

Lying and Misleading in Discourse^{*}

Andreas Stokke

— *penultimate draft, forthcoming in the Philosophical Review* —

Abstract

This paper argues that the distinction between lying and misleading while not lying is sensitive to discourse structure. It is shown that whether an utterance is a lie or is merely misleading sometimes depends on the topic of conversation, represented by so-called questions under discussion. It is argued that to mislead is to disrupt the pursuit of the goal of inquiry, i.e., to discover how things are. Lying is seen as a special case requiring assertion of disbelieved information, where assertion is characterized as a mode of contributing information to a discourse that is sensitive to the state of the discourse itself. The resulting account is applied to a number of ways of exploiting the lying-misleading distinction, involving conversational implicature, incompleteness, presuppositions, and prosodic focus. It is shown that assertion, and hence lying, is preserved from subquestion to superquestion under a strict entailment relation between questions, and ways of lying and misleading in relation to multiple questions are discussed.

1 Introduction

We sometimes have goals that urge us to mislead each other in conversation. Typically, we can meet them in two different ways, namely either by outright lying or by misleading while avoiding lying *per se*. There is a difference between lying and merely misleading. This distinction is important to us. We often take pains to stay on the right side of it in everyday matters. We build it into law codes, and it is a basic distinction in many religious systems of belief. There are famous cases of presidents and saints having exploited the difference dexterously, as do the rest of us with varying degrees of regret.

As the 1865 vote on the 13th Amendment was due to take place, Confederate representatives were traveling north for peace negotiations. On 31 January James Ashley wrote to

^{*}Thanks to Thomas Carson, Don Fallis, Torfinn Huvenes, Jed Lewinsohn, Salvador Mascarenhas, François Recanati, Jennifer Saul, and to audiences at Umeå University, Stockholm University, and Institut Jean Nicod (ENS, Paris) for helpful discussion. I am particularly indebted to Anders Schoubye and an anonymous referee for the *Philosophical Review* for extensive and very valuable suggestions.

Lincoln, “The report is in circulation in the House that Peace Commissioners are on their way or are in the city, and is being used against us. If it is true, I fear we shall loose the bill. Please authorize me to contradict it, if not true.” Lincoln wrote back, “So far as I know, there are no peace commissioners in the city or likely to be in it.”¹ In fact, the commissioners were on their way not to Washington, but to Fort Monroe, where Lincoln met them a few days later. Or consider the often cited case of Saint Athanasius who, when asked by pursuers sent by the emperor Julian to persecute him, “Where is the traitor Athanasius?”, replied, “He’s not far away.”²

This paper argues that the distinction between lying and merely misleading is sensitive to discourse structure. In particular, I argue that whether an utterance is a lie or is merely misleading depends on the topic of conversation, understood as the *question under discussion* (henceforth, QUD) in the sense of Roberts (2004), (2012). In the tradition from Stalnaker (1978), (1984), (1998), (2002), a discourse is taken to be a cooperative activity of information exchange aimed at the goal of inquiry, i.e., to discover how things are. In this setting a QUD is a subinquiry, that is, a strategy for approaching the goal of inquiry.

I propose that to mislead is to disrupt the pursuit of the goal of inquiry, that is, to prevent the progress of inquiry from approaching the discovery of how things are. On the view I will argue for, the difference between whether doing so counts as lying or as merely misleading depends on how one’s utterance relates to the QUD one is addressing.

The distinction between lying and merely misleading has traditionally been studied from the point of view of ethics.³ But recently, philosophers – e.g., Williams (2002), Carson (2006), (2010), Sorensen (2007), Fallis (2009), Stokke (2013), Saul (2012b) – have approached the distinction from within the philosophy of language in order to understand the linguistic difference between lying and merely misleading.

It is generally agreed that what distinguishes lies from utterances that are misleading but not lies is that the former convey misleading information with a certain kind of directness or explicitness. Accordingly, the standard approach has been to argue that, in lying, one *says* or *asserts* something misleading, as opposed to conveying it in some other, more indirect way.

This approach is well motivated by the kinds of cases that have traditionally been discussed in this area. The classic contrast between lying and merely misleading is the contrast between asserting or saying something one believes to be false (lying) vs. asserting or saying something one believes to be true in order to *conversationally implicate* something one believes to be false (merely misleading.)⁴ As an example, consider the following situation and the

¹See Lincoln (1953).

²See, e.g., Williams (2002, 102) who quotes a version on which Athanasius’s reply is, “Not far away.” I include the pronoun because it allows me to discuss ways of misleading with the person features of pronouns. See also MacIntyre (1995, 336).

³A central part of this literature stems from Kant’s (1797) notorious position that lying is always morally wrong. For discussion, see, e.g., Paton (1954), Isenberg (1964), Bok (1978), Kupfer (1982), Korsgaard (1986), MacIntyre (1995), Mahon (2006), (2009), Wood (2008), Carson (2010). More recently, Adler (1997), Williams (2002), Saul (2012a), (2012b) have questioned the moral significance of the distinction between lying and merely misleading.

⁴It is controversial whether lying requires believing the negation of the relevant content or merely not believing it. This issue is not under discussion in this paper. I assume the stronger view, since it makes things easier.

contrast between the dialogues in (1) and (2).⁵

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

- (1) Rebecca. Are you going to Paul's party?
Dennis. No, I'm not going to Paul's party.
- (2) Rebecca. Are you going to Paul's party?
Dennis. I have to work.

In both cases Dennis conveys the misleading information that he is not going to Paul's party. But while Dennis's utterance in (1) is a lie, his utterance in (2) is not a lie.

The standard approach explains this difference by pointing to the fact that, in (1), the misleading information is said or asserted, while, in (2), the same misleading information is merely implicated. In the former case Dennis says or asserts something he believes to be false. In the latter he says or asserts something he believes to be true, but thereby implicates something he believes to be false.

However, as I demonstrate in this paper, once we look beyond classic cases of this kind, we need discourse-sensitive notions of saying and asserting in order to capture the lying-misleading distinction. Some accounts of the lying-misleading distinction, like that of Saul (2012b), allow what is said to go beyond what is linguistically encoded by an utterance. Yet, as we will see, such views are still unable to account for the way the distinction depends on the state of the discourse.

I accept the view that lying requires assertion, and that assertion requires saying something, as opposed to conveying information in indirect ways, e.g., by conversational implicature. But I argue that what is said by an utterance, in a particular context, depends directly on the QUD that is addressed in that context. As a consequence the lying-misleading distinction is sensitive to the information structure of the discourse, that is, to the configuration of QUDs.

The view I defend agrees with Contextualist positions in the debate over the semantics-pragmatics distinction that what is said may go beyond linguistically encoded meaning. Yet I argue that what is said is nevertheless strictly constrained by a minimal kind of compositional meaning. As a result, on this view, what is said by a sentence in a context is determined systematically by the linguistic meaning of the sentence and the configuration of QUDs in the context.

In section 2 I present some initial evidence for the discourse-sensitivity of the lying-misleading distinction, and I show that this evidence presents a problem for Saul's (2012b) account. Section 3 details my own account based on the QUD framework. In section 4 I show how this account handles a number of different ways of exploiting the lying-misleading distinction in discourse. I apply the account to the classic contrast between assertion and implicature as well as to ways of committing oneself to misleading answers and of exploiting

Everything I will have to say can be transposed, *mutatis mutandis*, to fit the weaker view.

⁵Adapted from Davis (2010).

certain kinds of incompleteness for misleading purposes. In addition, I show that the account explains how one can use presuppositions to mislead. Finally, section 5 discusses two complications concerning QUDs. The first is that QUDs are often implicit. I discuss ways of using prosodic focus to indicate which QUDs are being addressed, and I show how this affects the possibilities of lying and misleading. Second, I discuss the consequences for the lying-misleading distinction of the fact that contexts typically contain multiple QUDs ordered by certain kinds of entailment relations. I show that assertion and lying are preserved in a particular way under a strict notion of question-entailment. Further, I demonstrate that entailment relations can explain that it is always an option to lie, whereas whether one can succeed in misleading while avoiding lying depends on which QUDs are in place.

2 Lying, Misleading, and Discourse-Sensitivity

In this section I present two initial kinds of evidence showing that the lying-misleading distinction is sensitive to discourse structure. The first comes from cases in which the difference arises depending on whether the utterance directly commits the speaker to a misleading answer to a QUD or not. The second kind concerns cases that exploit a particular kind of incompleteness in order to convey misleading answers.

2.1 Lying and Misleading

As illustrated by (1)–(2), the difference between lying and merely misleading is typically taken to be a difference between asserting disbelieved information and conveying it some other way. This difference carries strong weight in particular kinds of official discourse such as political and court room testimonies. Politicians are typically forbidden to lie. But in some situations they are not forbidden to carefully select what to say, even if doing so may be misleading, as long as their utterances are not outright lies.⁶ A similar strict implementation of the lying-misleading distinction is found in many legal systems. Saul (2012b) cites the following example pertaining to US law:⁷

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.

⁶See Williams (2002, 108–109) for discussion.

⁷Cf. *Bronston v. United States* 409 U.S. 352 (1973). See also Solan and Tiersma (2005) for discussion.

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 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.⁸

The difference between lying and other forms of misleading speech also plays a central role outside special contexts of this kind. A real estate agent is not allowed to lie to you about the state of the house she is trying to sell, but she may not be accountable for misleading utterances that are not lies. And even about trivial, everyday matters, many people will choose to say something true but misleading rather than to outright lie, all else being equal.

There are most likely many reasons for this attention to the difference between lying and merely misleading. One suggestion is that part of the importance placed on avoiding lying *per se* stems from the role assertion plays in the system for accumulating and pooling information that arguably underlies human development and cooperation.⁹ As I will argue, both lies and merely misleading utterances are ways of disrupting communal inquiry, i.e., the pursuit of truth. Yet lying is distinguished by exploiting a mode of contributing information to a discourse that involves a particular kind of commitment. Information conveyed by lying, as opposed to merely misleading, may thereby be more likely to be relied on by hearers in subsequent deliberation and action.¹⁰

But moreover, as we will see next, there is evidence that the difference between lying and merely misleading depends on the state of the discourse itself.

2.2 Committing to Misleading Answers

Consider the following story:

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 cut backs. But, despite this, Doris and Sean have remained friends. Sometime 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) Elizabeth. Why did Doris lose her job?
William. She insulted Sean at a party.

⁸Saul (2012b, 95).

⁹For relevant discussion, see Williams (2002).

¹⁰This is compatible with the views of, e.g., Adler (1997), Williams (2002), and Saul (2012a), (2012b) according to which there is no significant moral difference between lying and merely misleading, and as such is not necessarily a way of "fetishizing" assertion, to use Williams's (2002) phrase. That assertion plays a special role in the collective activity of information gathering does not imply that asserting disbelieved information is morally worse than other ways of misleading.

- (4) Elizabeth. How is Doris's relationship with Sean?
William. She insulted him at a party.

I think that William's utterance in (3) is a lie and that his utterance in (4) is not a lie, even though it is clearly misleading.¹¹ The reason for the difference seems clear. In both cases William's utterance provides a misleading answer to the question it is addressing. In the first case it provides the answer that the reason Doris lost her job was that she insulted Sean at a party. In the second case it provides the answer that her relationship with Sean is not good. Yet we have a strong sense that, whereas in the first case the answer is provided directly, or explicitly, in the second case the misleading answer is supplied indirectly, or implicitly.

At the same time, the only substantial difference between the dialogues in (3) and (4) is which question is being addressed. In other words, the same utterance (*modulo* the pronoun) is, in one case, a lie and, in the other case, merely misleading. This suggests that the difference between lying and merely misleading is sensitive to previous discourse structure, and in particular to the topic of conversation, or QUD. In other words, whether you lie or merely mislead depends on which question you are interpreted as addressing.

These observations can be corroborated by considering possible continuations of the discourses. In general, it is characteristic of utterances that are misleading but not lies that one can subsequently retreat from the misleading information one conveys. That is, one can deny that one intended to convey the relevant misleading information, while this does not involve retracting one's utterance completely. Correspondingly, when one is lying, one is typically committed to the misleading information one conveys in a particular sense.¹² Part of the reason the contrast between assertion and conversational implicature has typically been used to exemplify the difference between lying and merely misleading is arguably that it represents a contrast between a committing and a less committing way of communicating.

To illustrate, consider the continuations of our previous examples in (1') and (2').

- (1') Rebecca. Are you going to Paul's party?
Dennis. No, I'm not going to Paul's party.
Rebecca. Oh, don't you think he'll be disappointed?
Dennis. #No, I'm going to the party.

- (2') Rebecca. Are you going to Paul's party?
Dennis. I have to work.

¹¹I use "utterance" to mean what Stanley and Szabó (2000, 77–78) call "grammatical sentence", that is, roughly, disambiguated phonological form (or the equivalent, for written sentences.) Hence, an utterance, in this sense, does not include pragmatic enrichments or implicatures, and does not include saturation of indexicals such as the pronouns, demonstratives, or indexical adverbs of space and time. An utterance, in this sense, is not identical to what I call *minimal content* (see section 3.4), since the latter includes saturation of indexicals. In terms of the account I propose in section 3, utterances, on this terminology, determine minimal contents, in context, which in turn determine what is said, given a QUD. I allow myself to not strictly distinguish between utterance tokens and utterance types, and I do not always distinguish between "an utterance" in the sense of an act and in the sense of the object of such an act. By locutions such as *a's utterance u was a lie* I mean that *a* lied by making *u*.

¹²I am not suggesting that the impossibility of retreating in the way illustrated in the text is either necessary or sufficient for having lied, but merely that such impossibility is good evidence.

Rebecca. Oh, don't you think he'll be disappointed?

Dennis. No, I just meant I have to work now.

In both cases Dennis's initial utterance conveys the misleading information that he is not going to the party. When this is done by lying, as in (1'), the speaker is committed to this information, as witnessed by the infelicity of attempting to retreat from it in subsequent discourse. By contrast, when misleading while avoiding outright lying, as in (2'), the speaker is not so committed in that he can subsequently retreat from the misleading information.¹³

It is important to emphasize that we are not claiming that retreating subsequent to having been merely misleading is not, in an obvious sense, marked or infelicitous. For instance, by the retreat in (2'), Dennis openly signals having been uncooperative in the preceding discourse, so it is not surprising that his response is seen as annoying, exasperating, or the like. Yet the contrast with the corresponding response in (1') is clear.

As noted earlier, misleading while avoiding lying typically correlates with the possibility of subsequently claiming that one did not intend to convey the relevant disbelieved information – however obnoxious doing so may be.¹⁴ By contrast, lying typically correlates with a marked unintelligibility of this kind of subsequent retreat. I take this as evidence that lying involves commitment to the misleading information one conveys, whereas this type of commitment is avoided by utterances that are misleading but not lies.

This pattern is borne out by our case from above. Consider the continuations in (3') and (4').

(3') Elizabeth. Why did Doris lose her job?

William. She insulted Sean at a party.

Elizabeth. Oh, so he fired her because of that?

William. #No, that wasn't the reason.

(4') Elizabeth. How is Doris's relationship with Sean?

William. She insulted him at a party.

Elizabeth. Oh, so they're not on good terms?

William. No, they're still friends.

As before, the retreat in (4') is clearly seen as annoying due to its signaling that the speaker has previously been uncooperative. Yet, again, the contrast with (3') is clear. In (3') there is no possibility for the speaker to claim that he did not intend to convey the misleading information that Doris lost her job because she insulted Sean at a party.¹⁵ In other words, in (3'), William is

¹³Given that the contrast between (1) and (2) is the contrast between assertion and conversational implicature, this is just the standard observation that conversational implicatures are cancelable. However, the fact that misleading (vs. lying) in general allows for cancellation, or retreat, does not mean that misleading is always a matter of conversationally implicating disbelieved information. Although cancelability is generally thought to be a necessary condition for a proposition to count as a conversational implicature, few would argue that it is a sufficient one.

¹⁴I use "disbelieve", and cognates, such that "A disbelieves that *p*" means that *A* believes that not-*p*.

¹⁵The speaker can claim that he was not addressing the question that was asked, that is, that he was intending to contribute to a different topic of conversation. However, such a defense is tantamount to opting out of the

committed to the misleading information he conveys. By contrast, in (4'), he is not committed to the misleading information he conveys, i.e., that their relationship is not good.

Another difference exhibited by classic cases of lying vs. merely misleading concerns the possibilities for denials by hearers. A lie can be met with an explicit denial, while a merely misleading utterance does not permit such denials, but typically requires questioning the speaker's intentions instead. This difference is illustrated by (1'') and (2'').

(1'') Rebecca. Are you going to Paul's party?
Dennis. No, I'm not going to Paul's party.
Monica. Yes, you are!

(2'') Rebecca. Are you going to Paul's party?
Dennis. I have to work.
Monica. #Yes, you are!/Wait, are you trying to make her believe you're not going?

The same difference is exhibited by (3) and (4), as shown below.

(3'') Elizabeth. Why did Doris lose her job?
William. She insulted Sean at a party.
Garry. No, that wasn't the reason!

(4'') Elizabeth. How is Doris's relationship with Sean?
William. She insulted him at a party.
Garry. #No, their relationship is fine!/Wait, are you trying to make her believe they're not on good terms?

I think this is sufficient to group (3) with (1), and (4) with (2), with respect to the lying-misleading distinction. Dennis's utterance in (3) involves the same kind of committing mode of communicating as in classic cases of lying. By contrast, in (4), the speaker does not incur such a commitment.

I conclude that the case of (3)–(4) shows that the difference between lying and merely misleading sometimes depends on whether the speaker commits herself to a misleading answer to a question she is addressing. Next, we will see that this kind of question-sensitivity also occurs for cases involving incompleteness.

2.3 Exploiting Incompleteness

As pointed out by Saul (2012b), one can sometimes exploit particular kinds of incompleteness in navigating the lying-misleading distinction. Here is an example:

conversation altogether. In particular, if William claims, in (3), that he was not addressing Elizabeth's question, this is significantly not analogous to canceling the implicature in (4). In the latter case the reply is still claimed to have been intended to be relevant to the question, although what it was most naturally taken as contributing is claimed to have been unintended. By contrast, to claim, in (3), that one was not addressing the question is to claim that one was not taking part in the conversation at all. I am not concerned with such anomalous cases here.

Larry is keen on making himself seem attractive to Norma. He knows she's interested in logic – a subject he himself knows nothing about. From talking to her, Larry has become aware that Norma knows that he has just finished writing a book, although she doesn't know what it's about. In fact, the book Larry wrote is about cats. Recently, Larry also joined an academic book club where the members are each assigned a particular book to read and explain to the others. Larry has been assigned a book about logic. But he hasn't even opened it.

(5) Norma. What's the topic of the book you wrote?

Larry. My book is about logic.

(6) Norma. Do you know a lot about logic?

Larry. My book is about logic.

I think Larry's utterance in (5) is a lie, while his utterance in (6) is merely misleading. Moreover, it is natural to say that the reason is that, whereas in both cases Larry's utterance conveys the misleading information that he wrote a book about logic, only in the first case does this information constitute a direct answer to the question. As we will see, this is captured by the account I propose in section 3.

In support of these judgements we can consider continuations of the dialogues, as in (5') and (6').

(5') Norma. What's the topic of the book you wrote?

Larry. My book is about logic.

Norma. Oh, you wrote a book about logic?

Larry. #No, I just meant that the book I'm assigned to read is about logic.

(6') Norma. Do you know a lot about logic?

Larry. My book is about logic.

Norma. Oh, you wrote a book about logic?

Larry. No, I just meant that the book I'm assigned to read is about logic.

As this shows, these cases behave in a manner parallel to the ones examined earlier. In (5') the speaker cannot claim that he did not intend to convey the relevant misleading information. By contrast, this is possible in (6'), although the speaker is again seen as uncooperative. As with the previous cases, lying correlates with commitment to the misleading information, while merely misleading does not. And again, the contrast between (5) and (6) suggests that whether the speaker incurs commitment of this kind turns on which question is being addressed.

Similarly, these cases exhibit the behavior we should expect with respect to the possibilities for denials by hearers, as illustrated by (5'')–(6'').

(5'') Norma. What's the topic of the book you wrote?

Larry. My book is about logic.

Julie. No, the book you wrote is about cats!

(6'') Norma. Do you know a lot about logic?
Larry. My book is about logic.
Julie. #No, you don't know anything about logic./Wait, are you trying to make her believe that you know a lot about logic?

As before, this is evidence that (5) behaves like classic cases of lying, as in (1), while (6) behaves like classic cases of merely misleading, as in (2). In other words, the case of (5)–(6) demonstrates that one can sometimes rely on semantic incompleteness in order to mislead, while avoiding outright lying.

2.4 The Need for a Discourse-Sensitive Account

The cases we have examined suggest that an account of the lying-misleading distinction that relies on notions of saying or asserting that are not sensitive to discourse structure is likely to be inadequate. Importantly, this is so even given accounts on which what is said is allowed to go beyond linguistically encoded meaning.

Such an account is provided by Saul (2012b). Agreeing that lying requires saying something one believes to be false, Saul argues that what is said is constrained by the principle (NTE) below.

(NTE) A putative contextual contribution to what is said is a part of what is said only if without this contextually supplied material, *S* would not have a truth-evaluable semantic content in *C*.¹⁶

However, this principle is unable to account for examples such as (3). Williams's utterance is truth-evaluable without supplementation. It is true if and only if Doris insulted Sean at a party. Hence, (NTE) cannot count the utterance as *saying* that Doris lost her job because she insulted Sean at a party, and hence, Saul's account cannot agree with the judgment that the utterance is a lie. The reason is clear. Namely, (NTE) is not sensitive to discourse structure.

On the other hand a particular virtue of Saul's account is that it pays close attention to the semantics-pragmatics distinction, and as such is able to handle cases like (6).¹⁷ Arguably, in this case, the speaker exploits the fact that his utterance allows for different completions. The most salient interpretation of his utterance is that Larry wrote a book about logic. Clearly this is the completion Larry hopes the hearer will fasten on, since it is the one that will in turn furnish an answer to the QUD, i.e., that he knows a lot about logic. At the same time, Larry is not committed to this contribution, and moreover there is a possible completion that he believes to be true. It is natural to think that this feature of the example is what exempts Larry from lying. As Saul's account makes explicit, one may avoid lying if one or more available completions are believed to be true.¹⁸

¹⁶Saul (2012b, 57). For discussion of this principle, see also Recanati (1993, 242), Bach (1994, 160–162). Saul (2012b, x) is explicit that her discussion of what is said is does “not aim to show that any theory of saying is *the right theory* (or the wrong theory) of saying” but instead attempts to “discern which is the right theory of saying for a particular purpose – drawing the intuitive distinction between lying and misleading.”

¹⁷By contrast, e.g., neither Adler (1997) nor Williams (2002) considers examples beyond the classic contrast between assertion and conversational implicature.

¹⁸See Saul (2012b, 65).

So, although Saul's account makes explicit the fact that what is said, in the sense that is relevant for the lying-misleading distinction, may go beyond what is linguistically encoded in a sentence, her view retains the traditional conception on which the sentence is the basic unit of analysis. Yet, as we have seen this picture is unable to account for the way in which the lying-misleading distinction depends on discourse structure.

More generally, we can take from this section two desiderata for our account of the lying-misleading distinction. First, it must be sensitive to discourse structure, and in particular, to QUDs. Second, it must take into account the ways in which one can exploit incompleteness concerning what one can be construed as saying. In the following two sections, I provide an account that meets these criteria.

3 Questions under Discussion and What is Said

In this section I first explain how the notions of what is said and assertion figure in my account of the lying-misleading distinction. I then introduce the central components of the QUD framework, and I provide a rudimentary semantics for questions. Finally, I define a notion of what is said according to which whether a proposition counts as said or not by an utterance is determined directly by QUDs.

3.1 Saying and Asserting

I endorse the widely accepted, generic account of lying according to which to lie is to assert something one believes to be false.¹⁹ More precisely, for the purposes of this discussion, I propose to characterize lying as follows:

Lying

A lies if and only if

(L1) *A* asserts that *p*, and

(L2) *A* believes that not-*p*.

The challenge is to spell out a notion of assertion that will capture when an utterance is a lie and when not, and thereby provide an account of the lying-misleading distinction. Below, I outline the account of assertion I will be relying on.

In the tradition from Stalnaker (1970), (1978), (1984), (1998), (2002), we think of a discourse as proceeding against a background of shared information, called the *common ground* of the conversation, that is, the information taken for granted by the participants.²⁰ A discourse, on this view, is directed toward the goal of inquiry – to discover how things are, or what the actual world is like – by incrementally collecting true information in the common ground.

¹⁹This generic account is endorsed by, among others, Chisholm and Feehan (1977), Adler (1997), Williams (2002), Carson (2006), (2010), Sorensen (2007), Fallis (2009), Stokke (2013), (2014), Saul (2012b).

²⁰For a precise definition of common ground information, see Stalnaker (2002).

I assume that assertion requires *saying* something, and furthermore, in accordance with Stalnakerian orthodoxy, that assertion involves a bid to make what is said common ground. Hence, I will assume the following necessary conditions on assertion:

Assertion

In uttering a sentence S , A asserts that p only if

(A1) S says that p , and

(A2) by uttering S , A proposes to make it common ground that p .

So, in asserting that p , a speaker makes an utterance that says that p and thereby proposes to make it common ground that p .²¹

The reasons for distinguishing between what is said and what is asserted in this way are familiar. A standard consideration involves cases of irony. Consider, for example, the dialogue in (7).

- (7) Carol. Did you like the movie?
Ted. [Ironically] Yeah, it was great!

I take it that Ted's utterance says that the movie was great.²² (This will be a consequence of the account of what is said I propose.) But, uncontroversially, Ted is not asserting that the movie was great. Hence, saying that p is not sufficient for asserting that p . Nor is Ted asserting that the movie was not great – even though to communicate this content is the point of his utterance. Hence, proposing to communicate that p , or to make it common ground that p , is not sufficient for asserting that p . Rather, assertion requires making a bid to add information to the common ground by saying it.

3.2 Questions under Discussion

The proposal I want to put forward relies on the framework for understanding discourse structure developed by Roberts (2004), (2012). I now go on to set out the components of this framework.²³

Roberts's central insight is that, to approach the Stalnakerian goal of inquiry, the discovery of how things are, "we must develop strategies for achieving this goal, and these strategies involve subinquiries."²⁴ Such subinquiries are aimed at answering questions that have been accepted as being under discussion.

²¹I use "sentence" to mean grammatical sentence, in the sense of footnote 11.

²²On some ways of using the notion of *saying*, this assumption is false. E.g., as Neale (1992, 523) points out, "If U utters the sentence "Bill is an honest man" ironically, on Grice's account U will not have said that Bill is an honest man: U will have made as if to say that Bill is an honest man." The reason for this is that on a Gricean understanding, saying that p is more akin to how we are understanding asserting that p . Such an account uses different notions to make similar distinctions, e.g., "making as if to say" vs. "saying", or the like.

²³I am not concerned with giving an exhaustive summary of the details of this model of discourse structure, but just with providing the requisite background for the characterization of the notion of what is said that will generate my account of the lying-misleading distinction. Here I follow the presentation in Roberts (2012).

²⁴Roberts (2012, 4).

There are two kinds of questions, those that can be answered by a yes or a no, and those that cannot. The first are *polar questions* like those in (8), the latter are *wh-questions*, as in (9).

- (8) a. Is Mary working?
 b. Does Mary like peanuts?
 c. Is Mary in Rome?
- (9) a. Who is working?
 b. What does Mary like?
 c. Where is Mary?

I adopt the semantics for questions given by Roberts (2012), which in turn draws on the seminal works of Hamblin (1973) and Groenendijk and Stokhof (1984). As on the account in Hamblin (1973), a question is taken to denote a set of propositions – intuitively, the set of its possible answers. We call these the *alternatives* of the question. Roughly, the set of alternatives for a question q is the set of propositions obtained from first abstracting over all the wh-elements in q , if any, and then applying the resulting property to each entity of the appropriate sort in the domain of the model.²⁵

Consider, for example, the wh-question (9a). Abstracting over the wh-element *who* in (9a) yields the following property:

$\lambda x. x$ is working

Accordingly, the set of alternatives for (9a) is the set of propositions corresponding to

$[\lambda x. x$ is working] (x)

for each x of the appropriate sort in the domain. Suppose the domain only contains Mary, Kelly, and Jim. Then the alternatives of (9a) will be the set corresponding to

$[\lambda x. x$ is working] (Mary)

$[\lambda x. x$ is working] (Kelly)

$[\lambda x. x$ is working] (Jim)

In accordance with the familiar Stalnakerian view, we think of propositions as sets of possible worlds.²⁶ So the question in (9a) receives the following interpretation:

$\llbracket (9a) \rrbracket = \{ \{w : \text{Mary is working in } w\}, \{w : \text{Kelly is working in } w\}, \{w : \text{Jim is working in } w\} \}$

Now consider the polar question (8a). Since (8a) does not contain any wh-elements, its set of alternatives is just the singleton of the proposition corresponding to it:²⁷

²⁵See Roberts (2012, 10) for a precise definition.

²⁶See especially Stalnaker (1984).

²⁷See Roberts (2012, 10).

$$\llbracket(8a)\rrbracket = \{\{w : \text{Mary is working in } w\}\}$$

So a question denotes the set of its alternatives. A wh-question is analysed as abstracting into a property and in turn denoting the set of propositions obtained from applying that property to each of the relevant items in the domain. A polar question just denotes the singleton of the proposition corresponding to it.

The alternatives of a question can be seen as corresponding to a range of polar subquestions. Given the narrow domain of our toy example, the subquestions of (9a) are those in (10).

- (10) a. Is Mary working?
 b. Is Kelly working?
 c. Is Jim working?

Given this, we can distinguish between *partial* and *complete* answers. A partial answer to a question is an answer that constitutes a yes or a no to one or more of its subquestions. A complete answer is an answer that constitutes a yes or a no to each of its subquestions.

Following Groenendijk and Stokhof (1984), we can represent the set of complete answers to a question as a *partition* on logical space, i.e., the set of all possible worlds. A question partitions logical space into cells corresponding to each of its complete answers. A polar question determines a partition with two cells corresponding to its positive answer and its negative answer. Wh-questions determine partitions corresponding to all the possible distributions of truth-values to their alternatives.

For example, (9a) partitions logical space into eight cells. Let m , k , and j be the propositions that Mary is working, that Kelly is working, and that Jim is working, respectively. Then the partition set up by (9a) is represented by Table 1.

1 $m \ k \ j$	2 $m \ k \ \neg j$	3 $m \ \neg k \ j$	4 $m \ \neg k \ \neg j$
5 $\neg m \ k \ j$	6 $\neg m \ k \ \neg j$	7 $\neg m \ \neg k \ j$	8 $\neg m \ \neg k \ \neg j$

Table 1.

Cell 1 covers worlds in which Mary, Kelly, and Jim are all working, cell 2 covers worlds where Mary and Kelly, but not Jim, are working, and so on. Formally, a partition is therefore a set of sets of worlds.

Each cell represents a complete answer to (9a), that is, an answer that affirms or denies each of its subquestions. Correspondingly, we will say that a complete answer *rules in* a unique cell in the partition. In turn, a partial answer rules in at least one cell (but not all), and hence all complete answers are also partial answers. The cells that are ruled in represent the possibilities that are still live, given the answer to the question.

For convenience, let us henceforth use $[q?]$ to denote the partition determined by $q?$ We then characterize “answerhood” more formally, as follows:

Answerhood

An answer (partial or complete) to a question $q?$ is the union of a non-empty, proper subset of $[q?]$.

As above, intuitively, an answer is a statement that rules in at least one cell, but not all. We therefore specify that an answer to a question $q?$ must be the union of a *proper, non-empty* subset of the partition determined by $q?$, since neither a statement that rules in no cells nor a statement that rules in all cells of $[q?]$ is intuitively an answer to $q?$ ²⁸

To illustrate, consider the replies to (9a) in (11).

- (11) a. Mary is working.
b. Mary and Kelly are working.
c. Only Jim is working.

(11a) is a partial answer to (9a), since it is a positive answer to one of its subquestions, (10a), but remains neutral on the others. Correspondingly, (11a) rules in cells where Mary is working, i.e., cells 1–4. (11b) is a positive answer to two of the subquestions, but remains neutral on the third. So (11b) rules in cells where both Mary and Kelly are working, i.e., cells 1–2. Finally, (11c) is a complete answer in being an answer to each subquestion and rules in a unique cell, namely the cell where only Jim is working, i.e., cell 7.

3.3 Misleading and the Big Question

In the Stalnakerian model common ground information is represented by a set of possible worlds, called *the context set* – the set of worlds that are compatible with the information in the common ground.²⁹ The possibilities included in the context set are the possibilities that are live given what has already been collected in the common ground of the discourse. Roberts points out that the context set at any given time in a discourse can be thought of as the alternative set for the general question that corresponds to the Stalnakerian goal of inquiry, i.e., discovering what the actual world is like:

Stalnaker's goal of discourse can itself be viewed as a question, the Big Question, What is the way things are?, whose corresponding set of alternatives is the set of all singleton sets of worlds in the context set at a given point in discourse. (Roberts, 2012, 5)

In other words, inquiry can be understood as the pursuit of a complete, true answer to the Big Question – in Stalnakerian terms, of narrowing the context set to just the actual world. Accordingly, a QUD is a subinquiry, that is, a strategy towards the goal of answering the Big Question. In answering a QUD, one advances the discourse and thereby reduces the context set. That is, one effectively claims that the actual world is among the worlds one rules in.

Given this framework, to mislead is to disrupt the pursuit of the goal of inquiry, that is, to prevent the progress of inquiry from approaching the actual world. If one answers (9a) with (11a), one reduces the context set to a set from which worlds in which Mary is not working

²⁸I refrain from discussing this in detail here. See Schoubye and Stokke (2015) for details.

²⁹See in particular Stalnaker (1998), (2002).

are eliminated. (Here, as throughout, I focus for convenience on cases in which the relevant speech act is accepted by the other discourse participants.) If one effectuates this change while disbelieving that Mary is working, one reduces the context set to a set that, according to one's own beliefs, excludes the actual world. Thereby one steers the narrowing of the context set that is unfolding through the discourse away from the actual world.

There are many ways of contributing information to a discourse. Information can be contributed by assertion, conversational implicature, or by other means. There are equally many ways of contributing misleading information. One may do so by asserting something one believes to be false, by conversationally implicating it, or in some other way. Asserting that one is not going to a party and implicating that one is not going are both ways of contributing this information to the common ground of the discourse, and both effectuate a narrowing of the context set. If one believes that one is not going, both are ways of disrupting the pursuit of the goal of inquiry.

In the classic cases exemplified by (1)–(2), a speaker avoids lying, while still misleading, by contributing disbelieved information as a conversational implicature. According to the standard Gricean understanding, a proposition counts as a conversational implicature only if it is derived (or at least derivable) from information available to the hearer about what is said by the speaker.³⁰ In other words, on the orthodox Gricean conception, implicature-derivation must take as its input a content that the hearer recognizes as what is said. The derivation then proceeds along familiar lines appealing to the cooperative principle and the maxims and to available information about the context.

We will see that contributing disbelieved information by means of classic conversational implicatures of this kind is not the only way of being misleading, while avoiding lying.

3.4 Minimal Content

The central part of my account of the lying-misleading distinction is its characterization of the notion of what is said. On this characterization what is said by an utterance depends on QUDs. Briefly, what is said by a sentence in a context is the answer it provides to the QUD it is addressing. But moreover, this view also assumes that what a sentence can be used to say is constrained in a precise way by a species of *compositional* meaning.

For example, on the view I will argue for, while (12a), given the right QUD, can be used to say (12b), the same utterance cannot be used to say (12c) or (12d).

- (12) a. [Uttered by Larry] My book is about logic.
b. The book Larry wrote is about logic.
c. The book Larry wrote is about cats.
d. Larry knows a lot about logic.

To implement this, I assume that any grammatically complete, declarative sentence determines what I shall call a *minimal content* as a result of composition of the lexical meaning of its constituents, in context.

³⁰See in particular Grice (1989, 31). For relevant discussion and for a view according to which what is said is sometimes derived by way of Gricean maxims, see Soames (2008).

In many cases the minimal content of a declarative sentence is a proposition. In the terminology of Recanati (1989), Bach (1994), Cappelen and Lepore (2004), and others, such a proposition can be identified as a so-called *minimal proposition*, i.e., a proposition that is determined solely by composition of the constituents of the relevant sentence. For example, it is safe to assume that (13) determines the minimal proposition that Mary is working.

(13) Mary is working.

But moreover, even sentences that contain uncontroversial examples of context-sensitive expressions and constructions will often determine minimal propositions.

In particular, I follow Cappelen and Lepore (2004) in assuming that there is a *basic set* of context-sensitive expressions and constructions that are typically saturated before determination of minimal content. I take this set to include, at least, the pronouns (*I, you, he, she, it, we, they*, etc.), demonstratives (*that, this, those*, etc.), and the recognizable adverbs of space and time (*here, now, there, yesterday*, etc.), that is, roughly the set of expressions which Kaplan's (1989) formative treatment was designed to account for.³¹ So, for example, I assume that (14) determines the minimal proposition that Mary is working.

(14) [Demonstrating Mary] She is working.

Beyond the obvious candidates mentioned above, however, the issue of which expressions and constructions to include in the basic set is highly contentious and non-trivial.³²

Yet, for our purposes, it is not necessary to decide on a definite list of basic set elements. What is needed is the assumption that composition of lexical meaning, in context, determines constraints on what is said. For a range of context-sensitive expressions and constructions in natural language, it is plausible to think that these will very often be saturated prior to determination of this minimal content. This class of obvious context-sensitive elements, we may take it, largely coincides with Cappelen and Lepore's basic set. But we do not need to assume that such expressions and constructions are always assigned values before minimal content emerges.

In particular, on the view I will defend here, minimal content may fall short of complete propositionality. There are many ways in which this can come about. Two of these are, first, due to unclarity concerning how particular context-sensitive elements are to be saturated, e.g., in cases of vague demonstrations or where there is no determinate salient value for a particular context-sensitive parameter, and second, due to indeterminacies in the boundaries of the lexically encoded conditions of ingredient expressions.³³

³¹Note that I will not be assuming that if an expression or construction is a member of the basic set, its value must be determined in a Kaplanian manner, i.e., via characters operating on tuples representing linguistically relevant facts about the utterance situation. It may be that some basic set elements have their values determined by common ground information. All that is assumed is that this happens prior to determination of minimal content.

³²Cappelen and Lepore (2004, 1) provide a list that extends beyond the three basic categories mentioned in the text.

³³The notorious statement by Bill Clinton, "I did not have sexual relations with that woman", was claimed not to be a lie because, it was argued, Clinton was using the term *sexual relation* in a way such that its extension did not include the relevant acts. This is arguably an example in which indeterminacy arises due to lexical non-specificity. In this paper I focus on indeterminacy arising from contextual underdetermination.

But even in such cases, what a sentence can be used to say is constrained by its minimal content. In cases where the compositionally determined constraints on what is said fall short of propositionality, we can construe the minimal content as a range of candidate minimal propositions. As we will see, on my view, such cases allow for particular ways of exploiting the lying-misleading distinction.

As this suggests, my view disagrees with so-called *Radical Contextualists*, like Travis (1985), Sperber and Wilson (1986), Carston (2002), and Recanati (2004), (2010), who claim that compositional meaning always falls short of propositionality. On the other hand I refrain from adopting the strong position of so-called *Semantic Minimalists*, like Cappelen and Lepore (2004), according to which all grammatical, declarative sentences express propositional contents that are determined purely by lexical meaning and compositional procedures.³⁴ Instead, my view agrees with so-called *Moderate Contextualists*, like Bach (1994), in granting that compositional content, although sometimes fully propositional, may fall short of complete propositionality.³⁵

3.5 What is Said

I can now state the definition of what is said that my account of the lying-misleading distinction turns on. Using $\mu_c(S)$ to denote a minimal proposition expressed by a sentence S in a context c , I adopt the following definition of what is said:³⁶

What is Said

What is said by S in c relative to a QUD $q?$ is the weakest proposition p such that p is an answer (partial or complete) to $q?$ and $p \subseteq \mu_c(S)$.³⁷

So, what is said, according to this proposal, is the weakest answer to a QUD that entails a minimal proposition expressed by the utterance in question, given the context. As before, we understand an answer to a question $q?$ to be the union of a non-empty, proper subset of $[q?]$, or intuitively, a statement that rules in at least one cell, but not all.

Let me first explain the way this characterization of what is said works in cases where the minimal content expressed is propositional. On this view what is said by a sentence S

³⁴For a different type of of Semantic Minimalism, see Borg (2004).

³⁵The view I argue for also disagrees with Cappelen and Lepore's *Speech Act Pluralism* according to which "an utterance can assert propositions that are not even logical implications of the proposition semantically expressed. Nothing even prevents an utterance from asserting (saying, claiming, etc.) propositions incompatible with the proposition semantically expressed by that utterance." (Cappelen and Lepore, 2004, 4)

³⁶This proposal is defended in detail in Schoubye and Stokke (2015). See also Gauker (2012) for a view that has some similarities to this one.

³⁷A full account might take c to be a Kaplanian tuple consisting of speaker, time, location, and world, and include in such tuples a QUD (and perhaps other coordinates.) A sentence can then be defined as saying a proposition relative to such an extended Kaplanian tuple, the traditional coordinates of which are used to determine minimal content in the familiar way. (This does not rule out that some basic set elements may have their values determined in other ways, e.g., via common ground information.) On such an implementation, truth can then be defined relative to such contexts and worlds. Letting $\sigma_c(S)$ denote what is said by S relative to c (understood as just suggested), S is true relative to a context c and a w if and only if $w \in \sigma_c(S)$.

relative to a QUD $q?$ is always a proper subset of $[q?]$, that is, a proper subset of the complete answers to $q?$ Specifically, what is said is the weakest such proposition that entails the minimal proposition expressed by S , given the context. In other words, to determine what is said by S relative to $q?$, first find the proper subsets of $[q?]$ that entail the minimal proposition expressed by S , then pick the weakest of them. (By a proposition p being weaker than another proposition r , we mean that p is entailed by r , but not *vice versa*.)

To illustrate, consider our previous example of the reply to (9a) in (11a).

(9a) Who is working?

(11a) Mary is working.

The minimal proposition expressed by (11a) is the proposition that Mary is working. So we look for proper subsets of the partition set up by (9a) that entail this proposition. As seen from Table 1, there are several proper subsets of $[(9a)]$ that entail that Mary is working. In general, any proper subset that contains some m -worlds and no $\neg m$ -worlds entails the minimal proposition expressed by (11a).

For example, the subset covered by cells 1 and 2 entails that Mary is working. But if this subset counts as what is said, this amounts to the prediction that, relative to (9a), (11a) says that Mary and Kelly are working, which is clearly incorrect. However, this proposition is not the *weakest* of the propositions in $[(9a)]$ that entail the minimal proposition expressed by (11a). Rather, the weakest of them is just the set of all worlds where Mary is working, i.e., the worlds covered by cells 1–4. So, in this case, what is said is just the minimal proposition itself, i.e., the proposition that Mary is working.

What about cases in which minimal content is not fully propositional? As we said above, in such cases, the minimal content can be construed as a range of candidate minimal propositions. For example, as we will see, (12a) may fall short of determining a fully propositional minimal content, and instead be associated with a range of candidates. In such cases I suggest that the interpretation “looks through” the candidates and tries to find a suitable content for what is said, given the QUD. In each case what is said is evaluated in the same way as above, i.e., as the weakest proposition that is an answer to the QUD and entails the relevant candidate minimal proposition.

We will see that sentences can be used to say things that go beyond their minimal contents. Yet what is said is precisely constrained both by minimal, compositional meaning and by QUDs, and hence so is assertion. I take this to be a natural extension of the Stalnaker-Roberts view of information exchange. The goal of inquiry – to answer the Big Question – is approached via subinquiries, i.e., QUDs. When solving the task of which information to add to the common ground as the result of an utterance, the participants of a discourse will look at the minimal content of the utterance and at which QUD is being addressed in order to determine what the speaker said.

Consequently, one way of being misleading is to contribute disbelieved information to the discourse. To lie is to assert a disbelieved proposition p , which in turn requires uttering a sentence that counts as saying that p , given the context and the relevant QUD.

4 Accounting for the Lying-Misleading Distinction

In this section I demonstrate the consequences of the account of the lying-misleading distinction I proposed above. I start by showing how the account handles the cases we have looked at earlier. I then turn to consider some further phenomena involving presuppositions, negation, and prosodic focus.

4.1 The Classic Contrast

Consider first the classic contrast between assertion and conversational implicature, as illustrated by (1)–(2).

- (1) Rebecca. Are you going to Paul's party?
Dennis. No, I'm not going to Paul's party.
- (2) Rebecca. Are you going to Paul's party?
Dennis. I have to work.

We assume that, in (1), Dennis's utterance determines the minimal proposition that Dennis is not going to Paul's party. Rebecca's question is a polar-question, and hence its partition contains two cells corresponding to the positive answer and the negative answer, respectively. The negative answer entails (is identical to) the minimal proposition expressed by Dennis's utterance. According to our definition, this proposition is therefore said by the utterance. Moreover, since Dennis is obviously offering this proposition for common ground uptake, it is asserted. So because Dennis, in this case, is asserting something he believes to be false, we explain why he is lying.

Now look at (2). We assume that Dennis's utterance determines the minimal proposition that Dennis has to work. There is no answer to the question that entails this minimal proposition. Neither the proposition that Dennis is going to the party nor the proposition that he is not going to the party entails that he has to work. So our account does not predict that either the positive answer or the negative answer to Rebecca's question is said. Regardless, the classic observation is that Dennis, in this case, conversationally implicates that he is not going to the party. This implicature will typically be explained by appeal to the Maxim of Relation, "Be relevant."³⁸ Implicatures are offered for common ground uptake. Hence, Dennis is being misleading, since he is making a bid to update the common ground with disbelieved information.

As noted earlier, conversational implicatures are standardly seen as derivable from what is said. Intuitively, Dennis's utterance in (2) says that he has to work, and this content is uncontroversially the input to the derivation of the implicature that he is not going to the party. So our account should predict that what is said, in this case, is that Dennis has to work. However, it might not be obvious that we can meet this demand, since we do not predict that this content is said relative to the question in (2). Yet, as I now go on to explain, the QUD framework offers an elegant way of achieving this result.

³⁸See Grice (1989, 27) and Davis (2010, 7)

4.2 Defaulting to the Big Question

We have seen that, given the Stalnakerian understanding of the aim of inquiry, information exchanges are always, ultimately, directed toward resolving the Big Question. In this sense, therefore, the Big Question is always a background QUD (at least in serious conversation.) I suggest that, in cases where an utterance cannot be seen as saying something relative to a more immediate, local QUD, the interpretation typically defaults to seeing the utterance as addressing the Big Question directly. More precisely, when there is answer to the local QUD that entails any candidate minimal proposition, the utterance is interpreted against the Big Question. (I return to this in section 5.)

A complete answer to the Big Question is an answer that rules in one world and rules out all others. So the Big Question is the question that partitions logical space into one cell for each and every possible world, as illustrated by Table 2.

1	2	3	4	
w_1	w_2	w_3	w_4	...

Table 2.

The minimal proposition expressed by Dennis's utterance in (2) is the proposition that Dennis has to work. The weakest non-empty proper subset of the partition in Table 2, i.e., the weakest answer to the Big Question, which entails that Dennis has to work, is simply the set of all worlds where Dennis has to work. Hence, since Dennis's utterance in (2) addresses the Big Question, even if it does not address the local QUD raised by Rebecca, we predict that Dennis's utterance says that he has to work, and moreover, this proposition is asserted. In this way we predict the desired input to the derivation of the implicature that he is not going to the party.

Furthermore, this means that, on our account, if Dennis believes that he does not have to work, he is lying in addition to misleadingly implicating that he is not going to the party. Again, this is clearly the result we want.

4.3 Committing to Misleading Answers

Now let us consider the case of (3)–(4).

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

As we saw, discourse-insensitive accounts like Saul's (2012b) are unable to explain why William's utterance in (3) is a lie, while the same utterance, in (4) is merely misleading.

Williams's utterance, in each case, expresses the minimal proposition that Doris insulted Sean at a party. The weakest non-empty proper subset of the partition set up by Elizabeth's

question in (3), which entails this minimal content, is the proposition that Doris lost her job because she insulted Sean at a party. So, given our proposal, this proposition is said, and asserted, by William's utterance in (3). We therefore explain why William is lying.

By contrast, there is no subset of Elizabeth's question in (4) that entails that Doris insulted Sean at a party. In particular, the content that is most naturally taken to be implicated by William's utterance – that the relationship is not good, or the like – does not entail the minimal proposition expressed by the utterance, and therefore, it is not asserted, on our account.

Nevertheless, William's utterance in (4) is misleading because it contributes disbelieved information to the discourse, i.e., that the relationship is not good. This contribution is plausibly seen as a conversational implicature. Again, the most likely explanation appeals to a violation of the Maxim of Relation. As above, I take it that William's utterance in (4) is seen as addressing the Big Question, since it does not make an assertion with respect to the local QUD. Interpreted against the Big Question, William's utterance says that Doris insulted Sean at a party. On our account this proposition is asserted by the utterance and hence serves as input to implicature-derivation. This prediction accords with the intuitive judgment that if William believes that Doris did not insult Sean at a party, he is lying.

So we explain why one can lie by committing oneself to misleading answers to QUDs, namely because what one says, and hence what one asserts, depends directly on the QUD one is addressing.

4.4 Exploiting Incompleteness

The next type of case we looked at involved exploiting incompleteness, as in (5)–(6).

(5) Norma. What's the topic of the book you wrote?

Larry. My book is about logic.

(6) Norma. Do you know a lot about logic?

Larry. My book is about logic.

As we noted, Larry's utterance in (5) is seen as a lie, while in (6), he is being merely misleading. Intuitively, in both cases, Larry conveys that he wrote a book about logic, which is something he believes to be false. But only in (5) is this information a direct answer to the question. But at the same time, the way in which the information that Larry wrote a book about logic is provided turns crucially on his use of the possessive construction.

What minimal content is expressed by Larry's utterances in these cases? I take it to be sufficiently uncontroversial that the possessive construction requires contextual supplementation of a possession relation. Intuitively, we need to know what relation Larry is supposed to bear to the relevant logic book in order for the utterances to make sense.³⁹ This is arguably

³⁹Cappelen and Lepore (2004, 94–95) argue that the possessive construction is not included in the basic set. Since they take minimal content to determine truth-conditions, Cappelen and Lepore therefore do not think that contextual supplementation of the possessive is necessary for truth-conditional content. Yet Cappelen and Lepore are not committed to the view that their minimal truth-conditions should be intuitively recognizable as contents that are communicated.

a central part of the reason why one can use this kind of semantic incompleteness to mislead while avoiding lying.

But are possessives to be seen as belonging to the basic set of context-sensitive elements that are saturated prior to determination of minimal content? As I said earlier, I will not attempt to answer this question in this paper.⁴⁰ However, as we will see next, there are good reasons to think that, regardless of how this issue is resolved, our account will be able to offer a plausible explanation of this type of case.

Given the context as described, it is natural to say that the compositional meaning of Larry's utterance, in each case, falls short of determining a minimal proposition. Instead, it is indeterminate between a range of candidates, including, at least, those represented by (15a) and (15b).⁴¹ (We assume that the first person feature of *my* is saturated.)

- (15) a. The book Larry wrote is about logic.
b. The book Larry is assigned to read is about logic.

...

Moreover, note that there is no sense in which the utterance is compatible with (16), even though the latter is true.

- (16) The book Larry wrote is about cats.

So although it is insufficient to fix a fully propositional content, the compositional meaning of the utterance, in context, still specifies a concrete constraint on what it can be used to say.

In terms of the Stalnakerian framework, the compositional meaning of utterance rules out that it can be used to update the common ground with the information that the book Larry wrote is about cats, i.e., that it can be used to assert (16). However, the compositional meaning allows that it may be used to update the common ground with assertions like those in (15). Put differently, even though the incompleteness of the utterance means that it does not determine a propositional minimal content, it still maps out a region of logical space within which to locate what it can be used to say.

There are two main ways to motivate the claim that the utterance fails to determine a minimal proposition. First, it might be argued that the possessive is part of the basic set, but is not sufficiently fixed in this context. Second, it might be argued that the possessive is not part of the basic set. Both suggestions have a claim to being plausible. But we do not need to settle this matter here. With respect to how it exploits the lying-misleading distinction, the important feature of the case is the indeterminacy between different minimal propositions. For our purposes, it is less important to settle for an account of the source of this indeterminacy.

⁴⁰Recanati (2004) argues that possessives involve a variable that must be "saturated", but that they differ from Kaplanian indexicality, where the value in question is determined via a linguistic rule operating on what he calls "narrow context." Barker (1995) argues that the possessive in English involves an empty determiner that receives its value via a contextually salient possession relation. See also Stanley (2000, 489) and Carston (2002, 185–186).

⁴¹Some theorists, e.g., Heim (1992) and Barker (1995, 78–82), take possessives to involve definiteness, yet this view remains controversial, and I will not presuppose it. The paraphrases in (15) employ the definite article merely for convenience, and none of my arguments rely on any substantial assumption about the definiteness of possessives.

According to what we suggested earlier, the interpretation will attempt to find a suitable content for what is said in terms of the candidate minimal propositions and the space of answers to the QUD. In (5) there are several answers to the QUD that entail (15a). The weakest of these is the proposition that the book Larry wrote is about logic. So, given the QUD, this candidate minimal proposition yields a plausible content as what is said. Furthermore, none of the answers to the QUD in (5) entail (15b). The same is true for other possible completions. For example, consider the candidates in (17).

- (17) a. The book Larry is thinking about is about logic.
b. The book Larry is talking about is about logic.
c. The book Larry has been working on is about logic.
...

No answer to the question in (5) entails any of these. Rather, since (15a) yields a plausible outcome for what is said, in this case, we predict that Larry is interpreted as saying, and hence asserting, that the book he wrote is about logic. This also accords with our judgment that Larry is lying in this case.

Now consider (6). Neither of the candidate completions in (15) is entailed by either the positive or the negative answer to the QUD in (6). Nor is any other possible completion. So Larry's utterance in (6) cannot be construed as saying anything relative to the local QUD. So, according to our earlier proposal, the utterance is instead seen as addressing the Big Question. Notice, however, that in cases like this, where there are different candidate minimal propositions in play, the Big Question does not distinguish a particular content as what is said. We could say, therefore, that each of the candidates in (15) and (17) is said. But this is not right. Larry is not lying, in this case, and therefore (15a) cannot be said, or asserted. This conforms to our sense that Larry, in (6), is not asserting any particular completion.

Instead, I propose that, in cases where an utterance is indeterminate between a range of minimal propositions, and moreover cannot be construed as saying anything relative to the relevant local QUD, the interpretation pulls back even further and takes as what is said a particular kind of generalized proposition. This proposition is the union of all possible completions, i.e., not just the most salient ones. In Stalnakerian terms, since the interpretation cannot determine an assertion that situates the actual world within a particular completion, it settles for the observation that the utterance nevertheless locates the actual world somewhere in the region of logical space enclosed by the candidates taken together, i.e., the region of space determined by the compositional meaning of the utterance. In other words, the utterance is seen as making an assertion that is true if and only if at least one completion is true.

In our case, given that Larry's utterance cannot be interpreted as saying anything relative to the QUD, and since the Big Question does not distinguish a particular completion as said, the interpretation takes as what is said the union of all possible completions, i.e., like those in (15) and (17). There are different ways of paraphrasing the resulting generalized proposition, e.g., as in (18).

- (18) Larry bears some relation to a book about logic.

Intuitively, this is the content that would be understood by someone who had no access to the contextual information relevant for deciding between the candidate minimal propositions.

(18) is a plausible content for what is said in (6). And furthermore, Larry is naturally seen as proposing to add to the common ground the information it paraphrases. Part of the peculiarity of this case is that Larry makes an assertion that is, on the one hand, sufficiently specific to achieve his misleading purposes and, on the other hand, sufficiently non-committal to avoid lying. Given that (18) is asserted in (6), we predict that Larry is lying if he believes that (18) is false. This is the right result. In particular, we make the same prediction as on Saul's (2012b) account, according to which speakers in cases like this avoid lying if at least one possible completion is believed to be true.

It remains to be explained why Larry is being misleading in (6). Arguably, the reason is that his utterance succeeds in implicating that he knows a lot about logic. As we have seen, implicatures are derived from what is said. According to our account, the hearer, in this case, recognizes that Larry's utterance does not say anything relative to the local QUD, and moreover, that it does not say anything particular relative to the Big Question. Hence, finally, the generalized content in (18) is taken to be what is said. In turn, the implicature that Larry knows a lot about logic is inferred.

This inference is arguably drawn on the basis of, first, Norma's belief that Larry wrote a book about logic, and second, the Maxim of Relation. If Larry had not been aware of this belief of Norma's, he would not have thought that he could mislead her in this way. However, I will not attempt to provide a more precise account of this process here. Doing so is a challenge for any view on which utterances can give rise to implicatures even when their compositional meaning falls short of full propositionality.⁴² As before, our account can do no more than make sure it predicts a suitable input for implicature-derivation. I assume that, given the context and the beliefs of the participants, a plausible account can be given of how Larry succeeds in implicating that he knows a lot about logic by asserting (18) in response to the question in (6).

To further illustrate the way this account handles incompleteness, I want to briefly comment on a slightly different type of case. Saul (2012b) gives the following example:

Dave is lying in bed, and two nurses are discussing the treatment he needs. Ed holds up a bottle of heart medicine, points at it, and utters (19):

(19) Has Dave had enough?

Fred replies with (20):

(20) Dave's had enough.

As it turns out, Fred hates Dave, wants him to die, and plans to bring this about by denying him his much-needed heart medicine.⁴³

⁴²For examples of such views, see, e.g., Carston (2002), Recanati (2004).

⁴³Saul (2012b, 62). Example numbering altered.

As Saul notes, Fred's utterance is intuitively a lie. Accordingly, she takes this case to be one "in which there is no misunderstanding or disagreement about what is said."⁴⁴ Saul's principle (NTE) makes the right prediction for this example. As she observes, the most salient completion is the one in (21).

(21) Dave has had enough heart medicine.

Moreover, *pace* minimalists like Cappelen and Lepore (2004), (20) is not truth-evaluable without supplementation. Hence, given Saul's constraint on what is said (NTE), (21) is said.

Let me briefly explain how my account is able to agree with this. At least one central reason the completion involving heart medicine is salient is the prominence of Ed's question, and in particular, because his accompanying gesturing makes it obvious that he is asking about the heart medicine. Intuitively, Ed's question is seen as asking whether Dave has had enough heart medicine. In other words, Ed's utterance in (19) is interpreted as introducing the QUD in (22).

(22) Has Dave had enough heart medicine?

On the other hand, as in the case of (5)–(6), Fred's utterance in (20) is arguably not determinate with respect to a single minimal proposition. Rather, it suffices only to map out a range of candidates, like those in (23).

(23) a. Dave has had enough heart medicine.
b. Dave has had enough tea.
c. Dave has had enough bourbon.

...

(23a) is entailed by the positive answer to the QUD in (22). Hence, we predict that the proposition that Dave has had enough heart medicine is said, and in turn, asserted. This is the correct result, since the reason Fred is lying is intuitively that he asserts that Dave has had enough heart medicine.

Before moving on, I want to address a potential concern that might arise at this point. It might be asked, if we allow incomplete questions, like (19), to be filled out as completed QUDs, as in (22), why not just do the same for incomplete utterances, in the first place? Why not just argue that (20), in this context, gets completed as (23a), instead of going the roundabout way via QUDs?

In response, we should emphasize that the overall aim of our account has been to point out that the lying-misleading distinction is sensitive to discourse structure due to the relationship between what is said and which question is addressed. As we have seen, in the case of (3)–(4), one and the same, grammatically complete, declarative sentence can be used to say different things depending on the question it is addressing. This is evidence that any adequate account of what is said must be sensitive to discourse structure in this way. The suggestion made above concerning Saul's example of (19)–(20) demonstrates how such an account can handle cases

⁴⁴Saul (2012b, 61).

in which a question explicitly asked is incomplete. Moreover, as we will see in section 5, there are further reasons for thinking that QUDs can be introduced in other ways than by being explicitly asked in conversation. Below, I first turn to consider some further cases involving presupposition, negation, and prosodic focus.

4.5 Presuppositions and Pronouns

As noted earlier, since conversational implicature is not the only way of adding information to the discourse without saying it, conversationally implicating disbelieved information is not the only way of being misleading while avoiding lying. One of the most prominent sources of non-asserted information in natural languages are presuppositions. According to the Stalnakerian view, the hallmark of a presupposition is that it places a requirement on the common ground of a conversation. We can say that if a sentence S presupposes that p , S is felicitous, in a context c , only if p is either already common ground in c or is accommodated. As we will see below, this allows for two ways of being misleading with the use of presuppositions. First, if the common ground contains information one does not believe, one can sometimes serve one's communicative goals by presupposing the disbelieved common ground information. Second, one can offer presupposed information that one does not believe for accommodation.

To illustrate the first of these possibilities, let us return to the stock example of Saint Athanasius, mentioned earlier. Consider a modernized version:

Nathan is sitting in the office, when suddenly a henchman of a loan shark he owes money to bursts in. As he questions him, Nathan realizes that the henchman does not know that he is Nathan.

(24) Henchman. Where's Nathan?
Nathan. He's not far away.

There are arguably two ways in which Nathan is being misleading in this case. First, he is non-standardly using the 3rd person to refer to himself. And second, he conveys that Nathan is not in the office, while avoiding outright asserting that. Here I will discuss the first. (I will return to the second in section 4.6.)

According to one standard view, the so-called *phi-features* of pronouns (person, gender, and number) are presupposition triggers.⁴⁵ Specifically, the 3rd person feature is seen as triggering the presupposition that the referent of the pronoun is neither the speaker nor the hearer.⁴⁶

What role does this kind of presupposition play in a common ground model of conversation? It is obvious that at least one function of the features of pronouns is to facilitate reference-identification. Stalnaker (1998) is explicit that

⁴⁵See, e.g., Cooper (1983), Heim and Kratzer (1998), Heim (2008), Sauerland (2008).

⁴⁶There is debate about the source of this presupposition. On some views the 3rd person feature triggers the requirement that the referent be neither speaker nor hearer as a semantic presupposition. On others the 1st and 2nd persons encode semantic presuppositions, while the 3rd person feature is empty, and the presupposition concerning the participant role of the 3rd person is derived as an implicature. On these options, see Heim (2008), Sauerland (2008).

If certain information is necessary to determine the content of some speech act, then appropriate speech requires that the information be presumed to be shared information at the time at which that speech act is to be interpreted. (Stalnaker, 1998, 101)

Stalnaker observes that a use of the first person, singular pronoun *I* requires that the speaker "must be presuming that the information that she is speaking is available to her audience – that it is shared information." (ibid.)⁴⁷ If this is right, the analogous observation should hold of the 3rd person. When someone uses the 3rd person, they must be presuming that the information that the intended referent is neither speaker nor hearer is available to the audience. In particular, someone using the 3rd person is relying on common ground information about who the speaker and addressee are in order to allow the audience to identify the intended referent.

Furthermore, as Stalnaker (2002) emphasizes, common ground information need not be true, and need not even be believed by the participants:

Successful communication is compatible with presuppositions that are recognized to be false, but the information that they are being presupposed must be actually available, and not just assumed or pretended to be available. (Stalnaker, 2002, 716)

Along these lines, Stalnaker (1998) points out that

Sometimes the most effective way to communicate something true is to presuppose something false. For example, if you are presupposing something false but irrelevant, I may presuppose it as well, just to facilitate communication. (You refer to Mary's partner as "her husband," when I know that they are not married. But I might refer to him in the same way just to avoid diverting the discussion.) (Stalnaker, 1998, 100)

These observations make it possible to explain the way in which the use of the 3rd person pronoun in cases like (24) is misleading.

As in Stalnaker's example, Nathan is presupposing something false in order to serve his communicative purposes. (As we will see later, he is doing so in order to assert something true, and thereby avoid lying.) What Nathan is presupposing is that the referent of *he* is neither the speaker nor the hearer. Moreover, since it is obvious from the utterance and the context that *he* is intended to refer to Nathan, the presupposition of the pronoun is that Nathan is neither the henchman nor the person speaking. This presupposition is felicitous because it is obvious that the henchman thinks it is common ground that the person speaking is not Nathan.⁴⁸ Since it is advantageous to him, Nathan goes along with this.

On this analysis, then, Nathan's use of the 3rd person to refer to himself does not mislead by adding disbelieved information to the common ground. The disbelieved information that he is not Nathan is already common ground when he makes his utterance, and he is relying on

⁴⁷It is not required that this information be available before the utterance was made. See Stalnaker (1998, 101).

⁴⁸Well known complications concerning the representation of necessarily false propositions in terms of possible worlds will arise here. In particular, the set of worlds in which Nathan is not Nathan is the empty set. However, I assume that there are ways of accommodating the obvious fact that discourse participants may entertain false presuppositions about the identity of other participants, or of people in general. For relevant discussion, see Stalnaker (1978).

that fact in using the 3rd person. Rather, Nathan's use of the 3rd person is to be seen as the kind of misleading by which one allows others to continue to hold a false belief for disingenuous purposes.⁴⁹ Since one thereby deliberately prevents the common ground from approaching the truth, this constitutes another way of disrupting the pursuit of the goal of inquiry.

The same points apply to Stalnaker's own example of going along with the false presupposition concerning Mary's partner. Again, the speaker may later be accused of having been misleading, not by getting the hearer to accept false information, but by allowing the hearer to continue holding a false belief. As in all other cases, the perceived severity of this kind of breach will vary.

Presuppositions can also be used to mislead in the second sense distinguished above, that is, the direct sense of aiming for disbelieved information to become common ground as a result of one's utterance. One can offer presuppositions that one believes to be false for accommodation.⁵⁰

In the case of pronouns this is most obvious regarding the gender features of pronouns. Suppose, for instance, that Marion replies as in (25) in order to make Mick think that Tim's baby is a girl, even though she knows it is a boy.

- (25) Mick. Have you seen Tim's baby yet?
Marion. Yes, she's lovely.

The feminine feature of the pronoun triggers the presupposition that the referent is female. Since it is obvious that the referent is Tim's baby, Marion's use of the pronoun will – unless something out of the ordinary happens – have the effect of making it common ground that the baby is female. Hence, if Marion believes this is false, she is being misleading by contributing disbelieved information without asserting (or saying) it.

Both ways of misleading are also possible with other kinds of presuppositions. For example, if Allan knows that Kyle falsely believes that his wife cheated on him, Allan may exploit that fact in uttering (26).

- (26) Your wife really regrets that she cheated on you.

Allan is being misleading in allowing the false information that Kyle's wife cheated to remain common ground. But moreover, Allan is lying. Since *regret* is a factive verb, if Allan believes that Kyle's wife did not cheat, he must also believe that it is false that she regrets cheating. So, in this case, Allan is both presupposing and outright asserting disbelieved information. And moreover, if it was obvious that Kyle did not believe, or suspect, that his wife cheated, Allan would be offering disbelieved, presupposed information for accommodation while at the same time offering disbelieved, asserted information for common ground uptake. Finally, if Allan knew that the wife did cheat but that she does not regret it, this would be a case of exploiting a true presupposition in order to facilitate a lie.

In general, then, one can employ presuppositions for misleading purposes either by knowingly allowing disbelieved information to remain common ground, or by directly aiming for

⁴⁹This is what Chisholm and Feehan (1977) called omissive, positive deception *secundum quid*.

⁵⁰On accommodation, see Stalnaker (1998, 102–104).

disbelieved information to be accommodated. Since presuppositions are not said, both of these are ways of misleading while not lying, even though such maneuvers may be accompanied by outright lying if what is asserted is likewise disbelieved.

Moreover, note that in these cases, misleading while avoiding lying does not correlate with the same kind of possibility for retreat as we observed for earlier cases. For example, one cannot, in general, retract gender information presupposed by a use of a pronoun. Hence, these cases illustrate a different way of being misleading while avoiding full-blown assertion of disbelieved information. But more particularly, there is arguably a sense that this kind of merely misleading is morally worse. For example, Marion in (25) did something *prima facie* worse than the kind of misleading that is susceptible to subsequent retreat. It is reasonable to think that at least part of this difference is due to the effect being

4.6 Negation

We can now return to the example of Nathan and the henchman in order to explain the sense that Nathan contributes misleading information – i.e., that he is not in the office – to the discourse, while avoiding lying. This requires an explanation of what is said by utterances involving negations. As we will see, this is straightforward, given our proposal.

The henchman’s question partitions logical space into cells corresponding to its complete answers. A complete answer is an answer that constitutes a yes or a no to each place Nathan could be. Clearly, the list will be very long. Moreover, some of the places will exclude each other, while some will not. And furthermore, for some of the places, it will be vague whether they are far away or not. Let us assume that there are five salient places Nathan could be: the office, the mall, the bar, his home, and the garage. Moreover, assume that the bar is in the mall. So the question determines the partition in Table 3.

1 $o \neg b \neg m \neg h \neg g$	2 $\neg o b m \neg h \neg g$	
3 $\neg o \neg b m \neg h \neg g$	4 $\neg o \neg b \neg m h \neg g$	5 $\neg o \neg b \neg m \neg h g$

Table 3.

Suppose further that (i) Nathan’s home is nearby, (ii) the mall, and hence the bar, are far away, and (iii) the garage is neither near nor far. In other words, the places that are not far away are the office, home, and the garage. Consequently, there are many proper subsets of the question that entail the minimal content that Nathan is not far away – e.g., the subset covered by cell 1, the one covered by cells 1 and 5, etc. The weakest of these is the subset covered by cells 1, 4, and 5. So this proposition counts as what is said.

Given the way things are set up in this simplified example, this proposition can be paraphrased in more than one way, as in (27).

- (27) a. Nathan is neither at the bar nor in the mall.

- b. Nathan is either in the office, at home, or at the garage.
- c. Nathan is not far away.

These are all correct ways of describing what Nathan said, that is, of describing the proposition consisting of the worlds covered by cells 1, 4, and 5. But on any of them, what is said is clearly something Nathan believes to be true. Hence, he is not lying.

We can now ask why Nathan is seen as misleading – apart from the features we noted above concerning his use of the 3rd person. It is plausible to say that his utterance implicates that Nathan is not in the office. One way of suggesting how this implicature arises starts from the observation that what Nathan says, on any description, is weaker than the proposition that Nathan is in the office. The proposition that Nathan is in the office, i.e., the worlds covered by cell 1, entails the proposition that is said, i.e., the worlds covered by cells 1, 4, and 5. Correspondingly, on either of the descriptions in (27), what is said is entailed by the proposition that Nathan is in the office. Hence, the latter is a stronger statement than what Nathan asserted. Given a principle like Grice's (1989, 26) First Maxim of Quantity, "Make your contribution as informative as is required (for the current purposes of the exchange)", the speaker can be expected to contribute the strongest available content. According to this idea, then, since what was asserted is weaker than the proposition that Nathan is in the office, the utterance implicates that Nathan is not in the office.

5 Other Questions

We have seen that the ways in which one can mislead without lying depends on which QUD one is seen as addressing. Yet so far I have been assuming a simplified conception of how QUDs figure in discourses. In this section I look at two of these simplifications, and I comment on how a more realistic picture affects the ways in which one can lie vs. merely mislead. The first concerns the fact that QUDs are often implicit. The second concerns the fact that contexts typically contain multiple QUDs.

5.1 Implicit Questions and Prosodic Focus

The cases we have looked at until now have been cases in which the relevant QUDs are explicitly asked. Yet much communication does not proceed like this. We very often say things, even though there is no question explicitly asked. We can lie or be merely misleading in such cases, too. This might be thought to be a potential problem for my view, since its notion of what is said is only defined relative to QUDs. However, as I explain below, the fact that no question has been made explicit in the previous discourse does not mean that there are no QUDs.

Roberts (2012, 8) stresses that "questions are often only implicit, inferred on the basis of other cues." Specifically, Roberts's suggestion is that questions are often accommodated by the interlocutors in much the same way as presuppositions are. Retrieving the right question to accommodate, however, is arguably a much more complicated process than that of pre-

supposition accommodation.⁵¹ Even so, there are mechanisms that speakers use to indicate which questions they are addressing.

Roberts argues that one of these mechanisms, in English, is prosodic focus:

prosodic focus in English presupposes the type of question under discussion, a presupposition which enables the hearer, with some other contextually given clues, to reconstruct that question and its relation to the strategy being pursued. (Roberts, 2012, 8)

Roberts's discussion of this phenomenon is rich and highly detailed. My aim here is merely to sketch a rudimentary way of understanding the relation between focus and question accommodation as a way of suggesting one way in which QUDs can be in place even if they are not made explicit.⁵²

To take an example from Krifka (2011), note that (28) can be used to address either of the questions in (29).

(28) Fritz will go to Potsdam tomorrow.

- (29) a. Where will Fritz go tomorrow?
b. When will Fritz go to Potsdam?
c. Who will go to Potsdam tomorrow?
d. Who will go where tomorrow?
e. Who will go where when?

But, as Krifka points out, which question is being addressed will typically be indicated by intonational pattern. Representing focus by an *f*-subscript and stress using capitals, each intonational pattern in (30) is only felicitous relative to the corresponding question in (29).

- (30) a. Fritz will go [to POTsdam]_f tomorrow.
b. Fritz will go to Potsdam [toMORrow]_f.
c. [FRITZ]_f will go to Potsdam tomorrow.
d. [FRITZ]_f will go [to POTSDam]_f tomorrow.
e. [FRITZ]_f will go [to POTSDam]_f [toMORrow]_f.

In a case where there is no suitable QUD already in place, choosing one of the intonational patterns in (30) will indicate to the listener which of the questions in (29) is being addressed. In other words, prosodic focus – as expressed by intonational pattern – is one way in which speakers provide cues that allow other participants to accommodate QUDs. For example, uttering (30a) in a context in which no suitable question has previously been accepted as a QUD will typically have the effect that the participants will assume that the speaker is addressing (29a), and thereby accommodate that question as a QUD.

Focus can be used for misleading purposes in this way. Here is an example:

⁵¹Presupposition triggers differ in respect of the facility with which their presuppositions can be retrieved for accommodation. In some cases the process requires considerable contextual information. For example, this is notoriously the case with the iterative *too*, as noted by Kripke (2009).

⁵²As such, my discussion of focus and question accommodation diverges from Roberts (2012).

Melissa knows that Fritz is first going to Potsdam and then on to Rome tomorrow.

(31) Jack. Will Fritz go to Potsdam tomorrow?

Melissa. Fritz will go to [ROME]_f tomorrow.

Note that Melissa's utterance is not felicitous with a flat intonation. In other words, Melissa's utterance in (31) must be seen as addressing, not the question about Potsdam that was explicitly asked, but a different question, namely (29a). Relative to (29a), and ignoring the indexical, we predict that what is said is that Fritz will go to Rome tomorrow. Therefore, Melissa is not lying, since she is asserting something she believes to be true.

However, there is a clear sense that Melissa is being misleading in (31). We feel that her utterance conveys that Fritz is not going to Potsdam tomorrow, and this sense is particularly explicit due to the presence of Jack's question. I take it to be plausible that this effect can be explained on the basis of our prediction that what is said, in this case, is merely that Fritz is going to Rome tomorrow. Yet the issue of *how* to explain this effect is vexed and highly controversial. Even though I will not offer anything like a complete account of this issue here, it is worth gesturing at one way this kind of case might be explained within a framework of the kind we are presupposing.⁵³

On the Stalnaker-Roberts view of discourse, the goal of a discourse is to answer the Big Question – i.e., to reach the goal of inquiry – and this overarching goal is approached via the attempt to answer more local QUDs. Following Carlson (1982), Roberts uses the analogy of games, i.e., interactions involving goals, rules, moves, and strategies, to describe this endeavor. One type of rule that governs this kind of game, that is, the game of reaching the goal of inquiry via subinquiries, are conversational rules like the Gricean maxims, arising from rational considerations concerning the goal of the game. Along these lines, Roberts suggests that “Quantity 1 [follows] from the desire to maximize the payoff of a move in view of commitment to the ultimate goal [...]” (2012, 4) In particular, she argues that, when a QUD is raised, participants are generally expected to provide complete answers, if possible (see Roberts, 2012, 6.) So, if we are right that Melissa's utterance in (31) addresses (29a), there will be a presumption that she try to provide a complete answer to the latter question. Consequently, it will be presumed that Melissa wants to convey that Rome is the only place Fritz will go tomorrow. This might therefore be a way of explaining the inference that Fritz will not go to Potsdam as well. However, much more work is needed to flesh out this kind of account; let alone to determine its overall viability.

In other cases there is more room for maneuver concerning which question is to be accommodated, as in the following story:

In an office memo Barbara sees Doris's name on a list of recent lay-offs. Surprised she shows the memo to William and points to Doris's name on the list while making a quizzical, puzzled face.

(32) William. She insulted Sean at the Christmas party.

⁵³For other relevant discussion, see, e.g., Carston (1998), Levinson (2000), Sauerland (2004), Spector (2007).

Suppose, as in the previous scenario, that William knows that, while Doris did insult Sean, this is not the reason she got fired. Did he lie?

I think intuitions are unclear here. In particular, I think the degree to which one will judge William to be lying correlates with the degree to which one will judge that he can only be seen as addressing the question of why Doris got fired. To support this, suppose that the situation is as described above, but that Barbara, while pointing to the list, utters, “Why?”, “What?”, or the like. I think that, in these cases, intuitions that William is lying become stronger, and that the reason is that the degree to which his utterance can only be seen as addressing the question of why Doris got fired is strengthened.

One can avoid lying, while still being misleading, if one can exploit which question the hearers can be expected to accommodate, while remaining in a position to claim that one intended to address a different QUD.

5.2 Multiple Questions

The final issue I want to turn to concerns the fact that realistic contexts invariably contain multiple QUDs. Moreover, the different QUDs in a context are typically related to each other. One important way QUDs can be related is by specific kinds of entailment relations. I will look at two such relations and the way in which they influence ways of lying and misleading.

One notion of entailment between questions is the following.

Strict Question-Entailment

A question q_1 strictly entails a question q_2 if and only if completely answering q_1 yields a complete answer to q_2 .⁵⁴

Following Roberts (2012, 16), let us consider a simplified example of Strict Question-Entailment. Imagine a discourse concerning only two individuals, Hilary and Robin, and concerning only two kinds of foods, tofu and bagels. Then consider the following questions:

- (33) Who ate what?
- a. What did Hilary eat?
 - i. Did Hilary eat bagels?
 - ii. Did Hilary eat tofu?
 - b. What did Robin eat?
 - i. Did Robin eat bagels?
 - ii. Did Robin eat tofu?

Given the narrow domain described above, (33) strictly entails both (33a) and (33b), which in turn strictly entail the questions below them. A complete answer to (33) yields complete answers to both (33a) and (33b) and to their subquestions. This is easily illustrated by considering a fragment of the sixteen cell partition determined by (33). Consider, for example,

⁵⁴Cf. Roberts (2012, 7), Roberts (2012, 12), Groenendijk & Stokhof (1984, 16).

the three randomly selected cells from [(33)] in Table 4 (where *ht* means ‘Hilary ate tofu’, *rb* means ‘Robin ate bagels’, and so on.)

1	2	3
<i>ht</i> \neg <i>hb</i>	<i>ht</i> \neg <i>hb</i>	<i>ht</i> \neg <i>hb</i>
<i>rt</i> <i>rb</i>	<i>rt</i> \neg <i>rb</i>	\neg <i>rt</i> <i>rb</i>

Table 4.

If one answers (33) by specifying that Hilary ate tofu but not bagels and Robin ate both, one’s answer is represented by cell 1 of Table 4. As seen, this answer also yields complete answers to the subquestions, i.e., it also completely answers what Hilary ate, what Robin ate, and whether each of them ate tofu, bagels, both, or neither.

We can now see that our definition of what is said has the consequence that saying, and hence assertion, is preserved from subquestions to superquestions under Strict Question-Entailment. This can be spelled out as follows:

Upwards Preservation of What is Said

If a question q_1 strictly entails a question q_2 , then

If a proposition p is said relative to q_2 , p is said relative to q_1 .

To see that what is said is preserved in this way, observe that (34) says that Robin ate tofu relative to all of (33b_{ii}), (33b), and (33).

(34) Robin ate tofu.

However, even though (34) says that Robin ate tofu relative to (33), it says nothing relative to (33a), (33a_i), (33a_{ii}), and (33b_i). That is not a counterexample to Upwards Preservation of What is Said. Yet it shows that assertion is not preserved downwards from superquestions to subquestions under Strict Question-Entailment. That is, a proposition may be said relative to a question q_1 but not relative to a question q_2 it entails.

Earlier we assumed that when an utterance does not make an assertion with respect to a local QUD, the interpretation typically defaults to the Big Question. For example, if the only QUD available is (33b_i), (34) will be seen as addressing the Big Question, and thereby as saying that Robin ate tofu.

Since the Big Question is the most general question of all, it strictly entails all other questions.⁵⁵ The proposal that interpretations default to the Big Question can therefore be seen as a special case of a more general principle according to which, when there is no way of obtaining an assertion for an utterance from a local QUD, the interpretation moves up the hierarchy of strictly entailing QUDs until it reaches a QUD relative to which the utterance makes an assertion. Once this is achieved, since assertion is preserved upwards, the same proposition will be asserted relative to all strictly entailing super-QUDs, and ultimately the Big Question itself.

⁵⁵See Roberts (2012, 7).

This means that if one lies relative to a question, one lies relative to all questions that strictly entail it. Moreover, since the Big Question is always available, whether one can lie or not does not depend on which local QUDs are in place. As such, lying is always an option – namely, one can always utter a sentence that says something one believes to be false. Yet, as we will see next, the option of misleading without lying does depend on which QUDs are available.

If one conveys a misleading answer to a question, one will thereby also be conveying a misleading answer to all questions that strictly entail it. This is clear from (35).

- (35) Mark. Who ate what?
Kevin. Well, let's see – did Robin eat tofu?
Albert. Robin ate bagels.

Albert's response to Kevin's question is most naturally seen as conveying that Robin did not eat tofu.⁵⁶ Suppose Albert believes that Robin ate both bagels and tofu. In that case he is being misleading, but is not lying. He conveys disbelieved information without saying it. This misleading information is also a partial answer to Mark's superquestion, and hence Albert is also misleading with respect to this question.

However, this only holds *if* there is a subquestion against which one is being misleading. If, for example, the discourse in (35) only contained Mark's question, then Albert's utterance would not have been misleading – it would just have said that Robin ate tofu. In other words, whether or not one can succeed in misleading while avoiding outright lying depends on which QUDs are in place, and which QUD one is seen as addressing. But if one does contribute a misleading answer (even if one does not assert it) to a QUD, the same information will be a misleading answer to each strictly entailing QUDs.

It is of course implausible to think that QUDs are always (perhaps even often) related to each other by the relation of Strict Question-Entailment. Still, it is a natural thought that, in well functioning, truth-pursuing discourses, QUDs are related in ways that are similar to this. In particular, QUDs are arguably are often ordered such that answering them in turn will, step by step, be a way of progressing toward the goal of inquiry.

This idea is captured by Roberts's (2012) model of discourse structure. Roberts assumes that the QUDs of a discourse are totally ordered by a precedence relation $<$, which is simply the relation of one utterance having been made before another.⁵⁷ She then defines the following relation on the QUDs of a discourse:

⁵⁶In accordance with what was said in section 5.1, there is a clear preference for the intonational pattern in (i) for Albert's utterance in (35).

(i) Robin ate [BAgels]_f.

As such, the utterance might be taken to indicate that it is addressing (33b). This does not detract from the point I am making concerning the way lying and misleading behave under Strict Question-Entailment.

⁵⁷See clause (10d) in Roberts (2012, 14). To be sure, this is arguably too simple in that QUDs can be hierarchically ordered with respect to each other, independently of the order in which they were uttered, and furthermore. I ignore this here, though.

Contextual Question-Entailment

For all QUDs q, q' of a discourse, if $q < q'$, then the complete answer to q' contextually entails a partial answer to q .⁵⁸

One way in which Contextual Question-Entailment can be satisfied is if the QUDs of a discourse stand in the relation of Strict Question-Entailment. For example, a complete answer to (33b) yields a partial answer to (33), and so on. But Contextual Question-Entailment also allows for cases in which the relation between the QUDs of a discourse is more subtle than Strict Question-Entailment. To take an example from Roberts (2012, 19–20), consider the questions in (36).

- (36) a. What kinds of seafood will John eat?
b. Isn't he allergic to clams?

As Roberts (2012, 19) observes, completely answering (36b) does not straightforwardly entail a partial answer to (36a). Suppose that the answer to (36b) is *yes*. Even so, John might still eat clams for other reasons, e.g. out of politeness, or because he cannot resist. On the other hand, suppose the answer to (36b) is *no*. Even so, that does not directly entail that John *will* eat clams, since there might be other reasons why he will not. However, Roberts proposes that in many contexts in which this kind of discussion takes place, it will be common ground that (a) one will not eat anything one is allergic to and (b) one will eat something unless one has reason not to. (a) takes care of the case in which the answer to (36b) is *yes*. (b) takes care of the case in which the answer is *no*, given the further observation that an answer to (36b) is a partial answer to (37).

- (37) What reasons would John have for not eating clams?

In other words, Contextual Question-Entailment appears to be a promising way of explaining at least some of the intricate ways in which QUDs are related. Briefly, we try to order our QUDs such that those raised will deliver partial answers to the previous ones, given the common ground, and perhaps via bridge-questions.

Now consider the following situation:

Maria knows that John is allergic to clams, but that he will eat them anyway because he can't resist.

- (38) Kelly. What kinds of seafood will John eat?
Eric. Isn't he allergic to clams?
Maria. Yes, John is allergic to clams.

Intuitively, Maria is being misleading, since she conveys that John will not eat clams, and yet she is not lying because all she is asserting is that he is allergic to clams.

The reason Maria is not lying in (38) is clear. Her utterance expresses the minimal proposition that John is allergic to clams. The latter is entailed by the corresponding answer to Eric's question, and is therefore asserted by Maria. So we explain why she is not lying.

⁵⁸This is a slight simplification of clause (10g_{iii}) in Roberts (2012, 15).

How do we explain the sense that Maria is being misleading in (38)? That is, how do we explain that Maria succeeds in adding to the discourse the information that John will not eat clams? As noted, Roberts suggests that it is reasonable to assume that it is common ground that one will not eat something one is allergic to. It is natural to explain Maria's strategy for misleading as trading on a fact of this kind. The reason she can mislead Kelly and Eric into thinking that John will not eat clams by asserting that he is allergic to clams is because she can rely on a general assumption to the effect that people typically avoid eating what they are allergic to. In other words, the operation of Contextual Question-Entailment, and the configuration of the common ground, help explain the misleading nature of Maria's utterance.

As before, that Maria can succeed in misleading is due in part to which QUDs are in place – in this case, the misleadingness is due to the presence of Kelly's question. For example, suppose the question of what kinds of seafood John will eat is not a QUD, neither explicitly nor implicitly, e.g., as in (39).

- (39) Kelly. What ailments does John have?
Eric. Isn't he allergic to clams?
Maria. Yes, John is allergic to clams.

In this case Maria does not convey that John will not eat clams, or any other misleading information. So, whether one misleads or not depends on the constellation of QUDs in the context, and furthermore, on how they are related to each other.

As seen from (38), Contextual Question-Entailment is not sufficient to ensure that what is said is preserved from subquestions to superquestions, as Strict Question-Entailment is. Nevertheless, as demonstrated by (39), Contextual Question-Entailment does ensure something weaker. If one asserts a complete answer to a subquestion, in a discourse satisfying Contextual Question-Entailment, one will thereby be contributing a partial answer to its superquestions. Therefore, in such a discourse, one way of contributing a disbelieved partial answer to a QUD, while avoiding lying, is to assert a complete answer to a subquestion that one believes to be true but which, given the context, will be a way of achieving one's goals with respect to the superquestion.

6 Conclusion

To mislead is to disrupt the pursuit of the goal of inquiry, to prevent the discovery of how things are. To lie is to do so by assertion. Assertion is a mode of contributing information to a discourse that is sensitive to which question is under discussion. Therefore, whether one lies or merely misleads sometimes depends on the subinquiry one is engaged in, or which QUD one is addressing. In particular, one can lie by committing to an answer to a QUD that one believes to be false. Conversely, some misleading utterances avoid being outright lies by making sure they do not commit to disbelieved answers to QUDs.

As a collective system for pooling true information, inquiry is usually a prerequisite for informed action and decision making. In relying on the results of inquiry, we depend on the contributions of others. We are therefore naturally sensitive to the difference between when

a speaker commits herself to an answer to a subinquiry as opposed to merely suggesting an answer in more indirect ways. Contributions of the former kind are, at least typically, more likely to be drawn on in subsequent deliberation. This centrality of asserted information may be one source of the intuitive weight given to lying vs. merely misleading.

The central claim of this paper has been that what is said by a sentence, in context, depends systematically on QUDs, while being constrained by compositional meaning. On this view, QUDs play a central role in how what is said, or truth-conditional content, is determined, in context. As such, while the various Contextualist positions in the debate over the semantics-pragmatics distinction are right that what is said may go beyond linguistically encoded meaning, the view defended here provides a systematic theory of truth-conditional content that shows how what is said, in this sense, depends on compositional meaning.⁵⁹

References

- Adler, J. (1997). Lying, deceiving, or falsely implicating. *Journal of Philosophy*, 94, 435–452.
- Bach, K. (1994). Conversational implicature. *Mind and Language*, 9, 124–162.
- Barker, C. (1995). *Possessive descriptions*. Stanford, CA: CSLI Publications.
- Bok, S. (1978). *Lying: Moral choice in private and public life*. New York: Random House.
- Borg, E. (2004). *Minimal semantics*. Oxford: Oxford University Press.
- Cappelen, H., & Lepore, E. (2004). *Insensitive semantics – a defense of semantic minimalism and speech act pluralism*. Oxford: Blackwell.
- Carlson, L. (1982). *Dialogue games: An approach to discourse analysis*. Dordrecht, the Netherlands: Reidel.
- Carson, T. (2006). The definition of lying. *Noûs*, 40, 284–306.
- Carson, T. (2010). *Lying and deception: Theory and practice*. Oxford and New York: Oxford University Press.
- Carston, R. (1998). Informativeness, relevance, and scalar implicatures. In R. Carston & S. Uchida (Eds.), *Relevance theory – applications and implications* (Vol. 37, pp. 179–235). Amsterdam: John Benjamins.
- Carston, R. (2002). *Thoughts and utterances*. Oxford: Blackwell.
- Chisholm, R., & Feehan, T. (1977). The intent to deceive. *Journal of Philosophy*, 74, 143–159.
- Cooper, R. (1983). *Quantification and syntactic theory*. Dordrecht: Reidel.
- Davis, W. (2010). Implicature. *Stanford Encyclopedia of Philosophy*.
- Fallis, D. (2009). What is lying? *Journal of Philosophy*, 106, 29–56.
- Gauker, C. (2012). What Tipper is ready for: A semantics for incomplete predicates. *Noûs*, 46(1), 61–85.
- Grice, H. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Groenendijk, J., & Stokhof, M. (1984). *Studies on the semantics of questions and the pragmatics of answers*. PhD thesis, University of Amsterdam. <http://dare.uva.nl/record/123669>.
- Hamblin, C. (1973). Questions in Montague English. *Foundations of Language*, 10, 41–53.

⁵⁹See Schoubye and Stokke (2015) for a complete articulation and defense of this view.

- Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9, 183–221.
- Heim, I. (2008). Features on bound pronouns. In D. Harbor, D. Adger, & S. Bèjar (Eds.), *Phi-theory: Phi-features across modules and interfaces* (pp. 35–57). Oxford and New York: Oxford University Press.
- Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. Oxford: Blackwell.
- Isenberg, A. (1964). Deontology and the ethics of lying. *Philosophy and Phenomenological Research*, 24, 465–480.
- Kant, I. (1797). *On a supposed right to lie from philanthropy*. (M.J. Gregor, trans.), in M.J. Gregor (ed.), *Immanuel Kant, Practical Philosophy*, Cambridge and New York: Cambridge University Press, 1996, pp. 611–615. First published as ‘Über ein vermeintes Recht, aus Menschenliebe zu lügen’, in *Berlinische Blätter* 10, Blatt (6 September 1797), pp. 301–314.
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford and New York: Oxford University Press.
- Korsgaard, C. (1986). The right to lie: Kant on dealing with evil. *Philosophy and Public Affairs*, 15, 325–349.
- Krifka, M. (2011). Questions. In K. von Heusinger, C. Maienborn, & P. Portner (Eds.), *An international handbook of natural language meaning* (Vol. 2, pp. 1742–1785). de Gruyter.
- Kripke, S. (2009). Presupposition and anaphora: Remarks on the formulation of the projection problem. *Linguistic Inquiry*, 40, 367–386. (Edited transcript of a talk given in 1990)
- Kupfer, J. (1982). The moral presumption against lying. *Review of Metaphysics*, 36, 103–126.
- Levinson, S. (2000). *Presumptive meanings – the theory of generalized conversational implicature*. Cambridge, MA: MIT Press.
- Lincoln, A. (1953). To James M. Ashley. In R. Basler (Ed.), *Collected works* (Vol. 8). New Brunswick, NJ: Rutgers University Press. Retrieved from <http://name.umd1.umich.edu/lincoln8>
- MacIntyre, A. (1995). Truthfulness, lies, and moral philosophers: What can we learn from Mill and Kant? *The Tanner Lectures on Human Values*, 16, 307–361.
- Mahon, J. (2006). Kant and the perfect duty to others not to lie. *British Journal for the History of Philosophy*, 14, 653–685.
- Mahon, J. (2009). The truth about Kant on lies. In C. Martin (Ed.), *The philosophy of deception* (pp. 201–224). Oxford and New York: Oxford University Press.
- Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15, 509–559.
- Paton, H. (1954). An allegeded right to lie: A problem in Kantian ethics. *Kant-Studien*, 45, 190–203.
- Recanati, F. (1989). The pragmatics of what is said. *Mind and Language*, 4(4), 295–329.
- Recanati, F. (1993). *Direct reference – from language to thought*. Oxford: Blackwell.
- Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
- Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford and New York: Oxford University Press.

- Roberts, C. (2004). Context in dynamic interpretation. In L. Horn & G. Ward (Eds.), *The handbook of pragmatics* (pp. 197–220). Oxford: Blackwell.
- Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics*, 5, 1–69.
- Sauerland, U. (2004). Scalar implicatures in complex sentences. *Linguistics and Philosophy*, 27, 367–391.
- Sauerland, U. (2008). On the semantic markedness of phi features. In D. Harbour, D. Adger, & S. Bèjar (Eds.), *Phi-theory: Phi-features across modules and interfaces* (pp. 57–83). Oxford and New York: Oxford University Press.
- Saul, J. (2012a). Just go ahead and lie. *Analysis*, 72(1), 3–9.
- Saul, J. (2012b). *Lying, misleading, and what is said: An exploration in philosophy of language and in ethics*. Oxford & New York: Oxford University Press.
- Schoubye, A., & Stokke, A. (2015). *What is said?* MS.
- Soames, S. (2008). Drawing the line between meaning and implicature – and relating both to assertion. *Noûs*, 42, 440–465.
- Solan, L., & Tiersma, P. (2005). *Speaking of crime: The language of criminal justice*. Chicago: University of Chicago Press.
- Sorensen, R. (2007). Bald-faced lies! Lying without the intent to deceive. *Pacific Philosophical Quarterly*, 88, 251–264.
- Spector, B. (2007). Scalar implicatures: Exhaustivity and gricean reasoning. In M. Aloni, A. Butler, & P. Dekker (Eds.), *Questions in dynamic semantics*. Elsevier. (<http://www.semanticsarchive.net/Archive/WU2MGJkN/>)
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Oxford: Blackwell.
- Stalnaker, R. (1970). Pragmatics. In *Context and content* (pp. 31–46). Oxford and New York: Oxford University Press, 1999.
- Stalnaker, R. (1978). Assertion. In *Context and content* (pp. 78–95). Oxford and New York: Oxford University Press, 1999.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.
- Stalnaker, R. (1998). On the representation of context. In *Context and content* (pp. 96–114). Oxford and New York: Oxford University Press, 1999.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25, 701–721.
- Stanley, J. (2000). Context and logical form. In *Language in context – selected essays* (pp. 30–68). Oxford: Clarendon Press, 2007.
- Stanley, J., & Szabó, Z. G. (2000). On quantifier domain restriction. In *Language in context – selected essays* (pp. 69–110). Oxford: Clarendon Press, 2007.
- Stokke, A. (2013). Lying and asserting. *Journal of Philosophy*, CX(1), 33–60.
- Stokke, A. (2014). Insincerity. *Noûs*, 48(3), 496–520.
- Travis, C. (1985). On what is strictly speaking true. *Canadian Journal of Philosophy*, 15(2), 187–229.
- Williams, B. (2002). *Truth and truthfulness: An essay in genealogy*. Princeton, NJ: Princeton University Press.
- Wood, A. (2008). *Kantian ethics*. Cambridge and New York: Cambridge University Press.