

Non-literal lies are not exculpatory

Hüseyin Güngör*

To appear in *The Philosophical Quarterly*

Word count: $\approx 9,907$ (everything included)

June 13, 2024

Abstract

One can lie by asserting non-literal content. If I tell you “You are the cream in my coffee” while hating you, I can be rightfully accused of lying if my true emotions are unearthed. This is not easy to accommodate under many definitions of lying while also preserving the lying-misleading distinction. The essential feature of non-literal utterances is their falsity when literally construed. This interferes with accounts of lying and misleading, because such accounts often combine a literal construal of what is said by an utterance with a falsity requirement for lying. In the presence of non-literal lies such definitions struggle to make plausible predictions for non-literal lies and merely misleading utterances together. In this article I aim to fix this by extending Daniel Hoek’s pragmatic account of conversational exculpation to assertions in general. Since this mechanism is designed to compute the intended meanings of non-literal utterances, it straightforwardly predicts non-literal lies to be as such. The lying-misleading distinction is also preserved, because merely misleading utterances arise out of exploiting a different pragmatic mechanism—Gricean additive implicatures. Along the way I also draw some general lessons about assertion and implicatures.


Keywords: lying · misleading · implicatures · pragmatics · nonliteral content

*Thanks to Justin Bledin, Dan Hoek, Sonya Ringer and two anonymous reviewers for very helpful discussions on this material. Especially, two anonymous reviewers made some insightful and invaluable suggestions which reshaped my thinking about the whole dialectic. I cannot thank them enough.

1 Introduction

You are on a game show where you are shown a map of Italy on a touch-screen and asked to click on the region where the city Crotoné is located. The map of Italy is partitioned into 20 regions (see the image below). Crotoné is located in Calabria which includes the region shaped like the arch of the Italian boot on a political map, but not the heel-shaped region of Italy. You have no idea where Crotoné is, but know that your friend Zeynep, who lived in Italy for 10 years, must know. You phone her in for help. Zeynep resents you for not helping her in the past and is happy it is payback time. You ask Zeynep where Crotoné is and she says:

- (1) Crotoné is in the heel of the Italian boot.



You click on what you take to be the heel of the Italian boot and proceed to get eliminated from the show.

If you are like me, you must feel like you have been lied to. However, you are not lied to simply because Crotoné is not located in the literal heel of a literal boot and what Zeynep said was *literally* false. After all Zeynep cannot excuse herself by saying

“Why did you believe me? You know Crotone is not located in a literal heel!”. This is because the phenomenology of lying disappears if we imagine Zeynep saying:

(2) Crotone is in the arch of the Italian boot.

Hearing this, you would proceed to click exactly on where Crotone is and collect your prize. In other words, you would not fault Zeynep for saying something false with (2), whereas you would do so with (1), although both are *literally* false.

In this paper I propose an account of non-literal lies without losing the lying-misleading distinction.¹ In the next section I start discussing some prominent accounts of lying. I argue that none of the accounts can account for non-literal lies without losing the lying-misleading distinction. Such accounts oscillate between making uncomfortable predictions about the status of non-literal lies and predicting merely misleading utterances to be lies. Then I proceed to provide an account of assertion by generalizing Daniel Hoek’s account of non-literal expressions to an account of assertion in general (?). Armed with such an account, non-literal lies are predicted to be lies as desired.

Then I turn to the lying-misleading distinction. It is often noted that merely misleading utterances are not tantamount to lies, because the false content is not asserted, but merely implicated (e.g. ?; ?). However, since the meanings of non-literal utterances are also taken to be implicatures of their literal meanings, they interfere with merely misleading utterances. Gricean implicatures are *additive* contents on top of the assertion (?; 374), whereas the intended meanings of non-literal utterances are themselves the contents of assertions. The latter ensures that the speakers are committed to the intended meanings of non-literal assertions, whereas implicatures do not. So the mechanism which allows one to merely mislead without lying is distinguished from the mechanism which recovers the intended meanings of non-literal utterances. I make the case that Hoek’s mechanism underlying non-literal utterances are much

¹I will only focus on the case of metaphors with the game show example, but our explanation can be extended to all the non-literal linguistic phenomena *conversational exculpature* claims to capture such as loose talk (?; §4.1) and expressions involving fictional components (?; §4.3). Furthermore, I will only focus on the cases involving assertions of declarative sentences. It has been argued that one can lie with presuppositions (?), pictures (?) and other non-assertive means (?). I set aside such issues for the sake of bringing out the distinctive overlap of non-literal assertions and lying, namely, the essentially false literal content of non-literal assertions and how this interacts with the lying-misleading distinction. Thanks to a reviewer for discussion here.

more prevalent and hence part of the mechanism for assertion, whereas Gricean implicatures are merely computed from the intended meanings determined by Hoek's mechanism. This explains our choice for embedding Hoek's parochial looking mechanism into an account of assertion.²

2 Accounts of lying vs. non-literal lies

2.1 Saul's account

Although lying has been one of the canonical examples of moral wrongs, the interest in its linguistic aspects is relatively new.³ Such interest is doubly justified. Not only do accounts of assertion and saying have to do justice to our intuitions about whether someone lied or not on a given occasion, but our intuitions about lying can also inform theories of linguistic communication (a point forcefully emphasized by ?, §2). The heart of this debate concerns the hardship to find necessary and sufficient conditions for an utterance to be an instance of lying, while distinguishing lies from merely misleading utterances. The lying-misleading distinction is best exemplified by a real life case discussed by Jennifer Saul (?, p. 95):

Samuel Bronston had both personal and company bank accounts in several countries. At his company's bankruptcy hearing, the following exchange took place between Bronston and a lawyer...:

Lawyer: Do you have any bank accounts in Swiss banks, Mr. Bronston?

Bronston: No, Sir.

Lawyer: Have you ever?

Bronston: The company had an account there for about six months, in Zurich.

Because Bronston himself had earlier had a large personal bank account in Switzerland, he was charged with perjury. The basis of the perjury charge was that, while his second utterance above was literally true, it was deeply

²For a similar argument in the context of Frege's puzzle, see ?.

³The list is long, but some prominent works in this vein include ?, ?, ?, ? ? and ?. ? ties the moral hazard of lying to its linguistic analysis.

misleading in that it conveyed that Bronson had never had a personal Swiss bank account. The eventual verdict by the US Supreme Court was that a merely misleading statement is not perjury.

The US Supreme Court's verdict is not an odd one. One can easily imagine Bronston felicitously defending himself by arguing that he *just* said that the company had an account in Switzerland, *not* that he himself did not have an account. Compare this with a hypothetical scenario where Bronston says that he never had a bank account in Switzerland. It would be very easy in this case to convict Bronston of perjury. It is not an easy to find an account of what is communicated by an utterance, which can distinguish a case of lying from a case of merely misleading, especially when we add usual cases of non-literal utterances to the mix.⁴

There are many competing accounts of lying which also preserve the lying-misleading distinction. Some of them do not touch the issue of non-literal lies. For instance, Saul (? , p. 3) brackets non-literal lies altogether in her definition of lying:

Lying-Saul: *If the speaker is not the victim of linguistic error/malapropism or using metaphor, hyperbole, or irony*, then they lie iff (i) they say that *p*; (ii) they believe *p* to be false; (iii) they take themselves to be in a warranting context.

Saul's account cannot pass any verdict for the case of lying in the game show case, since non-literal expressions do not satisfy the antecedent of her definition. It is not hard to understand why she sets non-literal lies aside. Saul employs a notion of saying which is aimed to capture the literal meaning of an utterance in order to predict Bronston's case above to be a case of misleading without lying.⁵ Saul cannot allow what is *implicated* by an utterance to be what is said by that utterance, since this would take the implicature of Bronston's second utterance that he did not have a bank account in Switzerland to be what is said. This would wrongly predict Bronston's assertion to be a lie instead of being merely misleading. Since what is communicated by a metaphorical utterance is taken to be its *implicature* (e.g. ? , 34), Saul leaves such expressions out of her analysis altogether not to lose the lying-misleading distinction.⁶

⁴This is also mostly true of written text, though I will restrict my attention to utterances.

⁵Saul's definition of saying involves some context sensitivity to establish truth-evaluability of expressions, e.g. expressions involving pronouns or adverbial modifiers which can only be saturated contextually. See ? (§3) for more details.

⁶The correctness of a *verbatim* Gricean account of metaphors is suspect. One reason is that

2.2 Stokke's account

However, not everyone restricts their definition of lying to literal cases. For instance, ?? provides the following definition of lying:

Lying-Stokke: A speaker *A* lies by uttering a declarative sentence *S* iff (i) *S* says that *p*, (ii) *A* proposes to make *p* common ground and (iii) *A* believes that $\neg p$.

Here the definition does not bracket non-literal utterances, so it is fair game to ask whether non-literal lies are predicted to be lies while preserving the lying-misleading distinction. First, I will take for granted the Stalnakerian notion of assertion with Stokke. According to the Stalnakerian framework the assertion of *p* is a bid to make *p* common ground among the interlocutors (?). More relevant to our purposes is Stokke's conception of *saying*. In order to maintain sensitivity to the lying-misleading distinction, Stokke mostly cleaves to a literal conception of meaning where the meaning of a declarative sentence is determined as a function of the lexical meanings of its constituents (i.e. *minimal content*: ?, 102). Stokke's novel addition to the literal conception is to suppose that what is said or said by an utterance is sensitive to the subject matter of a conversation or question under discussion (*QUD* in the lingo: ?, §3.4).⁷

QUD-sensitivity helps Stokke with certain cases where Saul's account cannot get right (see especially ? §4.3). An example from Stokke explains the need for QUD-sensitivity in accounting for the lying-misleading distinction:

At an office Christmas party, William's ex-wife, Doris, got very drunk and ended up insulting her boss, Sean. Nevertheless, Sean took the incident lightly, and their friendly relationship continued unblemished. More recently, the company was sold, and Doris lost her job in a round of general

metaphorical utterances do not seem cancellable. There is a conflict in saying "You are the cream in my coffee. In fact, I do not like you at all", which does not exist in saying "I ate some apples. In fact, I ate them all". Uncancellability of non-literal speech dovetails with the reason why they are lies in the first place. See ? and ? for further reasons of doubting a classical Gricean account.

⁷Stokke's QUD-sensitive notion of saying is part of a recent trend in philosophy of language. See ? for QUD-sensitive beliefs, ? for QUD-sensitive desires, ? for QUD-sensitive decision-making and ? for QUD-sensitive intentions. In some sense, Stokke's position is more fundamental in that, if true, it would entail the QUD-sensitivity of all *propositional* attitudes, because such attitudes are taken to be relations between agents and propositions, the latter of which is partly determined by QUD. Some doubt whether such attitudes are uniformly propositional in the first place (e.g. ?, ?).

cutbacks. But, despite this, Doris and Sean have remained friends. Some-time later, William is talking to Elizabeth, who is interested in hiring Doris. However, William is still resentful of Doris and does not want Elizabeth to give her a job.

- (3) a. Elizabeth: Why did Doris lose her job?
b. William: She insulted Sean at a party.
- (4) a. Elizabeth: How is Doris's relationship with Sean?
b. William: She insulted him at a party.

Here although both (3-b) and (4-b) have the same literal meaning, only (3-b) is a lie. Despite being deeply misleading, (4-b) is not a lie, since it only suggests or implicates that their relationship is bad without explicitly saying so. This can only be predicted if what is said by (3-b) and (4-b) is sensitive to the questions (3-a) and (4-a). So Stokke proposes that someone lies when they *say* something they believe to be false and propose to make it common ground, while they merely mislead without lying if they say (in Stokke's QUD-sensitive sense) something they believe to be true, while implicating something false.

Despite the addition of QUD-sensitivity Stokke still relies on a literal conception of what is said in order to capture the lying-misleading distinction. Such loyalty to the literal meaning spells trouble for Stokke by failing to predict intuitive non-literal lies to be lies. For instance, go back to the game show case:

- (5) Crotona is in the heel of the Italian boot.

The literal meaning of (5) is false, because the compositionally determined meaning of (5) from its lexical constituents implies the proposition that Crotona is physically located in a piece of footwear with a literal heel. Of course, Italy is not a boot and does not have a literal heel. However, this literal and false content is not what Zeynep proposes to make common ground—she wants to make common ground the intended proposition that Crotona is located in the region Apulia, which is located in the metaphorical heel. Although the intended proposition is also false, what is more important is that what is said by (5) according to Stokke's conception and what Zeynep proposes to make common ground with the assertion of (5) are not the same proposition. So (i)

and (ii) in **Lying-Stokke** are not satisfied together and hence (5) is not an instance of lying.⁸

In addition to falsely predicting non-literal lies to be non-lies, Stokke's view do not predict non-literal merely misleading utterances to be so, either. Stokke deploys implicatures to explain why (4-b) is merely misleading. The idea is that, since what William said with (4-b) is literally true, but implicates something false, William is misleading Elizabeth without lying to her. If this is the recipe for merely misleading utterances, Stokke does not predict merely misleading non-literal utterances to be so, either:

- (6) [It is Fatma's birthday. Even though Ali loves Fatma, he was too busy and forgot to buy her a present.]
- a. Fatma: Did you buy me a present?
 - b. Ali: You are the cream in my coffee.

Even though (6-b) is not a lie, it is misleading, since Ali communicates something true, i.e. he loves Fatma, while implicating something false, i.e. he bought her a present. However, Stokke predicts that Ali literally says and implicates something false. So he does not predict (6-b) to be merely misleading.

Perhaps we can improve **Lying-Stokke** by either tweaking the notion of what is said (i) or what is asserted (ii). Let us try each option. First, we can carefully integrate some implicatures to substitute for what is literally said by some utterances. This is a natural expansion of Stokke's account, since Stokke leaves a complete account of implicatures for future work (?, 113). Implicatures are often divided into two: *additive* and *substitutional* implicatures (?). Additive implicatures enrich the literal meaning of an utterance by strengthening with additional content. Stokke's canonical examples of merely misleading utterances all involve additive implicatures, since additive implicatures preserve the commitment to the literal meaning. Substitutional implicatures supplant the literal meaning of the utterance altogether and remove the commitment to the literal meaning. Metaphors and irony are taken to be two canonical examples of substitutional implicatures (?). We can weaken Stokke's conception of what is said

⁸? charges Stokke's account with the prediction that non-literal assertions all end up being lies. ? defends against Keiser by proposing that the literal content of non-literal utterances are not the content proposed for common ground uptake, even if they are false and believed to be false. Our argument accepts Stokke's defense, but goes in the other direction. Stokke cannot predict non-literal lies to be lies.

by adding a disjunct to the effect that sometimes what is said can be the substitutional implicature of the utterance. For each utterance, we first determine which disjunct constitutes what is said, we feed it into **Lying-Stokke** and determine whether a certain utterance is a lie or not.

There are several problems with such weakening. First, it is too narrow. Non-literal utterances are not the only way to lie. One can also lie with loose speech:⁹

- (7) [Ali really cares everyone arrive at the party at roughly 6:00. Zeynep knows Salim arrived at the party at 6:01, but she wants Ali to yell at Salim.]
- a. Ali: When did Salim arrive?
 - b. Zeynep: He did not arrive at 6.

Intuitively, (7-b) is a lie, even though Zeynep said something literally true.¹⁰ The issue for the disjunctive proposal is that neither the literal content nor the substitutional implicature seem to yield the right results. Taking the literal content implies that Zeynep did not lie by saying that Salim did not arrive at 6, since this is literally true and she knows it. It is not also clear if any substitutional implicature is triggered by (7-b). How is (7-b) a case of Zeynep saying one thing, but meaning something completely different? Loose speech does not seem similar to metaphors or irony in completely overtaking the literal meaning of an utterance. Weakening what is said by substitutional implicatures does not get at the generality of the phenomenon of interest.

Second, there is no uniform mechanism for determining when substitutional implicatures are triggered. So we do not have a principled way of determining when substitutional implicatures take over the literal content. The most widely accepted proposal is that substitutional implicatures are triggered when the maxim of quality is overtly flouted (?, 34). Both metaphors and irony seem to fit this bill.

This generalization is not correct, though. Go back to the game show case. Suppose that instead of directly saying that Crotona is in the heel of the Italian boot, Zeynep starts thinking aloud in order to project a false semblance of deliberation. She starts by remarking that she has been around the arch of the Italian boot and then says:

- (8) Crotona is **not** in the arch of the Italian boot.

⁹Thanks to a reviewer for providing an analogue of this case and discussion here.

¹⁰Contextual standards for precision may also shift the intuitions about whether such assertions are lies and the final account will accommodate this fact.

Given the context, (8) is a lie, since Crotone is located at the region metaphorically described by the negatum of (8). Not only is it a lie, but it is a remarkable lie, come to think of it. (8) is literally true, because Crotone is not physically located in the arch of a piece of footwear, and this literal truth is believed by the speaker. But it still is a lie. Negated non-literal utterances enable one to lie by asserting a believed truth!¹¹ Moreover, since the literal meaning is true, the conditions for triggering substitutional implicatures do not obtain. Intuitively, we expect the substitutional implicature for (8) to substitute for the literal meaning of (8), but (8) does not violate the maxim of quality, so the substitutional implicature is not triggered. This shows that allowing substitutional implicatures to be sometimes what is said by an utterance does not help in general, because there is not a clear recipe for when they are triggered.

What if we instead try to modify the clause (ii) of what is asserted instead of (i) of what is said in **Lying-Stokke**? Similar challenges arise, so I will be brief. We cannot say that what is proposed for common ground uptake by uttering “p” cannot be what is implicated by the token utterance of “p” in general, since this would violate the lying-misleading distinction. We can perhaps say that either the literal proposition or its substitutional implicature of an utterance is proposed to make common ground. But then we can run through the problems above for (ii). Negated non-literal utterances do not trigger substitutional implicatures and hence they end up not as lies. In general, there does not seem to be a straightforward way of accommodating non-literal lies in Stokke’s framework.

2.3 Viebahn’s account

Before turning to my own account of lying, I would like to discuss Viebahn’s commitment-based account of lying (?), one of whose touted virtues is its ability to account for non-literal lies. Viebahn is globally pessimistic for the prospects of what he calls a *content-based* approach to non-literal lies. This approach attempts to locate the lying-misleading distinction in the content of an utterance and how this content is communicated, e.g. said vs. implicated. According to Viebahn such an approach is doomed when it comes to accounting for non-literal lies. Instead he proposes what he calls a *commitment-based* definition to lying:

Lying-Viebahn: *A lies to B if and only if there is a proposition p such*

¹¹Thanks to a reviewer for emphasizing the importance of this case in the dialectic.

that: (i) *A* performs a communicative act *C* with *p* as content; (ii) with *C*, *A* intends to communicate *p* to *B*; (iii) with *C*, *A* commits herself to *p*; and (iv) *A* believes that *p* is false.

Viebahn explains the notion of commitment to *p* in terms of speaker responsibility to justify whether she knows that *p*. Viebahn argues that with the commitment approach, literal and non-literal cases of lying and being merely misleading can be accounted for. This is not hard to see. In the cases of non-literal lies, the content the speaker is committed to are not the literal meanings expressed by the utterance in question, but their intended meaning. After figuring out the intended meaning, we then plug the committed content through the definition of lying. Viebahn accounts for the lying-misleading distinction by pointing out that misleaders avoid such a commitment, because their commitment is to the true, but misleading proposition.

I have several issues with the commitment-based approach. First, strictly speaking, cases of non-literal lies spell doom for content-based approaches *only if* it is shown that there is no mechanism by way of which we can extract the intended meanings from the literal utterances while also preserving the lying-misleading distinction. However, no such argument is given by Viebahn. So it is possible that there is a content-based account which does just that.

Second, the lying-misleading distinction seems sensitive to the components in content-based accounts. Our discussion of Stokke's account showed that some non-literal lies involving negated propositions turn out to be merely misleading exactly because of the interaction between negated non-literal content and Stokke's analysis of being merely misleading in terms of implicatures. Bypassing a content-based approach altogether cannot do justice to such content-based distinctions.

Third, Viebahn's account contains a very important untold story. For instance, Stokke's account goes wrong in very specific ways when we consider the intuitive cases of non-literal lies and misleads, as we displayed above. Viebahn's discussion does not tell us where exactly such content-based accounts go wrong other than not being commitment-based. This is unsatisfactory, because one would like to know, for instance, how the invariable falsity of non-literal propositions interacts with believe-falsity requirement of lying or how the truth of an utterance utilized in being merely misleading is related to the irrelevance of the utterance to the subject matter at hand. For instance, at a coarse-grained level, Bronston manages to avoid lying because of a particular strategy: he seems to say something literally taken true, but irrelevant to

the question at hand. However, he does this, while succeeding at saying something implausibly irrelevant. However, switching to a commitment-based approach washes out the prospects for discovering such nuances.

None of these issues may be a problem for those onboarded to a broadly Brandomian framework where the notion of commitment is fundamental and assertion is explained partly in terms of it (?). However, those of us who think the reason why one commits to a claim is because they assert it would love to see their content-based options exhausted before admitting defeat in accounting for non-literal lies. For the latter Viebahn's commitment-based approach does not offer the type of linguistic nuances one wants in accounting for the lying-misleading distinction. Viebahn's account simply expresses that *if* we already know the intended proposition in cases of non-literal content, then this is the proposition to which the liars commit and misleaders avoid commitment. However, it does not tell us anything about how speakers and hearers are able to figure out which proposition it is.¹²

3 Conversational exculpature and lying

In this section I provide a content- or saying-based definition of lying which cover literal and non-literal cases alike, while preserving the lying-misleading distinction. I first generalize Daniel Hoek's pragmatic mechanism of *conversational exculpature* (?) to a general definition of assertion and then employ it in a definition of lying. The components of this mechanism allow us to avoid the issues we have raised for various accounts of lying. Importantly, we do not touch the definition of what is said by an utterance and outsource all the work to what is communicated by a token use of it. So we can keep the mechanism to determine the literal meaning of utterances and do not suffer from the issues of determining the appropriateness of substitutional implicatures. Since the mechanism completely takes over the assertion part of the definition and is always active, we do not also face the problems of trigger-conditions. We can employ Hoek's mechanism to account for non-literal lies, while having room for additive implicatures as well. Additive implicatures can still be deployed to account for merely misleading utterances. This division of labor will yield a fruitful account of lying and misleading.

¹²See ? and ? for similar complaints about Viebahn's treatment.

3.1 Preliminaries

Daniel Hoek introduced a pragmatic mechanism called *conversational exculpature* to compute the intended meanings of utterances which fail Gricean maxim of relation in general (?, 26-30).¹³ Hoek observes that many interesting literal or non-literal utterances require a repair mechanism, if we are to take them as relevant in discourse (?, 153). The general idea is as follows. Conversational exculpature is a correction mechanism for an utterance, when the speakers flouts the Gricean maxim of relation. The mechanism works by taking the literal meaning of an utterance and leverages it against background assumptions and a subject matter to recover the intuitively intended meaning of an utterance.

The idea of background assumptions will be key to motivating the generality of the mechanism. Hoek borrows the what Mandy Simons calls “contextual presuppositions” as background assumptions (??). Contextual presuppositions serve to establish the relevance of a putatively irrelevant assertion to the subject matter at hand. An example should be helpful. Suppose you have a class starting at 2:00pm and all the participants in the class know this. The teacher walks in and utters as the clock hits 2:00pm:

(9) Okay folks, it is 2:00pm.

(9) manages to convey the message that the class starts now, even though its literal content has nothing to do with when the class starts. Simon’s idea is that the contextual presupposition that the class starts at 2:00pm supplements the literal meaning of (9) to establish the relevance of (9) to the issue at hand.


Hoek builds on this idea to compute the intended meanings of non-literal utterances. Go back to the game show case. When Zeynep asserts that Crotona is in the heel of the Italian boot, the reason we understand what she intends to communicate is because we compute the intended proposition with the aid of a contextual presupposition. Essentially, we participate in a make-believe where we pretend as if the landmass of Italy consists of a gigantic piece of footwear, whose parts coincide with the actual geographical locations of the regions in Italy. The make-believe consists in pretending to believe (roughly) the proposition that:

¹³See ?. I will mostly focus on metaphors, but his mechanism extends to loose speech, non-catastrophic presupposition failures and many others. Hoek’s account builds on Yablo’s relevance semantics (?, §9-10). ? also argues the mechanism is useful for a solution to Frege’s puzzle.

the region separating the Adriatic and the Tyrrhenian Sea, which is in actual fact the locus of the landmass of Italy, is instead occupied by a boot of vast proportions. (?, 184)

We strengthen the literal meaning of “Crotone is in the heel of the Italian boot” with this contextual presupposition. The literal meaning under the pretense of the Italian landmass made up of a literal boot consists of all and only the worlds where the location of Crotone is fixed in the literal heel. In none of these worlds is Crotone located in the literal arch, vamp or upper of the Italian boot.

However, note that this strengthened proposition is not the intended proposition yet, since it is true *only* in the worlds where Italy is a literal boot. We want the intended proposition to tell us something about Crotone’s location in the actual world or worlds like ours where Italy is not a literal boot. To generate the intended proposition that Crotone is in Apulia (metaphorically described as “the heel of the Italian boot”), we leverage the relevance of the conjunction of the literal meaning conjoined with the pretense to a subject matter or QUD. We assume that the subject matter CROTONE’S LOCATION IN ITALY is a partition of the logical space which groups worlds which agree on the location of Crotone in their respective cells (we ignore the worlds where Crotone is not in Italy).



Each cell packs all the worlds which host Crotona in that respective region, but differ otherwise. The key detail is that, given the pretense about the literal Italian boot, the worlds where Crotona is in the literal heel of the literal Italian boot shares the same cell as the worlds where Crotona is in the metaphorical heel of the metaphorical Italian boot. Why? Note that the literal meaning of “Crotona is in the heel of the Italian boot” says nothing about where Crotona is relative to the regional partition of Italy. Crotona can be in the heel and belong to any region, since the parts of the literal Italian boot can vary in which regions they correspond to. This variance goes away, once we bring in the contextual presupposition.¹⁴ The presupposition fixes the distribution of Italian regions as they are in the actual world to the parts of the literal Italian boot, because in our pretense we take the landmass of *actual* Italy to be replaced by a literal boot. The pretense of treating Italy as a literal boot helps us smuggle the actual distribution of regions in Italy into the boot version of Italy. Once this is fixed, the literal meaning of “Crotona is in the heel of the Italian boot” becomes relevant to the subject matter CROTONE’S LOCATION IN ITALY. It is true only in the worlds where Crotona is located in Apulia (the region metaphorically described as “the heel”), be it

¹⁴To a great degree at least, it still allows variance within the regions.

in literal or metaphorical boot worlds, and false elsewhere.¹⁵

Once the conjunction is relevant, we can generate the intended proposition from it by projecting the truth-value of the conjunction at each world in a cell to its cellmates. If the conjunction $p \wedge q$ of the literal proposition p and the contextual presupposition q is true in some world w and w' is in the same cell as w , then we assume $p \wedge q$ is also true in w' . If $p \wedge q$ is false at w , then $p \wedge q$ is also false at w' . Although the conjunction lacks a truth-value in worlds like ours where the presupposition is not satisfied, the truth-value projection assigns the conjunction a truth-value in worlds where the conjunction lacks a truth-value. This final step yields the intended proposition which takes on only one truth-value in each cell of the partition—a proposition relevant to the subject matter CROTONE'S LOCATION IN ITALY.

3.2 Formalism

Now let me illustrate how all of that complicated process comes in a fairly straightforward formalism. A **full proposition** p expressed by an utterance “ p ” is a set of possible worlds where “ p ” is true, i.e. $p \subseteq \Omega$ where Ω is the logical space. We also have the notion of a **partial proposition**. A partial proposition is a pair of sets of worlds $\langle t_p, f_p \rangle$ where t_p is the set of worlds where “ p ” is true and f_p is the one where it is false. “ p ” lacks a truth-value outside $t_p \cup f_p$. Given the notion of a partial proposition, we can define a full proposition p to be just $\langle p, \neg p \rangle$. Computing the intended

¹⁵A reviewer rightfully queries whether the overlap of the literal proposition and the contextual presupposition is sufficient to fix the actual location of Crotone. The reviewer provides a helpful example to think not. Suppose a phone case is lying on top of a chess board and the top half of the case overlaps with E4 on the board. Is placing a finger on the top half of the case sufficient to place a finger on E4? Intuitively, it seems not. Similarly, merely superimposing Crotone being in the heel of a boot and a make-believe boot Italy does not mean we have fixed the actual location of Crotone. Two points. First, in certain versions of the phone case we can also treat merely putting our finger on the phone case as placing our finger on E4. This would be just like the Crotone case where we pretend we did. In fact, we may never literally place our finger on E4 due to electrostatic forces, but in our everyday descriptions of the scenario we describe *as if* we do. Such examples further exemplify the need for bridging the literal meaning of our assertions with contextual presuppositions. Second, the location of Crotone example is somewhat different from the phone case example. The phone case does not physically occupy the surface of E4 and hence one's finger does not contact E4. The pretense of treating Italy as if it is a literal boot takes the landmass of Italy itself or its counterparts in the relevant worlds to be occupied by a literal boot. So the geographical location itself which is Italy in the actual world corresponds to the literal boot in the make-believe worlds. There may be some hard issues concerning the crossworld identity of Italy, geographical locations, cities and other notions establishing the close relation between actual geographical regions and their boot-counterparts, but they are orthogonal to the illustration of the linguistic phenomena of interest here.

propositions requires that we **restrict** the literal proposition to the contextual presuppositions. This comes about by conjoining components of a partial proposition with a contextual presupposition. We define an operation $p \upharpoonright q$ on a pair of propositions which takes two propositions p, q and returns $\langle p \cap q, \neg p \cap q \rangle$. As the reader might guess by now, this operation will formalize the idea of pretending as if the landmass of Italy were occupied by a literal boot, when interpreting “Crotone is in the heel of the Italian boot”.

Finally, we **complete** a proposition by a **subject matter** (or a question) to obtain a full proposition. A **subject matter** or **question** Q is a partition of Ω .¹⁶ Two worlds which agree on Q (written $w \sim_Q v$) is placed in the same cell of the partition. We say a proposition p is **about** or **relevant to** Q iff p is a union of Q . p is not about or irrelevant to Q iff there is a p -world in every Q -cell. A **completion** of a partial proposition $\langle t, f \rangle$ by a subject matter written $Q(\langle t, f \rangle)$ is the possibly partial proposition:

$$Q(\langle t, f \rangle) := \langle \{w : w \sim_Q v \text{ for some } v \in t\}, \{w : w \sim_Q v \text{ for some } v \in f\} \rangle$$

Completion of a partial proposition by a subject matter attempts to provide a full proposition which is the intended meaning of a non-literal expression. If $p \upharpoonright q$ is relevant to Q , then Q can complete the partial proposition to yield a full proposition, which is the intended meaning of the non-literal utterance. If, on the other hand, $p \upharpoonright q$ is not relevant to Q , then the completion is not well-defined. Ignore ill-definedness for now. We will come back to it later. The intuition behind completion is that we leverage the bearing of the literal meaning under the pretense on the subject matter to gauge the truth-value of the literal proposition if the make-believe actually obtained. The whole process can be pithily summarized in the following diagram:

¹⁶See ?, ? and ? for this conception of subject matters.




Figure 1. The process of conversational exculpature. The light-, dark-gray and white areas are where a proposition is respectively true, false and lacks a truth-value.

Let us work through our running example of Crotona to put it altogether. To obtain the intended proposition communicated by “Crotona is in the heel of the Italian boot”, we first take the literal meaning, true in all the worlds where there is actually a boot occupying the Italian peninsula and Crotona is located in the heel of this boot. Taken literally, this proposition is irrelevant to the subject matter CROTONE’S LOCATION IN ITALY, since in some worlds where Crotona is located at the heel of a Italy-sized boot, this location does not match the expressed location of Crotona and there are also some worlds where it does. When we restrict the literal proposition to the contextual presupposition of the existence of a literal boot occupying the landmass of Italy, we obtain a proposition which is true only at worlds where the location of Crotona (the heel of the Italian boot) matches the actual heel-shaped region of Italy (Apulia). When we complete this proposition with the subject matter, we obtain the full proposition which expresses the proposition that Crotona is in southeast Apulia, which corresponds to the region of Italy which has a bird-view appearance of the heel of a boot. This way we

take the literal meaning of the sentence “Crotone is in the heel of the Italian boot” to its intended meaning which expresses the region-based location of Crotone with the help of a pretense that Italian landmass is occupied by a large boot and the subject matter CROTONE’S LOCATION IN ITALY. After completion, we are absolved or *exculpated* of gratuitous presuppositions made by the literal meaning or contextual presuppositions. The intended proposition does not entail either.¹⁷

3.3 Assertion, exculpature and lying

I propose an assertion-based definition of lying similar to Stokke’s, but substitute the intended proposition as obtained by exculpature for the literal proposition as the content proposed for common ground uptake:¹⁸

Lying: A speaker A lies by uttering a declarative sentence S in a context c iff (i) S says that p , (ii) A proposes to make $Q_c(p \upharpoonright q_c)$ common ground and (iii) A believes $\neg Q_c(p \upharpoonright q_c)$.¹⁹

Our proposal consists in taking what is communicated by an utterance to be its exculpated content. (iii) also changes respectively, because we take the speaker to believe not the literal proposition expressed by an utterance, but its exculpated content. It is easy to check that **Lying** easily predicts cases of non-literal lies to be lies for the obvious reason that it is tailor-made to recover the intended meanings of non-literal utterances. In the game show case, Zeynep lies, because she proposes to make the intended proposition that she believes to be false common ground.

Importantly, this definition predicts the negated non-literal lies to be lies as well. Recall that negated non-literal lies tied up Stokke’s account in many knots, because the literal proposition expressed by such an utterance is not only true, but also believed to be true. Negated non-literal utterances also mess up the conditions under which substitutional implicatures are triggered and hence dim the prospects for Stokke’s account even further. Just to recall:

¹⁷The process resembles cooking a turkey breast between two slices of veal for flavor and throwing the veal away. Here the veal-flavored turkey breast is the intended proposition, whereas the slices of veal are the literal meaning and the contextual presupposition.

¹⁸Subject matters and contextual presuppositions are supplied by the context, which I denote with subscripts, e.g. Q_c and q_c .

¹⁹ \neg is defined on partial propositions as follows: $\neg(\langle t, f \rangle) := \langle f, t \rangle$.

(8) Crotone is **not** in the arch of the Italian boot.

The reason why we can easily predict (8) to be a case of lying is thanks to one of the most important features of conversational exculpature: its transparency to negation.²⁰ That is, $Q(\neg p \upharpoonright q) = \neg Q(p \upharpoonright q)$. Intuitively, this is not hard to see. When we negate a proposition, we just invert the colors in the logical space (as depicted in **Figure 2**). But this just reverts the colors in the restricted regions as well. When we complete this proposition, we obtain the negation of the non-negated intended proposition. Since the intended meaning of (8) just is the negation of the intended proposition that Crotone is in the Calabria region, this satisfies **Lying** and hence predicted to be a lie as desired.




Figure 2. Transparency of negation with respect to conversational exculpature.

We may think that **Lying** is too fine-tuned to handle cases of non-literal lies and does not generalize. After all, we took a specific pragmatic mechanism to explain the intuitive relevance, but literal irrelevance of certain utterances and embed it in an account of assertion in general. In defense, we may note that communication in general rarely, if ever, involves us asserting the literal meanings of our utterances.²¹ As Hoek exemplifies (? , 155), conversational exculpature pervades linguistic communication. Furthermore, even in the cases where the literal proposition is proposed for common

²⁰See ? (pp. 171-172) for the proof.

²¹See Yablo (? , §9.17) for further defense of the generality of conversational exculpature in the context of Frege's puzzle.

ground uptake, this can be recovered as a special case of exculpature where none of the literal commitments of the literal proposition are waived. Note that if the literal proposition is already relevant to the subject matter at hand, the process of exculpature yields the exact same proposition as long as the contextual presupposition is irrelevant to the subject matter (as shown in **Figure 3.** below).²²




Figure 3. The literal proposition recovered from exculpature.

So the account is not as parochial as it may seem, just because it relies on a mechanism initially proposed to capture certain classes of linguistic phenomena.

When it comes to merely misleading utterances, I follow Stokke with one minor tweak. I take someone to be merely misleading without lying, when they assert something true, but *additively* implicate something false just like Stokke. However, the

²²If we have literally no contextual presuppositions, we can take the contextual presupposition to be *contentless*, meaning that it is Ω itself. It is trivial to show that $Q(p \uparrow \Omega) = p$ for any Q -relevant p . There is an interesting curiosity which I will not pursue here regarding the interpretation of Ω as the universal contextual presupposition. It seems to hint at a truth norm for assertion (as in ?), since what is common to all worlds in Ω is that they make the proposition Ω true.

content relative to which additive implicatures are computed are not the literal proposition, but the exculpated one. Again consider:

- (10) a. Fatma: Did you buy me a present?
 b. Ali: You are the cream in my coffee.

Ali misleads Fatma here without lying to her, because Ali implicates with the true intended proposition of his utterance, i.e. that Ali loves Fatma, something false, i.e. that Ali bought her a present. As long as additive implicatures are computed with respect to the intended propositions, we easily explain how one can merely mislead with non-literal utterances as well. Importantly, even though Ali is not committed to the literal meaning of (10-b), he is committed to the intended meaning of (10-b). Similarly, Ali is not also committed to the additive implicature derived from the intended meaning of (10-b). This difference in commitment is also underlined by our definition where the intended proposition is asserted and hence committed to, whereas the additive implicature is *not* asserted and hence not committed to.

Note that the process of exculpature provides a completely different proposition than the literal one. The literal proposition expressed by a non-literal utterance is irrelevant to a subject matter, whereas the intended proposition is relevant. The literal proposition is invariably false in the actual world, whereas the intended one can be true or false. So in a way the process of exculpature is a case of substitutional implicature. However, it does not face any of the problems we have raised for it in our critique of Stokke's account. One reason is that we have defined assertion *in general* in terms of exculpature rather than seeking some trigger-conditions. Another is that it covers a wider range of linguistic phenomena than substitutional implicatures. As Hoek demonstrates in detail (?, 175-177), exculpature takes care of loose speech, presupposition failures, fictionalism about abstract objects and many other linguistic phenomena where the literal utterance falls short of the intuitive subject matter in discourse. So despite seeming similar to substitutional implicatures, conversational exculpature is superior to them in covering non-literal lies.

The proposal also improves on Viebahn's commitment-based approach, because it provides a precise recipe for determining which content is asserted with non-literal content and which content is committed to in virtue of it. It also reverses the order of explanation for assertion and commitment. According to our definition, one is com-

mitted to a proposition, *because* one asserted it rather than vice versa. So for those of us committed to content-based accounts of lying, the proposal gives us some reason to resist the commitment-based approach insofar as non-literal lies go.

3.4 Whence subject matters and ill-definedness

There are two serious and related hanging issues for our proposal. To tease the first, let us revisit the non-literal merely misleading case:

- (11) a. Fatma: Did you buy me a present?
 b. Ali: You are the cream in my coffee.

(11-b) is merely misleading, because the intended meaning is true, but the implicature is false. The intended meaning, however, is not computed with respect to the question (11-a). If it were, then the intended proposition that Ali loves Fatma would be irrelevant to (11-a) and hence the result would be ill-defined. Instead the intended proposition is computed relative to the question “Does Ali love Fatma?” to which it is relevant. So this raises the question of where the subject matters and questions to compute the intended propositions come from in general.

In **Lying** questions and contextual presuppositions are supplied by the context. The default framework for integrating subject matters and questions in discourse is Craige Roberts’ (??). Roberts takes conversation to be an issue-directed enterprise. Roberts models issues as a stack of questions under discussion hierarchically ordered by their importance to interlocutors. (?, §1.2) Importantly, the stack consists of questions which are “unanswered and answerable”. (?, 15) We require a slight adjustment to Roberts’ idea by keeping track of both unanswered and answered questions in discourse. Although the unanswered question is ordered by a hierarchy of discourse-importance, we can treat answered questions merely as a set of questions consisting of the resolved questions from the stack.

The reason for this expansion is that computing some of the intended propositions may require answered questions to be computed. For instance, suppose that it was common ground between Fatma and Ali that Ali loves Fatma, so this issue is no longer under discussion. If we had only questions *under discussion*, then we could not interpret the intended meaning of Ali’s metaphorical utterance, since the question “Does Ali love

Fatma?” would not be available.²³ So instead of merely having Roberts’ questions *under discussion*, we take questions *resolved or under discussion* in a context c to determine Q_c .

The last issue concerns the worst case where neither resolved nor at-issue questions are available to compute the intended proposition. These are the ill-defined cases which I promised to come back to. Everything we have said so far has presupposed that a given exculpated proposition $Q(p \upharpoonright q)$ is well-defined. If it is not well-defined, then our definition blows up and nothing is communicated with an utterance. This cannot in general be right. Just because a certain assertion is irrelevant to some subject matter at hand (with or without a contextual presupposition) does not mean that nothing is communicated by this assertion. So we have to say what happens when $Q(p \upharpoonright q)$ is ill-defined. Hoek can skip ill-definedness for his purposes, because it is possible that no intended meaning is recovered from certain pairs of questions and contextual presuppositions. But since we employ the mechanism of exculpature in a general definition of assertion, we need to assign an asserted content to an well-formed utterance, even if it falls undefined by the process.

Fortunately, we can do this in a principled way. One of Roberts’ central assumptions is that (non-degenerate) conversations is geared towards figuring out the way things are or pinpointing which world is the actual one (? , 6). This is implemented by taking the finest-grained question “What is the way things are?” (called the “the Big Question”) always as part of the stack of questions. This question assigns each world its own cell and any proposition which is not Ω is relevant to this subject matter. Any proposition rules out the cells where it is false and thus bears on this question. So even if there is no other question in the set of questions resolved or under discussion such that $Q(p \upharpoonright q)$ is well-defined for a proposition p , $Q_{Big}(p \upharpoonright q)$ will be automatically well-defined due to every proposition other than Ω bearing on Q_{Big} . Further, this proposition will be equivalent to the full proposition $\langle p \cap q, \neg p \cap q \rangle$, since the completion relative to Q_{Big} trivially returns the truth-value of $p \cap q$ and $\neg p \cap q$ in each one-world cell. This guarantees a question to well-define the procedure of exculpature.

Considering the function of the Big Question in the context of non-literal speech helps us make interesting distinctions between irrelevant literal and non-literal asser-

²³If the issue of whether Ali loves Fatma is settled, why is (11-b) not redundant and off in the first place? There can be multiple reasons for asserting common-ground information such as reminding and making it more salient (?, ?)

tions in a discourse. Since the literal content of non-literal utterances is by default false in the actual world, no non-literal utterance can express something true in the actual world through conversational exculpature. Depending on the truth of the contextual presupposition, it can either be truth-valueless or false. Questions other than Q_{Big} could then complete project the restricted proposition to make the proposition true in a world which otherwise assigned no truth-value. However, Q_{Big} places each world in its own cell, so there is no truth-value projecting among cellmates. In general, the intended proposition expressed by a non-literal utterance can be true or false in a world only if there is a question or subject matter *other than* the Big Question relative to which the intended meaning of the non-literal utterances are computed.

By contrast, this is not necessary for Simons' original cases for motivating contextual presuppositions:

(12) Okay folks, it is 2:00pm.

Unlike non-literal utterances like metaphors, the intended meaning of (12) can non-trivially vary in truth-value according to the truth-value of the contextual presupposition in a given world, i.e. whether the class starts at 2:00pm or not in a given world. This is because the truth-value of the literal meaning of (12) can covary with the truth-value of the context presupposition that the class starts at 2:00pm. There is more interesting work to be done here, but it will suffice to observe that in order for non-literal utterances to convey non-trivial and informative propositions, a question other than the Big Question is required.²⁴

4 Conclusion

In this paper I proposed a definition of lying which accounts for literal and non-literal lies alike, while preserving the lying-misleading distinction. First, I argued this is a much harder task than it is appreciated by showing how various accounts of lying

²⁴We mentioned that Hoek's account builds on ?. The biggest difference between them is Hoek's dependence on subject matters to compute the intended proposition. Hoek argues (?, 164-166) that subject matter-dependence is required to resolve certain ambiguities among possible intended propositions, even if there is a unique literal proposition and contextual presupposition. Our observation not only suggests subject matter-dependence, but specific subject matter-dependence for non-literal utterances. In order to obtain intended propositions which can be true in the actual world, one cannot use the Big Question—we need a more indirect question to get there.

and misleading cannot capture non-literal lies. Then I showed how my proposal does so. The lying-misleading distinction is preserved, because the pragmatic mechanism which interprets non-literal utterances is wholly distinct from Gricean *additive* implicatures. Along the way we have learned many lessons. In my opinion one of the big lessons is how hard it is to maintain a distinction between additive and substitutional implicatures, given that the trigger-conditions for the latter are utterly unclear. As far as I know, this is not noted in the literature. Also I think the negated non-literal lies are also never discussed in the literature and I take them to be the final tribunal for any putative general account of lying. The transparency of negation in Hoek's mechanism gets such lies exactly right. I think there is a lot more work to be done for the framework. The reader may also suspect it given the somewhat messy nature of our solutions to the problems in the previous section. I leave systematization of our remarks to future work. What I hope to have shown is that non-literal lies are not just a fringe case in the literature of lying and misleading—they not only adjudicate between accounts of lying, but in the process of doing so teaches us lessons about correct theories of implicatures.