



Intentionalism and bald-faced lies

Daniel W. Harris [†]

Hunter College, New York, NY, USA

ABSTRACT

In *Lying and Insincerity*, Andreas Stokke argues that bald-faced lies are genuine lies, and that lies are always assertions. Since bald-faced lies seem not to be aimed at convincing addressees of their contents, Stokke concludes that assertions needn't have this aim. This conflicts with a traditional version of intentionalism, originally due to Grice, on which asserting something is a matter of communicatively intending for one's addressee to believe it. I argue that Stokke's own account of bald-faced lies faces serious problems and give several responses on behalf of intentionalism. Some bald-faced lies are best understood as irrational attempts to deceive. Others are best understood as indirect speech acts of various kinds. Still, others are best understood as conventional speech acts, which differ from communicative acts like assertion in the ways that they must be embedded in social institutions or practices. An overarching theme of this essay is that we should not make theoretical decisions about how to classify speech acts by consulting ordinary usage.

KEYWORDS lying; bald-faced lies; Grice; intentionalism; speech acts

What is it to lie, or to be otherwise insincere? In *Lying and Insincerity*, Andreas Stokke defends detailed answers to these questions, but he also does much more than this. The answers that he defends are embedded in a broader theory of language use, and the answers that he rejects are embedded in rival pragmatic theories of their own. Stokke therefore treats his central topic as one battlefield on which larger disputes over the nature of speech acts and communication are to be fought.

This is also the spirit in which I will respond to Stokke. One of the rivals to his theory is my own favored account of speech acts and communication – a view that I will call 'intentionalism'. Stokke rejects this view on the ground that it doesn't fit with the ways that we normally think and talk about lying and insincerity. My aim here is to respond on behalf of intentionalism to this argumentative thread in Stokke's book.

The main phenomenon that fuels Stokke's objection to intentionalism is bald-faced lying, which Stokke describes as a kind of lying that is not

CONTACT Daniel W. Harris  danielwharris@gmail.com  Hunter College, New York, NY, USA

[†]Thanks to Andreas Stokke and Jessica Keiser for comments on a previous draft.

intended to deceive the addressee. After a brief overview of intentionalism (Section 1), I will outline Stokke's objection (Section 2). I will then explain Stokke's way of handling bald-faced lies (Section 3) and argue that it doesn't make sense of fleshed-out versions of some of his own cases (Section 4). I will then outline several strategies by means of which intentionalists should respond to Stokke's objection. First, I will argue that we should give little weight to whether a speech act counts as an assertion or a lie by the standard of ordinary usage, and that we should instead seek theoretical roles for these concepts to play (Section 5). I will then take a divide-and-conquer approach to putative examples of bald-faced lies, arguing that they should be divided between at least three categories, none of which poses a threat to intentionalism: some are irrational attempts to deceive (Section 6), some are indirect speech acts (Section 7), and some are conventional acts (like the act of testifying in court) rather than communicative acts (like assertion) (Section 8).

1. Intentionalism

Intentionalism is the view that humans communicate by revealing intentions to change each other's minds. To perform a communicative act is to produce an utterance with a communicative intention – an intention to have a psychological effect on one's addressee in part by revealing to them the intention to do so.¹ Successful communication occurs if the addressee recognizes what effect the speaker intended to have on them. Actually producing the intended effect is a further kind of success.

The term 'speaker' is somewhat misleading. Communication needn't involve the use of language.² I can perform a communicative act by means of any overt behavior that, I think, could lead my addressee to infer what I intend. Given the right circumstances, I can communicate with a wink or by throwing my roommate's possessions out the window. What language adds to this picture is a means for encoding and decoding richly structured evidence of communicative intentions, and so a way of getting much more elaborate intentions recognized. Of course, it normally wouldn't make sense for me to produce any utterance with any communicative intention I please. This follows from a broader

¹This is what Grice (1969) called performing an act of 'utterer's occasion meaning' and many others have called 'speaker(s) meaning'. The term 'communicative act' originates with Strawson (1964) and Bach and Harnish (1979). There are many variations on intentionalism, and I won't attempt to speak on behalf of all of them here.

²It is notable that seven of Grice's (1957) eight original examples of utterer's meaning – what I am calling 'performing a communicative act' – are nonlinguistic. On this point, see Harris (2016, Section 5).

fact about intentions, which is that it is irrational to intend something that one takes to be impossible (Bratman 1987; Broome 1999).

A variety of different communicative acts can be understood in these terms. I have elsewhere defended a simple version of intentionalism that is due to Grice (1968, 1969).³ On this view, asserting p is a matter of communicatively intending for one's addressee to believe p . To direct someone to ϕ is to communicatively intend for them to form an intention to ϕ . To ask someone a question is to communicatively intend for them to form an intention to answer it. According to intentionalists, performing and interpreting communicative acts requires only the ability to act with communicative intentions and recognize them in others. These abilities rest on cognitive capacities for intention formation and mindreading that are universal to normal human adults. Communicative acts are among the basic instruments of human social cognition.

However, not all speech acts are communicative acts. For example, the speech acts that figure in judicial and legislative proceedings, police confessions, wedding ceremonies, and other ritualized social interactions depend for their existence and success on much more than the cognitive capacities of those involved. Their performance is made possible by facts about the social conventions and institutions that are operative in the situations where they occur. In order for an utterance to count as a marriage vow, for example, it must abide by all of the rules that govern marriage vows in the relevant jurisdiction. In Nova Scotia, where I was married, my vow had to include certain words in a certain order, in either English or French (or, as it happened, a mix of both), uttered at the right point in the ceremony, under the supervision of one of the few locals who had been granted the legal power to officiate marriage ceremonies, and so on. My marriage vow was a 'conventional act' – one whose conditions of performance in a given jurisdiction are defined by that jurisdiction's social institutions and conventions. Communicative and conventional acts are fundamentally different, and mustn't be run together. This requires care because of their superficial similarity.

2. Bald-faced lies

Stokke's case against intentionalism revolves around bald-faced lies, which he takes to be lies that the liar knows won't deceive their addressees. Stokke (2018, 18) adapts the following example from Carson (2006, 290):

³See, for example, Harris (2014, 2017, 2019a, 2019b, 2020).

The Cheating Student

A student accused of cheating on an exam is called to the Dean's office. The student knows that the Dean knows that she did in fact cheat. But as it is also well known that the Dean will not punish someone unless they explicitly admit their guilt, the student says,

(1) I didn't cheat.

Stokke invites us to share his intuition that the student has lied in uttering (1). To lie, Stokke argues, is to assert something that one doesn't believe.⁴ It follows that in uttering (1), the student asserted that they didn't cheat. But, Stokke argues, the student knows that they won't convince the dean, and so doesn't intend to do so.

If Stokke's argument is sound, then we have an assertion that isn't intended to make the addressee believe its content. This appears to be a counterexample to the form of intentionalism that I wish to defend.

3. Stokke on lying, assertion, and common ground

Stokke's solution to the problem of bald-faced lies is to adopt a theory of assertion that, he thinks, allows him to classify bald-faced lies like (1) as assertions. The theory that he adopts is originally due to Stalnaker (1978). On this view, asserting p is a matter of proposing that p be added to the conversation's common ground.⁵ A proposition p is part of a conversation's common ground just in case the participants all accept p for the purpose of the conversation, and commonly believe that they do so (Stalnaker 2002, 716; Stokke 2018, 51). On Stalnaker's account, which Stokke accepts, it is possible to accept p for some purpose without believing p (Stalnaker 1984, 79–81). To accept a proposition for a given purpose is to treat it as true for that purpose, but not necessarily for others. If p is in the common ground, then the participants in the conversation may felicitously presuppose p within the conversation, for example, but they needn't act on p when the conversation is over. 'A person may accept something in one context, while rejecting it or suspending judgment in another', says Stalnaker (1984, 80): 'There need be no conflict that must be resolved when the difference is noticed, and he need not change his mind when he moves from one context to the other'.

On Stokke's view, all bald-faced lies are lies, and all lies are assertions. In lying to the dean, the student aims to add the proposition that they didn't

⁴This is a slight simplification of Stokke's view.

⁵I am ignoring some irrelevant details of Stokke's account.

cheat to the common ground. Succeeding at this does not entail making the dean believe the lie. It entails only that the dean accepts, for the purpose of the conversation, that the student didn't cheat – a much easier goal to achieve. Stokke (2018, 52) motivates this way of thinking about the case as follows.

The shared intuition about the case of the cheating student is that the reason the student makes her utterance – despite the fact that both she and the dean know full well that what she says is false – is that she wants to 'go on the record'. This idea lends itself readily to be explained in terms of the common ground. Namely, to say that the student wants to go on the record is just to say that the student wants it to be common ground that she did not cheat.

If this account works, then it allows Stokke to maintain that the student asserted that he didn't cheat, and thereby lied.

4. An objection to Stokke's account

Stokke's theory conflicts with the kind of intentionalism that I outlined in Section 1, since it gives up on the idea that genuine assertions are intended to change addressees' beliefs. I have elsewhere argued that common-ground-centric accounts of assertion and other speech acts don't work, and that the kind of intentionalism that I outlined in Section 1 is to be preferred.⁶ So I am not inclined to go along with Stokke's account.

I also think that there is a serious problem with Stokke's way of accounting for bald-faced lies. In the case at hand, Stokke conflates two different purposes for which the dean can accept the proposition that the student didn't cheat. It is one thing for the dean to treat this proposition as true for the purpose of the conversation, as he would have to in order for the proposition to be common ground. It is another thing for the dean to treat it as true for official disciplinary purposes. The latter – an extra-conversational matter – is what the student really cares about, and is what explains why they perform the speech act that they perform.

Consider one way that the student's conversation with the dean could play out:

⁶In Harris (2017), I argue that common-ground-centric theories of speech acts collapse categories of speech acts that should be distinguished. In Harris (2020), I argue that we can communicate without changing the common ground and without intending (or proposing) to change it, and that common ground should be thought of as an artifact of idealized models of conversation rather than as part of the phenomenon of communication itself.

STUDENT: I didn't cheat.

DEAN: That's false and you and I both know it. But I am not going to punish you since you haven't confessed. You'll just have to live with your guilt.

Let p be the proposition that the student didn't cheat. From the dean's response, we can see that the student's utterance not only hasn't convinced the dean to believe p ; it hasn't led the dean to accept p for the purpose of their conversation, either. He comes right out and says that p is false, after all. If their conversation continues, it would be infelicitous for either the student or the dean to presuppose p .

If we take the student to have been asserting p , then their speech act was a failure. But even in this extension of Stokke's scenario, the speech act *wasn't* a failure; it was a success. The student, we can reasonably suppose, doesn't care about how the conversation goes; they care about whether they will be punished. That is an extra-conversational goal – one whose accomplishment depends on how the student's utterance interacts with the dean's policy about punishment and confession, not on how the dean responds within the conversation. Put a different way, the student doesn't care whether the dean treats p as true *for the purposes of the conversation*; the student just wants the dean to treat p as true *for official disciplinary purposes*. And these two ways of accepting p as true can come apart. Achieving the former would, according to Stokke, be part of the aim of asserting p . But it is the achievement of the latter that is the student's intention. If so, then the student wasn't really performing an assertion at all, at least by Stokke's lights.

So, what speech act *did* the student perform, if not an assertion? I will consider two answers to this question in Sections 7–8. First, I will lay out some other parts of what I think the intentionalist should say about bald-faced lies in general.

5. What are we trying to accomplish here?

A crucial premise in Stokke's argument is that bald-faced lies are lies. One person who has rejected this premise is Dynel (2011), who contends that 'the folk term 'bald-faced lie' should not be mistakenly equated with the concept of a lie approached theoretically' (151). Stokke's response is that Dynel runs afoul of ordinary usage:⁷

⁷Stokke also responds to Dynel by pointing out that 'no sensible theories of lying propose to *equate* the terms "bald-faced lie" and "lie"', since 'any sensible theory must acknowledge the datum that many lies are deceptive lies'. True enough, but I doubt that Dynel can be charitably understood as accusing anyone of making *that* mistake.

Furthermore, it is implausible to claim that the student merely lied in some loose sense of *to lie*. Rather, it is the insistence that the student did not lie that relies on a non-standard sense of the word. If one wants to deny that bald-faced lies are lies, one will be forced to admit that one is confining one's notion of lying to a subclass of what we ordinarily take to be lies. (Stokke 2018, 21)

In other words, it is unintuitive, or contrary to the way that we would ordinarily use the verb, 'to lie', to say that bald-faced lies aren't lies.

I won't try to dispute Stokke's intuition. But I doubt that it should carry much weight. Doesn't the methodology of testing definitions against intuitive judgments presuppose that our aim is to analyze the ordinary concept of lying? Haven't all past attempts to analyze ordinary concepts ended in dismal failure?⁸ Why should we be particularly interested in our ordinary way of thinking about lying anyway? Couldn't our ordinary concept turn out to be inconsistent, or in need of significant revision, or at least missing some theoretically important nuance?

In the last section, I also suggested that some of the speech acts that ordinarily get classified as bald-faced lies aren't assertions, and I will be defending this claim later. One could try to respond to this claim in the same way that Stokke responds to Dynel – by saying that my suggestion flies in the face of our ordinary concept of assertion. But I think that adherence to ordinary usage should be even less appealing in this case. After all, the verb 'to assert' and its cognates barely show up in ordinary, non-philosophical English, and don't bare much relation to the various philosophical notions of assertion when they do. 'Assertion' is an unabashed technical term. And so the goal of capturing ordinary usage seems particularly out of place when it comes to assertion.

All of this raises a methodological question: If our goal when giving theories of lying and assertion should not be to analyze ordinary concepts, what should it be? And how should this affect our methodology?

In the case of assertion, I think that there is a relatively clear answer to these questions. On my view, assertion is a theoretical posit that earns its keep by playing a role in a theory of human communication. I take the best theory of at least much human communication to be the sort of intentionalism that I outlined in Section 1. In particular, I think that we often seek to change others' beliefs, and that a powerful

⁸For a more belabored version of this point, see Fodor (1998, Chap. 3–4). For an amusing illustration involving the concept *chair*, see Elbourne (2011).

and flexible way of accomplishing this is to offer them evidence that we intend to do so, banking on the assumption that this will give them a reason to change their mind accordingly. This is what I call 'assertion'. To justify this account of assertion, I would have to give reasons to think that the overall theory in which it is embedded is a good account of human communication. Of course, this does not preclude the possibility that there are other interesting concepts of assertion that might play valuable roles in other theories. Mathematicians sometimes call the lines of proofs that they endorse 'assertions', for example, and I have no problem with this additional technical usage. Homonymy and polysemy pose no problem so long as we don't let them confuse us. Whether or not either of these technical uses lines up with the ordinary way that we use the verb, 'to assert' (insofar as we use it at all) is beside the point.

In order to carry this methodology over to the concept of lying, we would need to know which theoretical enterprise(s) a theory of lying seeks to contribute to. There are various possibilities. Perhaps our ancestors' capacity for lying played an important role in the evolution of the human mind, and so we need to figure out the right concept of lying in order to articulate the biological functions of some of our cognitive capacities.⁹ Alternatively, we might need a concept of lying to play a role in an ethical theory that seeks to specify which actions are permissible. Perhaps we need concepts of lying for both of these purposes.

But some caveats are in order. First, I see no reason to assume in advance that all such inquiries will require the same concept of lying, any more than a theory of human communication and theory of mathematical proofs will converge on the same concept of assertion. Second, the concept of lying has to earn its keep in these inquiries. If it turns out that all of the important normative distinctions can be made using a broader concept of dishonesty, for example, then it may be that lying just isn't an interesting concept for the purposes of normative theorizing.¹⁰ Third, and most importantly for present purposes, it may turn out that the concept of lying that best plays a given theoretical role is one whose extension excludes some or all of what we would normally call 'bald-faced lies'. In fact, I think that this is quite plausible. Consider the concept of lying that we might want to play a role in normative ethics.

⁹See, for example, Trivers (2011).

¹⁰For an argument for this claim, see Barber (2020).

There seem to be important normative differences between deceptive lies and bald-faced lies. Successful deceptive lies spread falsehood and disinformation but successful bald-faced lies do not. We therefore have different reasons for disapproving of them, and different responses to them might be warranted. From the point of a consequentialist normative theory on which we have a particular interest in understanding which actions promote true belief and which promote misinformation, we might not want to lump deceptive lies and bald-faced lies together in a single category. Likewise, a Kantian ethicist who locates the wrong of lying in the intent to deceive might have theoretical use for a concept of lies that does not include bald-faced lies.

Again, the fact that these concepts of lying might run afoul of ordinary usage is beside the point. A concept of lying that does important theoretical work for us might turn out not to fit with ordinary usage, but then again it might not. Many concepts that have turned out to be indispensable for the purposes of scientific and normative theorizing bear little relation to ordinary usage. Is there any reason to expect that the concept of lying will be exceptional in this regard?¹¹

6. Some bald-faced lies are irrational attempts to deceive

Keeping the methodological detour of the last section in mind, let's turn to the question of what an intentionalist should say about various speech acts that we would normally classify as bald-faced lies. The first possibility that I want to explore is that some bald-faced liars intend to deceive even though they believe that they can't succeed.

Stokke rules out this possibility with the following principle.

Belief-Intention Constraint

A intends to ϕ only if *A* believes that she can ϕ . (Stokke 2018, 25)

If this principle is correct, then it is not possible to believe that one can't deceive someone and intend to do so anyway. Stokke presumably thinks that there is something about the nature of belief and intention that makes this combination of attitudes impossible.

¹¹I anticipate that some, having read this section, will accuse me of hypocrisy. Isn't intentionalism itself the project of analyzing concepts like *meaning* by testing definitions against intuitions and ordinary usage? I think this charge is accurate when directed against some of the founding documents of intentionalism, such as Grice (1957, 1969) and Schiffer (1972). However, I have elsewhere criticized early intentionalists' methodology and defended an alternative approach for precisely this reason (Harris 2019b). Intentionalism, I argue, should be construed as an empirical theory of the cognitive capacities underlying human communication, not as a project in conceptual analysis.

I think that Stokke is wrong about this, and I reject his principle. In its place, I think we should adopt the following principle:

Belief-Intention Constraint on Rationality

If A intends to ϕ and A 's beliefs rule out the possibility that A will ϕ , then A 's mental state is in this respect irrational.

On this version of the principle, it is not impossible but merely irrational to intend to do what one believes to be impossible. I am not alone in preferring to think of this constraint as a rational requirement rather than a metaphysical requirement. Versions stated in terms of rationality have been explicitly defended by Bratman (1987), Broome (1999, 2013), and Holton (2011), among others.

I won't try to defend this view, except to say that it makes good sense of some problem cases. For example, Wilson and Sperber (1988, 213) imagine someone who shouts 'Start, damn you!' at a car that isn't working. The speaker knows that their utterance won't change the car's intentions: cars don't have minds. So what is going on? We could try to use this case to motivate a radical revision to our theory of speech acts. But there is a much simpler solution, which is to say that the speaker was suffering from a momentary bout of irrationality when they shouted at their car. Presumably, if we asked the speaker about the episode later, they would agree that their shouting was irrational. So I recommend that we say of such cases that the speaker momentarily forms a communicative intention that conflicts with what they believe.¹²

I think that some bald-faced lies are irrational actions in the foregoing sense. They are intended to accomplish something that the speaker believes is impossible. This seems like a good way to describe a child who insists that they didn't hit their sibling immediately after their parents watched them do it, for example. It is also the best way to make sense of Stokke's (2018, 55) example of an obviously drunk person who insists that they haven't been drinking.¹³ However, I don't think it's best to think of all bald-faced lies in this way. In particular, there doesn't seem to be anything irrational about the student's lie to the dean: the student's goal is to avoid punishment, and they correctly believe that they can accomplish this goal by uttering (1).

¹²There is even some empirical evidence for this way of thinking about the case. Cognitive scientists have shown that we tend to anthropomorphize inanimate objects more often when they violate our expectations, such as when they malfunction (Epley, Waytz, and Cacioppo 2007). Mandelbaum and Ripley (2012, 364) argue, for example, that we are much more likely to anthropomorphize a car if it breaks down.

¹³Another possibility about the child and the drunk is that they have momentarily come to believe – presumably also irrationally – that their lie will be convincing.

7. Some bald-faced lies are indirect speech acts

How, then, should we characterize the student's speech act? I will consider two theoretical options, each of which, I think, helps us to understand some bald-faced lies. The first is that some of what we call bald-faced lies are actually indirect speech acts, wherein the speaker makes as if to assert something in order to indirectly accomplish some other conversational goal.

The idea that we sometimes make as if to perform one speech act in order to indirectly perform another should be familiar from Grice's (1975) theory of implicature. Grice distinguishes between two ways of implicating something. First, one can implicate p by *saying* q , in which case one means both p and q . As an example, Grice describes a case in which A says to B, 'I am out of petrol', and B replies, 'There is a garage around the corner', thereby implicating that the garage is open and has gas available (1989, 32). The speaker's implicature is not a replacement for what they appear to say in this case, it is an additional message. Second, one can implicate q by merely *making as if to say* p , in which case one only means q but not p . Grice's examples of this kind involve metaphor, hyperbole, and irony, such as his case in which a speaker describes a man who has recently betrayed them as 'a fine friend', thereby implicating that he is anything but (1989, 34). Here, the speaker does not mean that the man in question is both a fine friend and also a bad friend, but only the latter: the implicature is intended to supplant what the speaker makes as if to say, not supplement it.

The mechanism at work here is straightforward: The speaker makes an utterance that would be uncooperative if taken literally, and this leads the addressee to the conclusion that the speaker intended something other than what their words would normally indicate. If I were to tell you that X is a fine friend in a context where we both know that X has betrayed me, I would violate the maxim of quality, since I would be asserting something that we both take to be false. I anticipate that you will take me to be cooperative, and so assume that I am merely flouting the maxim of quality in order to hint at what I really mean, which is that X is a terrible friend. Here, the cooperative principle together with your extralinguistic evidence about my state of mind defeat the linguistic evidence that I give you, which is exactly my plan.

Most of Grice's examples of implicating by making as if to say, including those involving metaphor, hyperbole, and irony, involve flouting the maxim of quality. They all involve utterances that, if taken literally,

would indicate that the speaker is saying something that neither the speaker nor the addressee believes, and that the speaker therefore obviously has no chance of getting the addressee to believe. I think that a lot of what we normally call 'bald-faced lies' are like this too. By uttering something that is obviously false, and that would be obviously uncooperative if taken literally, the speaker manages to flout the maxim of quality and indirectly communicate something else.

What are these bald-faced liars *really* trying to communicate? I doubt that we can answer this question with any more generality than we can say, in general, what sorts of things are communicated ironically or metaphorically. We can communicate many things indirectly, and if bald-faced lies are mechanisms of indirect communication, we should expect them to be versatile. Nevertheless, we can consider some plausible examples.

Here is one common kind of case: A asks B a question, knowing that the answer is *p*. The reason that they do this is because the truth of *p* should prompt a certain further response by B. In response to A's question, B bald-facedly answers with something other than *p*. B does this as a way of indirectly refusing to respond in the way that A was expecting to follow from *p*.

Ann and Bob have divided up the chores. Bob's job is to mow the lawn at 2 p.m. At 2:05, Bob is still playing a video game. The following exchange unfolds:

(2) Ann: What time is it Bob?

Bob: Oh, it's about 1:45.

Let's suppose that it's completely obvious that it is after 2 p.m., and so Bob can't be (rationally) attempting to deceive Ann. And let's suppose that he's not irrationally attempting to deceive her. In that case, his most plausible real aim is to indirectly inform her that he's not going to start mowing the lawn quite yet (and perhaps to rub it in her face that she can't make him).

We can say something similar about Meibauer's (2014) example of a husband who says to his wife, 'I slept in my office last night', when it is obvious that he was with another woman. Assuming that the husband really does take it to be impossible to deceive his wife, and that he isn't speaking with an irrational intention to deceive her, we are left with the question of what he takes the point of his utterance to be. A plausible answer, at least on some ways of fleshing out the case, is that he wishes to communicate that he feels no need to apologize or admit guilt, that

he has no intention of talking about it further, and that he thinks that his wife has no way of holding him accountable.¹⁴

Finally, I think this is one plausible explanation of what the student is doing when he appears to deny cheating by uttering (1). Assuming that he really is certain that he won't deceive the dean, and that he thinks that the dean knows that he knows about the dean's unofficial policy of only punishing those who confess, the student might utter (1) as a way of indirectly communicating that he is not planning to confess and that the dean should not expect him to.

I hesitate to definitively conclude that this is the best explanation of the student's utterance, both because the case is not very thoroughly described and because I think there is one more possibility worth taking seriously (see Section 8), but I am convinced that many bald-faced lies are best understood this way.

8. Some bald-faced lies are conventional acts

It is a central claim of Stokke's theory that lies are always assertions (Stokke 2018, 31, Chap. 3). I think there are good reasons to disagree. Suppose that we ignore my protests in Section 5 and adhere to ordinary usage about what counts as a lie. Here, I want to argue that when we attend to the theoretically important properties of speech acts, we find that many things that ordinary usage deems to be lies are not assertions. My reason is that assertion is a communicative act but many intuitive examples of lies, including bald-faced lies, are done by means of conventional acts.

This point is similar to one made by Keiser (2016), and I will illustrate it by expanding on an example first used by her.¹⁵ In *The Godfather: Part II*, Frank Pentangeli has turned state's witness against his former boss, Michael Corleone. He has signed a sworn affidavit saying that Corleone is head of a crime syndicate and has ordered Pentangeli to commit murder. There is a good deal of other circumstantial evidence for these

¹⁴Meibauer's own explanation of this case is that the husband is violating the maxim of quality as a way of showing his wife that he is uncooperative, thereby insulting her. I think that this is a plausible account of some bald-faced lies as well and that Meibauer's example is too underdescribed to decide between my account and his.

¹⁵My terminology and some of my theoretical background differ from those of Keiser. Instead of distinguishing between communicative and conventional acts, Keiser distinguishes between 'illocutionary actions' and speech acts that are not genuine illocutionary actions. As she uses this terminology, I believe she cross-cuts the communicative-conventional distinction in some ways. However, Keiser and I agree on enough core cases that her distinction winds up figuring in her argument in much the same way that the communicative-conventional distinction figures in mine.

claims. The contents of Pentangeli's affidavit are, let us suppose, common knowledge among the authorities. Still, the law dictates that Pentangeli must actually testify before a senate hearing before Corleone can be charged with a crime. At the last moment, the Corleones manage to locate Pentangeli's brother, Vincenzo, and bring him from Sicily to America. Seeing his brother sitting with the Corleones in court, Pentangeli becomes ashamed and afraid for his family and decides to change his testimony, which then plays out as follows:

SENATOR KANE: Mr. Pentangeli, were you a member of the Corleone family? Did you serve under caporegime Peter Clemenza, under Vito Corleone, also known as the Godfather?

PENTANGELI: I never...I never knew no Godfather. I have my own family, Senator.

SENATOR KANE: Mr. Pentangeli, you are contradicting a sworn statement you previously made to me and signed. I ask you again, sir, here and now under oath: were you at any time a member of a crime organization headed by Michael Corleone?

PENTANGELI: I don't know nothing about that! Oh...I was in the olive-oil business with his father, but that was a long time ago, that's all.

SENATOR KANE: We have a sworn affidavit! We have it – your sworn affidavit, that you murdered on the orders of Michael Corleone. Do you deny this confession? Do you realize what will happen as a result of your denial?

PENTANGELI: Look, the FBI guys, they promised me a deal. So I made up a lot of stuff about Michael Corleone, cuz that's what they wanted. But it was all lies – uh – everything!...

Pentangeli utters his answers in an unserious tone, making it obvious that he is not testifying in good faith. The senators' reactions make it clear that Pentangeli's utterances have not changed anyone's beliefs about his role in the mafia. Nor do the senators go on to treat Pentangeli's testimony as true for the purposes of the conversation. One senator shouts 'I'm going to get to the bottom of this!', for example. It is implausible that Pentangeli's aim was to make anyone believe or accept any of what he said, and we do no harm to the film's plot by assuming that it was not.

Intuitively, Pentangeli's utterances are lies, and bald-faced lies. Should we conclude that his utterances are also assertions, and that asserting a proposition p therefore needn't be aimed at making anyone believe p , or even at making anyone accept p for the purposes of the conversation? I think that this would be the wrong conclusion to draw. Instead, we should recognize that legal testimony is a fundamentally different kind of speech act than assertion, but nonetheless one that can count as a lie by the standards of ordinary usage.

The very possibility of testifying in court depends on an elaborate system of institutions and procedural conventions that vary by jurisdiction. These institutions and conventions define the conditions under which testimony is possible in the relevant jurisdiction and the effects that it will have there. In some jurisdictions, a person can perform the act of testifying in court only if they have been sworn in as a witness, only if people occupying the right roles (judge, jury, prosecutor, etc.) are present, and only if a host of other background conditions obtain. The aim of sharing information or changing people's beliefs is often incidental to the purpose of legal testimony. We can imagine a witness who has previously asserted all of the contents of their testimony out on the courthouse steps, with all of the same people listening, immediately before testifying. The contents of their testimony could be common knowledge to everyone present. This would in no way diminish the purpose of their testimony. The point of testimony is to put some information on the court record, and one does this not by communicating with anyone but by playing a role in a social ritual – what Austin (1962) called a 'conventional procedure'.

By its nature, legal testimony always invokes a specific set of social institutions and conventions. Assertion does not work this way – not according to the intentionalist account of assertion, and not according to Stokke's own account of assertion either. Assertion needn't be performed within any particular jurisdiction, and its nature and performance conditions do not vary between jurisdictions. Assertion is part of the basic repertoire of human social interaction, not a move in a social ritual whose nature is constituted by localized conventions. Although testifying in court may involve making noises that are similar to those that one would make in performing assertions, they are fundamentally different speech acts.

I think that this point can be extended to the student's lie to the dean. The student, like Pentangeli, is not motivated by the aim of changing anyone's mind. Rather, the student's aim is to do whatever it takes to avoid punishment. The dean has a policy of not punishing anyone who doesn't confess. Knowing this, the student utters the magic words, with little concern for any other effects that their action might have. The student is doing essentially the same kind of thing as Pentangeli, albeit in a less formal or consequential setting.¹⁶

I anticipate the following objection: Although it's true that the dean has a policy of punishing students only if they explicitly confess, this policy is

¹⁶Cf. Keiser (2016, 471–472).

(let us suppose) an informal one, and not part of the school's official rules. This makes the situation different than Pentangeli's senate hearing, where the procedural intricacies are explicit and clear.

But I don't think that this difference matters. It is a well-worn point in the philosophy of law that an institution's rules needn't be written down anywhere in order to be binding. Common-law systems give us one example, as does the unwritten convention at the foundation of the U.S. legal system that the U.S. Constitution is a legally binding document but the Declaration of Independence isn't.¹⁷ Some societies have no formal means of making their rules explicit, but this does not mean that they lack laws, social institutions, or rituals in which conventional speech acts may figure. My argument in this section depends only on treating the dean's policy as a *de facto* institutional fact – one that the student can exploit by uttering the right words in the right context, thereby performing a conventional act of denying guilt. The student's knowledge of this social affordance is what explains their utterance. Like Frank Pentangeli's testimony, the student's speech act is intuitively a lie, and a bald-faced lie; but to infer from this that it is best understood as an assertion in anything like the usual sense is to ignore theoretical distinctions without which we have little hope of understanding the student's actions. To label the student's speech act an assertion is to commit what Austin (1962, 3) calls the 'descriptive fallacy'.

9. Conclusion

I have argued that some of what we normally call 'bald-faced lies' are irrational attempts to deceive and others are speech acts that superficially resemble assertions but are actually of fundamentally different kinds. Bald-faced lies in these categories pose no threat to intentionalism. By contrast, Stokke's own account of bald-faced lies struggles to make sense of elaborated versions of his own examples.

A theme in the foregoing has been that speech acts are too complex and varied to be entrusted to our ordinary ways of talking about them. Even if conceptual analysis were possible (which I doubt!), the recapitulation of our folk categories that would result would ignore deep and important distinctions. What's needed is careful theoretical sensitivity to the ways in which different speech acts fit into our broader aims and activities, and to the psychological and social mechanisms by which they play these

¹⁷These points are originally due to Hart (1994).

roles. This approach is likely to explode folk categories like lying, leaving us with a collection of successor concepts that fit into our different theoretical pursuits in different ways. So be it. This is one kind of philosophical progress.

Disclosure statement

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

ORCID

Daniel W. Harris  <http://orcid.org/0000-0002-4764-3703>

References

- Austin, J. L. 1962. *How to Do Things with Word*. Oxford: The Clarendon Press.
- Bach, K., and R. M. Harnish. 1979. *Linguistic Communication and Speech Act*. Cambridge, MA: MIT Press.
- Barber, A. 2020. "Lying, Misleading, and Dishonesty." *The Journal of Ethic* 24: 141–164.
- Bratman, M. 1987. *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Broome, J. 1999. "Normative Requirements." *Ratio* XII (4): 398–419.
- Broome, J. 2013. *Rationality Through Reasoning*. West Sussex: Wiley-Blackwell.
- Carson, T. L. 2006. "The Definition of Lying." *Nou* 40 (2): 284–306.
- Dynel, M. 2011. "A Web of Deceit: A Neo-Gricean View on Types of Verbal Deception." *International Review of Pragmatic* 3: 139–167.
- Elbourne, P. 2011. *Meaning: A Slim Guide*. Oxford: Oxford University Press.
- Epley, N., A. Waytz, and J. T. Cacioppo. 2007. "On Seeing Human: A Three-Factor Theory of Anthropomorphism." *Psychological Review* 114 (4): 864–886.
- Fodor, J. A. 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Grice, H. P. 1957. "Meaning." *The Philosophical Review* 66 (3): 377–388.
- Grice, H. P. 1968. "Utterer's Meaning, Sentence-Meaning, and Word-meaning." *Foundations of Language* 4 (3): 225–242.
- Grice, H. P. 1969. "Utterer's Meaning and Intention." *The Philosophical Review* 78 (2): 147–177.
- Grice, H. P. 1975. "Logic and Conversation." In *Syntax and Semantics 3: Pragmatics*, edited by P. Cole. New York: Academic Press.
- Grice, P. 1989. *Studies in the Way of Word*. Cambridge, MA: Harvard University Press.
- Harris, D. W. 2014. "Speech Act Theoretic Semantics." PhD Diss., City University of New York Graduate Center.
- Harris, D. W. 2016. "Intentionalism Versus the New Conventionalism." *Croatian Journal of Philosophy* 16 (47): 173–201.

- Harris, D. W. 2017. "A Puzzle About Context and Communicative Acts." *ProtoSociology* 34: 119–143.
- Harris, D. W. 2019a. "Intention and Commitment in Speech Acts." *Theoretical Linguistic* 45 (1–2): 53–67.
- Harris, D. W. 2019b. "Intention Recognition as the Mechanism of Human Communication." In *Sensations, Thoughts, Language: Essays in Honor of Brian Loar*, edited by A. Sullivan. London: Routledge.
- Harris, D. W. 2020. "We Talk to People, Not Contexts." *Philosophical Studie*.
- Hart, H. L. A. 1994. *The Concept of Law*. 2nd ed. Oxford: Oxford University Press.
- Holton, R. 2011. *Willing, Wanting, Waiting*. Oxford: Oxford University Press.
- Keiser, J. 2016. "Bald-Faced Lies: How to Make a Move in a Language Game Without Making a Move in a Conversation." *Philosophical Studie* 173 (2): 461–477.
- Mandelbaum, E., and D. Ripley. 2012. "Explaining the Abstract/Concrete Paradoxes in Moral Psychology: The Nbar Hypothesis." *Review of Philosophy and Psychology* 3 (3): 351–368.
- Meibauer, J. 2014. "Bald-Faced Lies as Acts of Verbal Aggression." *Journal of Language Aggression and Conflict* 2 (1): 127–150.
- Schiffer, S. 1972. *Meaning*. Oxford: Oxford University Press.
- Stalnaker, R. 1978. "Assertion." In *Syntax and Semantics* 9, edited by P. Cole, 315–332. New York: Academic Press.
- Stalnaker, R. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Stalnaker, R. 2002. "Common Ground." *Linguistics and Philosophy* 25: 701–721.
- Stokke, A. 2018. *Lying and Insincerity*. Oxford: Oxford University Press.
- Strawson, P. F. 1964. "Intention and Convention in Speech Acts." *The Philosophical Review* 73 (4): 439–460.
- Trivers, R. 2011. *Deceit and Self-Deception*. London: Allen Lane.
- Wilson, D., and D. Sperber. 1988. "Mood and the Analysis of Non-Declarative Sentences." In *Human Agency: Language, Duty, and Value*, edited by J. Dancy, J. M. E. Moravcsik, and C. C. W. Taylor, Chapter 10, 210–229. Stanford, CA: Stanford University Press.