



Inquiry

An Interdisciplinary Journal of Philosophy

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/sinq20

Expressive lies

Luise Mirow

To cite this article: Luise Mirow (06 May 2025): Expressive lies, Inquiry, DOI: [10.1080/0020174X.2025.2498037](https://doi.org/10.1080/0020174X.2025.2498037)

To link to this article: <https://doi.org/10.1080/0020174X.2025.2498037>



© 2025 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 06 May 2025.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)

Expressive lies

Luise Mirow 

Department of Historical, Philosophical and Religious Studies, Umeå University, Umeå, Sweden

ABSTRACT

In this paper, I argue that cases involving expressive language prompt us to reconsider the need of three key features figuring in recent definitions of lying: assertion, belief and truth-evaluability. I assume that sentences including expressive language should be given an expressivist treatment, and I argue that such sentences can plausibly be used to lie. The main aim of this paper is to propose a revised definition of lying. Based on the resources available to two different expressivist views, I propose two ways to revise recent definitions of lying. While the first proposal results in a minimally revised definition that merely expands the notion of content but otherwise largely preserves all three key features of recent definitions of lying, the second proposal diverges more radically from recent definitions of lying. This second proposal that I call ‘the attitude expression analysis’ shows us something unexpected about lying, namely that the crucial feature of lying is a mismatch between attitudes held by the speaker and how the speaker presents herself in speech rather than the assertion of a propositional content the speaker believes to be false.

ARTICLE HISTORY Received 2 October 2023; Accepted 9 April 2025

KEYWORDS Lying; expressivism; truth-evaluability; assertion; propositional content

1. Introduction

Lying is a phenomenon that severely complicates, if not undermines, successful and cooperative communication among humans. In many people’s minds, lies are closely connected to truth in the sense that a liar typically tries to conceal the truth in saying something she believes to be false. Studying an underexplored class of lies involving expressive language raises questions about this common way to think about lying.

CONTACT Luise Mirow  luise.mirow@umu.se  Department of Historical, Philosophical and Religious Studies, Humanisthuset, Umeå University, 90187 Umeå, Sweden

© 2025 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

While recent developments in the literature on lying have looked at lying from an increasingly broad range of perspectives,¹ no one has yet said anything about the relation between lying and expressivism. This paper addresses that lacuna. Considerations of cases involving expressive language prompt us to reconsider the need of three key features figuring in recent definitions of lying: assertion, belief and truth-evaluability. For the purpose of this paper, I assume without argument that certain sentences including taste predicates, normative predicates and epistemic modals should be given an expressivist treatment, and I argue that those sentences can be used to lie. The main aim of the paper is to propose a revised definition of lying. Based on two different expressivist views, I will propose two ways to revise recent definitions of lying. While the first proposal results in a minimally revised definition that largely preserves key features of recent definitions of lying, the second proposal diverges more radically from recent definitions of lying. This second proposal that I will call ‘the attitude expression analysis’ shows us something unexpected about lying, namely that the crucial feature of lying is a mismatch between attitudes held by the speaker and how the speaker presents herself in speech rather than the assertion of a propositional content the speaker believes to be false.²

According to recent definitions of lying, asserting a propositional content³ one believes to be false is necessary for lying. Such definitions of lying are problematic in view of examples like the following three cases that intuitively seem to be instances of lying and not merely misleading.

Soup

You have invited your friend over for dinner. You have cooked a fish soup with saffron. You know that your friend loves fish, but you do not know that your friend cannot stand the taste of saffron. At dinner, you ask your friend, ‘Do

¹See, e.g. Fallis (2014) and Viebahn (2017) on non-literal lies, Viebahn (2020) and Stokke (2024) on presuppositional lies, Viebahn (2019) and Dixon (2022) on pictorial lies, and Sorensen (2007) and Krstić (2022) on bald-faced lies.

²The idea that lying involves a mismatch between attitudes and presentation is reminiscent of a thought from Confucian philosophy that being trustworthy is being such that your presentation matches your attitudes (Sung 2023). Thanks to Ethan Nowak for making me aware of this connection.

³A few remarks about my use of ‘propositional content’ and ‘proposition’ in this paper: Strictly speaking, utterances or sentences in contexts have propositional content, not sentences. Henceforth, I will shorten this and sometimes just speak about sentences having propositional content. Moreover, I use ‘proposition’ and ‘propositional content’ interchangeably here. Sometimes, I will speak of sentences having propositional content and sometimes, I will speak of sentences expressing propositional content or sentences expressing propositions.

you like the soup?'. Although she does not find the soup tasty, your friend replies, 'The soup is tasty'.

Plagiarism

Your friend, a university teacher and researcher, considers plagiarism permissible, given that the competitive pressure in academia is so high and that several people plagiarise in their work. Although she does not think plagiarism is wrong, your friend tells you, 'Plagiarism is wrong'.

Mum

You are out for a walk with your friend. Coming from the other end of the street, a person is approaching the two of you. Your friend knows that your mother is not in town. However, although she knows that the person approaching you is not your mother, your friend points to the person and says, 'The lady over there might be your mother'.

The cases suggest that sentences like 'The soup is tasty', 'Plagiarism is wrong' and 'The lady over there might be your mother' can plausibly be used to lie. However, on certain expressivist views, these sentences do not have propositional content. If it is true that there are sentences that can be used to lie but that do not have propositional content, asserting a propositional content one believes to be false cannot be necessary for lying. And if asserting a propositional content one believes to be false is not necessary for lying, we have to reconsider the need of three key features figuring in recent definitions of lying: assertion, belief and truth-evaluability. For the purpose of this paper, I will assume an expressivist treatment of the cases above, and I will propose two ways to revise recent definitions of lying, based on the resources available to two different expressivist views.

The structure of the paper is the following. In section 2 and 3, I provide a brief background on recent definitions of lying and the two expressivist views relevant to this paper. Section 4 presents a diagnostic test for lying that supports the intuition that sentences like 'The soup is tasty', 'Plagiarism is wrong' and 'The lady over there might be your mother' can be used to lie. In section 5 and 6, I propose two different revised definitions of lying, based on the two expressivist views presented in section 3. While the first revised definition expands the notion of content figuring in lying but otherwise largely preserves the three key features figuring in recent definitions of lying, the second revised definition presents a more radical divergence from recent definitions of lying in dismissing all three key features, most importantly truth-evaluability.

2. Propositionalism and three key features of lying

Many philosophers think that one has to say or assert a proposition one believes to be false in order to lie (e.g. Saul 2013; Stokke 2018; Viebahn 2021). Call such views ‘propositionalism’ or ‘propositionalist definitions of lying’. Stokke (2018, 31) provides a representative formulation of propositionalism:

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

(LA1) A asserts that p to B, and

(LA2) A believes that p is false.

Stokke’s propositionalist definition of lying is a natural starting point for this paper because many other definitions of lying are very similar (see, e.g. Saul 2013) and Stokke’s definition has been used as a starting point for discussions about how to define lying by other philosophers (see, e.g. Viebahn 2021). The different versions of propositionalism mainly differ with respect to how they spell out the notion of assertion in (LA1).⁴ What they have in common is that the content of the speaker’s lie is propositional and thus *truth-evaluable*, that the speaker *asserts* this propositional content, and that the speaker *believes* this propositional content to be false. That is, three key features are at the core of propositionalism: assertion, belief and truth-evaluability.⁵

On propositionalist definitions of lying, propositional content provides the link between the different conditions for lying, namely asserting a proposition and believing that the proposition is false.⁶ It is not enough that the speaker believes one thing to be false and asserts another thing. We need to make sure that the thing the speaker believes to be false is the same thing the speaker asserts.⁷ In other words, the propositional content of the assertion and the propositional content believed

⁴While Saul’s (2013) definition of lying is formulated in terms of saying rather than asserting, Stokke (2018, 31 fn. 13) refers to Saul as one of the writers who endorses an assertion-based account of lying, presumably because Saul’s (2013) notion of saying in combination with her notion of warrant amounts to assertion.

⁵Stokke (2024, 1) even goes as far as to claim that it is close to consensus that lies are assertions. Pepp (2018a) calls attention to the fact that the exact details of the belief-condition for lying are more controversial than they appear to be.

⁶If one does not accept the existence of so-called bald-faced lies (see Carson 2006 and Sorensen 2007 on bald-faced lies), one might want to include a third condition: intending to deceive the hearer about the proposition. Since most recent definitions of lying (e.g. Saul 2013; Stokke 2018 and Viebahn 2021) do not include this condition in their definitions of lying, I do not include it here either. Nothing crucial for the purpose of this paper hangs on this choice.

⁷Saul (2013, 7–8) makes a similar point about linking the speaker’s belief and her intention to deceive the hearer.

to be false have to be the same.⁸ We would, for example, not judge a speaker, who believes that it is false that all mammals lay eggs and asserts that Jupiter is the largest planet in our solar system, to be lying – because there is no link between the two propositional contents. In section 5, we will see that an expanded notion of content can do the linking job just as well as propositional content.

3. Expressivism

While nearly everyone accepts that some sentences have propositional content, there are philosophers who think that there are also sentences that do not. Expressivists, for example, think that there are sentences which express non-cognitive states of mind rather than propositions. One way to develop an expressivist view is to say that normative expressions – such as ‘wrong’ – and non-normative expressions have different kinds of meaning (Schroeder 2010, 21–22). While sentences containing non-normative expressions typically express propositions, which are true or false, sentences containing normative expressions do not express propositions but rather non-cognitive states of mind, such as approval or disapproval. For instance, Ayer (1936 [1971], 142) describes the difference between normative and non-normative expressions as follows:

If now I [...] say, ‘Stealing money is wrong’, I produce a sentence which has no factual meaning – that is, expresses no proposition which can be either true or false. It is as if I had written ‘Stealing money!!’ – where the shape and thickness of the exclamation marks show, by a suitable convention, that a special sort of moral disapproval is the feeling which is being expressed. It is clear that there is nothing said here which can be true or false.

This is not to say that every expressivist would agree with Ayer’s characterisation,⁹ but many people think that sentences like ‘Plagiarism is wrong’ do not have propositional content. Other philosophers have made similar claims about sentences involving taste predicates or epistemic modals.¹⁰ Note that the most common way of using ‘expressivism’ is a

⁸This description of linking makes it sound as if the contents of our beliefs and the contents of what we assert are the same objects, that is, that they are identical. To be clear, I am not assuming that the link has to be identity. It might be that the contents of our beliefs and the contents of what we assert are related through, for instance, entailment rather than identity. However, whether the link is identity or some other relation does not matter for the purpose of this paper because identity would be one straightforward way to deliver linking.

⁹For a different way to develop an expressivist view, see, e.g. Gibbard (2003).

¹⁰E.g. Ninan (2024) defends, and MacFarlane (2014) discusses without endorsing it an expressivist view on taste predicates; Yalcin (2011) and Hawke and Steinert-Threlkeld (2021) defend an expressivist view on epistemic modals.

form of non-cognitivism that is different from Ayer's expressivism. Others, mainly people who do not work in metaethics, use 'expressivism' as a label for all non-cognitivist views.

In this paper, I distinguish between two different kinds of expressivism which I label 'minimal expressivism' and 'radical expressivism'. They represent a useful way of carving up the different expressivist resources and commitments for the purpose of this paper. Most importantly, although minimal expressivism and radical expressivism can be mapped to modern forms of non-cognitivism and classical expressivist views, respectively, they are in the context of this paper primarily expressivist views about lying. The common denominator between minimal and radical expressivism is that they both accept that sentences such as 'The soup is tasty', 'Plagiarism is wrong' and 'The lady over there might be your mother' can be used to lie, but they reject the propositionalist claim that sentences can only be used to lie if they have propositional content.¹¹ The differences between minimal and radical expressivism arise from their different stances on the three key features of propositionalism. Depending on which kind of expressivism one accepts and thus which kinds of resources are available, the extent of revisions made to propositionalist definitions of lying will vary.

In the spirit of modern forms of non-cognitivism such as Gibbard's (2003) and Yalcin's (2011), *minimal expressivism* assumes a broader notion of content than propositional content, but accepts expressivist-friendly notions¹² of assertion, belief and truth-evaluability. Minimal expressivism claims that what is crucial to lying is not propositional content but the three key features assertion, belief and truth-evaluability. Since there are many resources available to minimal expressivism, we will see in section 5 that the revisions made to the propositional definition of lying are minimal in the sense that minimal expressivism merely expands the notion of content but otherwise preserves the propositional definition of lying. As the label suggests, *radical expressivism*, however, has to go for more radical revisions. In line with classical

¹¹Note that expressivists in general do not have to accept that sentences like 'The soup is tasty', 'Plagiarism is wrong' and 'The lady over there might be your mother' can be used to lie, but in this paper I will only discuss such expressivist views that accept this claim.

¹²When speaking about 'expressivist-friendly notions', I have in mind notions that fit with the overall expressivist framework. For example, an expressivist-friendly notion of truth-evaluability might be merely deflationary in the sense that there is nothing metaphysically or explanatorily deeper or more robust to truth than that 'The soup is tasty' is true if and only if the soup is tasty. For different versions of deflationism, see e.g. Field (1987) and Horwich (1998).

expressivist views such as Ayer's (1936 [1971]), radical expressivism dismisses all three key features of propositionalism, most importantly truth-evaluability. As we will see in section 6, the revised definition of lying based on radical expressivism shows us something unexpected about lying, namely that the crucial feature of lying is none of the three key features embraced by propositionalism, but a mismatch between attitudes held by the speaker and how the speaker presents herself in speech.

4. Expressive lies

In this section, I will argue that a diagnostic test for lying indicates that sentences like 'The soup is tasty', 'Plagiarism is wrong' and 'The lady over there might be your mother' can be used to lie. Defending this claim requires that sentences like 'The soup is tasty', 'Plagiarism is wrong' and 'The lady over there might be your mother' can be used to lie in a narrow sense of lying that singles out lying in a broader category of *insincerity*. In particular, the insincerity might consist in misleading without lying.¹³ For example, Gibbard (2003, 75) claims that one is insincere if one is not in the state of mind one expresses. Since expressing a state of mind is 'to purport to have it, whether or not one does' (Gibbard 2003, 75), one is insincere if one does not have the state of mind one purports to have. For instance, when a speaker utters 'Plagiarism is wrong' when she is not in a state of mind of disapproval towards plagiarism, she is insincere.

Usually, we assume that there is a difference between lying and insincerity. As Stokke (2018, 5) describes, insincerity is generally considered to be a broader category than lying:

Insincere speech encompasses many other forms of discourse than simple forms of lying to deceive others. We have various ways of deceiving each other with language while avoiding outright lying. We are skilful in navigating the landscape between lying and merely misleading. (...) We feign attitudes and emotions we do not have by using non-declarative language such as questions, imperatives, or exclamations. And we sometimes speak insincerely without aiming to deceive anyone.

¹³As Saul (2013, 20) points out, the distinction of interest is more precisely 'between lying and (intentionally) *merely* misleading'. I will use 'misleading' short for 'intentionally merely misleading' in this paper. Even more precisely, Saul (2013, 28) notes that 'we really have a three-way distinction – between lying, misleading, and accidentally saying something false'. I will ignore accidental falsehoods and focus on the distinction between lying and misleading. Pepp (2018b) contrasts lying with 'untruthfully implicating' rather than with 'misleading'.

Along the lines of Gibbard's thoughts, one might want to object to the claim that sentences like 'The soup is tasty', 'Plagiarism is wrong' and 'The lady over there might be your mother' can be used to lie, claiming that utterances of such sentences are not lies but merely insincere – in particular, cases of misleading rather than lies.

Conveniently, a diagnostic test that is used when one is uncertain whether an utterance is a lie or misleading can help us with the question whether utterances of sentences like 'The soup is tasty', 'Plagiarism is wrong' and 'The lady over there might be your mother' are lies or not. Among others,¹⁴ Viebahn (2021, 293–294) presents such a test to distinguish lying from misleading. The criterion for the test is 'whether speakers can consistently deny accusations of lying following their original utterances' (Viebahn 2021, 293). The speaker passes the test and thus counts as having misled if she can offer a consistent denial, and she fails the test if she cannot offer a consistent denial.¹⁵ It is not quite clear what exactly counts as a consistent denial, but I take it that the denial must not contradict the speaker's original utterance.¹⁶

Let us take a look at how the test works with a standard example of lying and misleading. Imagine a conversation between Anna and Ben. Anna had baked a chocolate cake and put it in the fridge, but now she notices that all that is left of the cake are a few crumbs. She asks her friend Ben, 'Did you eat the chocolate cake?'. Ben did eat the chocolate cake, but he replies, 'No, I did not eat the chocolate cake'. Later, Anna finds out that Ben ate the cake, and she confronts him by saying, 'You lied to me!'. Ben replies, 'No, I did not lie. I did not say that I did not eat the chocolate cake. I was just saying that I did not eat the chocolate cake'. That is not a consistent denial because it clearly contradicts Ben's original utterance. As a result, Ben fails the test, which indicates that he was lying.

Imagine now a slightly different version of the scenario. Anna asks Ben, 'Did you eat the chocolate cake?', and he did in fact eat it, but this time he replies, 'I do not like chocolate'. Again, later Anna finds out that Ben ate the cake, and she accuses him of lying. This time, Ben replies, 'No, I did not lie. I did not say that I did not eat the chocolate cake. I was just saying that I do not like chocolate' – which is true. This denial might

¹⁴See, e.g. van Elswyk (2020), Viebahn (2020), Stokke (2018) and Viebahn (2017).

¹⁵For the purpose of this paper, it only matters whether the speaker fails the test and thus is lying. What exactly the speaker does when passing the test – whether she is misleading, speaking insincerely in another way, or whatever else that is not lying – is not important.

¹⁶I am here using the term 'utterance' to make room for both assertions and expressions of attitudes such as approval or disapproval.

appear a bit nit-picky, but even if it is annoying, it does not contradict Ben's previous utterance. All he said was that he does not like chocolate. He merely implicated that he did not eat the chocolate cake because of his dislike of chocolate. As a result, he passes the test, which indicates that he was not lying but misleading.

Before applying the test to one of the three cases from above, a few notes are in order. Since I am discussing lying in relation to expressivism, I need to take into account whether someone who thinks that the semantics for sentences like 'The soup is tasty', 'Plagiarism is wrong' and 'The lady over there might be your mother' is expressivist has particularly good resources to reject that such sentences can be used to lie. Regarding the lying test, I will not try to answer the question what convinced expressivists would reply if they were accused of lying by uttering sentences like 'The soup is tasty', 'Plagiarism is wrong' and 'The lady over there might be your mother', but rather what kind of resources expressivists would say that ordinary language users have to reply to accusations of lying. For instance, a denial along the following lines would not make sense in ordinary language settings: 'No, I did not lie. I did not claim that the soup was tasty. In fact, I did not make a claim at all'.

Moreover, it is not clear how to think about consistency in the expressivist case in which we cannot rely on notions of assertion and truth-evaluability.¹⁷ Some might prefer to talk about coherence instead of consistency. I will assume that there is a way of making sense of a notion of consistency or coherence in expressivist-friendly terms,¹⁸ and for the purpose of the present discussion, I will stick with the term 'consistency'. Keeping these considerations in mind, let us apply the lying test to one of the three cases from above.

Remember, in *Soup*, your friend uttered:

(1) The soup is tasty.

Finding out that your friend did not find the soup tasty, you accuse her of lying. What can your friend respond to your accusation? She could reply:

(2) No, I did not lie. I did not say that the soup was tasty. I was just expressing my gustatory approval of the soup.

¹⁷Thanks to an anonymous referee of this journal for raising this issue.

¹⁸One way to spell out consistency in an expressivist way is to follow Gibbard (2003, chapter 3).

However, (2) is inconsistent with your friend's original utterance in (1). We are evaluating the responses to the lying test in a context in which we assume that to say 'The soup is tasty' *is* to express one's gustatory approval of the soup. So, claiming that she just expressed her gustatory approval of the soup will not help your friend to defend herself against the accusation of lying because expressing her gustatory approval of the soup *is* to say that the soup was tasty.

Note that I leave aside so-called exocentric cases which Lasersohn (2005, 670) describes as 'contexts in which someone other than ourselves is specified as the judge'. For example, imagine that you and your friend are buying dog food for your dog, and you say, 'Let's buy this dog food – it is delicious'. Later, your friend serves you the dog food for dinner. You refuse to eat it, and your friend says, 'But you said this dog food is delicious'. What you meant was of course that the dog food is delicious *for your dog*, but *not for you*. Excluding exocentric cases here is justified because all that I need for the purpose of this paper is *a* reading on which 'The soup is tasty' is a lie. It does not matter if one sometimes does not lie with sentences like 'The soup is tasty' as long as one sometimes *does*.

Since I am excluding exocentric cases, the speaker cannot use the following denial instead of denial (2):

- (3) No, I did not lie. I did not claim that the soup tastes good to me. I just meant that it would taste good to people who do not have a weird aversion to saffron.¹⁹

(3) is an exocentric case because in her denial, the speaker refers to other people than herself as the judge of whether the soup is tasty or not. However, in *Soup*, you have explicitly asked your friend whether *she* likes the soup. In that case, it is clear that the question under discussion concerns *your friend's* taste, and not taste in some exocentric sense. Denial (3) might at first sight look more plausible than denial (2), but it is still inconsistent because it contradicts your friend's original utterance (1).

Contrast your friend's original utterance (1) with a slightly different utterance: 'Fish is delicious'. Finding out that your friend did not find the soup tasty, you accuse her of lying. This time, your friend can reply:

- (4) No, I did not lie. I did not express my gustatory approval of the soup. I was just expressing my gustatory approval of fish.

¹⁹Thanks to Neri Marsili for raising this kind of objection.

Given that your friend loves fish, but the saffron ruined the soup for her, it makes sense for her to deny your accusation that she was lying. All your friend said was that fish is delicious. She merely implicated that the soup is tasty – because it contains fish that she finds delicious. This time, your friend passes the test, which means that she was not lying, but misleading, making you believe that she found the soup tasty.

To sum up, the lying test indicates that sentences like ‘The soup is tasty’, ‘Plagiarism is wrong’ and ‘The lady over there might be your mother’ can be used to lie, even on a narrow understanding that contrasts lying with misleading.²⁰ Therefore, I assume that there are expressive lies such as described in *Soup*, *Plagiarism* and *Mum*.

5. Minimal expressivism: lying with expressive content

In this section, I present a way to revise the propositional definition of lying based on the resources available to minimal expressivism. The proposal builds on an expanded notion of content that can play the role of propositional content in definitions of lying. I call the expanded notion of content based on work by Gibbard (2003), Yalcin (2011) and MacFarlane (2014) ‘expressive content’. Since minimal expressivism accepts expressivist-friendly notions of assertion, belief and truth-evaluability, one can simply replace propositional content with expressive content in the revised definition of lying. This move expands the notion of content figuring in definitions of lying, but otherwise preserves the three key features of propositionalist lying.

For the purpose of this paper, I assume that propositional content is content that has a truth-value that only varies across worlds.²¹ That is, propositional content determines a function from a possible world to a truth-value. That is on the face of it incompatible with expressivism because it is part of the expressivist views that interest me here that sentences do not express propositions. In contrast to propositional content, expressive content determines a function from a pair of a possible world and something else²² to a truth-value and is therefore a broader notion of

²⁰Due to space restrictions, I only walk through the lying test for *Soup* here, but one could spell out the test similarly for *Plagiarism* and *Mum*.

²¹Two remarks about propositional content: First, propositional content can also be understood in broader ways than I do here. Second, broader ways of understanding propositional content will give notionally different but substantively equivalent ways of formulating the proposal in this section.

²²Henceforth, I will use ‘something else’ as a generic term for whatever else the second component in the function from a world and something else to a truth-value is in expressive content. For instance, it might be a plan (in the normative case), a taste (in the case of taste judgements) or an information state (in the epistemic modal case).

content than propositional content. Due to space limits, I will here focus on a minimal expressivist view about taste²³ judgements, but something similar could be said about normative or epistemic modal judgements.²⁴

While propositional content is just a set of possible worlds, expressive content is a set of possible worlds *and* tastes. Some sentences (such as 'There is salt in the soup') have a content that is only sensitive to the world component. Some other sentences (such as 'The soup is tasty') have a content that is only sensitive to the taste component. Yet other sentences (such as 'The soup is tasty and there is salt in the soup') have a content that is sensitive both to the world component and to the taste component (MacFarlane 2014, 172).²⁵

Expressivists are not alone in appealing to contents that vary in truth-value across more than possible worlds. For instance, relativists like MacFarlane (2014) also make use of such contents. Therefore, expanding propositional content to expressive content is not sufficient for an expressivist view. To get a complete expressivist view, one needs more than an expanded notion of content. For instance, Gibbard (2003) offers a metatheory of content according to which contents are psychological states and MacFarlane (2014) suggests a post-semantics that can account for the difference between a relativist and an expressivist view about taste judgements. For the purpose of this paper, the details of what one adds to the expanded notion of content to make the view truly expressivist do not matter. Hence, I will focus on the notion of expressive content and the reader should feel free to complement it

²³When speaking about 'taste', I follow MacFarlane (2014, 143–144) in meaning a standard of taste, more specifically, a gustatory standard. This meaning of 'taste' as a standard of taste is different from objective taste – or what we might call 'flavour' – that is, different ways food can taste (e.g. salty, sweet, sour, bitter or umami). A standard of taste is a scale in the sense that one can evaluate different experiences of flavour more or less positively or negatively on that scale of tastiness. The flavours one likes are evaluated positively according to one's standard of taste, while the flavours one dislikes are evaluated negatively according to one's standard of taste. That is, one cannot dislike the flavour of chocolate but evaluate it positively according to one's standard of taste. Most often, we think of a standard of taste being an individual's standard of taste, but sometimes we can think of a group's standard of taste. Unless stated differently, I will mean a gustatory standard when speaking about 'taste' in the remainder of this paper.

²⁴What I say about expressive content is derived from Gibbard (2003), Yalcin (2011) and MacFarlane (2014). MacFarlane (2014, 167–175) presents, without endorsing it, a characterisation of expressivism about taste judgements that is inspired by Gibbard's (2003) expressivism about normative judgements. Yalcin (2011) defends a way of adapting Gibbard's view to epistemic modal judgements.

²⁵The taste component is here to be understood as affected by one's standard of taste, while objective taste (or flavour) affects the world component. For example, if one adds more salt to the soup, the objective taste of the soup alters, which ultimately affects the world component. Potentially, altering the objective taste of the soup might also affect the taste component, for example, if the soup gets so salty that one would evaluate it more negatively according to one's standard of taste. Consequently, 'The soup is tasty' could have been false if the standard of taste had been different in the context of utterance, but it could also have been false if the objective taste had been different (e.g. if the soup had been saltier) while the standard of taste stayed the same.

with whatever specific expressivist metatheory or post-semantics they find suitable.

What is the expressive content of 'The soup is tasty', then? The expressive content of 'The soup is tasty' is the set of all world-taste pairs at which the sentence 'The soup is tasty' is true. When the speaker utters 'The soup is tasty', she expresses a state of mind that agrees with all the states that judge the soup to be tasty and that disagrees with all the states that judge the soup not to be tasty (MacFarlane 2014, 171).

Expressive contents can be true or false relative to, for instance, a world and a taste. The important question here is whether minimal expressivism can make sense of what it is to believe that an expressive content is false. Making sense of what it is to believe that an expressive content is false requires a monadic truth predicate (i.e. 'is true' and 'is false'), as opposed to a relative truth predicate (i.e. 'is true relative to' and 'is false relative to'). Minimal expressivism can have such a monadic truth predicate that applies to expressive contents. To believe that p is true is basically just to believe p . It is to be in a state of mind that rules out every world-taste pair that p is false relative to. To believe that p is false is to be in a state of mind that rules out p being true. In other words, it is to be in a state of mind that rules out every world-taste pair that p is true relative to. This is not any more mysterious than being in a state of mind that rules out every possible world that p is true relative to.

Note that for minimal expressivism, the difference between different states of mind is a difference in content, not in attitude. That is why minimal expressivism accepts a notion of belief. One can believe that it is snowing outside, and one can believe that the soup is tasty or that plagiarism is wrong. The only difference between those beliefs is that they have different types of contents. There is no difference in attitude between the three beliefs.²⁶ The notion of belief accepted by minimal expressivism is deflationary or 'thin' in the sense that it is not reserved for states of mind with mind-to-world direction of fit only (see MacFarlane 2014, 167). Such states of mind with mind-to-world direction of fit are usually classified as 'thick' beliefs as opposed to states of mind like

²⁶This is not to suggest that states of mind never differ with respect to attitudes but only with respect to contents. One can, for instance, both hope and believe that it is snowing outside, which would be an example of difference in attitude but not in content. The point made on behalf of minimal expressivism is rather that what makes the state of mind different in the expressivist case compared to the non-expressivist case is the content. It is the content, not the attitude that differs between states of mind involving expressive contents and propositional contents. Minimal expressivism can therefore accept a notion of belief. One can believe an expressive content just as one can believe a propositional content.

tastes that have a world-to-mind direction of fit. As MacFarlane (2014, 173) points out, '[f]or the expressivist, then, the surface language of 'believing' masks a deep difference in psychological kind'. Minimal expressivism can accept that one can believe expressive contents like *the soup is tasty*, but this notion of belief is thinner than the one usually attributed to states of mind with mind-to-world direction of fit.

Minimalist expressivism can provide a story about how expressive content can be asserted, for instance with the help of a notion of assertion inspired by Stalnaker's (2002; 2014) notion of assertion.²⁷ According to Stalnaker (2014, 36), participants in a conversation share a pool of information that is updated throughout the conversation. They presuppose and take for granted certain information, they accept certain things for the purpose of the conversation, they add new information, and they remove information that is incompatible with what they commonly believe and know in the conversation. This information is the *common ground* among the conversational participants (Stalnaker 2002, 701). The common ground is a set of propositions, where propositions are understood as sets of possible worlds. The set of possible worlds that are compatible with the propositions in the common ground is the *context set* (Stalnaker 2014, 36). For example, if the proposition *it is snowing outside* is in the common ground, then it is snowing outside in every possible world in the context set.

According to Stalnaker (2014, 51), the characteristic way an assertion changes the context set is by adding information to the common ground. For example, when you make the assertion 'It is snowing outside', you propose to update the common ground with the piece of information that it is snowing outside (at the time and place of your utterance). Updating the common ground with information about the world means reducing the context set,²⁸ that is, eliminating the possibilities that are not compatible with the information gained from your assertion. As long as I go along with your assertion, the common ground will be updated in the way you propose. The effect of that update will be that we get a new, updated context set – the set of possible worlds that are compatible with the information that it is snowing outside. Possible worlds in which it is not snowing outside are eliminated from the

²⁷This is not a new move. Stokke's (2018) common ground definition of lying is also based on Stalnaker's notion of assertion.

²⁸Reducing the context set is not the only way to update the common ground. For example, when conveying permissions or using epistemic modals like 'might', one does not reduce but rather expands the common ground. See Rothschild and Yablo (2023) for an account of permissive updates.

context set. The updated background information of our conversation includes the information that it is snowing outside.

Something similar could be happening when one utters the sentence 'The soup is tasty'. The context set is then not just a set of possible worlds, but a set of world-taste pairs. It does not only include information about what the world is like, but it also includes tastes. For example, when uttering 'The soup is tasty', the speaker proposes to eliminate from the context set all the world-taste pairs at which the soup is not tasty. If the other participants of the conversation accept the speaker's assertion,²⁹ it becomes common ground that the soup is tasty. That is, they accept for the purpose of the conversation that the soup is tasty. In that sense, expressive content can be asserted.³⁰

Having established that expressive content can be truth-evaluable and believed to be false as well as asserted, one can simply plug in expressive content into propositionalist definitions of lying. Expressive content allows us to account for sentences like 'The soup is tasty', 'Plagiarism is wrong' and 'The lady over there might be your mother' being used as lies on a minimal expressivist view. Given that propositional content is a function from a world to truth-value and expressive content is a function from a pair of a world and something else to a truth-value, on the proposed revised definition of lying, expressive content replaces propositional content because we now have a different kind of object as content. The old kind of object, that is, propositional content, only varies in truth-value across worlds. The new kind of object, that is, expressive content, can vary in truth-value across worlds and something else (for instance, plans, tastes, or information states).

Among those new objects, some will be interesting, and some will be boring in the sense of Egan (2007, 11). The boring objects correspond to what we previously thought of as propositional content. They are boring because they only vary in truth-value across worlds, but not across other parameters such as plans, tastes, or information states. That is, the additional parameter is sometimes idle. For instance, let us say that expressive contents are true or false relative to worlds and tastes. That

²⁹A possible problem with this idea of accepting assertions like 'The soup is tasty' might be that aesthetic judgements are assumed to come with an experiential requirement that one has to have experienced the object one judges to be, for instance, beautiful or tasty (see, e.g. Franzén 2018). In that sense, it is not clear whether the other participants of the conversation can accept the speaker's assertion of 'The soup is tasty' if they have not tasted the soup themselves. However, accepting for the purpose of the conversation that something is tasty might not require all participants of the conversation to have tasted the soup.

³⁰See Yalcin (2018) for a more detailed account of asserting expressive content.

also goes for the expressive content *that grass is green*. However, in this case, the taste parameter is idle, and the truth-value only varies across worlds. The interesting objects are interesting in the sense that they vary in truth-value across both worlds and something else (for instance, plans, tastes, or information states). Since we have both boring and interesting objects among the new contents, the revised definition of lying can account for all the lies that propositionalist definitions of lying can account for as well as for expressive lies.

Based on the resources available to minimal expressivism, I suggest the following *minimally revised definition of lying* (where c_e stands for expressive content):

A is lying to B if and only if there is a content c_e such that A asserts c_e to B, and A believes that c_e is false.

The minimally revised definition of lying looks very much like Stokke's (2018, 31) propositionalist definition of lying, just with a broader notion of content. The difference to propositionalist definitions of lying is that when we consider whether A believes that c_e is false, we do not only need to look at worlds, but also at something else (for instance, plans, tastes, or information states).

How could the minimally revised definition of lying be spelled out for an example of a lie, say for 'The soup is tasty' uttered by your friend who does not find the soup tasty? By uttering 'The soup is tasty', your friend proposes to eliminate from the context set all the world-taste pairs at which the soup is not tasty. Your friend believes that it is false that the soup is tasty, but she proposes to update the context set with the expressive content that the soup is tasty. Since she believes that the content of her assertion is false, your friend is lying.

The minimally revised definition of lying expands the notion of content but preserves the three key features of propositionalist lying. Thus, as the label suggests, the minimally revised definition is very similar to propositionalist definitions of lying.

6. Radical expressivism: lying without truth-evaluability, assertion and belief

What about going further than expanding the notion of propositional content and also questioning all three key features of propositionalist lying, most importantly truth-evaluability? In this section, I propose a revised definition of lying based on the resources available to radical

expressivism, that is, a definition of lying that does not require truth-evaluability, assertion and belief. The revised definition that I will call ‘the attitude expression analysis’ accounts for lying in terms of directly and conventionally expressing an attitude one does not take oneself to have rather than in terms of asserting a propositional content one believes to be false.

How could one revise the propositionalist definition of lying based on the resources available to radical expressivism? One possibility is to take inspiration from Gibbard’s (2003, 75) idea mentioned in section 3, namely that insincerity means purporting to have a state of mind that one does not have.³¹ However, ‘purporting to have a state of mind one does not have’ is too broad for our purposes. Since we are interested in preserving the distinction between lying and misleading, Gibbard’s suggestion will not do the job. The lying-misleading distinction would collapse and leave us with mere insincerity.³² Nevertheless, there is something to Gibbard’s suggestion that one can make use of. The question is whether Gibbard’s suggestion can be restricted enough to preserve the lying-misleading distinction.

One possible way to restrict Gibbard’s idea is to replace ‘purporting to have a state of mind one does not have’ with ‘directly expressing an attitude one does not have’. As a starting point, ‘directly expressing’ means not to express something by expressing something else.³³ For instance, making an implicature is a paradigm example of expressing something by expressing something else. Adding ‘directly’ thus ensures that we can distinguish between lying and misleading because implicatures are not directly expressed. With an implicature, the speaker expresses one thing by expressing another. For example, when uttering ‘Fish is delicious’ in a context as described in section 4, the speaker expresses her gustatory approval of the soup *by* expressing her gustatory approval of fish. Implicatures used for misleading are thus not expressed directly, which means that they do not mistakenly count as lies on the revised definition.

The case of implicatures shows that ‘directly expressing’ means to express something in an unmediated way. However, the question is whether that is sufficient as a condition for lying. As we will see, the case of facial expressions illustrates that it is not a sufficient condition

³¹This notion of insincerity already figures in Searle’s (1969) work on speech-acts.

³²That is of course no real surprise since Gibbard (2003, 75) only speaks of insincerity, not of a narrower understanding of lying.

³³See Barber (2019) on a different notion of ‘expressing’ in the context of lying.

for lying to express an attitude one does not have in an unmediated way. Consider the following case:

Anger

Carla has very expressive facial expressions. Whenever Carla is angry with her partner David, she scowls in a highly recognisable way. This scowl directly, in an unmediated way expresses Carla's feeling of anger towards David, and David is easily able to recognise Carla's scowl as an expression of anger. One day, Carla is planning a surprise party for David. Carla does not want David to know about the party. To deceive David, Carla scowls at David despite not feeling any anger towards him. David recognises the characteristic scowl and is deceived. He comes to believe that Carla is currently angry with him.

This case intuitively feels like a case of somebody, here Carla, directly expressing an attitude that she does not have with the intention to deceive her audience. However, it does not seem to be a lie.³⁴ The revised definition of lying should be able to exclude cases like *Anger* from counting as lying.

Anger shows that it is not sufficient to directly – in the sense of being unmediated – express an attitude one does not have in order to lie. There are different options to further restrict the revised definition of lying. One option would be to add that the attitude in question does not only need to be expressed directly but also *linguistically*, for instance by uttering a sentence like 'The soup is tasty' or by holding up a sign with the words 'The soup is tasty' written on it. If one were to choose this option, one could exclude facial expressions such as in *Anger* from counting as lies. However, one would also exclude the possibility of lying with pictures because pictures do not express attitudes linguistically. It is controversial whether one can lie with pictures.³⁵ For the purpose of the present discussion, I would prefer to leave open the possibility of pictorial lies. Therefore, something different than 'linguistically' needs to be added to restrict the revised definition of lying.

A different option would be to restrict the revised definition of lying in terms of conventionality. That is, one would add that the attitude in question needs to be expressed *conventionally* as well as directly. There are different ways to spell out conventionality. Here, I have Lewis' (1969) notion of conventionality in terms of regularities which are based on certain expectations and preferences in mind. In this sense, an utterance

³⁴Thanks to Andrew Peet for presenting this example to me.

³⁵See e.g. Viebahn (2019) for an account of lying that includes pictorial lies, and see e.g. Stokke (2018) for an account of lying that excludes pictorial lies.

of the sentence ‘The soup is tasty’ is conventionally used to express the speaker’s gustatory approval of the soup. Adding ‘conventionally’ rather than ‘linguistically’ to the revised definition of lying, one leaves open the possibility to express attitudes in ways that do not require linguistic elements. In particular, one leaves open the possibility of lying with pictures, insofar as pictures can be used conventionally to express attitudes. Importantly, since facial expressions express emotions in a non-conventional way, adding ‘conventionally’ ensures that facial expressions like in *Anger* above can be excluded from the revised definition of lying.³⁶

Finally, one needs an insincerity condition in the revised definition of lying to make sure that the speaker misrepresents her attitudes and not just happens to express an attitude she unknowingly does not have. I therefore suggest replacing ‘directly and conventionally expressing an attitude one does not have’ with ‘directly and conventionally expressing an attitude one does not take oneself to have’.³⁷

Based on the resources available to radical expressivism, I propose the following *radically revised definition of lying* that I henceforth call ‘the attitude expression analysis’:

A is lying to B if and only if A directly and conventionally expresses to B an attitude that A does not take herself to have.

The attitude expression analysis of lying can account for both propositionalist lies and expressive lies. A’s utterance of ‘I did not eat the chocolate cake’ is a lie because A directly and conventionally expresses to B an attitude – in this case, the belief that she did not eat the chocolate cake – which she does not take herself to have. A’s utterance of ‘The soup is tasty’ is also a lie because A directly and conventionally expresses to B an attitude – in this case, her gustatory approval of the soup – which she does not take herself to have. As seen above, implicatures are excluded from the definition because the attitude in question is not directly expressed and thus the speaker is misleading but not lying.

Having presented the attitude expression analysis, in what remains, I will address what I take to be a serious challenge for this definition of

³⁶A third option would be to restrict the revised definition of lying in terms of commitment. That is, in order to lie one would need to directly express an attitude that one is committed to but does not have. This is a possible path to take, and following this path, one could for instance develop Viebahn’s (2021) commitment-based account of lying. However, as Pepp (2022) illustrates, it is quite difficult to develop an account of the relevant sense of commitment present in lying. Therefore, I have here chosen the conventionality path rather than the commitment path.

³⁷See Pepp (2018a; 2018b) for a nice overview of the complications that arise once one tries to untangle the details of untruthfulness in lying.

lying. The challenge arises from the common idea that non-declarative sentences cannot be used to lie. While this challenge is difficult, I will argue that there are some prospects to deal with it. My proposal is that the best way to deal with the challenge is to take a two-pronged approach. While one cannot lie with imperative and interrogative sentences because they do not directly express the relevant attitudes, it is hard to deny that exclamative sentences do express attitudes directly. However, I will argue that this is not an implausible consequence. While it is clear that one cannot lie with imperative and interrogative sentences, that is not clear in the case of exclamative sentences.

Non-declarative sentences such as imperative (e.g. 'Eat your soup!'), interrogative (e.g. 'Would you like more soup?') and exclamative expressions ('Yum!') are potentially problematic for the attitude expression analysis because they seem to express attitudes directly, but it is widely agreed that one cannot lie with non-declarative sentences.³⁸ Why would it be so problematic to include non-declarative sentences in the definition of lying? Could one not simply bite the bullet and accept that the attitude expression analysis comes with the cost that non-declarative sentences can be used to lie? Consider the imperative sentence 'Eat your soup!'. At least at first sight, it seems not entirely unreasonable to claim that by uttering the imperative sentence, the speaker expresses an attitude – for instance, the desire for the addressee to eat her soup – and that in case the speaker does not take herself to have that attitude, she is lying.

However, while most philosophers would agree that non-declarative sentences can be insincere, they cannot be used to lie.³⁹ For instance, it is clear that one can mislead by asking a question or issuing a command with which one intends to convey something one believes to be false, but one would not be lying in such cases. For example, when asking 'Is it snowing outside?', the speaker might convey that she is interested in finding out the answer to the question, and in case she

³⁸See, e.g. Stokke (2018, 199) on non-declarative sentences that can be insincere and used to mislead, but not to lie. Stokke's (2014) account of insincerity is inspired by Searle's (1969) insincerity conditions for different speech acts that also inspired Gibbard's (2003) account of insincerity.

³⁹Condoravdi and Lauer (2012, 43) even claim that '[t]he speaker of an imperative cannot be taken to be *insincere* with respect to the desire he communicates with an imperative'. This is shown by the infelicity of accusations of lying following utterances of imperative sentences. For example, one cannot respond to 'Eat your soup!' with 'You are lying, you do not really want me to eat my soup'. However, it is not quite clear how Condoravdi and Lauer's claim that it is impossible for a speaker to *lie* about her desire by using an imperative can establish the claim that the speaker cannot be *insincere* with respect to the desire. Remember that one can be insincere without lying, for instance by misleading without lying.

was in fact not interested in finding out the answer, she would mislead and thus be insincere, but not lying.

If non-declarative sentences cannot be used to lie, but seem to express an attitude the speaker does not take herself have, how can they be excluded from the attitude expression analysis of lying? One possible route would be to argue that when uttering a non-declarative sentence, the speaker does not *directly* express the attitude in question after all. What else would then be expressed directly by non-declarative sentences? There are different options. For the purpose of the present discussion, I will focus on imperative sentences, but something similar could be said about interrogative sentences.

Along the lines of Condoravdi and Lauer (2012, 38), one could say that what utterances of imperative sentences directly express is ‘a certain content related to the addressee’s future actions’, and that the speaker’s attitude – for instance, the desire for ‘the content to become reality’ – is merely conveyed.⁴⁰ When uttering ‘Eat your soup!’, the speaker would directly express the content of the addressee eating her soup, and the speaker would merely *convey* an attitude, namely the desire for the addressee to eat her soup. On such a view, imperative sentences would not be included in the attitude expression analysis because the attitude that the speaker does not take herself to have would not be expressed directly.

Along similar lines, a different option would be to take inspiration from Portner (2007) and say that when uttering an imperative sentence, the speaker directly expresses an item that is proposed to be added to the addressee’s To-Do List, and the speaker only indirectly expresses her attitude. When uttering ‘Eat your soup!’, the speaker would then directly express the item of eating the soup as proposed to be added to the addressee’s To-Do List, and only indirectly express her attitude of desire for the addressee to eat the soup. Again, on such a view, imperative sentences would not be included in the attitude expression analysis because the attitude that the speaker does not take herself to have would not be expressed directly.

The point is that when uttering an imperative sentence like ‘Eat your soup!’, there is no attitude that is directly and conventionally expressed. There are many attitudes that *could* be expressed, for instance the speaker’s belief that the addressee should eat her soup, or the speaker’s desire

⁴⁰Condoravdi and Lauer (2012, 38) note that this characterisation is true of directive and wish-type uses of imperative sentences, but not of all types of imperatives.

for the addressee to eat her soup, but those attitudes are not expressed directly and conventionally in the way as 'The soup is tasty' directly and conventionally expresses the speaker's gustatory approval of the soup.

Given the above, there are options to exclude non-declarative sentences from the attitude expression analysis of lying. At least, that applies to imperative and interrogative sentences. However, the explanation given for imperative and interrogative sentences is difficult to apply to exclamative sentences as they *do* express attitudes directly and conventionally. Having said that, I take it not to be implausible that one can lie with exclamative expressions. Consider the following case:

Chef

Your friend is an avid amateur chef and wants to apply to a famous cooking show on television. However, she is not sure whether her cooking skills are good enough. Before applying to the show, she therefore invites you over for a test round. She has prepared her signature dish and asks you to be her judge so that she can decide whether she should take the courage and apply to the cooking show. You taste the dish and think it tastes horrible. Despite the fact that you do not find the dish tasty, you utter 'Yum!'. Your friend comes to believe that the dish was tasty and applies to the cooking show but fails in the first round. Crushed, she realises that you did not find the dish tasty. So, she gets back to you and says, 'You lied to me! You said 'Yum!' although you did not find the dish tasty.'

I take it not to be implausible that you have lied by uttering 'Yum!'. If that is right, the attitude expression analysis can easily account for this exclamative lie: You have directly and conventionally expressed an attitude – your gustatory approval of the soup – that you do not take yourself to have, and thus you have lied.

Remember that we are assuming that expressivism is true about taste predicates, that is, that when uttering 'The soup is tasty', one is expressing one's gustatory approval of the soup, and we are assuming that one can lie in cases like *Soup*. Given that we already accept these two assumptions, the step to accepting exclamative lies like in *Chef* is not that big. If we accept that one can lie by uttering 'The soup is tasty', we should also find it plausible that one can lie by uttering 'Tasty!' if asked 'Is the soup tasty?'. After all, both utterances express the same attitude – namely the speaker's gustatory approval of the soup – and they do so directly and conventionally. Along similar lines, we should then find the step from 'This is yummy' to 'Yummy!' or 'Yum!' as plausible as the step from 'The soup is tasty' to 'Tasty!', especially given an expressivist

treatment which makes these utterances very similar. Again, the utterances express the same attitude. In fact, all of these utterances directly and conventionally express the speaker's gustatory approval of the soup.

To sum up, while one cannot lie with imperative and interrogative sentences because they do not directly express the relevant attitudes, it is hard to deny that exclamative sentences do express attitudes directly. However, I have argued that this is not an implausible consequence. While it is clear that one cannot lie with imperative and interrogative sentences, it is not implausible that one can lie with exclamative sentences. The attitude expression analysis can account for the difference between imperative and interrogative sentences on the one hand and exclamative sentences on the other hand.

7. Conclusion

In this paper, I have presented three cases involving expressive language that prompt us to reconsider the need of three key features figuring in propositionalist definitions of lying: assertion, belief and truth-evaluability. I have assumed that sentences including expressive language should be given an expressivist treatment, and I have argued that such sentences can plausibly be used to lie. The main aim of this paper has been to propose a revised definition of lying. Based on the resources available to two different expressivist views, I have proposed two ways to revise recent definitions of lying. While the first proposal results in a minimally revised definition that merely expands the notion of content but otherwise largely preserves all three key features of propositionalist definitions of lying, the second proposal diverges more radically from propositionalist definitions of lying. This second proposal that I have called 'the attitude expression analysis' shows us something unexpected about lying, namely that the crucial feature of lying is a mismatch between attitudes held by the speaker and how the speaker presents herself in speech rather than the assertion of a propositional content the speaker believes to be false.

I have identified a challenge for the attitude expression analysis that arises from the common idea that non-declarative sentences cannot be used to lie. While the challenge is difficult, I have argued that there are some prospects to exclude imperative and interrogative sentences from counting as lies. However, and maybe surprisingly, it turns out not to be implausible that exclamative sentences can be used to lie. If that is right, the attitude expression analysis can account for such exclamative lies.

Acknowledgements

Earlier versions of this paper were presented at the Higher Seminar in Philosophy at Umeå University (2021 and 2022), the Berlin Language Workshop (2022), the Umeå Winter Language Workshop (2022) and the Philosophy Seminar at KTH Royal Institute of Technology (2023). I am grateful to the participants for helpful discussion and comments on these occasions. I would also like to thank Emanuel Viebahn, Jessica Pepp, Nils Franzén, Andrew Peet as well as the anonymous reviewers and editors of this journal for helpful comments. I am particularly grateful to Torfinn Huvenes, Pär Sundström and Ethan Nowak for their extensive comments on earlier drafts of this paper.

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Luise Mirow  <http://orcid.org/0009-0001-9310-1462>

References

- Ayer, Alfred Jules. 1936 [1971]. *Language, Truth and Logic*. Harmondsworth: Penguin Books.
- Barber, Alex. 2019. "Lying, Misleading, and Dishonesty." *The Journal of Ethics* 24 (2): 141–164.
- Carson, Thomas L. 2006. "The Definition of Lying." *Noûs* 40 (2): 284–306.
- Condoravdi, Cleo, and Sven. Lauer. 2012. "Imperatives: Meaning and Illocutionary Force." *Empirical Issues in Syntax and Semantics* 9: 37–58.
- Dixon, Daisy. 2022. "Lies in Art." *Australasian Journal of Philosophy* 100 (1): 25–39.
- Egan, Andy. 2007. "Epistemic Modals, Relativism and Assertion." *Philosophical Studies* 133: 1–22.
- Fallis, Don. 2014. "Review of the Book *Lying, Misleading, and What is Said: An Exploration in Philosophy of Language and in Ethics*, by Jennifer Mather Saul." *European Journal of Philosophy* 22 (S1): e17–e22.
- Field, Hartry. 1987. "The Deflationary Conception of Truth." In *Fact, Science and Morality*, edited by Crispin Wright, and George McDonald, 55–117. New York: Blackwell Publishing.
- Franzén, Nils. 2018. "Aesthetic Evaluation and First-Hand Experience." *Australasian Journal of Philosophy* 96 (4): 669–682.
- Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Hawke, Peter, and Shane Steinert-Threlkeld. 2021. "Semantic Expressivism for Epistemic Modals." *Linguistics and Philosophy* 44: 475–511.
- Horwich, Paul. 1998. *Truth*. Oxford: Blackwell.
- Krstić, Vladimir. 2022. "On the Connection Between Lying, Asserting, and Intending to Cause Beliefs." *Inquiry* 68 (2): 643–662.
- Lasersohn, Peter. 2005. "Context Dependence, Disagreement, and Predicates of Personal Taste." *Linguistics and Philosophy* 28: 643–686.

- Lewis, David. 1969. *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.
- MacFarlane, John. 2014. *Assessment Sensitivity – Relative Truth and its Applications*. Oxford: Clarendon Press.
- Ninan, Dilip. 2024. "An Expressivist Theory of Taste Predicates." *Philosophers' Imprint* 24 (13): 1–23.
- Pepp, Jessica. 2018a. "Truth Serum, Liar Serum, and Some Problems About Saying What You Think is False." In *Lying: Language, Knowledge, Ethics, and Politics*, edited by Eliot Michaelson, and Andreas Stokke, 42–62. Oxford: Oxford University Press.
- Pepp, Jessica. 2018b. "Assertion, Lying, and Untruthfully Implicating." In *Oxford Handbook of Assertion*, edited by Sanford Goldberg, 829–850. Oxford: Oxford University Press.
- Pepp, Jessica. 2022. "What is the Commitment in Lying." *Journal of Philosophy* 119 (12): 673–686.
- Portner, Paul. 2007. "Imperatives and Modals." *Natural Language Semantics* 15: 351–383.
- Rothschild, Daniel, and Stephen Yablo. 2023. "Permissive Updates." In *Outstanding Contributions to Logic: Kit Fine on Truthmakers, Relevance, and Non-classical Logic*, edited by Federico L. G. Faroldi, and Frederik Van De Putte, vol. 26, 615–662. Cham: Springer.
- Saul, Jennifer Mather. 2013. *Lying, Misleading, and What is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford: Oxford Scholarship Online.
- Schroeder, Mark. 2010. *Noncognitivism in Ethics*. London: Routledge.
- Searle, John. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Oxford: Oxford University Press.
- Sorensen, Roy. 2007. "Bald-Faced Lies! Lying Without the Intent to Deceive." *Pacific Philosophical Quarterly* 88 (2): 251–264.
- Stalnaker, Robert. 2002. "Common Ground." *Linguistics and Philosophy* 22 (5): 701–721.
- Stalnaker, Robert. 2014. *Context*. Oxford: Oxford University Press.
- Stokke, Andreas. 2014. "Insincerity." *Noûs* 48 (3): 496–520.
- Stokke, Andreas. 2018. *Lying and Insincerity*. Oxford: Oxford Scholarship Online.
- Stokke, Andreas. 2024. "Lies are Assertions and Presuppositions are Not." *Inquiry*: 1–24.
- Sung, Winnie. 2023. "A Confucian Account of Trustworthiness." In *Perspectives on Trust in the History of Philosophy*, edited by David Collins, Iris Vidmar Jovanović, and Mark Alfano, 15–33. Lanham: Lexington Books.
- van Elswyk, Peter. 2020. "Deceiving Without Answering." *Philosophical Studies* 177 (5): 1157–1173.
- Viebahn, Emanuel. 2017. "Non-Literal Lies." *Erkenntnis* 82: 1367–1380.
- Viebahn, Emanuel. 2019. "Lying with Pictures." *British Journal of Aesthetics* 59 (3): 243–257.
- Viebahn, Emanuel. 2020. "Lying with Presuppositions." *Noûs* 54 (3): 731–751.
- Viebahn, Emanuel. 2021. "The Lying-Misleading Distinction: A Commitment-Based Approach." *The Journal of Philosophy* 118 (6): 289–319.
- Yalcin, Seth. 2011. "Nonfactualism about Epistemic Modality." In *Epistemic Modality*, edited by Andy Egan, and Brian Weatherson, 295–332. Oxford: Oxford University Press.
- Yalcin, Seth. 2018. "Expressivism by Force." In *New Work on Speech Acts*, edited by Daniel Fogal, Daniel W. Harris, and Matt Moss, 400–430. Oxford: Oxford University Press.