

Forthcoming in: A. Wiegmann (ed): *Lying, Fake News, and Bullshit*. Bloomsbury: Advances in Experimental Philosophy.

What does it take to tell a lie?

Emanuel Viebahn

Freie Universität Berlin

Abstract

Lying requires asserting a disbelieved proposition, that much is widely accepted in the debate on how to define lying. But what else is required? Does lying require a particular linguistic manner of expression, such as saying? Does the proposition asserted have to be false (and not merely disbelieved)? And does lying require an intention to deceive? The aim of this chapter is to provide an opinionated introduction to the debates on these questions that takes into account both theoretical considerations and empirical data.

1. Introduction

What does it take to tell a lie? That is the question of a vibrant philosophical debate stretching back at least to Augustine's *De mendacio* ([395]/2013). Naturally, there is no agreement on how it should be answered – there is no undisputed definition of lying. But there are three assumptions that most participants to the debate accept. Firstly, it is widely agreed that the target concept of lying does not cover all insincere communicative acts, and that it contrasts with the concept of misleading. Secondly, it is standardly held that telling a lie requires asserting something one disbelieves. And, thirdly, there is a consensus that, insofar possible, definitions of lying should respect the intuitions of ordinary speakers on what it takes to tell a lie.

In this chapter, I will take these three basic assumptions as starting point to introduce three of the most hotly contested questions concerning the concept of lying: Does lying require a particular linguistic manner of expression, such as saying? Does the proposition asserted have to be false (and not merely disbelieved)? And does lying require an intention to deceive? My aim is to provide an opinionated introduction to the debates on these questions that takes into account both theoretical considerations and empirical data. I will begin by introducing and illustrating the shared

assumptions (Section 2), before devoting a section each to the questions about manner of expression (Section 3), actual falsity (Section 4) and deceptive intent (Section 5).

2. Three widely shared assumptions in the debate on how to define lying

To get started, let us consider the three basic assumptions in a bit more detail. While none of them is uncontested, they are widely accepted, and can thus provide a stable footing for the discussion of the more controversial issues of the following sections.

The first assumption concerns the target of the debate. Which notion is it that philosophers and linguists are trying to capture? While the expression ‘lying’ is sometimes used in a broad sense that applies to many different kinds of insincerity (cf. Saul 2012: 1), it is widely agreed that the target concept of lying does not cover all insincere communicative acts. In particular, it is commonly held that the target concept of *lying* contrasts with the concept of *misleading* or, to be more precise, the concept of *performing an intentionally misleading utterance*. (Henceforth I will pick out the latter concept with the simple label *misleading*). Here is an example that brings out this contrast (cf. Saul 2012: 70, Viebahn 2021: 289):

A dying woman asks a doctor whether her son is well. The doctor saw the son yesterday, when he was fine, but knows that he was killed shortly afterwards. The doctor wants to spare the dying woman the news of her son’s death. She utters:

Version A: (1) He’s fine.

Version B: (2) I saw him yesterday and he was fine.

Most theorists (and ordinary speakers) would say that the doctor lies by uttering (1), while her utterance (2) is misleading, but not a lie. Accordingly, the target notion of the debate is the narrow concept of lying exemplified by (1), which contrasts with the concept of misleading exemplified by (2). While some theorists have dissented on this point (see, e.g., Bolzano [1810]/2007, Meibauer 2014), I will henceforth assume that the relevant concept of lying is one that is narrow enough to contrast with the concept of misleading.

The second common assumption is that lying requires asserting a disbelieved proposition. This assumption gives us two ingredients generally thought to be necessary for lying: *the communicative act of assertion* and *disbelief of the proposition asserted*. There are different ways of spelling out these ingredient concepts. Does asserting a proposition (in the sense relevant for defining lying) require uttering a sentence that semantically expresses that proposition? And when does an agent count as disbelieving a proposition? Must they take the proposition in question to be *false*, or does it suffice that

they *do not believe it to be true*? These questions shall be left open here, in order to state the second assumption as a requirement on lying that most theorists in the debate would accept:

Lying requires asserting a disbelieved proposition.

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

(L1) A asserts *p* to B, and

(L2) A disbelieves *p*.

Returning to our initial example, it is easy to see that utterance (1) fulfils both clauses of this requirement: the doctor asserts a proposition she disbelieves. With utterance (2), by contrast, the doctor puts forward a proposition she believes to be false (namely, that the woman's son is fine), but does not assert this proposition.¹ And the doctor believes the proposition that she does assert (namely, that she saw the son yesterday and he was fine) to be true. Utterance (2) thus involves no proposition that fulfils both (L1) and (L2). Again, there are some theorists who disagree even with this fairly minimal requirement (cf. Fox Krauss 2017), but I will take it as given in what follows, in order to focus on the more contentious issues.

The third point of agreement is methodological. How should we go about finding an adequate definition of lying? Given that the debate is mostly conducted by philosophers, uncontroversial methods of investigation involve theoretical arguments, thought experiments and examples, as well as expert intuitions about these examples. But the debate on how to define lying has also been one of the first debates in contemporary philosophy in which the intuitions of ordinary speakers have been widely accepted as important data points. Thomas Carson made this attitude explicit as early as 2006:

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. (Carson 2006: 301)

And since then, there has been little dissent and plenty of agreement when it comes to the importance of ordinary intuitions.² Of course, the intuitions of ordinary speakers are not everything.

¹ The expression *putting forward a proposition* is meant to cover different ways of trying to communicate a proposition, including stronger communicative acts, such as asserting, and weaker communicative acts, such as suggesting.

² For explicit agreement, see e.g. Fallis (2009: 32) and Saul (2012: vii). For dissent, see Dynel (2011: 151) and Harris (2020: 7).

On the one hand, there are often different ways to interpret experimental findings, and theoretical considerations and expert intuitions play a role in guiding such interpretation. On the other hand, there may be areas in which ordinary intuitions turn out to be inconsistent, in which case a decision must be made on which of these intuitions to respect in proposing a definition of lying. As is typical for empirically informed debates in philosophy, the aim is thus to arrive at a reflective equilibrium between empirical data (intuitions) and theoretical considerations. In line with this, I will discuss both theoretical considerations and empirical data points in the next sections. Empirical data is of particular importance in cases in which expert intuitions diverge, and we will see below that there are several questions in the debate where folk intuitions can help to resolve expert disagreements.

3. Does lying require a particular manner of expression?

Lying, we are assuming, requires asserting a disbelieved proposition. But what does it take to assert a proposition? That is one of the main points of contention in the debate on how to define lying. It is also the question that separates views with respect to how they hold apart utterances that are lies and utterances that are merely misleading, given the plausible assumption that misleading contains all the same ingredients as lying *sans* asserting.

The most common answer to the above question makes use of the concept of *saying*. Lying requires *saying* a disbelieved *proposition*, and *saying a proposition*, in turn, requires uttering a sentence that has the proposition in question (or a very closely related proposition) as its semantic content. Views of this kind have been defended by Fallis (2009), Carson (2010), Saul (2012) and Stokke (2018), amongst others. Given that these definitions tie lying to semantically expressed *content*, they can be labelled *content-based definitions* (cf. Viebahn 2021: 290).

Let us briefly focus on Andreas Stokke's definition of lying as a concrete example:

Stokke's definition of lying (Stokke 2018: 31)

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

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

Stokke's view of assertion has two components, given by (L1) and (L2). The first component is the saying-account of assertion just introduced, where what is said is 'strictly constrained by a minimal kind of compositional meaning' (Stokke 2018: 81). The second component is Stalnaker's (1999, 2002) view that the essential effect of asserting is a proposal to update the common ground of the conversation. Importantly, Stokke holds that conversational implicatures do not belong to what is

said. This means that, by uttering (2), the doctor does not count as *saying* that the woman's son is fine (though she puts this proposition forward as a conversational implicature). As a result, Stokke's definition correctly classifies the doctor's utterance of (2) as a non-lie and her utterance of (1) as a lie.

While the other theorists defending content-based definitions of lying disagree with the second part of Stokke's view of assertion (and thus would phrase (L2) differently), they concur with the first part and with spelling out what is said narrowly so that conversational implicatures do not count as said. Accordingly, they arrive at the same results for the doctor's utterances of (1) and (2).

By tying the concept of lying to the concept of saying, and by spelling out saying in a narrow way, definitions in this group entail that lying requires a particular manner of expression: the content of a lie has to be semantically expressed through the use of linguistic expressions. This precludes the possibility of lying with other manners of expression; in particular, it precludes (i) the possibility of lying with conversational implicatures and (ii) the possibility of lying with non-linguistic communicative acts. However, it has recently been argued that these are two ways in which content-based definitions of lying are too narrow, as I now want to indicate. First, I will consider whether conversational implicatures can be used to lie, before then turning to the possibility of lying with pictures.

3.1 Lying beyond what is said?

Does the content of a lie have to be *said* (in the sense just outlined)? Paradigmatic cases of lying, such as (1), suggest a positive answer to this question. And indeed, it is plausible that in many cases of lying the content of the lie does align with the semantic content of the sentence uttered. But it has been argued that this need not always be the case. Interestingly, one of the central cases that has been used to argue for the possibility of lying beyond what is said was introduced by Jennifer Saul, who defines lying in terms of what is said (cf. Saul 2012: 16; Viebahn 2021: 296):

Metaphor

Ada is a keen gardener but has had an exceptionally bad crop of tomatoes. Ada wants Bill to think that her crop was in fact great, so when she meets Bill and he asks how her crop of tomatoes has been, she utters (3):

(3) I've got tomatoes coming out of my ears.

Conversational implicature: Ada has had a great crop.

Ada's utterance (3) is metaphorical and thus non-literal: there is a mismatch between the proposition that is said (on a narrow view of what is said) and the proposition Ada puts forward with her utterance. While Ada says that she has tomatoes coming out of her ears, she puts forward that she

has had a great crop. On Grice's (1975: 53–4) influential account of metaphor the proposition put forward is analysed as a conversational implicature.

Nonetheless, several theorists have remarked that Ada is intuitively lying, and indeed lying about her crop. This verdict is supported by a diagnostic test Stokke (2018: 89–91), Viebahn (2021: 293–4) and others employ to decide whether or not an utterance is a lie.³ The test considers how speakers can respond to accusations of lying following an insincere utterance. Returning to our initial example once more, we can suppose that a colleague accuses the doctor of lying to the dying woman. If the doctor has uttered the misleading (2), she can consistently respond to the accusation along the following lines:

- (4) I didn't lie. I didn't claim that the dying woman's son was fine. I merely claimed that I had seen her son the day before and that he was fine then, which was indeed the case.

By contrast, a consistent response of this kind is not possible if the doctor utters (1) and is accused of lying. The diagnostic test says that the possibility of such consistent responses to accusations of lying can be seen as indicating that the utterance in question is not a lie (but merely misleading), while the absence of this possibility can be seen as indicating that the utterance in question is a lie. In the case of (3), the diagnostic test suggests that Ada's utterance is indeed a lie, as a response of the following kind would clearly be inconsistent:

- (5) I didn't lie. I didn't claim that I had a great crop. I merely claimed that I had tomatoes coming out of my ears.

So we have a first case that suggests that the content of a lie does not have to align with the content of the sentence uttered.⁴

Theorists have discussed other cases of apparent non-literal lies that suggest that the possibility of lying beyond what is said is not restricted to metaphorical utterances. For instance, I have discussed the following cases involving hyperbole and loose use, respectively (Viebahn 2017: 1368):

³ See García-Carpintero (2023), however, for criticism of this diagnostic test.

⁴ Note that while the case does involve a disbelieved proposition (namely that Ada has tomatoes coming out of her ears), this is not a viable candidate for being the content of the lie, as Ada clearly does not put forward this proposition.

Hyperbole

Carl desperately wants Daisy to come to his party, which is in full swing. When Carl calls Daisy, she says she'll come, but only if there is some food. Carl is aware that all the food has been eaten, but nonetheless utters:

(6) There's tons of food left.

Conversational implicature: There is lots of food left.

Loose use

Fred is meant to start work at the office at about nine o'clock. Recently, however, he has developed a habit of arriving considerably late, and on several occasions he has not arrived until after ten o'clock. His boss has heard about this and asks Fred's colleague Emily when he arrived at the office today. Emily knows that Fred did not arrive until about ten o'clock, but she wants to protect him and utters:

(7) Fred arrived at nine o'clock.

Conversational implicature: Fred arrived at about nine o'clock.⁵

In both cases, there is a mismatch between the proposition that is said (on a narrow view of what is said) and the proposition put forward, which the speaker disbelieves. In both cases, the proposition put forward is standardly treated as a conversational implicature.⁶ And both cases are intuitively cases of lying, with the diagnostic test supporting this verdict (as consistent denials appear to be unavailable). So there is a range of cases that (at least initially) suggest that it is possible to lie beyond what is said.

Apparent examples of non-literal lies have not gone unchallenged. Saul (2012: 18) has argued that utterances such as (3) are only counted as a lie 'when we are thinking of lying as general deceptiveness'. And García-Carpintero (2018: 208) notes that at least some of the cases in question need not be analysed as non-literal utterances. For example, he notes that the "tomatoes out of the ears" metaphor is pretty dead' (208), in which case (3) might semantically express (and thus say) that Ada has had a great crop.

However, responses of both kinds face challenges. Saul's observation does not seem to fit with the outcomes of the diagnostic test, which indicates that the cases involve *lies rather than misleading utterances*. And the diversity of apparent non-literal lies might make it difficult to argue that *all* apparent cases of non-literal lies are in fact literal. In view of these challenges, it is not very surprising that there has been comparatively little pushback overall on the possibility of non-literal lies.

⁵ Also see Simpson (1992: 630), Adler (1997: 444, fn. 27) and Keiser (2016: 469, fn. 17).

⁶ See e.g. Grice (1975: 53–4) and Davis (2007) on hyperbole and Lasersohn (1999) and Davis (2007) on loose use.

So theoretical considerations and expert intuitions can be taken to suggest that lies need not be said. What do the intuitions of ordinary speakers have to say on the matter? Until recently, empirical data offered a mixed view of the issue.

On the one hand, Weissman & Terkourafi (2019) conducted a much-discussed study that seemed to indicate that speakers generally consider utterances featuring disbelieved conversational implicatures *not* to be lies. Weissman & Terkourafi presented subjects with fifteen vignettes featuring disbelieved conversational implicatures of different kinds and asked them whether the speaker lied (indicating their answers on a 7-point Likert scale). Subjects considered the speaker to be lying in only two of these cases, with mean ratings for most vignettes far below the midpoint (and thus toward the non-lie end of the scale).

On the other hand, studies by Reins & Wiegmann (2021) and Wiegmann et al. (2021), which followed a similar approach, found that subjects mostly considered speakers to be lying in cases involving disbelieved conversational implicatures. For instance, Reins & Wiegmann (2021) tested four cases involving *generalised* conversational implicatures, all of which received mean lie-ratings above the midpoint of the scale, and four cases involving *particularised* conversational implicatures, three of which received mean lie-ratings above the midpoint of the scale.

The contours of this mixed picture were recently clarified through a study by Wiegmann (2023a), which indicates that for some of Weissman & Terkourafi's vignettes lie-ratings were low because the speaker's deceptive intent was not made clear to the subjects. Wiegmann (2023a: 110–3) replicated the results of the original study by Weissman & Terkourafi, but also ran a version in which the vignettes made explicit that the speaker intends the hearer to come to believe the disbelieved proposition (put forward as a conversational implicature). This increased lie-ratings for all vignettes, and led to lie-ratings clearly above the midpoint for nine out of fifteen cases.

Overall, existing studies thus indicate that ordinary speakers judge at least some cases involving disbelieved conversational implicatures to be lies. This is a first piece of evidence that supports the possibility of lying beyond what is said (where what is said is construed narrowly and excludes conversational implicatures). At the same time, ordinary speakers do not judge all utterances involving disbelieved conversational implicatures to be lies; and even when mean lie-ratings for disbelieved conversational implicatures are above the midpoint, they tend to be significantly lower than mean lie-ratings for paradigmatic cases of lying, such as (1). This supports the first of the three basic assumptions (Section 2), according to which the concept of lying is distinct from the concept of misleading. Finally, the huge variety of conversational implicatures calls for further empirical investigation of which *kinds* of (utterances involving) disbelieved conversational implicatures are judged to be lies by ordinary speakers.

3.2 Lying without words?

We have seen that non-literal lies present a first way in which lying can arguably go beyond what is said. But all of the cases considered so far are compatible with the idea that lying is a *linguistic* phenomenon: that we have to use language to tell a lie. The view of lying as a linguistic phenomenon is still predominant, and it is contained in many content-based definitions of lying. After all, these definitions tie lying to the notion of what is said, and if saying is a linguistic phenomenon (as seems initially plausible), then so is lying. Accordingly, Carson states that lying requires the production of ‘a linguistic token’ (2006: 287), and Fallis says that liars ‘have to use language to express a proposition’ (2009: 37).⁷

However, the idea that lying is a linguistic phenomenon has been challenged: it has been argued that it is possible to tell lies with the help of pictures, and thus without using language. Some of these challenges are situated outside the contemporary debate on how to define lying. For instance, Winfried Nöth (1997) and Daisy Dixon (2022) argue that it is possible to lie with pictures, but their arguments are not focussed on the narrow notion of lying that contrasts with misleading. Other challenges, however, *are* concerned with the concept of lying as it is at issue in the current debate: Viebahn (2019) and García-Carpintero (2023) argue that, even if the concept of lying is construed narrowly so as to contrast with misleading, it is possible to lie with pictures.

Here is an example I have used to support the possibility of pictorial lies (Viebahn 2019: 243):

- (P1) Martha wants Nora to think that Oscar and Paula kissed. Martha knows that in fact Oscar and Paula never kissed. She carefully manipulates a photograph in such a way that it shows Oscar and Paula kissing and messages it to Nora with no further comment. When Nora receives the photo, she comes to believe that Oscar and Paula kissed.

In my view, there are three reasons to take (P1) to involve a pictorial lie (cf. Viebahn 2019: 244–6). Firstly, it would be natural to report Martha’s communicative act by saying that Martha lied about Oscar and Paula kissing. Secondly, (P1) displays a contrast with cases of mere pictorial misleading, such as (P2):

- (P2) Martha wants Nora to believe that Oscar and Paula are a couple (which is false, although Oscar and Paula have moved to the same flat). Martha sends Nora accurate pictures of Oscar and Paula standing outside their flat and carrying boxes into the flat to misleadingly suggest that they are a couple.

⁷ Also see Saul (2012: 1) and Stokke (2018: 31).

The contrast between (P1) and (P2) seems to parallel that between the lie (1) and the misleading utterance (2) in the linguistic domain. This suggests that (P1) is not merely misleading, but a lie. And thirdly, the diagnostic test suggests that Martha lied in (P1), as she apparently cannot consistently dismiss accusations of lying.

If (P1) indeed involves a pictorial lie, then there are many instances of pictorial lies. A real-life example is provided by the front page of *The Mirror* on 9 August 1997, which showed a manipulated photo of Lady Diana and Dodi Fayed as they are about to kiss, although at the point the photo was taken they were actually facing in different directions (as the unedited photo reveals). It seems fair to say that the makers of *The Mirror* lied to their readers. While (P1) and the newspaper example feature manipulated images, unedited, accurate pictures can also be used to lie if they are presented as being about something else than they are actually about. For instance, a son might pictorially lie to his mother about being in school by sending yesterday's (unedited, accurate) photo of himself in school.⁸ While pictorial lies involving manipulated photos can be labelled *content-lies*, pictorial lies involving photos presented as being about something else than they are actually about can be labelled *target-lies*.⁹

There has been little disagreement with these purported examples of pictorial lies – though this might in part be due to the fact that the possibility of pictorial lies has only recently entered the debate. García-Carpintero (2023: 524–5), however, has argued that the diagnostic test of consistent responses does not seem to work with pictorial lies. While he agrees that (P1) and the other examples discussed are lies, he remarks that, having sent the manipulated photo, Martha could consistently respond as follows:

- (8) I didn't lie about Oscar and Paula kissing. I was only resending you a picture someone else had sent me.

I do not quite share the intuition that response (8) would be consistent, given Martha's previous actions (after all, she did manipulate and circulate the photo herself). But it is worth investigating further whether the diagnostic test can be applied to pictorial cases in the same way as it can be applied to linguistic cases, or whether changes are required.

Do intuitions of ordinary speakers allow for the possibility of pictorial lies? There is initial empirical data indicating a positive answer to this question. Viebahn & Wiegmann (2022) conducted a study in which participants were confronted with examples of insincere linguistic and

⁸ See Viebahn (2019: 247–8) for further discussion of these and other examples.

⁹ On the mechanisms of content-lies and target-lies, see Lewerentz & Viebahn (2023).

pictorial communication. These examples included standard cases of linguistic lying and misleading, on the one hand, and candidates for pictorial lying and misleading, on the other. The candidates for pictorial lies were structurally similar to (P1), while the candidates for pictorial misleading were structurally similar to (P2). In each case, the vignette included the picture that was used to pictorially communicate.

What were the results? Participants clearly rated purported cases of pictorial lies as lies, even though they were given the possibility to indicate that they were misleading, but not lies. Indeed, Viebahn & Wiegmann found no significant difference in the lie-ratings concerning the linguistic cases and the pictorial cases (and no difference between lie-ratings for purported content-lies and target-lies). Furthermore, judgements about purported pictorial lies differed significantly from judgements about purported pictorial misleading. However, judgements about candidates for pictorial misleading were closer to judgements about candidates for *pictorial* lying than to judgements about candidates for *linguistic* misleading. These latter judgements may suggest that, when it comes to ordinary usage, the lying-misleading distinction is less pronounced in the case of pictorial communication.

3.3 Where to go from here?

We have seen that there are both theoretical and empirical reasons to think that lying does not require saying, where saying is understood narrowly to exclude conversational implicatures and pictorial contents. To end this section, I now want to outline two reactions to this outcome, and thus two ways in which the debate might move forward.

Firstly, it is possible to retain the idea that lying should be defined in terms of what is said, while broadening the notion of what is said so as to cover some or all of the purported cases of lying discussed above. García-Carpintero (2021, 2023) pursues this route by defending the following content-based definition of lying:

García-Carpintero's definition of lying (García-Carpintero 2023: 515)

A lies in communicating proposition p to B if and only if:

- (L1) A assertorically commits to p,
- (L2) A's utterance says/makes explicit p,
- (L3) A believes p to be false.

This definition differs from e.g. Stokke's definition in several respects, but presently the most important aspect is that it retains the notion of what is said in (L2). As we have seen, García-Carpintero accepts the possibility of pictorial lies, and thus includes propositions semantically (explicitly) expressed by pictures in his notion of what is said (2023: 525–8). By contrast, he disagrees

about the possibility of non-literal lies, thus not extending his notion of what is said to conversational implicatures.¹⁰ As a result, the definition can account for pictorial lies, though it seems to be in conflict with the possibility of linguistically lying beyond what is said. Further theoretical and empirical research will hopefully clarify whether this is a feature or a bug of the theory.

Secondly, theorists have responded to the possibility of non-literal and pictorial lies by moving away from the content-based approach and instead defining lying in terms of communicative commitment (Viebahn 2021, Reins & Wiegmann 2021).¹¹ The underlying idea here is that the difference between lying and misleading does not concern the manner in which content is expressed, but rather the *force* with which it is advanced. Lying requires speakers to commit to a disbelieved proposition, while misleading utterances involve a disbelieved proposition that is put forward in a non-committal way.¹² For example, I have proposed the following commitment-based definition of lying:

Viebahn's definition of lying (Viebahn 2021: 300)

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

- (L1) A performs a communicative act C with p as content;
- (L2) with C, A intends to communicate p to B;
- (L3) with C, A commits herself to p; and
- (L4) A believes that p is false.

As I spell out the notion of commitment, communicative agents can commit themselves not only to semantically expressed content, but also to content that is conversationally implicated or pictorially expressed; they can avoid commitment through certain *additive* conversational implicatures (though not through *substitutive* conversational implicatures).¹³ The definition can therefore account for non-literal lies as well as pictorial lies, without counting every case involving a disbelieved conversational implicature as a lie.¹⁴

The commitment-based approach is supported by a study reported in Reins & Wiegmann (2021), which collected not only lie-ratings on vignettes featuring various insincere utterances, but

¹⁰ García-Carpintero (2023: 514) also includes conventional implicatures and linguistically triggered presuppositions in what is said; his view of what is said is thus broader than that of Stokke (2018), who does not include presuppositions in what is said.

¹¹ Marsili (2021) defends a definition of lying that requires both saying and commitment.

¹² The commitment-based approach takes its cue from the influential idea that assertion is tied to commitment, defended by Peirce (1934), Brandom (1983), Kukla & Lance (2009) and many others.

¹³ While *additive* conversational implicatures are put forward *in addition to* the semantic content of the sentence uttered, *substitutive* conversational implicatures are put forward *instead of* the semantic content of the sentence uttered. The examples of non-literal lies discussed above feature substitutive implicatures, while the initial example of misleading (2) features an additive conversational implicature. See Dinges (2015: 56) and Viebahn (2021: 311) for discussion.

¹⁴ For a related but different way of spelling out the notion of commitment, see Reins & Wiegmann (2021).

also asked subjects whether the speakers in the vignettes had taken on a commitment to a disbelieved content. Reins & Wiegmann found a strong relationship between commitment-ratings and lie-ratings, though of course it is an open question how exactly subjects interpreted the commitment question.

Criticism of the commitment-based approach has been voiced by Pepp (2022), who has objected that there are cases in which mere misleaders seem to take on a commitment (of the relevant kind) to a disbelieved proposition. According to Pepp, commitment-based definitions of lying can thus only succeed if further work is invested to find a notion of commitment that applies only to cases of lying.¹⁵

4. Does lying require actual falsity?

The second basic assumption states that lying requires asserting a disbelieved proposition. But must this proposition also be actually false, or does it suffice that the speaker disbelieves it? That is a second point of contention in the debate on how to define lying, and the topic of the current section.

The traditional view on this matter is *the subjective view of lying*: it entails that actual falsity of the proposition asserted is not required for lying. Subjectivists thus allow for the possibility of true lies. Subjectivism was advocated by Augustine, who mentions the possibility of true lies in the following passage:

So, he lies who holds one thought in his mind and gives expression to another through words or any other signification. [...] And so one can tell the truth by lying, if one takes something to be false and puts it forward as true, although what is put forward is in fact true. (Augustine [395]/2013: 62–4, my translation)

Nowadays, subjectivism is probably the majority view in the debate, with proponents including Chisholm & Feehan (1977: 152), Williams (2002: 96), Fallis (2009), Saul (2012), Stokke (2018), Maitra (2018), García-Carpintero (2018), Pepp (2019) and Wiegmann & Viebahn (2021).

However, a number of theorists have recently advanced theoretical and empirical arguments in favour of *objectivism* about lying, the view that lying requires *both* disbelief *and* actual falsity of the proposition asserted. Objectivists thus rule out the possibility of true lies and add a third clause to the basic requirements for lying mentioned in Section 2:

¹⁵ Also see Marsili & Löhr (2022) for criticism along other lines.

Lying requires asserting a disbelieved and false proposition.

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

- (L1) A asserts p to B,
- (L2) A disbelieves p , and
- (L3) p is false.

Objectivism has been defended by Carson (2006: 285, 2010: 15–17), Turri & Turri (2015), Benton (2018), Holguín (2021) and Kallestrup (2023).

Theoretical arguments for objectivism have rested on expert intuitions about examples or sentences. For instance, Carson intuitively feels that the following case does not involve a lie, but an attempted lie, as the proposition asserted turns out to be true:

I go fishing on a boat with a friend, John. He and I both catch a fish at the same time. Although we don't realize it, our lines are crossed. I have caught a very big fish and John has caught a little one, but we mistakenly believe that I caught the small fish and John caught the big one. We throw the two fish back into the water. I go home thinking that I caught a small fish. When I return, my father, an avid fisherman, asks me how I did. I say that I caught a very large fish and threw it back into the water, thereby intending to deceive him about [the] size of the fish that I caught. (Carson 2006: 285)

Carson holds that this intuition speaks against subjectivism and supports objectivism about lying. Taking a slightly different approach, Holguín (2021: 5354) argues for objectivism by considering whether there are situations in which the following sentences have true readings:

- (9) Although we know everything Bill Clinton is saying is true, what we don't know is whether he is lying.
- (10) Even if what Lexie says is true, there is a good chance it is a lie.

If subjectivism about lying were true, Holguín points out, there should be situations in which the sentences have true readings. But, in his view, sentences (9) and (10) 'sound terrible' (2021: 5354), regardless of the situation.

These theoretical arguments for objectivism are controversial. With respect to Carson's example, Stokke comments that 'many will be attracted to the opposite verdict on this case', as 'the mere fact that the fisherman's story is true—even though no one is aware of this—is not sufficient to show that no lie has been told' (2018: 33). And Stokke presents the following case that he intuitively takes to be a lie:

In a foregone age in which people believed in spontaneous generation, Robert is asking his friend Thomas for advice before taking an exam. Robert asks, “Where do rats come from?” Thomas believes that rats are generated spontaneously by wet dirt, and he knows that this is the answer that will be taken as correct on the exam. But he wants Robert to fail the exam and so he tells him, “They are born by other rats.” Robert trusts Thomas’s response and gives it as his answer on the exam, which he thereby fails. (Stokke 2018: 33–4)

I share Stokke’s verdicts on both cases, as well as his impression that most theorists will have intuitions that fit with subjectivism. If this is right, expert intuitions about examples do not present a strong case for objectivism.

Can Holguín’s sentence-based argument for objectivism do any better? That is not at all clear, either: if we know that everything Bill Clinton is saying is true but do not know whether he believes it to be true, then (9) does not sound terrible at all, and indeed exactly right (at least to me). Likewise, if we suspect that Lexie believes what she says to be false, then (10) does not sound terrible and might well be true. But I acknowledge that this is a case of expert intuition vs. expert intuition, and so it would be good to know how the intuitions of ordinary speakers pan out with respect to this question.

Fortunately, empirical studies provide helpful findings on whether ordinary speakers take lying to require actual falsity, though the route to these findings includes some twists and turns.¹⁶ A first data point was provided by Coleman & Kay’s (1981: 31) pioneering and influential study on the expression ‘lie’, which included the following case:

One morning Katerina has an arithmetic test she hasn’t studied for, and so she doesn’t want to go to school. She says to her mother, ‘I’m sick.’ Her mother takes her temperature, and it turns out to Katerina’s surprise that she really is sick, later that day developing the measles.

A large majority of subjects judged Katerina’s utterance to be a lie (despite the truth of the proposition asserted), which is a first data point in favour of a subjective view of lying among ordinary speakers.

However, a recent study by Turri & Turri (2015) called these findings into question. In the first part of the study, Turri & Turri confronted subjects with a case structurally similar to that of Coleman & Kay and simply asked them whether the speaker had lied. The results seemed to confirm

¹⁶ See Wiegmann & Meibauer (2019) and Wiegmann (2023b) for helpful overviews of this development.

the findings of Coleman & Kay, with more than 80 per cent of subjects judging the speaker to have lied. But Turri & Turri wanted to be sure that these results were not impacted by confounding factors, such as the wish to attribute blame to the speaker. Thus they conducted a second survey, in which they asked subjects which of the following two options better describes the speaker in the vignette:

- (A) He tried to lie and actually did lie.
- (B) He tried to lie but only thinks he lied.

In this case, 90 per cent of subjects chose option (B), according to which the speaker did not lie. According to Turri & Turri, this result confirmed the suspicion that the findings of the first part were influenced by confounding factors, and it provided a reason to assume an objectivist view of lying among ordinary speakers.

The apparent tension between the findings of Coleman & Kay (1981) and Turri & Turri (2015) was resolved by a study by Wiegmann et al. (2016). Wiegmann et al. suspected that the wording of (A) and (B) led subjects to interpret their task as being about the possible falsity of the statement, rather than about a possible lie. To test for this possibility, they ran the study with the following modified response options, which specify that the statement in question is true:

- (A1) He tried to lie and actually did lie although what he said turned out to be true.
- (B1) He tried to lie but only thinks he lied because what he said turned out to be true.

The findings confirmed their suspicion: a clear majority of subjects opted for (A1) and thus judged the true but disbelieved assertion to be a lie.¹⁷

As it stands, empirical findings thus support the view that ordinary speakers accept a subjective view of lying that allows for the possibility of true lies. Given that these intuitions are taken to be of importance for the debate on how to define lying, this outcome can help to settle the expert disagreement about putative examples of true lies, tipping the balance at least somewhat in the direction of subjectivism about lying.

¹⁷ See Turri & Turri (2021) for a further empirical challenge to the subjective view of lying, and Wiegmann (2023b) for a further empirical response.

5. Does lying require an intention to deceive?

A third contested question concerning the concept of lying asks whether lying requires an intention to deceive the addressee. This is another question that has a long history and goes back at least to Augustine ([395]/2013). Though it is not entirely clear which answer Augustine preferred, the traditional answer is that of *deceptionism*, according to which lying does indeed require a deceptive intent. On one common way of spelling out this deceptive intent, accepted by e.g. Isenberg (1964), Chisholm & Feehan (1977), Williams (2002), it consists in intending the addressee to come to believe the disbelieved proposition asserted. This leads to the addition of clause (L3) to the basic requirements for lying:

Lying requires asserting a disbelieved proposition and an intent to deceive.

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

- (L1) A asserts p to B, and
- (L2) A disbelieves p, and
- (L3) by asserting p, A intends to deceive B into believing p.¹⁸

Deceptionism was long the default view about lying, but in the last two decades there has been a considerable shift towards *non-deceptionism*, the view that lying does not require an intention to deceive. Non-deceptionism has been advocated by Carson (2006, 2010), Fallis (2009, 2010), Saul (2012), Stokke (2013, 2018), Krstić (2019, 2020) and Viebahn (2021), among others.

The shift towards deceptionism started in the theoretical part of the debate with the discussion of *bald-faced lies* by Carson (2006) and Sorensen (2007). A bald-faced lie is a lie that is told in a situation in which it is mutually believed by the speaker and the audience that the speaker is lying. Here is Carson's (2006: 289) first example of a bald-faced lie, which involves a threatened witness:

Bill has seen Bob commit a murder. The investigator knows that Bill has witnessed the murder (there is clear video evidence), and Bill knows that the investigator knows this. To convict Bob, an eyewitness statement is needed, so the investigator asks Bill: 'Did you see the murder?' Bob has threatened to kill Bill unless he claims not to have seen any murder. Bill replies:

- (11) I didn't see any murder.

¹⁸ The deceptive intent has been spelled out differently, for instance by Lackey (2013: 237, 241). Lackey argues that lying merely requires an intention *to be deceptive* in performing a disbelieved assertion, where being deceptive can consist in bringing about a false belief or in concealing information.

Among theorists in the debate, the most common reaction to this case is that Bill is lying, and that he has no intention to deceive the investigator with respect to not having seen any murder. If Bill is indeed lying, then this is a bald-faced lie, as it is mutually believed by Bill and the investigator that he is lying. And there are two reasons to hold that Bill does not have an intention to deceive the investigator. Firstly, Bill performs the utterance to protect his life. What matters to him is to do as Bob tells him to, and any possible deception is an unintended ‘side effect’ (Carson 2006: 89). Secondly, the fact that Bill knows that the investigator knows that he is lying rules out any rational intention to deceive the investigator: Bill cannot rationally intend to achieve what he takes to be impossible.

As a second example of a bald-faced lie, Carson (2006: 290) discusses a case of a cheating student:

Having been caught cheating in an exam, a student faces the dean, who has unsailable evidence of the student’s offence. For fear of potential lawsuits, the dean only punishes students who admit their offence. The student knows this, and thus utters:

(12) I did not cheat in the exam.¹⁹

Here, too, it seems that the student is lying, as is mutually believed by the student and the dean. And the student apparently does not have an intention to deceive the dean about the content of his assertion.

In addition to bald-faced lies, theorists have recently discussed other kinds of apparent lies without an intention to deceive. The general idea has been to find lies with ‘an alternative kind of motivation at work compared with an intent to be deceptive’ (Sneddon 2021: 48). One such alternative motivation can be for the addressee to deduce a proposition the speaker believes to be true, as in the following example introduced by Krstić (2020: 758–9):

A vicious murderer, Tony, is hiding from the police in Pinocchio’s house. In search of Tony, the police knock on Pinocchio’s door asking whether Tony is hiding in his house. Pinocchio wants to give Tony away but he is afraid that, if he gives any indication of this to Tony, Tony will hurt him. Luckily, Pinocchio knows that the police know that his nose starts to grow at the very instant he forms the intention to lie and that they know that he knows that they know how his nose behaves, etc.,

¹⁹ I have slightly changed both of Carson’s examples to ease discussion.

and he also knows that Tony does not know anything about this. Therefore, he asserts to the police:

(13) Tony definitely isn't in my house.²⁰

As Krstić notes, Pinocchio's aim is not to deceive the police, but rather for them to come to believe a proposition he believes to be true (namely the negation of the proposition asserted).

While many theorists hold that examples of bald-faced lies and alternative motivation lies are a reason to abandon deceptionism, there has also been criticism of this reaction. According to a first objection, the speakers in such cases are not in fact lying, as they fail to make an assertion. For example, Ishani Maitra (2018: 75–8) argues that the student in Carson's second example is *acting as if* he makes an assertion, without actually making an assertion.²¹ A second objection alleges that the speakers in such cases do have an intention to deceive, after all. This objection has been advanced by Jennifer Lackey (2013: 241–2), who acknowledges that Carson's student is making an assertion, but argues that he does intend to be deceptive by withholding his admission of wrongdoing and thereby concealing evidence needed for his punishment. While a full treatment of these objections is beyond the scope of this chapter, I do want to discuss some reasons to doubt that they are fully successful.

The first objection may be right about some alleged examples of undisguised lies in the debate, but it is less convincing with respect to the examples I have just discussed. In these examples, the speakers do seem to assert the content they put forward. This point can be supported by the observation that the utterances license reactions that are possible with assertions ('Are you sure you didn't see any murder?', 'Are you sure you didn't cheat?', 'Are you sure Tony isn't in your house?'). Such reactions are not licensed in cases in which speakers merely act as if they are making an assertion.

For the second objection, it matters how the intention to deceive is spelled out. It seems right that bald-faced lies and alternative-motivation lies can involve deceptive intentions of *some* kind. For instance, it can be argued that there is some sense in which the student is acting deceptively in uttering (12). But, firstly, standard formulations of the intention to deceive require an intention to get the addressee to come to believe the disbelieved proposition asserted, and such an intention does not seem to be present in any of the current cases. And, secondly, at least some alternative-motivation lies, such as Pinocchio's utterance of (13), appear to involve no deceptive intent whatsoever towards the addressees.

²⁰ This example is in some respects similar to Augustine's ([395]/2013: 66) example of the benevolent travel adviser.

²¹ Similar arguments have been put forward by Leland (2015), Keiser (2016) and Hawley (2018).

Finally, empirical work suggests that ordinary speakers accept that there are lies without an intention to deceive, though once again the route to this finding has not been straightforward. A first step was taken by Arico & Fallis (2013), who presented subjects with six cases of bald-faced lying similar to those involving (11) and (12). Virtually all subjects judged the speakers in these cases to be lying. However, subjects were not asked whether the speakers had an intention to deceive.

In a subsequent study by Meibauer (2016), putative bald-faced lies were again clearly judged to be lies, but subjects were also asked about deceptive intent and indeed ascribed an intention to deceive to the speakers; the same pattern emerged in a study by Rutschmann & Wiegmann (2017). These findings undercut the potential support judgements by ordinary speakers about bald-faced lies could have offered for non-deceptionism.

However, the most recent empirical findings do seem to speak in favour of non-deceptionism. Focussing on alternative-motivation lies, Krstić & Wiegmann (2022) tested whether ordinary speakers consider Pinocchio to be lying in uttering (13) and whether they take him to have an intention to deceive the police. A clear majority judged Pinocchio to be lying, and a clear majority judged him *not* to have an intention to deceive the police.

So, empirical studies have shown that ordinary speakers take both bald-faced lies and alternative-motivation lies to be lies. But they have also brought out that they only take some of these lies to lack an intention to deceive. Still, if this latter finding proves robust, that would be enough to indicate that ordinary speakers have a concept of lying that does not require an intention to deceive, which in turn could provide a push toward non-deceptionism about lying.

6. Conclusion

This chapter has provided an opinionated overview of three contested questions about the concept of lying. We have encountered arguments for a view that places fairly minimal demands on lying. According to this view, lying requires asserting a disbelieved proposition, but this proposition need not be said (as it can be conversationally implicated or pictorially expressed), it need not be actually false (in addition to being disbelieved), and the speaker need not intend the addressee to come to believe the asserted proposition (so there need not be a deceptive intent). While there are objections to each of these three points, we have seen that empirical studies suggest that they are in accordance with the concept of lying as it is held by ordinary speakers.²²

²² I would like to thank Alexandra Pop for helpful comments on a previous version of this chapter.

References

- Adler, J. (1997), 'Lying, Deceiving, or Falsely Implicating', *Journal of Philosophy*, 94: 435–52.
- Arico, A. and D. Fallis (2013), 'Lies, damned lies, and statistics: An empirical investigation of the concept of lying', *Philosophical Psychology*, 26: 790–816.
- Augustine ([395]/2013), '*De mendacio*', in *Opera/Werke*, Vol. 50, edited and translated by A. Städele, Paderborn: Schöningh.
- Benton, M. (2018), 'Lying, belief, and knowledge', in J. Meibauer (ed), *The Oxford Handbook of Lying*, Oxford: Oxford University Press: 120–33.
- Bolzano, B. ([1810]/2007), 'Über die Wahrhaftigkeit' in: E. Morscher and K. Strasser (eds), *Erbauungsreden der Studienjahre 1804/05 bis 1819/20*, Stuttgart: Frommann-Holzboog.
- Brandon, R. (1983), 'Asserting', *Noûs*, 17: 637–50.
- Carson, T. (2006), 'The definition of lying', *Noûs*, 40: 284–306.
- Carson, T. (2010), *Lying and Deception: Theory and Practice*, Oxford: Oxford University Press.
- Chisholm, R. M. and T. D. Feehan (1977), 'The intent to deceive', *The Journal of Philosophy*, 74: 143–59.
- Coleman, L. and P. Kay (1981), 'Prototype semantics: The English word lie' *Language*, 57: 26–44.
- Davis, W. (2007), 'Knowledge claims and context: Loose use', *Philosophical Studies*, 132: 395–438
- Dinges, A. (2015), 'Innocent implicatures', *Journal of Pragmatics*, 87: 54–63.
- Dixon, D. (2022), 'Lies in Art', *Australasian Journal of Philosophy*, 100: 25–39.
- Dynel, M. (2011) 'A Web of Deceit: A Neo-Gricean View on Types of Verbal Deception', *International Review of Pragmatics*, 3: 139–67.
- Fallis, D. (2009), 'What is Lying?', *The Journal of Philosophy*, 106: 29–56.
- Fallis, D. (2010), 'Lying and Deception' *Philosophers' Imprint*, 10: 1–22.
- Fox Krauss, S. (2017), 'Lying, risk and accuracy', *Analysis*, 77: 726–34.
- García-Carpintero, M. (2018), 'Sneaky Assertions', *Philosophical Perspectives*, 32: 188–218.
- García-Carpintero, M. (2021), 'Lying vs. misleading: The adverbial account', *Intercultural Pragmatics*, 18: 391–413.
- García-Carpintero, M. (2023), 'Lying versus misleading, with language and pictures: the adverbial account', *Linguistics and Philosophy*, 46: 509–32.
- Grice, H. P. (1975), 'Logic and Conversation', *Syntax and Semantics*, 3: 41–58.
- Harris, D. W. (2020), 'Intentionalism and bald-faced lies', *Inquiry*: Online First.
- Hawley, K. (2018),: 'Coercion and Lies', in E. Michaelson and A. Stokke (eds), *Lying. Language, Knowledge, Ethics, and Politics*, 229–45, Oxford: Oxford University Press.
- Holguín, B. (2021), 'Lying and knowing', *Synthese*, 198: 5351–71.

- Isenberg, A. (1964), 'Deontology and the Ethics of Lying', *Philosophy and Phenomenological Research*, 24: 463–80.
- Kallestrup, J. (2023), 'The myth of true lies', *Theoria*: Online First.
- Keiser, J. (2016), 'Bald-faced lies: how to make a move in a language game without making a move in a conversation', *Philosophical Studies*, 173: 461–77.
- Krstić, V. (2019), 'Can you lie without intending to deceive?', *Pacific Philosophical Quarterly*, 100: 642–60.
- Krstić, V. (2020), 'On the nature of indifferent lies, a reply to Rutschmann and Wiegmann', *Philosophical Psychology*, 33: 757–71.
- Krstić, V. and Wiegmann, A. (2022), 'Bald-Faced Lies, Blushing, and Noses that Grow: An Experimental Analysis', *Erkenntnis*: Online First.
- Kukla, R. and M. Lance (2009), *Yo! and 'Lo!': The pragmatic topography of the space of reasons*, Cambridge, MA: Harvard University Press.
- Lackey, J. (2013), 'Lies and deception: an unhappy divorce', *Analysis*, 73: 236–48.
- Lasersohn, P. (1999), 'Pragmatic Halos', *Language*, 75: 522–51.
- Leland, P. (2015), 'Rational responsibility and the assertoric character of bald-faced lies', *Analysis*, 75: 550–4.
- Lewerentz, L. and Viebahn, E. (2023), 'Truth and Directness in Pictorial Assertion', *Linguistics and Philosophy*: Online First.
- MacFarlane, J. (2011), 'What is assertion?', in: J. Brown and H. Cappelen (eds), *Assertion*, Oxford: Oxford University Press: 79–98.
- Maitra, I. (2018), 'Lying, Acting, and Asserting', in E. Michaelson and A. Stokke (eds), *Lying. Language, Knowledge, Ethics, and Politics*, Oxford: Oxford University Press.
- Marsili, N. (2021), 'Lying, speech acts, and commitment', *Synthese*, 199: 3245–69.
- Marsili, N. and G. Löhr (2022), 'Saying, commitment, and the lying-misleading distinction' *The Journal of Philosophy*, 119: 687–98.
- Meibauer, J. (2014), *Lying at the Semantics-Pragmatics Interface*, Boston/Berlin: De Gruyter.
- Meibauer, J. (2016), 'Understanding bald-faced lies. An experimental approach', *International Review of Pragmatics, Special Issue: Empirical Approaches to Lying and Deception*, 8: 247–70.
- Nöth, W. (1997), 'Can pictures lie?', in W. Nöth (ed), *Semiotics of the Media. State of the Art, Projects, and Perspectives*, 133–46, Berlin, Boston: De Gruyter Mouton.
- Peirce, C. S. (1934), 'Judgment and assertion', in *Collected papers of Charles Sanders Peirce*, Vol. V: 385–87, Boston: Harvard University Press..

- Pepp, J. (2019), ‘Assertion, Lying, and Untruthfully Implicating’, in S. Goldberg (ed), *The Oxford Handbook of Assertion*, Oxford: Oxford University Press, 828–58.
- Pepp, J. (2022), ‘What is the commitment in lying’, *The Journal of Philosophy*, 119: 673–86.
- Reins, L. and A. Wiegmann (2021). ‘Is lying bound to commitment? Empirically investigating deceptive presuppositions, implicatures, and actions’, *Cognitive Science*, 45: e12936.
- Rutschmann, R. and A. Wiegmann (2017), ‘No need for an intention to deceive? Challenging the Traditional Definition of Lying’, *Philosophical Psychology*, 30: 438–57.
- Saul, J. M. (2012), *Lying, Misleading, & What is Said*, Oxford: Oxford University Press.
- Simpson, D. (1992), ‘Lying, Liars and Language’, *Philosophy and Phenomenological Research*, 52: 623–39.
- Sneddon, A (2021), ‘Alternative motivation and lies’, *Analysis*, 81: 46–52.
- Sorensen, R. (2007), ‘Bald-faced lies! lying without the intent to deceive’, *Pacific Philosophical Quarterly*, 88: 251–64.
- Stalnaker, R. (1999), *Context and Content*, Oxford: Oxford University Press.
- Stalnaker, R. (2002), ‘Common ground’, *Linguistics and Philosophy*, 25: 701–21.
- Stokke, A. (2013), ‘Lying and Asserting’ *The Journal of Philosophy*, 110: 33–60.
- Stokke, A. (2018), *Lying and Insincerity*, Oxford: Oxford University Press.
- Turri, A. and J. Turri (2015), ‘The truth about lying’, *Cognition*, 138: 161–8.
- Turri, A. and Turri, J. (2021), ‘Lying, fast and slow’, *Synthese*, 198: 757–75.
- Viebahn, E. (2017), ‘Non-literal lies’, *Erkenntnis*, 82: 1367–80.
- Viebahn, E. (2019), ‘Lying with pictures’, *British Journal of Aesthetics*, 59: 243–57.
- Viebahn, E. (2021), ‘The lying-misleading distinction: a commitment-based approach’, *The Journal of Philosophy*, 118: 289–319.
- Viebahn, E. and A. Wiegmann (2022), ‘An empirical perspective on pictorial lies’, OSF Preprint: <https://osf.io/wp7va/> .
- Weissman, B. and M. Terkourafi (2019), ‘Are false implicatures lies? An empirical investigation’, *Mind & Language*, 34: 221–46.
- Wiegmann, A. (2023a), ‘Lying with deceptive implicatures? Solving a puzzle about conflicting results’, *Analysis*, 83: 107–118.
- Wiegmann, A. (2023b), ‘Does lying require objective falsity?’, *Synthese*, Online First.
- Wiegmann, A. and J. Meibauer (2019), ‘The folk concept of lying’, *Philosophy Compass*, 14: e12620.
- Wiegmann, A., J. Samland, and M. R. Waldmann (2016), ‘Lying despite telling the truth’, *Cognition*, 150: 37–42.

Wiegmann, A. and Viebahn, E. (2021), 'True lies and Moorean redundancy', *Synthese* 199: 13053–66.

Wiegmann, A., P. Willemsen, and J. Meibauer, (2021), 'Lying, Deceptive Implicatures, and Commitment', *Ergo*, 8: 709–40.

Williams, B. (2002), *Truth and Truthfulness*, Princeton: Princeton University Press.