COPiL CAMBRIDGE OCCASIONAL PAPERS *in* LINGUISTICS
Volume **17**, Article 5: pp. 68–93, 2025 | ISSN: 2050-5949

Lying is the Most Fun a Girl Can Have With Presupposition: An Experimental Investigation into the Role of Presuppositions in Gradient Judgements of Lying and Misleading*

Oliver Cooney
University of Cambridge

Abstract The pursuit of an adequate definition of 'lying' – one which excludes instances of mere 'misleading' has garnered much attention in recent years due to rising tension between theoretical approaches and experimental evidence. While Stokke's (2013, 2016, 2018) influential framework defines lies as disbelieved assertions, evidence from Viebahn, Wiegmann, Engelmann & Willemsen (2021), Reins & Wiegmann (2021) and Reins, Wiegmann, Marchenko & Schumski (2021) has shown that native speakers consider deceptive non-asserted content such as presuppositions to be instances of lying. Yet the empirical research incorrectly construes presupposition as a homogenous category. Presuppositions have been shown to vary in (a) their logical relationship to an utterance, (b) their treatment in instances of failure (both explored by Glanzberg (2005) and (c) their level of projection and entailment (Degen & Tonhauser 2022). The current study addresses this mis-construal by demonstrating that presuppositional strength has a significant impact on native speaker judgements; weak presuppositions are less likely to be considered lies than strong presuppositions. Importantly, this variation arises in a gradient rather than in discrete categories. The results are used to propose that assertion-based accounts of lying are inadequate, while Viebahn's (2020, 2021) emerging commitment-based view naturally accommodates the present findings. This is paired with insights from lexical semantics to advocate a nuanced understanding of lying.

1 Introduction

The distinction between 'lying' and 'misleading,' while intuitive, boasts an illusive character that has intrigued both linguists and philosophers for many years. The pursuit of a definition of lying – one which excludes mere misleading – has gained particular traction in the last two decades due to rising tensions between theoretical and experimental work.

Stokke's (2013, 2016, 2018) theoretical framework defines lies as disbelieved assertions, thus excluding deceptive presuppositions from being considered lies by virtue of not being asserted. However, experimental evidence from Viebahn et al. (2021), Reins & Wiegmann (2021), and Reins et al. (2021) has illustrated that native

* Special thanks are given to my supervisor, Prof. Napoleon Katsos, for his invaluable guidance in this research.

English speakers consider the deceptive use of presuppositions to be clear instances of lying. These findings form the basis of Viebahn's (2020, 2021) proposal that a lie is any deceptive act, the content of which a speaker is committed to.

Yet the empirical research incorrectly construes presupposition as a homogeneous category. Presuppositions vary in their logical relationship to an utterance's main content, giving rise to Glanzberg's (2005) separation of strong and weak presupposition. Alongside this, Degen & Tonhauser (2022) identified variation in the level of projection and entailment across various factive presupposition triggers. Those triggers which are less projective and less closely tied to the main content have been ignored in research on presuppositional deception, thus limiting the generalisation of research conclusions.

The current study investigates whether these variable properties give rise to previously unidentified variation in lying judgements. Native English speakers are presented with deceptive presuppositions and asked to rate how far they consider these to be lies. The results illustrate a significant impact of presuppositional strength on lying judgements such that weaker presuppositions are less likely to be considered lies. Importantly, these judgements are organised on a continuous gradient, rather than in binary categories.

These results are used to propose that Stokke's (2013, 2016, 2018) assertion-based approach is inadequate given its inability to account for gradience. By contrast, Viebahn's (2020, 2021) commitment-based view naturally accommodates such evidence. Based on this, a broader account of lying is proposed which also incorporates the insights of prototype theory within lexical semantics. From this, a nuanced understanding of both lying and presuppositional strength is developed which sees both as continuous, frequently mediated by both linguistic and pragmatic context.

2 Presupposition

2.1 Defining presupposition

Stalnaker (1973, 1974, 1999) defines a presupposition as a proposition which interlocutors take for granted during conversation. This is exemplified in example (1), where the use of a definite description (DD), *Dave's girlfriend,* in (1a) indicates that the speaker assumes both that (1a) and that their interlocutor knows that (1b). Thus, (1b) is a presupposition of (1a).

(1) a. Even Dave's girlfriend drives a Honda.

b. Dave has a girlfriend.

c. Dave's girlfriend is unlikely to drive a Honda.

Stalnaker (1973) embeds his account in a broader understanding of assertion as an attempt to update the common ground, or the knowledge shared between interlocutors. This is formalised as a set of possible worlds, dubbed the 'context set'. Each interlocutor holds their own context set and presuppositions are those

propositions which hold for all worlds within it. An asserted proposition, by contrast, holds only of some worlds. Thus, an assertion aims to update an interlocutor's context set by removing those worlds for which it does not hold.

Under this approach, the focus-sensitive particle (FSP), *even*, in (1a) also acts as a presupposition trigger, indicating that in all worlds in the speaker's context set, it holds that (1c). Moreover, (1a) can only be considered felicitous if it is the case that (1c) also holds for all worlds in the interlocutor's context set. Construed in this way, presupposition is understood as a pragmatic restriction on the contexts in which a trigger can be used felicitously. The ability for this account to accommodate a presupposition like *even* is Stalnaker's (1974) motivation for moving away from the semantic approach to presupposition given by Strawson (1952) and expanded by van Fraassen (1966). These authors define presupposition through a logical relationship to an assertion such that presuppositions are necessary for propositions to obtain a truth value. (1a) presupposes (1b) because without the existence of a girlfriend of Dave, there can be no one to drive or not drive a Honda, thus (1a) would yield no truth value. While this may suffice for (1b), Stalnaker indicates that (1c) cannot be understood in this way, as *even* contributes nothing to the truth conditions of (1a). If an interlocutor were to be unaware or even explicitly disbelieve that (1c), (1a) would be infelicitous, but it would not be incomprehensible; both truth conditions and truth values could be obtained. Therefore, this infelicity cannot be understood as a semantic failing. A comprehensive theory of presupposition must account for all examples, thus Stalnaker's (1973) recognition of this truth-conditional diversity privileges his account over others. This diversity drives a deeper understanding of presupposition explored in Subsection 2.2.

2.2 Presupposition failure and the taxonomy of presupposition

Much of the presupposition literature has focused on instances of presupposition failure, where a presupposition trigger is used infelicitously. Under Stalnaker's (1973) view, this arises due to differences in the context sets of the interlocutors. When using a trigger, the speaker indicates they hold a presupposition, but this presupposition does not hold of all worlds in the listener's set. Given they arise solely from differences in knowledge, presupposition failures are ubiquitous, and interlocutors must have some mechanism to mitigate their damage. Lewis (1979) suggests this is accommodation. In cases of failure, an interlocutor may begin to assume a presupposition which they did not previously assume, accommodating it into their context set.

Yet as previously noted, not all presuppositions must be accepted for successful communication to occur, meaning accommodation is not always necessary. Glanzberg (2005) utilised this variation to divide presuppositions into two major categories, shown in Table 1. So-called 'strong' presuppositions like (1b) are necessary for the assignment of a truth value, and therefore are obligatorily accommodated. By contrast, 'weak' presuppositions like (1c) are inessential to utterance comprehension, thus are only optionally accommodated.

Within theoretical linguistic literature, the strong-weak dichotomy has been repeatedly suggested, featuring in work by Abusch (2009), Zevat (1992), Kadmon (2001), Simons (2001). All authors contend that presuppositions differ in their logical relationship to the assertion, but they differ in finer details, such as Zevat's (1992) inclusion of change-of-state verbs (COSs) as strong triggers.

Weak Triggers	**Strong Triggers**
Focus-sensitive Particles	Demonstratives
Iteratives	Factive Verbs
	Clefts

Table 1 Taxonomy of presupposition triggers by strength as devised by Glanzberg (2005).

However, the experimental literature, despite supporting the existence of variation, does not support such a strict dichotomy. Limited evidence that the classes of triggers are distinct was provided by Cummins, Amaral & Katsos (2019), who investigated the acceptability of responses to polar questions. Stimulus questions included either a strong or weak presupposition, while responses varied on two parameters: the acceptance or denial of the presupposition and of the main content. Responses which affirmed the question's main content but denied the presupposition were far more acceptable when the presupposition was weak. The denial of a strong presupposition instead required the denial of the main content. This illustrates that the triggers pattern distinctly in their relation to the main content. However, the authors also found that the denial of a weak presupposition was sufficient grounds to deny the main content, as in (2). If weak presupposition denial can also cause the main content to fail, then perhaps weak presuppositions are not so strictly distinguished from strong ones.

(2) Did Brian lose his wallet again?
No, because he never lost it before.

(adapted from Cummins et al. 2019: 208)

Domaneschi, Carrea, Penco, & Greco. (2013) presented participants with examples of sentences using strong (DDs, factives, COSs) and weak (FSPs and iteratives) triggers and asked participants about the presupposed content. They found that participants were able to provide more correct answers on questions about strong presuppositions than weak ones. Yet the authors suggested this only indicated a greater attentiveness to strong presuppositions and could not confirm that weak presuppositions were not accommodated. A direct attempt to assess the accommodation of a weak trigger *weider* (German *again*) was undertaken by Tiemann (2014) and Tiemann, Schmid, Bade, Rolke, Hertrich, Ackermann, Knapp & Beck (2011). Participants were presented with *weider*-failures and questioned on the presupposition. Results indicated that the presupposition was not accommodated, with

participants answering 'only one' to (3) despite *again* triggering the presupposition that Linda had received a pink lamp previously.

(3) Linda received a pink lamp again.
How many pink lamps did Linda receive?

Yet this methodology was criticised by Schwarz (2016), who suggested participants may have interpreted the question as referring only to the immediate event. Under this interpretation, Linda did receive 'only one' lamp in the event described by (3), regardless of whether she had received more in the past. Thus, the results cannot conclusively indicate that weak triggers are not accommodated, leaving the existence of a sharp division still contested.

To complicate matters still, Domaneschi et al.'s (2013) results for COSs fell in between the two classes, reflecting a wider debate concerning their status as strong triggers. Geurts (1994) illustrated that COSs do not consistently trigger a presupposition that the original state occurred, as in (4).

(4) a. You haven't brought your lighter with you. Have you stopped smoking?
b. You seem irritable. Have you stopped smoking recently?

The use of *stop* in (4a) indicates that the speaker assumes the interlocutor has smoked previously, while this is not necessarily the case of (4b). Simons (2001) utilised this difference to suggest that COSs and factives do not encode their presupposition in their lexical entry, while other triggers do.

Simons' (2001) claim echoes a sentiment found elsewhere in the sceptical literature on factive presuppositions. Factive verbs are claimed to presuppose the truth of their complement clause, yet Karttunen (1971) identified a sub-class of factives who lose this property within polar questions and conditional antecedents. *Realise* is one example of these 'semi-factives' as it triggers the presupposition in (5a) but can take a non-factive reading in (5b) and (5c). Contrastively, *regret* in (6) consistently carries a factive presupposition.

(5) a. Dave realised he was wrong.
b. If Dave realises that he is wrong, he will apologise.
c. Did Dave realise he was wrong?

(6) a. Dave regrets that he broke the vase.
b. If Dave regrets that he broke the vase, he will apologise.
c. Did Dave regret that he broke the vase?

Degen & Tonhauser (2022) investigated this variability in presupposition projection experimentally, presenting participants with polar questions using various factives and asking them to rate the speaker's certainty of the presupposition. They also investigated whether the utterances entailed their presuppositions by asking whether the content of a clause 'followed naturally' from its presentation within a factive predicate. The results indicated a wide range of projection and entailment across the class of factives. For example, *know* and *see* ranked higher than *reveal* and *announce*, with (7c) only following naturally from (7a) and (7b) when using the former two verbs. This replicated the results of Tonhauser, Beaver & Degen (2018) who found speakers were seen to be more certain of presuppositions triggered by *know* than *stop*.

(7) a. Sandra knew/saw/revealed/announced that Julian stole a cupcake.

b. Did Sandra know/see/reveal/announce that Julian stole a cupcake?

c. Julian stole a cupcake.

This research illustrates that presupposition triggers are not a homogeneous class, varying in their logical relationships, their salience to the interlocutor, and their projectivity. Yet, experimental evidence indicates that dividing triggers into binary categories may be less appropriate than a gradient approach, shown in Table 2.

Strength	**Examples**
Stronger	Definite Descriptions (DD)
	Factives like *know*, *see*
↓	Factives like *reveal*, *announce*
	Change-of-State predicates (COSs)
↓	Iteratives
	Additives
Weaker	Focus-sensitive particles (FSPs)

Table 2 New taxonomy of presupposition triggers by gradient strength.

Stronger triggers like DDs or *know* have been consistently shown to be logically necessary for an utterance's main content, and experimental evidence has supported this in projectivity and acceptability ratings. By contrast, the literature has shown weaker triggers like *too* to be less connected to the main content. Between these anchor points lies a diverse array of triggers who do not definitively belong to either category. While theoretical literature deems them strong (i.e., they are essential for truth valuation), experimental evidence shows their projection, entailment, and acceptability during failure is weaker than the canonically strong triggers. These 'middle triggers' therefore comprise a gradient stage between the strong and weak.

3 Lying

3.1 The failing of traditional definitions

A similar friction arises in the literature on lying whereby theoretical work often defines lying to the exclusion of phenomena that naïve native speakers in the experimental literature consider to be lies.

Stokke's (2013, 2016, 2018) semantic definition argues that lying can be defined through assertion: a speaker is said to have lied if they have asserted something they believe to be false. In this, Stokke (2013) employs Stalnaker's (1999) notion of assertion outlined in Subsection 2.1. A speaker lies if they do not believe a proposition, p, to be true, yet intend to make p common ground.

Given the mechanism of accommodation, this makes presupposition an excellent vehicle for deception. In using a presupposition trigger, the speaker acts as if they hold this presupposition, leading an interlocutor to accommodate it. Yet, for Stokke, presuppositions are not part of the assertion. Given they hold of all possible worlds in the context set, they cannot update the common ground by eliminating incompatible worlds. Consequently, deceptive presuppositions do not constitute lies for Stokke, but rather mere 'misleading'. A similar sentiment permeates other assertion-based definitions of lying, such as Saul (2012).

However, experimental evidence illustrates that ordinary speakers do consider lying to include some disbelieved non-assertions. Viebahn et al. (2021) presented participants with deceptive questions which made use of both presuppositions and implicatures believed to be false by the speaker. This included examples like (8a) which presupposes (8b) which the speaker believed to be false. All cases were judged to be lies.

(8) a. Are you going to Claire's party?

b. Claire is having a party.

Importantly, Viebahn et al. (2021) also presented participants with clear cases of deceptive behaviour, such as the use of a fake accent, and found these to not be considered lies but were misleading. Similar results were demonstrated by Reins & Wiegmann (2021) and Reins et al. (2021) finding that both English and Russian speakers viewed disbelieved presuppositions to be instances of lying.

The importance of folk concepts of lying for definitional literature has been noted by Carson (2006a), Fallis (2009), Arico & Fallis (2013). While all authors note that it is unreasonable to expect a definition to account for all intuitions – especially where ordinary speakers disagree – it remains a key desideratum that a definition should aim to be as comprehensive as possible. As such, it becomes necessary to revise Stokke's (2013, 2016, 2018) definition.

3.2 The commitment-based definition

Considering their results, van Fraassen (1966) proposes that rather than assertion, it may be more accurate to define lying through commitment, sketching a definition developed into a fuller account by Viebahn (2020, 2021).

Viebahn (2020, 2021) proposes that the distinction rests on the fact that liars commit themselves to the truth of what they communicate, while misleaders do not, as elaborated in (9). This definition resembles the one offered by Carson (2006b), who proposes liars warrant the truth of what they communicate and believe themselves to be warranting.

(9) A speaker, *A*, lies iff:

A performs a communicative act, *C*, with *p* as content

With *C*, *A* intends to communicate *p* to an interlocutor

With *C*, *A* commits themself to *p*

A believes that *p* is false (adapted from Viebahn 2021: 300)

Commitment is characterised by Viebahn (2020) as evoking 'justificatory responsibility': a liar takes on the responsibility to defend *p*, while a misleader does not. This responsibility arises when a speaker cannot felicitously respond to an audience challenge by denying they ever claimed *p*. For example, if, upon uttering (10a), a speaker was asked "how do you know Mary has a brother?" they could not respond that they never claimed she did. The same would be true if the speaker had uttered (10b), thus indicating that justificatory responsibility holds for both presuppositions and main content.

(10) a. Mary has a brother.

b. Mary is going shopping with her brother.

c. I think Mary has a brother.

d. I know Mary has a brother

If, however, a speaker were to utter (10a) in a fake French accent, and were challenged on their nationality, they could feasibly respond that they "never claimed to be French." Viebahn argues this can explain why the use of a false accent is misleading, but not a lie: the speaker is not committed to the proposition that she is French. Thus, commitment, as identifiable through justificatory responsibility, can successfully delineate the two.

While Viebahn (2020, 2021) does not commit himself to a specific communicative act (lying can occur with assertion, presupposition, or implicature), he does recognise that different linguistic forms indicate greater or lesser degrees of commitment. For example, Viebahn demonstrates that (10c) permits the denial of the audience challenge by saying one never claimed (10a). However, they would still need to

justify why they believe it. In uttering (10d), however, the speaker cannot say they did not claim (10a). For Viebahn, this indicates (10c) carries a weaker level of responsibility, but justificatory responsibility, nonetheless. The impact of this gradience will become important in Subsection 4.1.

Support for this definition can be found in Reins & Wiegmann (2021) who found a correlation between judgements on lying and various metrics of commitment. Participants were presented with the use of deceptive presuppositions, implicatures and non-verbal actions and asked rate whether they considered them lies and whether the speaker could deny the content of the deception. They found that lie ratings negatively correlated with ratings of commitment. The content of non-verbal acts could easily be denied, but these were not considered lies, while presuppositions and implicatures showed the inverse. This indicates that commitment may play a role in lying and accounts which incorporate commitment may better characterise speaker intuitions.

3.3 Returning to presuppositional lies

However, these speaker intuitions are yet to be comprehensively investigated. The experimental evidence which showed deceptive presuppositions to be lies made use only of strong presuppositions, relying heavily on the use of DDs and a limited set of factives. Without evidence of how weak presuppositions are regarded, the potential arises for these not to be considered lies. This becomes especially tenable if one recalls that presuppositions are optimal deceptive tools through accommodation. Weaker presuppositions are those which less strongly necessitate accommodation, and thus may be less successful in deceiving. In turn, this may diminish their categorisation as lies.

This conclusion is also perfectly plausible given the definitional literature. In recognising that commitment is gradient, Viebahn (2020) introduces the possibility that some presuppositions may indicate less commitment than others. Committing to the main content of an utterance requires one to commit to its strong presuppositions, given this content can only be true if they are true. The absence of this logical dependency for weak presuppositions means commitment to them does not follow naturally from commitment to the main content. If lying is seen as dependent on commitment, then this predicts deceptive weak presuppositions will less strongly be considered lies.

Though variable judgements are not incompatible with assertion-based frameworks, though they are predicted to arise differently. For Stokke, the truth conditions of an assertion update the common ground by eliminating those worlds in which they are not met. Given that weak presuppositions do not contribute to truth conditions, they do not aid this process and are therefore not asserted. However, Stokke (2016) characterises an assertion as a 'minimal proposition' which updates the common ground. Given that strong presuppositions are logically necessary for this common ground update, it is possible that these are included in the minimal proposition. Under this interpretation,Stokke's (2013, 2016, 2018) framework would eliminate only weak presuppositions from being considered lies.

The following study investigates this research gap, assessing the predictions of the two theories for the under-investigated weak presuppositions.

4 The Current Study

4.1 Research aims and hypotheses

The present study aims to uncover whether deceptive presuppositions are treated heterogeneously in lie judgements. In particular, the study seeks to uncover if variable presupposition strength causes some deceptive presuppositions to not be considered lies alongside those which are. While previous research has illustrated that speakers do consider deceptive presuppositions to be lies, these studies have neglected to assess weak presuppositions, which the theoretical literature has suggested form a distinct class. Both assertion- and commitment- based approaches to lying suggest this class will be less strongly considered lies, though they differ in exactly how.

For Stokke's (2013, 2016, 2018) approach, the hypothesis stems from lies having to be asserted. An assertion, in Stokke's view, is a strictly minimal proposition which can include strong presuppositions but sharply excludes weak presuppositions. Therefore, this predicts participants will sharply divide presuppositions into those which are, and those which are not considered a lie based on strength.

For Viebahn's (2020, 2021) commitment-based approach, however, the boundary is more blurred. Lying is instead related to commitment to disbelieved information, and commitment is a continuous variable. Thus, speakers may not classify all triggers as definitively lies or not lies, but instead place some in between these groupings due to their ability to display middling levels of commitment. This finding would contradict previous theoretical work on presuppositional strength, which suggested a sharp binary between weak and strong presuppositions in line with Stokke's (2013, 2016, 2018) approach.

Experimental work on presupposition, however, has shown that triggers continuously, rather than discretely, vary in their accommodation, projection, and entailment. This work therefore predicts Viebahn-type result, with 'middle triggers' being given middling lie ratings. Under Stokke's (2013, 2016, 2018) approach, these triggers are suggested to be concretely considered lies.

4.2 Participants

46 participants were recruited on Prolific Academic and completed an online survey implemented in Qualtrics. They received £2.25 in exchange for their participation.

The average age interval of participants was 18-24 years, though ages spanned between 18 to 44 years across 12 males, 23 females, and 5 non-binary participants. Participants were required to be at least 18 years old and a native English speaker. In post-survey demographic questions, 6 participants reported that they did not have "at least one parent/caregiver speak English to [them] from birth" and their results were subsequently excluded.

4.3 Materials

4.3.1 Vignettes and questions

A series of vignettes were constructed involving an interaction between two interlocutors in which the communicative intention of the speaker was made clear by the narrative. Examples of this are given in (11).

(11) a. Mike and Alex are co-workers. Mike wants Alex to believe he owns a car.
In fact, Mike does not own a car.
Mike tells Alex: "Someone tried to break into my car yesterday."
Alex comes to believe that Mike owns a car.

b. Mike and Alex are co-workers. Mike wants Alex to believe he owns a car.
In fact, Mike does not own a car.
Mike tells Alex: "I have a car now."
Alex comes to believe that Mike owns a car.

(11a) illustrates the use of a deceptive presupposition, as Mike uses the DD "my car" to communicate that he owns a car by presupposing it. This is contrasted with the use of a deceptive assertion in (11b) to convey the same information. 18 examples were constructed in this fashion, with each scenario having a presupposition and an assertion variant. The latter variants were used as controls, as false assertions are examples of canonical lies. The former were constructed using nine different presupposition triggers, whose selection is explained in section 4.3.1. Nine filler items were constructed which involved asserting true information.

Participants were then presented with a statement describing the speaker as having lied or misled the interlocutor and asked to rate their agreement with the statement on a Likert scale, given in Figure 1. The scale varied from 'Strongly Agree' to 'Strongly Disagree' and utilised 6 points to force a definitive choice on the critical sentence while also recording the strength of this judgement. Ethical approval for this methodology was obtained from the Ethics Committee of the Faculty of Modern and Medieval Languages and Linguistics at the University of Cambridge.

Figure 1 Sample question taken from survey.

4.3.2 *Presupposition triggers*

Triggers were selected from the gradient of presupposition strength proposed in Subsection 2.2, repeated in Table 3 adding the example triggers used in the study.

Strength	
Stronger	DD
	Factives: *know* and *see*
	Factives: *reveal* and *announce*
	COSs: *stop*
	Iteratives: again
Weaker	Additives: *too*
	FSPs: *even*

Table 3 Presupposition triggers used in the study organised by gradient strength.

While the taxonomy is to be understood as a gradient, for ease of analysis the scale was split into three sections: strong, middle, and weak. Section 5 discusses the results of each individual trigger to ensure the variability within these sections is taken under consideration.

The make-up of this gradient is based primarily on the taxonomy of triggers outlined by Glanzberg (2005), Zevat (1992) and Tiemann (2014) which deem DDs and *know* to be strong triggers while *again, too,* and *even* are deemed weak. The experimental literature surveyed in Subsection 2.2 supports this as Domaneschi et al. (2013) found DDs and factives are mandatorily processed during conversation, while FSPs and iteratives are not. Similarly, the acceptability ratings collected by Cummins et al. (2019) categorised *again* and *too* as weak in comparison to *stop.*

However, under this taxonomy, *stop* is grouped with the middle triggers, given the conflicting results attested in the literature. While theoretically considered a strong trigger, Tiemann et al. (2011) found acceptability ratings for *stop* failures patterned with weak triggers like *again* and *too.* Interestingly, Domaneschi et al.'s results for *stop* patterned with neither strong nor weak triggers, instead falling between them. This makes *stop* an intriguing edge case which appears to neither be conclusively strong nor weak.

Also within this middle category are the factives *reveal* and *announce,* based on the results of Degen & Tonhauser (2022). The authors demonstrated these factives rank very low in projection and entailment compared to others. *Know* and *see* were amongst those found instead to strongly project and entail, suggesting they are strong triggers.

4.4 Design

The study employed a Latin Square design meaning that no participant saw the same scenario with both assertion and presupposition. Participants were randomly assigned to one of two groups. Where Group 1 saw a scenario in its presupposition variant, Group 2 saw the assertion variant. Both groups saw a mixture of presuppositions and assertions. The presentation of the scenarios was randomly ordered, though the ordering of questions was consistent, with lie-ratings presented before misleading-ratings.

5 Results

5.1 Lie ratings

Given the format of questions reported in section 4.3.1, 'Strongly Agree' indicated participants strongly believed the trigger to be a lie, while 'Strongly Disagree' indicated they believed it not to be. Lie ratings were quantified by converting the Likert scale to a numerical scale ranging from 1–6, where 1 corresponds to 'Strongly Agree.' Thus, the lower the number, the more strongly participants believed the use of the trigger to be a lie. This was done to calculate a modal average for each trigger and for each broad class of trigger (strong, middle, and weak). These are reported in Figure 2, where the lefthand bar indicates the result for the assertion condition and the righthand bar indicates the result for the presupposition condition.

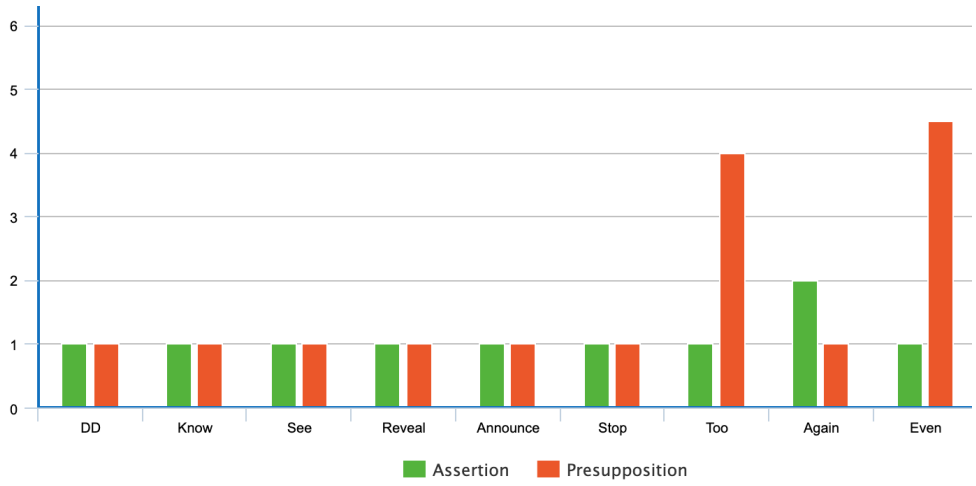

Figure 2 Mean lie ratings by trigger across both assertion and presupposition conditions.

All strong and middle triggers (DDs, *know, see, reveal, announce, stop*) reported a mode of 1 for both the assertion and presupposition conditions. This indicated participants strongly agreed that these were lies. By contrast, weak triggers had an overall mode of 3, though the mode for the individual triggers varied: *too* carried a mode of 4, *even,* 4.5, and *again,* 1.

The two conditions were compared to assess whether deceptive presuppositions are considered lies in the same way deceptive assertions are. During numerical analysis, the converted Likert scale ratings were interpreted as ordinal data, allowing for the implementation of a Kruskal-Wallis test for significance. This indicated no significant difference between ratings when using an assertion and when using strong triggers (p = 0.7409) or middle triggers (p = 0.68997). However, there did arise a significant difference in ratings when using a weak trigger (p = 0.0001). Deeper analysis of this class revealed a significant difference when using *too* (p = 0.00001) and *even* (p = 0.05) but not when using *again* (p = 0.12982). Pairing this with the mode, it becomes clear that *even* (mode: 4.5) is not considered a lie. While a mode of 4 for *too* also suggests it is not a lie, it is important to note that this reflects a bimodal distribution across 3 and 5, which lie either side of the definitive judgement. Moreover, the aggregate agreement rating for *too* (calculated using all scores between 'Strongly Agree' and 'Somewhat Agree'), was equal to the aggregate disagreement rating, suggesting results are inconclusive but remain distinct from those obtained for the assertion. The remaining triggers are strongly considered lies, indistinct from a deceptive assertion.

To assess the validity of strength-based groupings and detect gradience effects, the trigger classes were compared against one another. This revealed no significant difference between the strong and middle triggers (p = 0.3445) but did reveal one between middle and weak triggers (p = 0.00001). Furthermore, a significant difference was detected between *again* and *stop* (p = 0.01383), while no significant difference was found between the middle triggers (p = 0.248). This supports the separation of *again* from the middle class, despite being considered a lie. However, a significant difference also prevailed between *again* and *even* (p = 0.00123) and *again* and *too* (p = 0.04164), suggesting that a homogeneous weak class cannot accommodate *again* either.

In terms of dispersion, the range of answers for presuppositions was low for strong triggers (DD: 1, *know*: 0, *see*: 2) and middle triggers, except for *stop* (*reveal*: 1, *announce*: 1, *stop*: 5). This was reflected in the interquartile range (IQR) which was 0 for all strong and middle triggers besides *stop,* which had an IQR of 1. For weak triggers, range was much higher (*too*: 4, *again*: 4, *even*: 5) as was the IQR (*too*: 2, *again*: 2, *even*: 2).

The assertion condition had a much lower range, with all scenarios carrying a range of 0–3, besides *see* and *even* (DD: 0, *know*: 1, *see*: 5, *reveal*: 3, *announce*: 0, *stop*; 1, *too*: 1, *again*: 3, *even*: 5). However, these two higher ranges were not reflected in the IQR, with *see* and *even* both only giving 0.5, while the rest of the triggers gave 1 or 0.

5.2 Misleading ratings

The modal misleading rating for each trigger was 1, or 'Strongly Agree,' with the aggregate agreement rating between 18–20. The same result was found for the assertion conditions, indicating that deceptive presuppositions and deceptive assertions are considered equally misleading, irrespective of their differing lie ratings.

Once again, numerical equivalents of the Likert ratings were used to conduct a Kruskal-Wallis test for significance. This revealed no significant difference between ratings between the two conditions for any trigger (p = 0.66758 for overall comparison; p = 0.65171 for strong triggers; p = 0.52366 for middle triggers; p = 0.71332 for weak triggers). There also arose no significant difference between the classes of triggers (p = 0.452) indicating they were all considered equally misleading.

6 General Discussion

6.1 Discussion of individual results

The primary aim of this paper was to assess whether deceptive presuppositions are universally considered lies. The results reveal this is not the case. Contrary to the conclusions of previous research, lie judgements vary with respect to the strength of the deceptive presupposition. However, where theoretical work has proposed presuppositions be divided into two discrete categories (the strong and weak presuppositions), the results of the present study instead support the view that strength is a continuous gradient, in line with the conclusions of empirical studies. To illustrate this, it is necessary to explore the findings for individual triggers, which will be used to inform a broader perspective on lying in Subsection 6.2.

6.1.1 Strong and middle triggers

In line with previous research, all deceptive strong triggers were considered to be instances of lying, and there arose no significant difference between their ratings and the ratings given for deceptive assertions. This result was also replicated for the middle triggers, which were conclusively deemed to be lies. With no significant difference detected between the strong and middle triggers, the distinction between them appears less relevant to lying research. The motivation for dividing triggers came from evidence that certain strong triggers behaved differently in terms of projection, entailment, and acceptability ratings. These included *stop, reveal*, and *announce*. However, it appears the properties of these 'middle triggers' did not prevent them from being considered lies.

An important caveat to this, however, is that the differences identified for middle triggers are context-dependent, with certain contexts allowing projection and others precluding it. However, the present study only assessed each trigger in a single context. Thus, it may be the case that other linguistic contexts prompt different lie ratings for these triggers. This is an important area for future research, the effect of which will be explored in subsection 6.2.3.

Both strong and weak triggers were also considered significantly misleading, just as all deceptive assertions were. This suggests that misleading is not an alternative categorisation to lying, but instead a prerequisite; all lies are misleading, but not all instances of misleading are lies. Theories of lying are often construed such that the definitions of lying and misleading are in stark opposition to one another. This finding suggests instead that theories should seek to define lying as a phenomenon

building upon the properties of misleading, rather than contradicting them. This possibility will be opened in section Subsection 6.2.

6.1.2 Even

By contrast, while it was also considered misleading, the results demonstrate that native speakers do not consider the deceptive use of *even* to be a lie. The modal rating of 4.5 (falling between 'Strongly Disagree' and 'Somewhat Disagree') significantly differs from the modal lie-rating of 1 ('Strongly Agree') for the same scenario using a deceptive assertion.

This strongly contradicts the previous conclusions of Viebahn (2020), Reins & Wiegmann (2021), and others that all deceptive presuppositions are considered lies. While the current study did replicate their results on strong triggers, the treatment of *even* vindicates the critique that their conclusions are too general. By failing to include weak triggers in previous lying research, conclusions that span the entire class of presuppositions have inappropriately ignored the diversity of the phenomenon.

6.1.3 Too

Yet it also appears inaccurate to suggest that all weak triggers are not considered lies, as the results for *too* are inconclusive. While the modal lie-rating of 4 ('Strongly Disagree') would indicate it is at least considered weakly to not be a lie, this is the average of a bimodal distribution across 'Disagree' and 'Somewhat Agree.' Not only do these modes fall on either side of the lying threshold, but so does the overall frequency, with the aggregate frequency of disagreement answers (from 'Somewhat Disagree' to 'Strongly Disagree') being equal to that of the agreement answers (from 'Somewhat Agree' to 'Strongly Agree'). This equivalence suggests that native speakers do not share a conclusive view as to whether this trigger is considered a lie or not.

As will be discussed further in section 6.2.1, construing 'lying' as a concept with a definitive threshold may be inaccurate. Thus, it would be unfair to propose that results lie on either side of a strict binary. It is fair, however, to claim that too is considered less of a lie than other deceptions. This is supported by the significant difference found in lie ratings for the use of *too* and for the same scenario using a deceptive assertion. *Too* is different from a canonical lie and from a strong presuppositional lie, but one cannot claim that this difference is in not being a lie at all.

6.1.4 Again

The results for the final trigger, *again*, complicate this further as its deceptive use was considered a lie, contrary to this study's hypothesis. This stems from both the higher aggregate agreement (15 compared to an aggregate disagreement rating of 5) and from the mode of 1 ('Strongly Agree'). Even more striking is the lack of a significant difference between the lie-ratings for *again* and for the same scenario

with a deceptive assertion (p = 0.12982), which indicates that *again* is considered to be just as much of a lie as a deceptive assertion, similar to the results found for strong and middle triggers.

At first glance, this may indicate that *again* has been misclassified as a weak trigger, yet results show reclassification would be inappropriate. The results for *again* were shown to be significantly different from all strong and middle triggers, yet no significant difference was detected within or between these groups. This is bolstered by the range data for the triggers. For *again,* every option on the scale received at least 2 responses, with an overall range of 4 and an IQR of 2.*Too* and *even* carried a range of 4 and 5, respectively, and IQR of 2. Range scores like this are highly uncharacteristic of the strong and middle triggers, who all had a range between 0–2 and an IQR of 0. The only exception to this is *stop,* which had a range of 5 and IQR of 1, yet closer inspection of the raw data reveals this was caused by a single response of 'Strongly Disagree,' while all other participants responded in agreement. This suggests that weak triggers may be better characterised as causing deceptions which participants are unsure about, rather than those which are not considered lies.

Yet it appears incorrect to suggest that *agree* is conclusively weak. A significant difference was found between the results for *again* and those for *too* and *even,* yet no such difference was detected between the latter two. Moreover, both *too* and *even* were treated significantly differently from a deceptive assertion, while *again* was not. It appears, then, that *again* does not belong to either the weak nor the middle or strong classes, but instead falls somewhere in between. This finding has incredibly important implications for our understanding of presuppositional strength and the theory of lying, as explored in Subsection 6.2 and Subsection 6.3.

6.2 Implications for the theory of lying

6.2.1 Gradience in lying

Taken together, these results indicate that lying is not a discrete category, but rather a gradient, whereby certain forms of deception are taken to more strongly be lies than others. In the context of presupposition, the results reveal the gradient in Table 4. Given the conflicting results obtained for *too*, it is unclear at exactly which point presuppositions are no longer considered lies. Yet this serves only to further emphasis the gradient nature of the category which incrementally decreases membership, rather than definitively prohibiting it.

The experimental research surveyed in Subsection 2.2 suggested that gradience would exist for lie judgements within the set of factives, given that they vary in their projection and entailment. Though the current study only assessed 4 of the 14 verbs Degen & Tonhauser (2022) proved to vary in projection and entailment, these were taken from either end of their scale, reflecting a significant difference in their behaviour. However, the present study detected no significant differences between the factives. A potential explanation for this is given in section 6.2.3. Setting this

Strength	
More strongly lies	DD, Factives, COSs
	Iteratives: *again*
	Additives: *too*
Less strongly lies	FSPs: *even*

Table 4 Indiscrete gradient for presupposition triggers by strength.

aside, the current results do not support positing finer gradience within the category of factives.

The identification of gradience within the category of lying has important theoretical implications, revealing the inadequacy of the assertion-based approach. Subsection 3.3 demonstrated that Stokke's (2013, 2016, 2018) view is not incompatible with the existence of strong presuppositional lies, so long as one includes strong presuppositions within the assertion. However, this cannot be said for weak presuppositions. Weak presuppositions are, by definition, not relevant to the truth conditions of an utterance, thus are inessential for the utterance to update the common ground. For Stokke (2013), assertions are minimal propositions, thus inessential content is not seen to be asserted. If a lie must be an assertion, then weak presuppositions cannot be lies. The current study presents two major challenges to this view.

The first, and perhaps most crucial, is that the majority of participants considered *again* to be a lie, while at least half of participants considered *too* to be a lie. In both being weak presuppositions, this renders Stokke's (2013, 2016, 2018) view untenable, as his definition inaccurately excludes weak presuppositions where speaker judgements include them. The second challenge comes from the evidence that *again* is considered more of a lie than *too* which, in turn, is considered more of a lie than *even.* This gradience is incompatible with the assertion-based approach, which relies upon a harsh binary between those aspects of meaning which are and are not contained within an assertion. Even if one were to consider weak presuppositions to be contained within an assertion, all weak presuppositions would have to be found here. This would mean that universally, weak presuppositions must all be considered lies, and there is no room for gradience. Conversely, if they are cast out of the assertion, they must all be cast out, yet research has illustrated no conclusive view of weak presuppositions as resting on either side of the lying threshold.

The commitment-based approach, however, naturally permits gradience in lie judgements by recognising that commitment is a gradient property. Viebahn (2020) notes that linguistic expressions carry varying levels of commitment to their content which cannot be dichotomously classified as entirely committed or entirely uncommitted. Rather, speakers illustrate a continuum of commitment when using different expressions. Where an utterance carries greater commitment, a speaker

can be said more strongly to have lied if they believe this utterance to be false. Thus, the commitment-based approach paves the way for deceptive presuppositions to vary in how strongly they are considered to be lies, as is shown to be the case in the findings of the present study.

For the predictions of the commitment-based view to be fully realised, the findings of this study must be correlated with findings on how commitment varies across weak triggers. This has yet to be studied directly, yet previous research by Reins & Wiegmann (2021) provided some insight. The researchers assessed the perceived commitment levels and lie status of a number of presuppositions and found that where each presupposition was considered a lie, it was also considered to carry a high degree of commitment. As characteristic of presupposition research, they considered only strong presuppositions, meaning their findings cannot be tied to the variable lie ratings found for *again, too,* and *even.* Yet their methodology can be preliminarily applied to critical sentences used in this study, given in (12).

One key metric Reins & Wiegmann (2021) assessed was deniability: whether a speaker can feasibly deny they ever claimed a presupposition. If a presupposition can be denied, a speaker is less committed to it.

(12) a. Olivia: Even Mike passed the maths exam!

b. Stephen: I passed my maths exam, too.

c. Liam: I know Chris has been sleeping with other women.

If Olivia were criticised for saying Mike is bad at maths, it seems perfectly plausible for her to say she never claimed so, despite this presupposition being triggered by *even* in (12a). By contrast, it would be entirely infelicitous for Liam to say he never claimed it to be true that Chris has been sleeping with other women, as this is presupposed by *know* in (12c). Between these extremes lies the semi-infelicity of a retort by Stephen that he never claimed to have passed other exams due to the presupposition triggered by *too* in (12b). This doesn't appear as infelicitous as Liam's denial, but it is certainly not as plausible as Olivia's defence. While only preliminary, this indicates that deniability does correlate with the findings of the current study in that (12a) is less of a lie than (12b) which is, in turn, less of a lie than (12c). It is crucial that deniability, along with other metrics of commitment, are re-assessed in an experimental setting with a broader audience before concluding that Viebahn's view is accurate. However, the results of the current study do suggest that a commitment-based approach may be more fruitful in defining lying than an assertion-based view, given the existence of gradience.

6.2.2 A potential caveat: prototype theory

Yet one potential source for gradience remains undiscussed: prototype effects. Prototype theory, pioneered by Rosch (1973, 1975) and Rosch & Mervis (1975) argues that lexical meaning cannot be understood as a discrete list of necessary and sufficient conditions. Instead, it argues that word meanings correspond to

cognitive concepts which bear prototypical idealisations. The real-world examples that can be referred to by a given word need not be identical to this prototype but will share some characteristics with it. The exact number of shared characteristics can vary, giving rise to the typicality effects which motivated the development of the approach. Rosch (1975), Rosch, Simpson & Miller (1976) (amongst others) found that in an experimental setting, entities were not dichotomously categorised as belonging to a category or not, but their membership would vary as a function of how many properties they shared with an abstract prototype. The more they resembled the prototype, the more strongly and unanimously they were identified as an appropriate referent for the word. If one understands lying as a cognitive concept construed in this way, then the gradience effects in the present study can be understood as yet another example of prototypicality.

In fact, Coleman & Kay (1981) investigated prototype effects with regards to lying by presenting participants with vignettes of lies. These vignettes varied in the features they shared with a prototypical lie, given in (13).

(13) A lie is an asserted proposition, *p*, such that…

p is false

The speaker believes *p* to be false

The speaker intends to deceive the addressee

They found that where scenarios included a greater number of these features, participants were more likely to consider the speaker to have lied. Furthermore, instances with a greater number of prototypical properties were also considered more strongly to be lies than those with fewer.

Coleman & Kay (1981) do not claim their list of prototypical properties to be exhaustive, thus there remains the potential to include some property relevant to presupposition. To account for the apparent distinction between strong and weak presuppositions, this would have to be a property more pertinent in strong presuppositions. One candidate for this is assertion. As explored previously, strong presuppositions can be understood as those contained in an assertion, while weak presuppositions are not. Thus, if the human prototype of a lie is asserted, deceptive strong presuppositions will more closely resemble it. Deceptive weak presuppositions, on the other hand, will share fewer features, making them less likely to be considered lies. Crucially, though, this does not preclude them from being part of this category, as referents with fewer prototypical features can still be referents. This can therefore accommodate the result that strong and weak presuppositions are judged differently while also permitting some deceptive weak presuppositions to be lies, as demonstrated by *again*.

While this appears to be a way to resurrect Stokke's (2013, 2016, 2018) definition, once again the issue of gradience persists: The proposed prototype would not allow judgements on weak presuppositions to vary. If all weak presuppositions are equally un-asserted, they should equally resemble (or not resemble) the prototype. This predicts, then, that they should be considered equally in lie judgements, be that as a

weak lie or not a lie at all. Yet the study revealed that *again* is considered a lie, while *too* and *even* are not. If the prototype account is to explain the effects observed in this study, the prototypical lie must also hold some gradient property which varies across the weak presuppositions. As noted above, one potential property is commitment.

If high commitment is part of the prototype of a lie, this would account for the observed gradience, as it is already established that commitment is gradient. The gradience effects which arise in lying judgements, then, are not just a function of whether or not a referent shares prototypical properties, but also to what extent it shares these properties if they are gradient. Even where the insights of prototype theory can account for native intuitions on lying, this requires one to recognise the potential insights of Viebahn's (2020, 2021) commitment theory and view assertion-based approaches as inadequate in isolation.

6.2.3 Contextualism in the taxonomy of lying

Thus far, the results the present study have illustrated that lie ratings vary as a function of the linguistic means used, i.e. the presupposition trigger. The particularities of this variation have been used to suggest that this stems from the variable commitment carried by different triggers. Yet as noted throughout discussions on strong triggers, presuppositions vary also as a function of the context in which they are used.

In particular, certain strong presuppositions can be seen to not project or entail when used in certain constructions such as questions. The present study did not assess the effect of these conditions, instead opting to present all triggers in a simple, declarative context. However, it may be the case that this could bear a strong influence on the judgements collected.

Although *announce* was classified as a lie in this study, Kiparsky & Kiparsky (1970) note that *announce* does not always project. The critical context used in the present study is given in (14a) which strongly presupposes the truth of the complement clause. However, presenting announce in a polar question, as in (14b), permits a non-factive reading.

(14) a. John wants Tracy to think her cousin is getting a divorce. In fact, her cousin is happily married and doesn't intend to divorce. John tells Tracy: "Your cousin announced that she is getting a divorce." Tracy comes to believe her cousin is getting a divorce.

b. John wants Tracy to think her cousin is getting a divorce. In fact, her cousin is happily married and doesn't intend to divorce. John asks Tracy: "Did your cousin announce that she is getting a divorce?" Tracy comes to believe her cousin is getting a divorce.

If (14b) is taken by the interlocutor to not project, then there is potential for it to be judged less of a lie than (14a), where the presupposition cannot be denied. If this is the case, this presents a potential distinction between middle and strong triggers

which was not identified in the present study. This distinction, though, would only arise in certain linguistic contexts.

There is also potential for the pragmatic context to influence lie decisions too. In defining weak presupposition, Glanzberg (2005) proposes that these presuppositions are optionally accommodated, rather than obligatorily so, due to their lack of influence on the truth conditions. However, when a weak presupposition is used in certain discourse contexts, its accommodation can become mandatory. In the present study, *even* was judged strongly to not be a lie, based on the scenario given in (15a). Here, the presupposition that Mike is bad at maths does not contribute to the truth conditions that Mike passed the maths exam. However, (15b) presents another usage of *even* where the failure of its relevant presupposition precludes an appropriate response to the preceding discourse.

(15) a. Olivia wants Paula to believe her brother, Mike, is bad at maths. In fact, Mike is very good at maths. Olivia tells Paula: 'Even Mike passed the maths exam!' Paula comes to believe that Mike is bad at maths.

b. Olivia wants Paula to believe her brother, Mike, is bad at maths. In fact, Mike is very good at maths. Paula asks Olivia: 'Was the maths exam difficult?' Olivia responds: 'Even Mike passed the maths exam!' Paula comes to believe that Mike is bad at maths.

In (15b), Olivia's response can only answer Paula's question through the implicature that it was easy. This implicature can only arise if one presupposes that Mike is bad at maths: if Mike, who is very bad at maths, passed the exam, it must have been an easy exam. The presupposition triggered by *even* is therefore obligatory here, as one cannot glean a relevant contribution from Olivia's response unless one takes the presupposition to be true. Though obligatory accommodation is not typically a feature of weak presuppositions, this example demonstrates that it can be prompted in certain pragmatic contexts. It may be the case that this influences lie judgements, with (15b) being judged more of a lie than (15a).

Though the examples given were not assessed in the current study, future research on contextual influence on lie judgements may reveal their importance. Where the given speculations prove fruitful, the gradient given in Table 3 must be rethought as a function of the experimental context, rather than an absolute ranking of triggers. It may be the case that triggers in alternative contexts appear much higher or much lower in this scale.

6.3 Implications for presupposition theory

In Glanzberg's (2005) original conception of a strength-based taxonomy of presupposition, triggers were divided into two distinctive groups. However, the present research has illustrated that, just as lying should be conceived of as a gradient, so too should presuppositional strength. In fact, the gradients follow naturally from one another.

The relation of presuppositional strength to lying stems from the definition of strong presuppositions as those which must be accommodated in cases of failure. A deceptive presupposition is merely another instance of failure, and therefore prompts accommodation. The difference here is that the failure is intentional. The speaker aims to trigger accommodation, allowing the lie to be successful. Considering this, it was hypothesised that strong presuppositions would be more successful than weak presuppositions in lying because they prompt obligatory accommodation, i.e., they force the lie to be successful. Given that weak presuppositions are only optionally accommodated, the likelihood of successful deception is lower.

Assuming this likelihood of success is related to the prototype of lying, the current research has illustrated that one cannot identify a definitive cut-off point between the more and less successful. Edge cases like *again* and *too* illustrate that weak presuppositions are not a homogeneous class which is conclusively distinct from the strong triggers. As such, it becomes more appropriate to conceive of their strength as a continuum, rather than a binary.

This reconceptualisation is entirely plausible given the definition of a weak trigger: they prompt *optional* accommodation, not a lack thereof. Proposing that weak triggers are inconsistently accommodated does not stand them in stark opposition to those which are always accommodated, as accommodation can and does still occur. The definition does, however, categorise strong presuppositions in an all-or-nothing fashion: if a trigger is not always accommodated, its accommodation is not obligatory, thus it is not a strong trigger. Yet as illustrated in Subsection 2.2, even purported strong triggers do not always require accommodation. Certain constructions can permit a reading without a presupposition (as in Example 5) or even preclude the projection of the presupposition entirely (as in example Example 7). Not only does this threaten the homogeneity of the strong trigger class, but it also damages the view that strong presupposition triggers always require accommodation. Given the absence of a significant difference in the treatment of strong and middle triggers, it seems incorrect, at least in the case of declarative sentences, to suggest these triggers constitute discrete classes of strength. Once again, the current findings corroborate previous empirical research in illustrating a heterogeneity within classes and a gradience between them. This refutes the conception of presuppositional strength as distinct classes and suggests in its place a continuum.

7 Conclusion

In summary, the results of the present study confirm that presuppositions do not behave as a homogeneous class in judgements of lying. Instead, the weak presuppositions previously absent from lying research are shown to be less likely considered lies. The strong presuppositions, though, conform to previous findings that deceptive presuppositions are considered lies. Crucially, this variability does not arise as a strict binary, but rather as a gradient where weak presuppositions can be considered lies, not lies, or somewhere in between. This variability, however, did not extend upward to the so-called 'middle triggers', which were deemed theoretically strong

but behaved differently in experimental settings. Though it remains necessary to study this further, this indicates preliminarily that, in some contexts, middle triggers need not be distinguished from the strong. Instead, all triggers can be considered on a continuum. This gradience was shown to be incompatible with the assertion-based approach to lying, but is entirely plausible under the commitment-based view. While the results do not directly confirm the commitment approach, they provide support for it as a more fruitful approach to the lying-misleading distinction.

References

Abusch, D. 2009. Presupposition triggering from alternatives. *Journal of Semantics* 27. 37–80. doi:10.1093/jos/ffp009.

Arico, A. J. & D. Fallis. 2013. Lies, damned lies, and statistics: An empirical investigation of the concept of lying. *Philosophical Psychology* 26. 790–816. doi:10.1080/09515089.2012.725977.

Carson, T. L. 2006a. The definition of lying*. *Nous* 40. 284–306. doi:10.1111/j.0029-4624.2006.00610.x.

Carson, T. L. 2006b. The definition of lying. *Nous* 40(2). 284–306. doi:10.1111/j.0029-4624.2006.00610.x.

Coleman, L. & P. Kay. 1981. Prototype semantics: The english word lie. *Language* 57. 26–44. doi:10.1353/lan.1981.0002.

Cummins, C., P. Amaral & N. Katsos. 2019. Backgrounding and accomodation of presuppositions: an experimental approach. *Proceedings of Sinn und Bedeutung* 17. 201–218. https://ojs.ub.uni-konstanz.de/sub/index.php/sub/article/view/341.

Degen, J. & J. Tonhauser. 2022. Are there factive predicates? an empirical investigation. *Language* 98(3). 552–591. doi:10.1353/lan.0.0271.

Domaneschi, F., E. Carrea, C. Penco, & A. Greco. 2013. The cognitive load of presupposition triggers: mandatory and optional repairs in presupposition failure. *Language, Cognition, and Neuroscience* 29(1). 136–146. doi:10.1080/01690965.2013.830185.

Fallis, D. 2009. What is lying? *The Journal of Philosophy* 106. 29–56. http://www.jstor.org/stable/20620149.

van Fraassen, B. C. 1966. Singular terms, truth-value gaps, and free logic. *The Journal of Philosophy* 63. 481. doi:10.2307/2024549.

Geurts, B. 1994. *Presupposing*: Universität Osnabrück dissertation.

Glanzberg, M. 2005. Presuppositions, truth values, and expressing propositions. In G. Preyer & G. Peter (eds.), *Contextualism in philosophy: knowledge, meaning, and truth*, 349–396. Oxford University Press.

Kadmon, N. 2001. *Formal pragmatics*. Blackwell.

Karttunen, L. 1971. Some observations on factivity. *Research on Language and Social Interaction* 4. 55–69. doi:10.1080/08351817109370248.

Kiparsky, P. & C. Kiparsky. 1970. Fact. *Progress in Linguistics* doi:10.1515/9783111350219.143.

Lewis, D. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8. doi:10.1007/bf00258436.

Reins, L. M. & A. Wiegmann. 2021. Is lying bound to commitment? empirically investigating deceptive presuppositions, implicatures, and actions. *Cognitive Science* 45. doi:10.1111/cogs.12936.

Reins, L. M., A. Wiegmann, O. P. Marchenko & I. Schumski. 2021. Lying without saying something false? a cross-cultural investigation of the folk concept of lying in russian and english speakers. *Review of Philosophy and Psychology* 14. doi:10.1007/s13164-021-00587-w.

Rosch, E. 1973. Natural categories. *Cognitive Psychology* 4. 328–350. doi:10.1016/0010-0285(73)90017-0.

Rosch, E. 1975. Cognitive representations of semantic categories. *Journal of Experimental Psychology: General* 104. 192–233. doi:10.1037/0096-3445.104.3.192.

Rosch, E. & C. B. Mervis. 1975. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology* 7. 573–605. doi:10.1016/0010-0285(75)90024-9.

Rosch, E., C. Simpson & R. S. Miller. 1976. Structural bases of typicality effects. *Journal of Experimental Psychology: Human Perception and Performance* 2. 491–502. doi:10.1037/0096-1523.2.4.491.

Saul, J. M. 2012. *Lying, misleading, and what is said.* Oxford University Press.

Schwarz, F. 2016. Experimental work in presupposition and presupposition projection. *Annual Review of Linguistics* 2. 273–292. doi:10.1146/annurev-linguistics-011415-040809.

Simons, M. 2001. On the conversational basis of some presuppositions. *Semantics and Linguistic Theory* 11. 431. doi:10.3765/salt.v11i0.3099.

Stalnaker, R. 1973. Presuppositions. *Journal of Philosophical Logic* 2. doi:10.1007/bf00262951.

Stalnaker, R. 1974. Pragmatic presuppositions. In M. K. Munitz & P. Unger (eds.), *Context and content: Essays on intentionality in speech and thought* Semantics and Philosophy, 197–214. New York University Press.

Stalnaker, R. C. 1999. *Assertion* 78–95. Oxford University PressOxford. doi:10.1093/0198237073.003.0005.

Stokke, A. 2013. Lying and asserting. *Journal of Philosophy* 110. 33–60. doi:10.5840/jphil2013110144.

Stokke, A. 2016. Lying and misleading in discourse. *Philosophical Review* 125. 83–134. doi:10.1215/00318108-3321731.

Stokke, A. 2018. *Lying and insincerity.* Oxford University Press.

Strawson, P. F. 1952. *Introduction to logical theory.* London, Methuen; New York, Wiley.

Tiemann, S. 2014. *The processing of 'wieder' ('again') and other presupposition triggers*: Universitat Tübingen dissertation. http://hdl.handle.net/10900/57169.

Tiemann, S., M. Schmid, N. Bade, B. Rolke, I. Hertrich, H. Ackermann, J. Knapp & S. Beck. 2011. Psycholinguistic evidence for presuppositions: On-line and off-line data. *Proceedings of Sinn und Bedeutung* 15. 581–596. https://ojs.ub.uni-konstanz.de/sub/index.php/sub/article/view/401.

Tonhauser, J., D. Beaver & J. Degen. 2018. How projective is projective content? gradience in projectivity and at-issueness. *Journal of Semantics* 35. doi:10.1093/jos/ffy007.

Viebahn, E. 2020. Lying with presuppositions. *Noûs* 54. 731–751. doi:10.1111/nous.12282.

Viebahn, E. 2021. The lying-misleading distinction: A commitment-based approach. *The Journal of Philosophy* 118. 289–319. doi:10.5840/jphil2021118621.

Viebahn, E., A. Wiegmann, N. Engelmann & P. Willemsen. 2021. Can a question be a lie? an empirical investigation. *Ergo* 8. doi:10.3998/ergo.1144.

Zevat, H. 1992. Presupposition and accommodation in update semantics. *Journal of Semantics* 9. 379–412. doi:10.1093/jos/9.4.379.

Oliver Cooney
University of Cambridge
ojc42@cam.ac.uk