). Choosing referring expressions in belgian and netherlandic dutch: Effects of animacy. *Lingua*, 145, 104–121. <https://doi.org/10.1016/j.lingua.2014.03.007>
- Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior Research Methods*, 45(4), 1191–1207. <https://doi.org/10.3758/s13428-012-0314-x>
- Young, L., & Saxe, R. (2009). Innocent intentions: A correlation between forgiveness for accidental harm and neural activity. *Neuropsychologia*, 47(10), 2065–2072.
- Zarcone, A., van Schijndel, M., Vogels, J., & Demberg, V. (2016). Saliency and attention in surprisal-based accounts of language processing. *Frontiers in Psychology*, 7, 1–17. <https://doi.org/10.3389/fpsyg.2016.00844>
- Zeijlstra, H. (2012). There is only one way to agree. *The Linguistic Review*, 29(3), 491–539.
- Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2), 162.
- Zwaan, R. A., Radvansky, G. A., Hilliard, A. E., & Curiel, J. M. (1998). Constructing multidimensional situation models during reading. *Scientific Studies of Reading*, 2(3), 199–220.

Cite this article: Glauch, K. (2024). The Culprit-Hypothesis: how blame attribution influences the interpretation of expressive adjectives, *Language and Cognition* 16: 2169–2187. <https://doi.org/10.1017/langcog.2024.56>