

Lying, Deceptive Implicatures, and Commitment

Alex Wiegmann, Pascale Willemsen & Jörg Meibauer

Forthcoming in *Ergo*

Abstract

Deceptive implicatures are a subtle communicative device for leading someone into a false belief. However, it is widely accepted that deceiving by means of deceptive implicature does not amount to lying. In this paper, we put this claim to the empirical test and present evidence that the traditional definition of lying might be too narrow to capture the folk concept of lying. Four hundred participants were presented with fourteen vignettes containing utterances that communicate conversational implicatures which the speaker believes to be false. We further collected several potential proxy measures of lying, to get a better understanding of when a deceptive implicature is considered a case of lying. The results indicate that most implicatures (ten out of fourteen) were evaluated as lies and that lie ratings were closely tracked by the degree to which speakers were considered to have committed themselves to the truth of the content conveyed by their deceptive implicatures.

Keywords

definition of lying, conversational implicature, deceptive implicature, commitment

Author Note

Alex Wiegmann, Institute of Philosophy II, Ruhr University Bochum

Pascale Willemsen, Institute of Philosophy, University of Zurich

Jörg Meibauer, Institute of Linguistics, University of Mainz

Correspondence concerning this article should be addressed to:

Alex Wiegmann, E-mail: Alexander.Wiegmann@ruhr-uni-bochum.de

1. Introduction

Kevin is shocked when he finds an empty pack of cigarettes in the trash. He confronts his wife, Sally. Sally does not want Kevin to know that she secretly smokes from time to time, so she says, “Evelyn was here yesterday, and she is a smoker.” In fact, Evelyn visited Sally yesterday and she is a smoker. The empty pack is not Evelyn’s pack but the one Sally smoked.

Did Evelyn lie to Kevin? According to a widely held view in the philosophical literature (henceforth: the standard view), lying entails that, in order to lie, speakers have to assert something they believe to be false.¹ This requirement may be spelled out as follows (cf. Stokke 2018; Viebahn 2019a):

A lies to B if and only if there is a proposition *p* such that

1. A asserts that *p* to B, and
2. A believes that *p* is false.

Applying these conditions to the case at hand leads to the verdict that Sally did not lie, because she did not *assert* anything that she believes to be false. Sally does indeed believe that Evelyn was there yesterday, and she also believes that her friend is a smoker. Had Sally replied by saying “The empty pack of cigarettes in the trash belongs to Emily”, her utterance would have counted as a lie. However, by replying “Evelyn was here yesterday, and she is a smoker” Sally merely implicated that the cigarettes belonged to Evelyn – she did not assert it. Sally’s utterance, it is often claimed, might be misleading but does not amount to lying. In this paper we challenge this widely held view theoretically and test its empirical adequacy. The results of our experiment indicate that the folk do consider some deceptive implicatures as instances of lying. We further suggest that a commitment-based definition of lying that also include certain deceptive implicatures might be promising.

The paper is structured as follows. In Section 2, we first discuss the importance of empirical studies for philosophical discussions about the concept of lying. Section 3 introduces conversational implicatures and describes how they can be understood as a means of lying. In Section 4, we elaborate on the role of commitment for lying and argue that, at least in some cases, the speaker seems to be committed to the content of a deceptive implicature and therefore might be considered to be lying. In Section 5, we briefly summarize the sparse empirical literature on lying

¹ This general account has been endorsed by many authors, e.g. Chisholm & Feehan (1977), Adler (1997), Carson (2006, 2010), Sorensen (2007), Fallis (2009), Saul (2012), Stokke (2018). However, the proposals often differ from each other in relying on different accounts of assertion. Stokke (2018), for instance, bases assertions on the notion of common ground, which results in the following definition:

A lies to B if and only if there is a proposition *p* such that

- (L1) A says that *p* to B, and
- (L2) A proposes to make it common ground that *p*, and
- (L3) A believes that *p* is false.

(L2) might be considered to be a replacement condition for a prominent but contested requirement, namely the intention-to-deceive condition (cf. Mahon 2016).

by means of deceptive implicature. In Section 6, we present our own experiment in which we tested whether fourteen cases of deceptive implicatures are judged to be lies. We further collected data on several factors that might explain under what conditions a deceptive implicature is considered to be a case of lying. We found that most deceptive implicatures were judged to be cases of lying, and that lying ratings were strongly correlated with the speaker's commitment to the content of the deceptive implicature. The folk concept of lying, so we conclude in Section 7, is broader than most prominent definitions assume. We discuss whether a commitment-based definition of lying might be promising.

2. The Relevance of Empirical Research

Lying is a concept that is deeply anchored in our social life. For this reason, it does not come as a surprise that some philosophers explicitly state that a good definition of lying should capture the ordinary use of this concept.² For instance, Carson (2010: 33) writes

Lying is a concept used in everyday language, and moral questions about lying arise in people's everyday experience. There are no compelling reasons to revise or reject the ordinary language concept of lying -- at least the burden of proof rests with those who would revise or reject it. Therefore, consistency with ordinary language and people's linguistic intuitions about what does and does not count as a lie is a desideratum of any definition of lying.

In a similar vein, Fallis (2012: 974) highlights that the debate is about the folk concept of lying and says

The goal of the philosophers [...] is, apparently, to find the definition that best captures the ordinary usage of the term 'lying'. That is, the participants in this debate are not disagreeing about the best way to define a highly technical notion [...]; nor is each theorist merely laying out his own idiosyncratic conception of lying.

If the aim is to develop a definition of lying that is consistent with ordinary language, it seems difficult to deny the importance of empirical studies on people's concept of lying. We can only evaluate whether the definition is appropriate, if we know how the folk concept works. Against this background, an empirical examination of whether a definition of lying should include the possibility to lie with (certain) deceptive implicatures seems in order (for an overview of empirical studies on people's concept of lying, see Wiegmann & Meibauer 2019).

Of course, there might be interesting philosophical projects to define lying that do not (primarily) aim to capture people's intuition about this concept. Such projects will not be touched by the evidence we present in this paper. The aim of this paper is to provide an empirical basis for definitions of lying which aim to capture the folk concept, and we remain silent on the implications of our findings for other definitory projects. We also wish to emphasize that empirical studies have

² To the best of our knowledge, no philosopher claims that ordinary people's intuition about lying are of little or no relevance for the definition of lying

their limitations (cf. Gerken 2018) and their findings might be overridden by theoretical considerations. Empirical studies cannot and must not replace philosophical theorizing and conceptual analysis. But there is currently little reason to doubt they have potential to advance the philosophical discussion.

3. Conversational Implicatures

In this section, we consider some of the tenets of the theory of implicatures, as put forward by Grice (1989a, b). A central focus of linguistics and philosophy of language is the study of the meanings of linguistic entities. Meaning is traditionally researched in semantics and pragmatics. Semantics deals with literal meaning (which remains relatively stable across contexts), while pragmatics deals with non-literal and context-dependent meaning (see Borg 2012; Carston 2002; Levinson 2000; Recanati 2004).

A semantic notion related to sentence meaning is *entailment* (or *implication*). For instance, the sentence *Peter and Mary went shopping* entails that Mary went shopping. If the sentence *Peter and Mary went shopping* is true, then it must also be the case that the sentence *Mary went shopping* is true. Therefore, *Peter and Mary went shopping* entails *Mary went shopping*. In many situations, however, upon hearing this utterance, we would think that Peter and Mary went shopping *together*. This extra meaning is not conveyed by the semantics of the sentence, but is rather a pragmatic inference, namely an *implicature*. The conversational implicature arises as a hypothesis about intended meaning without being expressed (or coded) literally. This can easily be seen by comparing the sentence *Peter and Mary went shopping* with *Peter and Mary went shopping but not together*. Cancelling the implicated content explicitly is possible without yielding a contradiction. Therefore, it is adequate to distinguish between entailments and implicatures, as Grice (1989a) famously pointed out. Take, as another example, the sentence *Children are children*. From a purely semantic point of view, this sentence is utterly uninformative (being a tautology). However, due to pragmatic inferences, the addressee might derive the additional meaning that the speaker is suggesting that it is pointless to worry about children's bad behavior (see Meibauer 2008). Conversational implicatures therefore not only make a significant contribution to communication, but also constitute, as Grice argued, another building block of what he calls "speaker meaning".

Conversational implicatures, arrived at on the basis of pragmatic inference, have several distinctive properties: they are cancellable (e.g., by adding a cancelling phrase in the ensuing discourse; for pertinent problems see Haugh 2013, Zakkou 2018), they vary according to specific contexts, and are derivable via Gricean reasoning (see the general pattern described in Grice 1989a: 31). In addition, since conversational implicatures have certain contents, they can be evaluated in the true/false dimension. Suppose someone says *Peter and Mary went shopping*, and the addressee

derives the implicature that the two of them went shopping together. Here we can ask whether it is true that they went shopping together. Since implicatures can be true or false, they can be used to deceive. The sentence about Peter and Mary might be used in a context in which Tom intends Peter's wife Sally to believe that her husband has gone shopping together with Mary, and make further inferences about her husband's infidelity.

While a widely accepted view is that one cannot lie by using deceptive implicatures, some researchers such as Bok (1999), Faulkner (2007), Meibauer (2005, 2011, 2014a, 2019), and Ortony and Gupta (2019) have argued that lying by using deceptive implicatures is possible. According to such broader and more inclusive notions of lying, Tom's deceptive implicature, intended to create a false belief in Sally, would not only count as a case of deception, but also constitute a case of lying. Although Grice did not endorse this view, one might, also from a Gricean perspective, count deceptive implicatures as lies by pointing to the idea of the "total signification of an utterance" (Grice 1989b: 41), i.e., the complete meaning that a speaker conveys and a hearer computes, which includes implicated content. Deceptive implicatures, after all, are used to deceive, can be evaluated in the true/false dimension, and, as we will argue, can involve a high degree of speaker commitment, which might suffice to count as a lie according to the folk concept of lying.

4. Conversational Implicatures and Commitment

According to virtually all definitions of lying in the philosophical literature, one cannot lie by means of deceptive implicatures. This claim is often taken for granted, or it justified by reference to examples which are assumed to reliably trigger certain intuitions. In the philosophical literature, it is widely accepted that cases of deceptive implicatures are misleading, but they are clearly not cases of lying (e.g., Davis 2005; Saul 2012; Stokke 2018; Viebahn 2019a).³ However, if the aim is to develop a definition of lying that captures ordinary people's use and understanding of this concept, the intuitions of philosophers might not be decisive (cf. Section 2) and might not even conform to the lay concept. Hence, we should probably wait for empirical support before we take it for granted that deceptive implicatures are clearly not cases of lying.

In addition to pumping intuitions about concrete examples, some philosophers argued for a general feature that deceptive implicatures intuitively share. Stokke (2016) and also Viebahn (2019a) argue that we can exclude deceptive implicatures from our definition of lying because someone who accused of lying by deceptive implicatures can convincingly reject the allegation. This, so they argue, is not possible when a speaker lies by asserting something false. Consider a dialogue of the following form:

³ It should be noted that since Stokke's view is sensitive to the question under discussion, some implicatures (according to the characterization in Section 3) could be classified as lies (see Stokke 2018: 82-83, for two examples).

Rebecca: You lied to me!

Dennis: No, why do you think so?

Rebecca: You claimed that p but believed/knew p to be false!

A very convincing way to ward off this accusation would be for Dennis to demonstrate that he did not, in fact, assert that p. For instance, let us assume that Rebecca finds out that Dennis went to a party and is angry because she mistakenly thinks that Dennis texted her to say that he did not go to the party. Dennis can easily dismiss this allegation by showing her that, in fact, he wrote that he did go to the party.

Some authors argue that, in cases of deceptive implicatures, a defense similar to the one just described is available to the person accused of lying by means of deceptive implicatures. For instance, Viebahn (2017, 2019a; see also Stokke 2016) has recently used the following scenario (Davis 2005) to claim that the accused person can respond by denying having claimed that p.

Party:

Dennis is going to Paul's party tonight. He has a long day of work ahead of him before that, but he is very excited and can't wait to get there. Dennis' annoying friend Rebecca comes up to him and starts talking about the party. Dennis is fairly sure that Rebecca won't go unless she thinks he's going, too.

(1) Rebecca: Are you going to Paul's party?

Dennis: No, I'm not going to Paul's party.

(2) Rebecca. Are you going to Paul's party?

Dennis: I have to work.

Viebahn (2019a) argues that after (2), Dennis could offer the following sincere response to Rebecca when she accuses him of lying after finding out that he went to the party:

I did not lie. I didn't claim that I wasn't going to go to Paul's party. I merely claimed that I had to work, which I did.

By contrast, Dennis cannot sincerely reply in this way after giving his answer in (1). In a nutshell, the argument for cases in which p is believed to be false seems to be: if a speaker asserts that p, then she is committed to the truth of p and cannot plausibly deny having claimed that p (cf. Michaelson, 2016). Therefore, she lied. By contrast, if p is only implicated, then the speaker is not – or at least not sufficiently strongly – committed to the truth of the content of the implicature (in short: committed to the implicature) and can thus plausibly deny having claimed that p. Therefore, she did not lie. In the following, we discuss whether this argument is plausible.

For our present purpose and the discussion to come, it will not be necessary to spell out exactly when speakers are committed to an implicature and what exactly it means to be committed (for an example of a spelled-out account, see Section 7.2). We assume that ordinary speakers have at least a rough understanding of this notion, according to which speakers who are committed to a certain content cannot plausibly deny having communicated this content and have to take responsibility for it (for an overview of the latter view in the context of assertion, see Pagin 2016). Furthermore,

it should be noted that the view that a speaker can be committed to the truth of a proposition conveyed by means of an implicature is not unheard of in the philosophical literature (e.g., Camp 2011; Garcia-Carpintero 2018; Geurts, 2019; Green, 2006; Marsili, 2014; Meibauer, 2014b; Wilson & Sperber, 1986).

Let us now take a closer look at the Party case. Rebecca and Dennis are friends and have a casual conversation in an everyday context. Hence, we should assume that both conversational partners are cooperative and interesting in meaningfully participating in the conversation. Rebecca now asks a very specific question, namely whether Dennis is going to Paul's party. Dennis could answer with yes or no, but he does not. Instead he tells Rebecca that he has to work. In this context, there seems to be only one reasonable interpretation of Dennis' answer. By saying that he has to work, Dennis is communicating in a very efficient way that he will not attend the party, and that the reason for his absence is that he has to work. Assuming cooperativeness, this is the only meaningful interpretation of the statement. If we denied that Dennis claimed that he will not go to the party, then we should be baffled by his response, because it would not be related to the question, let alone answer it. In more drastic words: Dennis's answer would not make any sense at all.

To illustrate this point, compare the Party case with the following, as discussed by Saul (2012). A dying woman asks the doctor whether her son is well. The doctor saw him yesterday, when he was well, but knows that he was later killed in an accident. Naturally, the doctor intends to keep this information from the dying woman in comfort her. So, he says *I saw your son yesterday and he was fine*, believing she will infer that her son is fine. In this case, and in contrast with the Party case, the doctor's utterance makes sense even if we do not derive the implicature from it. To see this, imagine that the doctor actually saw the son only before the accident (and does not know about the later accident). The dialogue *A: How is my son? – B: I saw your son yesterday and he was fine* seems perfectly fine and not odd at all. By contrast, Dennis's reply in the Party case would not address Rebecca's question if we do not derive the implicature that Dennis will not attend the party because he has to work. In such cases, in which the interpretational space left for the response is extremely narrowed down by both the yes/no format of the question and the response being complete uninformative on the level of what is said, it seems reasonable to hold that Dennis committed himself to the implicated claim that he will not go to the party, as no other meaning seems to answer the question at issue.

Rejecting the analysis above would make countless successful communications from everyday life appear very mysterious. Probably, most of us have witnessed or participated in conversations similar to the one Rebecca and Dennis were having – be it about attending a party, the next training session in one's sport club, or whatever. Excluding special circumstances – such

as people wanting to demonstrate that they have mastered the difference between implicating and asserting – it seems very unlikely that in such a context, *I have to work* will be considered anything but a conclusive answer. If, having received this reply, the questioner were to ask what this answer is supposed to mean, we would probably be baffled, and might even doubt her cognitive abilities. Moreover, our analysis of this case seems to be vindicated by understanding commitment to p as taking responsibility for the truth of p and making oneself liable to social sanctions in case p is false (cf. Peirce 1934). If, after responding *I have to work* the person shows up at the party and is criticized, then it seems at least questionable whether responding to the accusation with *I didn't claim that I wasn't going to go to Paul's party. I merely claimed that I had to work, which I did* would be a promising way to reduce social sanctions.

So perhaps a good case can be made for considering Dennis as having committed himself to not attending Paul's party. Although already hinted at above, it might be worth being more explicit about what we do and do not want to claim here. We neither claim that a speaker is always committed to an implicature to the same degree as in cases of assertion, nor that the speaker is always committed to an implicature to a similarly high degree as in the Party case. For instance, if a speaker says that she *believes* that p, the implicature seems to be that she does not *know* that p (Horn, 1984; see also our third vignette in the experiment). However, her commitment to not knowing that p might appear intuitively weaker than asserting that she does not know that p (and weaker than Dennis' commitment to not attending the party). What *we do claim* is that at least in some cases it seems reasonable to consider the speaker to be committed to an implicature to a degree that is sufficiently high to make the corresponding utterance a case of lying.

As indicated above, and in line with some authors (e.g., Marsili 2014, 2019; Mazzarella et al. 2018; Moeschler 2013) we assume that commitment comes in degrees. When a speaker asserts that p, she seems clearly committed to p, although even between different assertions the strength of commitment might vary (compare “p” versus “certainly/obviously/undoubtedly p”). When implicating p, the strength of commitment is intuitively (and on average) weaker than in cases of asserting p and it may vary strongly across different cases of deceptive implicatures. Crucially, we take it that in some cases the degree of commitment is sufficiently high for considering the speaker as having claimed that p. One kind of such cases might exhibit the structure of the Party case, in which a yes/no-question is asked and the answer would make sense without deriving the implicature, but there might be (many) other cases.

Besides addressing the more general question of whether ordinary speakers consider deceptive implicatures to be cases of lying, our experiment is designed to test the hypothesis that lie judgments for deceptive implicatures go hand in hand with the degree to which the speaker is considered to be committed to the implicature. Before we describe our experiment in more

detail, we will first briefly summarize the findings of previous studies that have empirically investigated lying by means of implicatures.

5. Empirical Studies on Deceptive Implicatures

Weissman & Terkourafi (2018) argue against the possibility of lying by deceptive implicature. They presented 20 participants with 15 cases of deceptive implicatures. The test question was whether the agent lied, and participants could indicate their judgment on a 7-point Likert-item ranging from 1 (“definitely not a lie”) to 7 (“definitely a lie”). Except for four cases, participants’ mean ratings were below the midpoint of the scale (labeled “neither a lie nor not a lie”), indicating that they did not judge the target utterances to be lies.

The results of Or et al. (2017) point in the other direction. In their experiment, 28 participants were asked to evaluate assertions, both with regard to whether the speaker told the truth, and also with regard to whether the speaker lied. Some assertions were literally true but carried a deceptive implicature, while others were just true. Participants could indicate their disagreement with the claims that the speaker told the truth and that the speaker lied on a 7-point Likert-item. The mean rating for the literally true but deceptive-implicature-including sentences was 3.2 and below the midpoint of the scale (4), indicating that these were considered to be lies, but not full-fledged lies.

Antomo et al. (2018) found similar results. They presented 30 participants with uncontroversial cases of lying (intentionally false assertions) and deceptive implicatures. Intentionally false assertions were judged to be clear cases of lying (median of 1 on a 5-point-scale), and to a greater extent than deceptive implicatures. However, participants also tended to evaluate deceptive implicatures as cases of lying (with a median of 2). The results of a second experiment testing children led to similar results.

Wiegmann & Willemsen (2017) obtained results that speak rather strongly for the view that deceptive implicatures count as lying. They ran four experiments focusing on violations of Grice’s maxims (Grice, 1989a). Overall, they tested eight cases of deceptive implicatures, each of which was presented to about 100 participants. Each case ended with the claim that the speaker lied to the hearer, and participants could either agree or disagree with this claim. It was consistently found that the majority of participants considered the speaker who used a deceptive implicature to be lying.

To sum up, there have only been a few studies that have empirically investigated lying by means of deceptive implicatures. The sample sizes are mostly small and the results differ widely. Consequently, we are a long way from having a clear picture of what is going on.

6. Experiment

The main goal of the current experiment is to clarify whether ordinary speakers believe that lying by means of deceptive implicatures is possible. To this end, we presented 200 participants with fourteen vignettes involving different kinds of deceptive conversational implicatures, and asked them to judge whether the protagonist lied.⁴ On the basis of previous findings (as summarized in the previous section), we expected lying ratings to differ between the fourteen cases. To better understand these differences, we additionally asked participants questions concerning potential proxy measures of lying, such as the probability of the addressee ending up with a false belief, and whether the utterance was factually false. We predicted that the best proxy measure would be the degree to which the speaker was perceived to be committed to the implicature.

6.1 Participants

In this experiment, participants were recruited on Prolific Academic and completed an online survey implemented in Unipark. All participants were required to be at least 18 years old, English native speakers, and to have an approval rate from previous studies on the platform of at least 90%. 402 participants started the survey, and the data of 371 were included in the analysis. 31 participants were excluded for finishing the experiment in less than three minutes, failing the attention check, or indicating severe technical issues when doing the survey. The mean age was 32.3 years, 43% were male, 56% female, and 1% were non-binary or preferred not to disclose their gender. Participants received £1.50 for an estimated 15 minutes of their time (£6/h).

6.2 Design, Procedure and Materials

We implemented a 4 (kind of deception: Lie vs. Deceptive Implicature_a vs. Deceptive Implicature_b vs. Truth) * 14 (vignettes: fourteen short stories modeled after popular and different kinds of false implicature) mixed design. The kind of deception was tested as a between-subjects factor, while the fourteen different vignettes were tested within subjects. Participants were randomly assigned to one of the four between-subjects conditions. In the Lie condition, participants evaluated prototypical cases of lying; in the Truth condition, they evaluated

⁴ One might object that our short vignettes do not guarantee that a deceptive conversational implicature is present, because they do not explicitly state that all the conditions necessary for a conversational implicature to arise (Grice 1989a: 30-31) are met. (See Saul 2002 for an illuminating discussion of these conditions, and whether they are entirely under the speaker's control.) Nevertheless, we believe that it is plausible to assume that our participants understood them as containing the deceptive implicature that we had in mind. First, directly before the speaker makes the crucial utterance, we describe why the speaker says what she says and what she wants to achieve with this (in contrast to Weissman & Terkourafi 2018, who do not make the speaker's motive explicit). Second, we make the implicature we have in mind explicit in our questions about the proxy measures (and the lie ratings do not significantly differ between the condition in which we first ask the question about the proxy measures and the condition in which the lie judgments were collected first). We thank an anonymous reviewer for raising this worry.

prototypical cases of telling the truth. The two Deceptive Implicature conditions differed in the order in which the relevant test questions were asked (see below for further details).

Participants first read general instructions to familiarize themselves with the task and the response formats. They were then presented with the fourteen scenarios in randomized order. Afterwards, they were presented with an attention check in the form of a vignette featuring a person who accidentally said something that turned out to be false, although the agent had very good reason to believe what he said to be true (we excluded participants who judged this agent to be lying). On the final page, participants were asked about their gender, age, and whether any technical or other issues occurred.

The fourteen vignettes were modeled after cases of conversational implicatures discussed in the pertinent literature (e.g., Grice 1989a and Levinson 2000; see Meibauer 2006 for an overview). We can classify these cases (i) along the GCI/PCI (generalized conversational implicature/particularized conversational implicature) distinction and (ii) by the Gricean maxims or principles involved.

Table 1: Classification of the stories used in the experiment.

	<i>Quantity</i>	<i>Relation</i>	<i>Manner</i>
<i>GCI</i>	1, 2, 3, 14	5	10, 11, 12
<i>PCI</i>	4, 13	6, 7, 8, 9	

Here is one of the fourteen cases we employed (for the wording of the remaining vignettes, see the appendix):

(1) Quantity—GCI—scalar <his, a> (Grice 1989a: 37-38)

Ben and Pia have been married for 10 years and are on their way to a party in Paris. In their hometown, word on the street is that their marriage is in a bad place. Coincidentally, Robert is in Paris for business at the same time. Robert detests Ben, as Ben got elected mayor of their town: a job Robert wanted to have. Robert happens to meet Ben and Pia one evening in the hotel bar. When Robert returns to their home town, he sees an opportunity to get Ben in trouble. So he says to his friends in the tennis club, “In the hotel bar I saw Ben kissing a woman!” In fact, Ben was kissing his wife Pia, and Robert knew that.

[Test prompt:]

Robert lied to his friend.

Please indicate the extent to which you disagree or agree with this claim.

completely disagree (1)-(2)-(3)-(4)-(5)-(6)-(7) completely agree

In the Lie condition, the deceptive implicatures were made explicit to transform the utterance into a prototypical lie. Hence, in the Lie condition the speaker in vignette (1) says:

“In the hotel bar I saw Ben kissing a woman—and the woman was not his wife!”

and in the Truth condition he tells his friends:

“In the hotel bar I saw Ben kissing a woman—it was Pia, his wife!”

In the two Deceptive Implicature and the Lie conditions, participants were presented with five additional questions:

Commitment: To what extent do you disagree or agree with the following claim? Robert committed himself to the claim that Ben was kissing a woman who was not his wife.

Suggestion: To what extent do you disagree or agree with the following claim? Robert suggested that Ben was kissing a woman who was not his wife.

Probability: After hearing what Robert said, how likely is it that Robert’s friends will come to believe that Ben kissed a woman who was not his wife?

Falsity: To what extent do you agree with the following claim? Robert’s statement was factually false.

Morality: How would you morally evaluate Roberts’s behavior in the story you just read?

These five questions were added in the two Deceptive Implicature conditions, to test which of these five measures would best predict lie ratings for the deceptive implicatures. As already hinted at in Section 4, we expected commitment ratings to do best. The other proxy measures were included for the following reasons. The *Suggestion* question is, so to say, a weaker version of the Commitment question and would track lie ratings more closely if lying only required that a speaker’s utterance suggested – rather committed herself to – a believed-false claim. The *Probability* question takes into account that causing a false belief in the addressee seems to be a main function of lying. However, even though intended, the speaker cannot know for sure that they will cause the false belief. The *Falsity* question tracks the idea that cases of deceptive implicatures might amount to lying because people consider a literally true utterance carrying a deceptive implicature as overall false. The *Morality* question was included because morality has been shown to influence a wide range of judgments (see Knobe 2010 for an overview). It is plausible that people are more likely to judge those utterances to be lies that they consider morally wrong.

The two Deceptive Implicatures conditions (a & b) differed only with respect to the order in which the lie-question and the five proxy measure-questions were presented. In Deceptive Implicatures_a (and in Truth and Lie), participants first saw the vignette together with the lie-question. After they provided their lie rating, the vignette was again presented together with the five proxy measure-questions. In Deceptive Implicatures_b, participants first saw the vignette together with five proxy measure questions (order randomized), and after they provided their judgments to the proxy measures questions, they saw the vignette together with the lie-question on the next page.

We implemented these two orders of presentation because we could not rule out that the participants’ proxy measure-ratings would influence their lie ratings (and vice versa). For instance, judging that a certain assertion was factually false might increase the probability of considering this assertion to be a lie; or first judging that a behavior was morally good might decrease the probability

of judging the corresponding assertion to be a lie. Alternatively, once participants have judged the assertion to be a lie, this might increase the probability of the five proxy measure-questions being answered in support of the lying judgment.

6.3 Results⁵

6.3.1 Deceptive Implicature_a versus Deceptive Implicature_b

We first tested, by means of a mixed ANOVA, whether lie ratings in the two Deceptive Implicature conditions were affected by the order in which the lie-question and the proxy measure-questions were presented, with the order of presentation as a between-subjects factor, and the vignettes as a within-subject factor. This was not the case. There was neither a main effect of order of presentation $F(1, 185)=1.94$; $p=.0165$; $\eta_p^2 = .011$, nor any interaction of order of presentation and vignette ($F(13, 2379)=1.61$; $p=.074$; $\eta_p^2 = .009$). Hence, we collapsed the data for the two Deceptive Implicature conditions (combined $n = 185$) in order to analyze the lie ratings.⁶

6.3.2 Lie ratings

The lie ratings for the three conditions and the fourteen cases are summarized in Figure 1. Unsurprisingly, lie ratings in the Lie condition were all significantly above the midpoint of the scale (all $p < .0001$) and close to the upper limit of the scale. Lie ratings in the Truth condition were significantly below the midpoint of the scale (all $p < .0001$) and close to the lower limit of the scale. The lie ratings for each of the fourteen deceptive implicatures lay between their corresponding prototypical variants. Ten out of the fourteen Deceptive Implicatures vignettes received mean lie ratings significantly higher than the midpoint of the scale (4), three were not significantly different from the midpoint, and one was significantly below it (see Table 2). Looking at the median ratings instead of the means, the result pattern (see Figure 2) is similar, with eleven medians greater than the midpoint, two being equal to the midpoint, and one being below it. Furthermore, in ten out of the fourteen cases the most frequently chosen option was the one indicating highest agreement with the claim that the agent lied (see Figure 2).

⁵ The data can be accessed here: https://osf.io/qma3x/?view_only=a23eb956cf4b4169aa380661735ed41c

⁶ An anonymous reviewer raised the worry that presenting fourteen cases to each participant could artificially increase lie ratings. To address this worry, we analyzed the data of participants' first rating in the deceptive implicature condition (in which the lying question was asked before the proxy measures; $N=93$), because these very first ratings are surely not affected by the demand characteristic mentioned by the reviewer. The average of all first ratings, which included all fourteen cases of deceptive implicatures due to the randomized design, was very similar to the overall average in this condition (including not only the first but all fourteen cases), namely 4.98 vs 4.91. We are thus confident that lie ratings were not artificially inflated by the within-design of our experiment.

Table 2: Summary statistics for the lying and commitment ratings in the fourteen Deceptive Implicature vignettes. T-values, p-values, and effect size (ES) refer to the test against the midpoint of the scale (4). Reported p-values are not adjusted for multiple comparisons, (adjusting them would not change the inferences; adjusted alpha-level = 0.05/14=0.004).

Scenario	Means	Std. Dev.	N	95% CI	t-value	p-value	ES (d)
1	5,05	1,95	185	[4,77; 5,33]	7,33	<.001	0,54
2	6,51	1,06	185	[6,36; 6,67]	32,14	<.001	2,36
3	3,14	1,88	185	[2,87; 3,41]	-6,21	<.001	-0,46
4	5,41	1,87	185	[5,14; 5,68]	10,28	<.001	0,76
5	5,18	1,67	185	[4,94; 5,42]	9,58	<.001	0,70
6	4,08	1,87	185	[3,81; 4,35]	0,59	.556	0,04
7	4,55	1,99	185	[4,26; 4,84]	3,78	<.001	0,28
8	4,46	2,03	185	[4,17; 4,75]	3,08	.002	0,23
9	3,77	2,00	185	[3,48; 4,06]	-1,58	.115	-0,12
10	4,21	1,97	185	[3,92; 4,50]	1,45	.147	0,11
11	5,26	1,66	185	[5,02; 5,50]	10,39	<.001	0,76
12	5,08	1,74	185	[4,83; 5,33]	8,43	<.001	0,62
13	5,30	1,60	185	[5,07; 5,54]	11,05	<.001	0,81
14	5,35	1,63	185	[5,11; 5,58]	11,20	<.001	0,82




Figure 1: Lie judgments for the fourteen vignettes in the three conditions. Error bars represent 95% CI. The scale ranged from 1 to 7, with “1” indicating strong disagreement with the claim that the agent lied and “7” indicating strong agreement.




Figure 2: Violin plot of the lie ratings for the fourteen vignettes in the Deceptive Implicature condition. White circles represent medians, black circles represent modes.

6.3.3 Proxy measures

To assess how strongly the mean lying ratings for the fourteen scenarios are associated with the mean ratings of each of the five proxy measures for the fourteen scenarios, we correlated the mean lying ratings from the condition in which the lie-question was always asked first (Deceptive Implicature_a; $n=93$) with the mean ratings of the five proxy measures in the condition in which the proxy measure-questions were always asked first (Deceptive Implicature_b, $n=92$), so that the resulting correlation was based on two independent samples, i.e. each participant either contributed a lying rating or a proxy measure rating. The mean ratings for both the five proxy measures and the lying ratings from these two conditions are summarized in Table 3. All except for the Falsity proxy measure exhibit a significant correlation with the lying ratings. The strongest correlation was found for Commitment ratings (see Figure 3), the agreement or disagreement with the claim that the agent committed herself to the implicature ($r = .87$, $p < .001$); followed by Suggestion ($r = .79$), Probability ($r = .68$), Morality ($r = -.58$) and Falsity ($r = .46$). Furthermore, the mean squared error (MSE) was lowest for Commitment (MSE = 0.41), followed by Morality (MSE = 0.60; inverted ratings), Probability (MSE = 0.93), Suggestion (1.00), and Falsity (4.48).⁷ Unsurprisingly,

⁷To assess how robust these rankings were, we bootstrapped one million samples (for lying $n = 93$, for proxy measures $n = 92$, with replacement) to see how often commitment obtains the highest correlation and the lowest MSE, respectively. Commitment had the highest correlation in 98.6% of all samples (suggestion in 1.4%) and in 91.1% the lowest MSE (Morality in 8.9%). The data and the script for the bootstrap procedure can be accessed here: https://osf.io/qma3x/?view_only=a23eb956cf4b4169aa380661735ed41c

commitment ratings for the cases of straightforward lies (i.e. believed-false assertions) were—similarly to lie ratings—consistently close to the ceiling (never below 6).

Table 3: Summary statistics (mean and standard deviation) for the lying and proxy measure ratings in the fourteen Deceptive Implicature vignettes. Correlation values refer to the correlation of the respective proxy measure ratings with the lying ratings for the fourteen Deceptive Implicature cases. Reported p-values are not adjusted for multiple comparisons.

Scenario	Lie		Commitment		Suggestion		Probability		Falsity		Morality	
	mean	std	mean	std	mean	std	mean	Std	mean	std	mean	std
1	5,29	1,91	5,80	1,37	6,29	1,20	6,21	0,91	2,11	1,82	1,60	0,98
2	6,57	1,00	6,28	1,08	6,35	1,40	6,37	0,98	5,84	1,82	2,50	1,43
3	2,98	1,82	3,52	1,82	3,70	1,80	4,10	1,78	2,83	1,91	4,57	1,22
4	5,35	1,61	5,64	1,50	5,86	1,48	5,45	1,63	2,97	1,82	2,97	1,45
5	5,62	1,74	6,12	1,19	6,54	0,75	6,07	1,26	1,70	1,40	2,86	1,12
6	3,90	1,99	5,35	1,37	5,63	1,24	5,30	1,22	2,91	1,91	2,72	1,23
7	4,73	2,02	5,37	1,61	5,77	1,55	6,30	0,93	1,92	1,93	3,88	1,24
8	4,46	2,13	4,71	2,05	5,14	1,99	4,91	1,97	1,62	1,38	2,60	1,45
9	3,97	2,07	4,87	1,76	5,46	1,54	5,66	1,34	2,16	1,63	2,73	1,25
10	4,25	1,98	5,25	1,69	5,89	1,38	5,64	1,37	3,59	2,26	3,87	1,25
11	5,27	1,67	5,89	1,46	6,29	1,10	5,91	1,08	4,73	1,91	3,34	1,11
12	5,37	1,67	5,62	1,56	6,26	1,23	6,24	1,24	3,02	2,06	2,86	1,19
13	5,43	1,56	5,30	1,74	5,65	1,36	5,26	1,40	4,39	1,96	2,34	1,31
14	5,56	1,58	5,43	1,43	5,68	1,37	5,36	1,46	3,70	1,98	2,05	1,26
Correlation with lying			0,87		0,79		0,68		0,46		-0,58	
p-value for correlation			<.001		.001		.007		.094		.029	




Figure 3: Lie and Commitment ratings for the fourteen scenarios ($r=.87$). The lie-statement had the form ‘[Speaker] lied to [addressee]’ and the commitment-statement ‘[Speaker] committed herself to the claim that [implicature]’.

To further assess the robustness of commitment as the best proxy measure, we zoomed in to the participant's level.⁸ Commitment was confirmed to have the strongest relationship with the lying ratings ($r = .74$), followed by Suggestion ($r = .62$), Morality ($r = -.55$), Falsity ($r = .52$), and Probability ($r = .36$). Moreover, using – mutatis mutandis – the same method at the vignette level, Commitment was again the best proxy measure ($r = .5$), followed by Falsity ($r = 0.47$), Morality and Suggestion ($r = -.37$ and $r = .37$, respectively), and Probability ($r = .16$).

6.3.4 Individual response pattern

An exploratory analysis of lying ratings at the level of individual participants reveals that agreement ratings differed quite a lot across the fourteen cases for most participants. In fact, not a single participant provided the same rating for all cases. On the contrary, the difference between the highest and the lowest rating (henceforth: range) was 6 for many participants (78 out of 185, 42%), i.e. they strongly agreed (“7”) and strongly disagreed (“1”) at least once. Moreover, 28% (51/185) had a range of 5, and for 19% (36/185) the range was 4. By contrast, only 11% (20/185) had a range of 3 or lower (6% (12/185) had a range of 3, 4% (7/185) had a range of 2, and only one participant had a range 1).⁹

7. Discussion

In our experiment, deceptive conversational implicatures were predominantly judged to be lies (ten out of fourteen cases). The best proxy measure for lying ratings was the degree to which speakers were perceived to have committed themselves to the truth of the conversational implicature. The stronger they do, the more likely their utterance is to be considered a lie. These results indicate that ordinary speakers consider it possible to lie by means of deceptive implicatures, and they offer a potential explanation for the diverse findings in the literature (see Section 5). Moreover, the findings suggest that a definition of lying based on the notion of commitment might be promising.

7.1 Limitations and Objections

Before we turn to the implications of our study for the debate on the definition of lying, we would like to address some potential objections to and limitations of our study. First, in our experiment, as in most experimental approaches, rating scales were used. However, it is perhaps unclear whether a 7-point scale is the right device for representing participants' intuitions about lying. A

⁸ Here, the correlation for each participant ($n=185$) across the fourteen scenarios was calculated and z-transformed. The average of the z-transformed values was then transformed back to get the average correlation on the level of participants. Correlations on the level of participants are—as in this case—usually lower on the level of participants than on the level of mean ratings, because in the latter idiosyncrasies carry less weight and can cancel each other out.

⁹ The individual response pattern of all participants can be seen in the form of a short video here: https://osf.io/qma3x/?view_only=a23eb956cf4b4169aa380661735ed41c

scale seems to suggest or even presuppose that lying comes in degrees, rather than being a yes/no decision between “The utterance was a lie” and “The utterance was not a lie” (Rutschmann & Wiegmann 2017). Since there are arguments for the scalar nature of lying (Marsili 2014, 2019; see also Coleman & Kay 1981, who consider lying to be a prototype concept), a scalar measure seems appropriate. Moreover, a recent study (Viebahn et al. 2020) found that binary and degree-measures of lying judgments correlate virtually perfectly, suggesting that they both measure the same aspects in an experimental study.

Second, one might worry that the participants who judged a speaker using a deceptive implicature to be lying only did so because we did not offer – in addition to “The speaker lied to the addressee” – more fine-grained response options such as “The speaker misled the addressee”. This is a reasonable worry that we cannot refute with the data we collected in the current experiment. Fortunately, though, data from a study (unpublished author ms) that included such options still found that most cases of deceptive implicatures were considered to be cases of lying (see Viebahn et al. 2020, who found that including this option slightly lowered lie attributions).¹⁰

Third, the vignettes that we used all represent cases discussed in the Gricean, neo-Gricean and post-Gricean literature in which they serve as examples for various kinds of implicatures. However, it is not clear whether, or to what extent, these single vignettes are representative of an implicature phenomenon. We can safely say, however, that for a certain cover story that belongs to a certain class of implicatures, we find a specific effect. Nevertheless, it would be premature to conclude that any other cover story belonging to that class would yield the same results.

7.2 Implications for the definition of lying

The results of our experiment suggest that the traditional definition of lying is too narrow to capture ordinary speakers’ judgments that classify some deceptive implicatures as cases of lying. If one agrees that a crucial aim of the definition of lying is to capture the folk concept, one possible reaction to such findings would be to extend the traditional definition to include all deceptive implicatures (cf. Meibauer 2005; 2014a: 125). This would lead us to the following broader definition of lying (the “or” is inclusive):

A lies to B if and only if

there is a proposition *p* such that

1. A asserts that *p* to B, and
2. A believes that *p* is false

¹⁰ As a very helpful reviewer (no irony here) noted that it is still possible that had we included an additional question asking whether the speaker misled the addressee, the lying ratings in our study would be significantly lower. While we consider this possibility not to be very likely, we cannot rule this out and just want to add that we chose to not include the additional question in order to keep this study simple and comparable to previous studies (e.g., Weissman & Terkourafi 2018).

or there is an implicature q such that

1. A implicates that q to B , and
2. A believes that q is false

However, this definition might be too broad, because of the empirical finding that some cases of deceptive implicatures were not judged to be instances of lying by laypeople. Furthermore, it would not be able to account for a salient feature of our results. While participants overall tended to agree that speakers can lie by means of deceptive implicatures, the agreement ratings were always significantly higher for the corresponding deceptive assertions. This gap is mirrored by the degree to which participants perceived the speaker to have committed herself to the believed-false claim. As briefly mentioned above, when the speaker lied in a straightforward way (i.e. asserted a believed-false claim), commitment ratings were – as one would expect – consistently close to the ceiling of the rating scale and always higher than in corresponding cases of deceptive implicatures. Furthermore, commitment ratings also tracked the finding that the lying ratings for the fourteen cases of deceptive implicatures differed to quite some extent, as indicated by the result that commitment exhibited the highest correlation with the lying ratings and the lowest MSE. In light of these findings, understanding lying as commitment to a claim that is believed to be false seems promising. Perhaps a definition along the lines of Viebahn’s (2019a: 13) is worth considering¹¹:

A lies to B if and only if there is a proposition p such that:

1. A performs a communicative act C addressed to B with the content p ;
2. by performing C , A commits herself to p ; and
3. A believes that p is false.¹²

Obviously, in order to establish a commitment-based definition of lying, several issues need to be addressed. The two most pressing questions are the following. What exactly does it mean to be committed to the truth of a proposition? And what are the features of an implicature that decide whether or not a speaker is committed to its content? The first question has been extensively

¹¹ It is worth noting that although this definition is compatible with our findings, Viebahn (2019) considers neither Dennis (in the party scenario) nor Sally (in the smoking scenario) to be lying.

¹² To incorporate the finding that participants’ agreement ratings for the claim that a speaker lied by means of deceptive implicatures differed to quite some extent between the fourteen cases (see Section 6.4.3), one could add a clause allowing for more or less clear lies, depending on how strongly the speaker committed herself to the deceptive content of her utterance. Moreover, the insincerity condition (3), which is not our focus here, might need to be adapted in light of empirical findings (see Marsili 2019, for how the insincerity condition could be modified).

discussed in the context of assertion (see Pagin 2016 for an overview), and has been the focus of recent discussions (see the special issue in *Theoretical Linguistics* 45 (1-2), 2019). Given our empirical approach, it will not come as a surprise that we believe that answering the question of what commitment means in this context requires further empirical studies, which is why we will refrain for now from proposing a specific account.

Nevertheless, it is probably worth considering how commitment has been spelled out in the literature.¹³ We have chosen to present Viebahn's (2021) recently proposed account, for the following reasons. First, it is a variant of a view endorsed by several authors, namely the idea that commitment to *p* entails a certain justificatory responsibility to defend *p* if challenged (see Pagin 2016, for an overview of such accounts). Moreover, since Viebahn's account has been developed in the course of defending a commitment-based definition of lying that includes certain implicatures, it might help us to address the second question (about the features of an utterance that determine whether the speaker is committed to the implicated content). This account is described as follows:

The notion of commitment [...] is to be understood as the commitment in assertion [...]. By requiring liars to commit themselves to a proposition they believe to be false, [...] [this notion of commitment] requires them to take on a justificatory responsibility to defend that they *know* the proposition put forward. And this justificatory responsibility, in turn, is based on the absence of consistent dismissals to knowledge-challenges regarding the proposition in question.

Viebahn's (2021) account not only provides a characterization of what it means to be committed to *p*, but additionally offers a criterion for determining whether a speaker committed herself to a certain proposition.

The Commitment in Lying

In performing a communicative act *C* with a proposition *p* as its content, a speaker *A* commits herself to *p* (in the sense relevant for the lying-misleading distinction) iff *A* cannot consistently dismiss an audience challenge in response to *C* to defend (or justify) that she knows that *p*.

While it is beyond the scope of this paper to closely examine Viebahn's criterion, it seems fair to note two things about it. First, it seems to make correct predictions for several kinds of implicatures. For instance, it seems plausible that a speaker cannot consistently dismiss an audience challenge in response to *C* to defend (or justify) that she knows that *p* in cases of substitute implicatures, such as saying *I've got tomatoes coming out of my ears* to communicate that one has

¹³ In putting forward Viebahn's (2021) we neither commit ourselves to his account and neither do we claim that it is able to account for all of our findings (thanks to an anonymous reviewer prompting us to make these points explicit).

obtained a good crop (Viebahn 2021). Second, it seems that the question of whether the speaker can consistently dismiss the challenge is not easy to answer. For instance, let us consider again the case of Robert saying *In the hotel bar I saw Ben kissing a woman!* and assume that someone challenges Robert by asking *Do you really know that he did not kiss his wife?* Can Robert consistently dismiss this challenge by saying something like *I did not claim that he kissed someone who wasn't his wife. I merely claimed that he kissed a woman?* The answer to this question does not seem to be straightforward (and Viebahn himself does not deal with GCIs). This need not be a bad thing, because it might reflect the finding that participants' judgments for deceptive implicatures are not as clear as they are for the corresponding assertions. However, there might be easier-to-identify features of an utterance that determine whether the speaker committed herself to an implicature. Again, further empirical studies seem the way to go, but let us make some educated guesses here. First, as discussed in Section 4, one kind of cases in which a speaker is considered to have committed herself to the truth of an implicature might be if her utterance would not be relevant at all without deriving the implicature. Without deriving the implicature, replying *I have to work* to the question *Are you coming to the party?* would not be relevant to the question at all (Party case). The same holds for the answer *Evelyn was here yesterday, and she is a smoker* to the question of who the cigarettes belong to (Cigarette case). By contrast, consider case (3). Here, *answering I believe that he has a girlfriend* seems relevant to the question of how chances are for getting together with the man in question, even without deriving the implicature that the speaker does not know it. Second, a speaker might commit herself to the truth of an implicature if the implicature corresponds to a strong default interpretation. For instance, the empirical evidence (in our study but also in Weissman & Terkourafi 2018) that deceptive cardinals are cases of lying might be explained by the fact that, in contexts where the exact number of something is important, utterances like *have two children* seem to be naturally understood as *exactly two* (cf. relevance theorists, Carston 2002; Sperber & Wilson 1986, who consider GCIs as explicatures and belonging to the explicit content expressed). A further candidate for when speakers are committed to implicatures might be cases in which the implicature is a substitute for what is said, as in metaphors and hyperbole (cf. Viebahn 2017). For instance, Carl saying, in (11), that the kids were bouncing and bouncing and bouncing seems to be just a substitute for saying that the kids were bouncing for a long time.

Let us conclude the discussion with addressing a potential worry that could arise at this point. One might be skeptical that without specifying commitment in the context of lying, any real progress has been made by the current study. However, such an assessment would neglect important findings and features of the study. First, the high number of participants and vignettes in the current study arguably provides the strongest evidence that ordinary people consider some cases of deceptive implicatures as cases of lying. Second, concerning the relations of lying and

several proxy measures, the finding that commitment best tracks lying ratings is not trivial or obvious, even though the view that commitment is important for lying can already be found in the literature (Meibauer 2014b; Viebahn 2017, 2019a). The other proxy measures we investigated were arguably not any less plausible from the start. For instance, if falsity had been the best proxy measure, there would have been a very natural explanation of why one can lie with deceptive implicatures: they can affect truth evaluations – and the more strongly they render an utterance factually false, the more likely the utterance is to be considered a case of lying. Moreover, assuming that the point of lying is to cause a false belief in the addressee, it would not be surprising if the probability of doing so would be the best proxy measure of lying. And taking into account the seminal work in experimental philosophy done by Knobe (2003, 2010), it would not have been shocking if our findings had revealed that moral evaluations most closely track lie judgments. Hence, although there is still a long way to go, the findings of the current study might point towards the right direction.

To sum up, the current study provides strong evidence that ordinary speakers consider a wide range of deceptive implicatures to be cases of lying and suggests a link between lying and being perceived to be committed to a believed false claim. The latter finding suggests that defining lying in terms of commitment might be promising. However, before a commitment-based definition of lying can be established further theoretical and empirical research is needed to clarify the concept and its determinants.

References

- Adler, Jonathan E. (1997). Lying, Deceiving, or Falsely Implicating. *The Journal of Philosophy*, XCIV, 435–452. doi: 10.2307/2564617
- Antomo, Mailin, Susanne Müller, Katharina Paul, Markus Paluch, and Maik Thalmann (2018). When Children Aren't More Logical Than Adults: An Empirical Investigation of Lying by Falsely Implicating. *Journal of Pragmatics*, 138, 135-148. doi: 10.1016/j.pragma.2018.09.010
- Bok, Sissela (1999). *Lying. Moral Choice in Public and Private Life*. Vintage Books.
- Camp, Elizabeth (2012). Sarcasm, Pretense, and The Semantics/Pragmatics Distinction. *Noûs*, 46(4). doi: 10.1111/j.1468-0068.2010.00822.x
- Carson, Thomas L. (2006). The Definition of Lying. *Noûs*, 40(2), 284–306. doi: 10.1111/j.0029-4624.2006.00610.x
- Carson, Thomas L. (2010). *Lying and Deception: Theory and Practice*. Oxford University Press.
- Carston, Robyn (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Blackwell.
- Chisholm, Roderick M. and Feehan, Thomas D. (1977). The Intent to Deceive. *The Journal of Philosophy*, 74(3), 143–159. doi: 10.2307/2025605
- Coleman, Linda and Paul Kay (1981). Prototype Semantics: The English Word *lie*. *Language*, 57, 26–44. doi: 10.1353/lan.1981.0002
- Davis, Wayne (2005). Implicatures. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/sum2005/entries/implicature/>
- Fallis, Don (2009). What is Lying? *The Journal of Philosophy*, 106(1), 29–56. doi: 10.5840/jphil200910612
- Fallis, Don (2012). Lying as a Violation of Grice's First Maxim of Quality. *Dialectica*, 66(4), 563–581. doi: 10.1111/1746-8361.12007
- Faulkner, Paul (2007). What is Wrong with Lying? *Philosophy and Phenomenological Research* 75(3), 535–557. doi: 10.1111/j.1933-1592.2007.00092.x
- García-Carpintero, M. (2018). Sneaky assertions. *Philosophical Perspectives*, 32(1), 188–218. doi: 10.1111/phpe.12116
- Gerken, Mikkel (2018). Truth-Sensitivity and Folk Epistemology. *Philosophy and Phenomenological Research*. Advance online publication.
- Geurts, Bart (2019). Communication as Commitment Sharing: Speech Acts, Implicatures, Common Ground. *Theoretical Linguistics* 45(1-2), 1–30.
- Green, Stuart P. (2006). *Lying, Cheating, and Stealing. A Moral Theory of White-Collar Crime*. Oxford University Press.

- Grice, Paul (1989a). Logic and Conversation. In Paul Grice (1989), *Studies in the Way of Words* (22–40). Harvard University Press.
- Grice, Paul (1989b): Further Notes on Logic and Conversation. In Paul Grice (1989), *Studies in the Way of Words* (41-57). Harvard University Press.
- Horn, Laurence (1984). Toward a New Taxonomy for Pragmatic Inference: Q-Based and R-Based Implicature. In Deborah Schiffrin (Ed.), *Meaning, Form, and Use in Context: Linguistic Applications* (11–42). Georgetown University Press.
- Knobe, Joshua (2003). Intentional Action and Side Effects in Ordinary Language. *Analysis*, 63(3), 190–194. doi: 10.2307/3329308
- Knobe, Joshua (2010). Person as Scientist, Person as Moralist. *Behavioral and Brain Sciences*, 33(4), 315–329. doi: 10.1017/S0140525X10000907
- Levinson, Stephen (2000). *Presumptive Meanings. A theory of Generalized Conversational Implicature*. MIT Press.
- Mahon, James E. (2016). The Definition of Lying and Deception. *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/win2016/entries/lying-definition/>
- Marsili, Neri (2014). Lying as a Scalar Phenomenon: Insincerity along the Certainty-Uncertainty Continuum. In Sibilla Cantarini, Werner Abraham, and Elisabeth Leiss (Eds.), *Certainty-uncertainty – and the Attitudinal Space in Between* (154–173). Benjamins.
- Marsili, Neri (2019). Lying and Certainty. In Jörg Meibauer (Ed.), *The Oxford Handbook of Lying* (170–182). Oxford University Press.
- Mazzarella, Diana, Robert Reinecke, Ira Noveck, and Hugo Mercier (2018). Saying, Presupposing, and Implicating: How Pragmatics Modulates Commitment. *Journal of Pragmatics*, 133, 15–27.
- Moeschler, Jacques (2013). Is a Speaker-Based Pragmatics Possible? Or How Can a Hearer Infer a Speaker's Commitment? *Journal of Pragmatics*, 48(1), 84–97.
- Meibauer, Jörg (2005). Lying and Falsely Implicating. *Journal of Pragmatics*, 37, 1373–1399. doi: 10.1016/j.pragma.2004.12.007.
- Meibauer, Jörg (2006). Implicature. In Keith Allan (Ed.), *Encyclopedia of Language and Linguistics* Vol. 5 (2nd ed., 568–580). Elsevier.
- Meibauer, Jörg (2008). Tautology as Presumptive Meaning. *Pragmatics & Cognition*, 16(3), 439–470. doi: 10.1075/pc.16.3.02mei
- Meibauer, Jörg (2011). On Lying: Intentionality, Implicature, and Imprecision. *Intercultural Pragmatics*, 8(2), 277–292. doi: 10.1515/iprg.2011.013
- Meibauer, Jörg (2014a). *Lying at the Semantics-Pragmatics Interface*. De Gruyter Mouton.

- Meibauer, Jörg (2014b). A Truth That's Told with Bad Intent. Lying and Implicit Content. *Belgian Journal of Linguistics*, 28, 97–118. doi: 10.1075/bjl.28.05mei
- Meibauer, Jörg (2019). Lying, Implicating, and Presupposing. In: Jörg Meibauer (Ed.), *The Oxford Handbook of Lying* (193–202). Oxford University Press.
- Michaelson, E. (2016). The Lying Test. *Mind & Language*, 31(4), 470–499.
- Or, Shirly, Mira Ariel, and Orna Peleg (2017). The Case of Literally True Propositions with False Implicatures. In Innocent Chiluba (Ed.), *Deception and Deceptive Communication: Motivations, Recognition Techniques and Behavioral Control* (67–108). Nova Science.
- Ortony, Andrew and Swati Gupta (2019). Lying and Deception. In Jörg Meibauer (Ed.), *The Oxford Handbook of Lying* (149–169). Oxford University Press.
- Pagin, Peter (2016). Assertion. *The Stanford Encyclopedia of Philosophy*.
<https://plato.stanford.edu/archives/win2016/entries/assertion/>
- Peirce, Charles S. (1934). *Collected papers of Charles Sanders Peirce*. Harvard University Press.
- Rutschmann, Ronja and Alex Wiegmann (2017). No Need for an Intention to Deceive? Challenging the Traditional Definition of Lying. *Philosophical Psychology*, 30(4), 434–453. doi: 10.1080/09515089.2016.1277382
- Saul, Jennifer M. (2002). Speaker Meaning, What Is Said, and What Is Implicated. *Nous* 36(2), 228–248.
- Saul, Jennifer M. (2012). *Lying, Misleading, and What Is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford University Press.
- Sorensen, Roy (2007). Bald-Faced Lies! Lying without the Intent to Deceive. *Pacific Philosophical Quarterly*, 88(2), 251–264. doi: 10.1111/j.1468-0114.2007.00290.x
- Sperber, Dan and Deirdre Wilson (1986). *Relevance: Communication and Cognition*. Harvard University Press.
- Stokke, Andreas (2016). Lying and Misleading in Discourse. *Philosophical Review*, 125(1), 83–134. doi: 10.1215/00318108-3321731
- Stokke, Andreas (2018). *Lying and Insincerity*. Oxford University Press.
- Viebahn, Emanuel (2017). Non-Literal Lies. *Erkenntnis*, 82(6), 1367–1380. doi: 10.1007/s10670-017-9880-8
- Viebahn, Emanuel (2019a). Lying with Presuppositions. *Nous*. Advance online publication. doi: 10.1111/nous.12282
- Viebahn, Emanuel (2019b). Lying with Pictures. *The British Journal of Aesthetics*, 59(3), 243–257. doi: 10.1093/aesthj/ayz008

- Viebahn, Emanuel (2021). The Lying-Misleading Distinctions: A Commitment-Based Approach. To appear in: *The Journal of Philosophy*.
- Viebahn, Emanuel, Alex Wiegmann, Nele Engelmann, and Pascale Willemsen (2020). Can a Question be a Lie? An Empirical Investigation. Preprint available at: <https://osf.io/jfyn8>.
- Weissman, Benjamin and Marina Terkourafi (2018). Are False Implicatures Lies? An Empirical Investigation. *Mind & Language*, 1–26. doi: 10.1111/mila.12212
- Wiegmann, Alex and Jörg Meibauer (2019). The Folk Concept of Lying. *Philosophy Compass*, 2019, e12620. doi: 10.1111/phc3.12620
- Wiegmann, Alex and Pascale Willemsen (2017). How the Truth Can Make a Great Lie: An Empirical Investigation of Lying by Falsely Implicating. *CogSci* 2017, 3516–3521.
- Wilson, Deirdre and Dan Sperber (1986). Inference and Implicature. Reprinted in Steven Davis (Ed.), 1991, *Pragmatics: A Reader* (377–394). Oxford University Press.