



Lying, speech acts, and commitment

Neri Marsili¹

Received: 13 April 2020 / Accepted: 21 October 2020 / Published online: 18 December 2020
© The Author(s) 2020

Abstract

Not every speech act can be a lie. A good definition of lying should be able to draw the right distinctions between speech acts (like promises, assertions, and oaths) that can be lies and speech acts (like commands, suggestions, or assumptions) that under no circumstances are lies. This paper shows that no extant account of lying is able to draw the required distinctions. It argues that a definition of lying based on the notion of ‘assertoric commitment’ can succeed where other accounts have failed. Assertoric commitment is analysed in terms of two normative components: ‘accountability’ and ‘discursive responsibility’. The resulting definition of lying draws all the desired distinctions, providing an intensionally adequate analysis of the concept of lying.

Keywords Definition of lying · Speech act theory · Assertion · Commitment · Performative utterances · Deception · Insincerity

1 Introduction

Dishonest communication plays an important role in the spread of misinformation, often with dramatic consequences: recent, blatant examples are the false promises that supported the Brexit campaign (see e.g. Chappell 2016; Watson 2018), and the falsehoods (spread by Twitterbots and fake news websites) that plagued the US presidential elections in 2016 (Silverman 2016; Allcott and Gentzkow 2017) and 2020 (Ferrara et al. 2020). Given the social and moral significance of lying, it is not surprising that disciplines as diverse as sociology, linguistics, and psychology have displayed an increasing interest in its analysis. A fundamental philosophical question that cuts across these disciplines concerns how to define lying.

The author would like to thank Jennifer Saul, Paul Faulkner, Don Fallis, Guido Löhr, Eliot Michaelson, and two anonymous referees for helpful comments.

✉ Neri Marsili
nerimarsili@ub.edu; nerimarsili@gmail.com; neri.marsili@unibo.it

¹ Dipartimento di Filosofia e Comunicazione - FILCOM, Università di Bologna, Via Zamboni 38, 40126 Bologna, Italy

Several authors have attempted to offer an analysis of the concept of lying in terms of necessary and sufficient conditions. A variety of different proposals have emerged in the literature, sparking a lively debate about which definition best captures our intuitions (for an overview, see Mahon 2015). This paper presents a puzzle for existing accounts of lying, showing that they are all unable to track our intuitions about whether a given utterance is a lie, and puts forward a definition of lying that is able to solve it.

With some approximation, extant definitions of lying can be grouped into three families: deceptionist definitions (according to which all lies are intended to deceive) assertion-based definitions (according to which all lies are assertions), and hybrid accounts (which incorporate both requirements). Let us briefly familiarise ourselves with each view.

According to deceptionist definitions (Isenberg 1964; Primoratz 1984; Mahon 2008; Lackey 2013), lying consists in saying (as opposed to implying) what you believe to be false, with the intention of deceiving your audience into believing what you said. More formally:

Deceptionist definitions:

S lies to A iff:

- (a) *S utters a declarative sentence with content p* ¹
- (b) *S believes that $\neg p$*
- (c) *S intends to deceive A about p*

The distinctive feature of deceptionist definitions is the ‘intention to deceive’ requirement (c) (which can be phrased in slightly different ways, see Mahon 2008; Fallis 2018). Beyond the mere intuition that lying is a form of intentional deception, a key theoretical motivation for including this requirement is its ability to differentiate between genuine lies and other believed-false declarative utterances that are not lies, such as ironic, metaphorical, and fictional utterances, which are not meant to deceive the audience about their literal content.

In recent years an impressive case has been mounted against deceptionist accounts (Carson et al. 1982, p. 17; Carson 2006; Sorensen 2007, 2010; Arico and Fallis 2013; Fallis 2015, 2018; Krstić 2018, 2019; Marques 2020), prompting several authors to abandon condition (c). Scholars who reject (c) acknowledge that a definition featuring only (a) and (b) would be too broad, as it would include ironic, metaphorical, and fictional utterances. Typically, their solution is to replace (c) with a condition requiring that the speaker *genuinely asserts* that p . More formally:

¹ Condition (a) can be formulated in slightly different ways: some authors phrase it as “S says that p ” (e.g. Saul 2012; Stokke 2013a), others as “S states that p ” (e.g. Chisholm and Feehan 1977; Mahon 2015). I adopted this formulation because it is neutral about the semantics of performative utterances, a topic discussed at length in the next section (§2.1). Different formulations aside, condition (a) tracks the requirement that a *locutionary* act with content p must be performed, as opposed to the requirement (set by condition (d), cf. p. 3) that a specific *illocutionary* act (i.e. assertion) is performed. My phrasing of (a) is not meant to rule out subsentences (“For you!” indicating a letter) and elliptical signs (nodding in response to a question); I am leaving aside these complications merely for ease of exposition, as it is customary in the literature.

Assertion-based definitions:

S lies to *A* iff:

- (a) *S* utters a declarative sentence with content *p*
- (b) *S* believes that $\neg p$
- (d) In making the utterance, *S* is asserting that *p*

Scholars who endorse assertion-based definitions of lying² tend to agree that a speaker lies iff she *asserts something insincerely*, but disagree on what to count as an assertion for the purpose of defining lying.³ In other words, assertion-based definitions of lying differ depending on how the ‘assertion-condition’ (d) is formulated. **Hybrid accounts** (the third family of definitions) incorporate both condition (c) and condition (d) in their definition of lying.⁴

The next section (§2) introduces a new puzzle for definitions of lying: distinguishing between speech acts that can be lies and speech acts that cannot. It shows that deceptionist definitions are unable to make the right distinctions in this respect. The subsequent sections will review the most prominent assertion-based definitions (Stokke 2013a, b, 2018; Fallis 2012, 2013; Carson 2006), showing that these proposals are either similarly unable to draw the required distinctions (§3.1–3) or vulnerable to further counterexamples (§3.4). Where these accounts have failed, I argue that a definition based on the notion of assertoric commitment can succeed. After introducing a novel account of assertoric commitment (§4), I show that the resulting definition of lying avoids the difficulties affecting other accounts, and provides an adequate analysis of the concept of lying (§5).

2 The puzzle of explicit performatives

One of the main contentions of this paper is that a good definition of lying should be able to draw a distinction between the speech acts that are ‘lie-apt’ and those that are not. I will argue that some *explicit performative sentences* can be used to lie (§2.1), while others can be used to deceive, but not to lie (§2.2).⁵ The importance

² This label was first introduced by Stokke (2013a). Proponents of this view include Carson (2006, 2010); Sorensen (2007, 2010); Fallis (2009, 2012, 2013); Stokke (2013a, 2018).

³ Carson (2006, 2010) and Saul (2012) suggest that a further condition might be required, namely that the asserted proposition be actually false—but neither commits to this further requirement (for compelling empirical reasons not to include this condition, see Wiegmann et al. 2016). Also, different authors take (d) to have different significance. Some (e.g. Chisholm and Feehan 1977, p. 142; Fallis 2009, p. 33; Meibauer 2014) take their proposed phrasing of (d) to be a *definition* of assertion. Others do not wish to “[commit themselves] to a view of the final analysis of the phenomenon of assertion” (Stokke 2013a, b, p. 46, cf. Carson 2006, p. 300).

⁴ The label ‘hybrid’ is mine. Defenders of this view include Simpson (1992); Mannison (1969); Chisholm and Feehan (1977); Kupfer (1982); Newey (1997); Williams (2002); Meibauer (2005, 2014); Faulkner (2007, 2013). Many of these authors are motivated to endorse both (c) and (d) by Gricean considerations about the nature of communicative acts and testimony (cf. fn 10).

⁵ In what follows, my discussion will inevitably be limited to a few examples, since it is practically impossible to discuss every performative verb of the English language. The chosen linguistic sample, however, is significant: my token utterances are representative of classes of speech acts (assertives, commissives, directives) on which we have straightforward intuitions. I will not consider other classes, such as *declarations*

of this becomes apparent once we realise (§2.3–3) that most existing definitions are inaccurate, precisely because they are unable to draw this distinction.

2.1 Lying with explicit performatives

Explicit performative sentences (‘explicit performatives’ for brevity) are declarative sentences of the form “I (hereby) [performative verb] that Φ ”, in which the speaker performs a given illocution (promising, asserting, betting, etc.) by declaring that they are performing that illocution. Utterances (1) to (3) are examples of explicit performatives that can be lies. To simplify the discussion, I have marked the *content* of each speech act (what the speaker is promising, asserting, swearing, etc.) with an asterisk:

- (1) I assert that (1*) I received expressed consent from the patient
 (2) I promise that (2*) I will wear a blue dress at the wedding
 (3) I swear that (3*) I saw the defendant at the crime scene

Intuitively, (1), (2)⁶, and (3) can be lies under the right circumstances—whenever the speaker believes, respectively, that (1*), (2*) or (3*) is false (and aims to convince the interlocutor that these propositions are true). To put the same point differently: the fact that you are explicitly asserting, promising, or swearing that something is the case does not render you immune from the accusation of having lied.

It could be argued, however, that performative utterances can *never* be lies. Since assuming the opposite (i.e. that some performatives *can* be lies) is crucial to the main argument delivered in this paper, I will begin by reconstructing and dismissing the case against performative lies. The reader who already shares the intuition that (1–2–3) are genuine lies can jump to §2.2, where I proceed to expose the rest of my argument.

Let us call the view that performative utterances can never be lies the ‘*No-Performatives View*’. This view maintains that (1–2–3) *cannot* be lies, despite our pre-theoretical, naïve intuitions about them, and can be motivated by a ‘descriptivist’ semantic theory of the content of explicit performatives. A descriptivist semantics is one that identifies the propositional content of our explicit performatives with the full sentences (1, 2, 3), rather than the embedded that-clauses (1*, 2*, 3*).⁷ On this view, if you utter (3), you *assert that you are swearing* that you saw the defendant at the crime scene. If we interpret performatives in this literal way, it becomes apparent that it is virtually impossible to lie by uttering them (cf. Searle 1989, p. 539; Marsili 2016, pp. 275–277).

To appreciate this point, recall that lying requires insincerity: you must believe that the content of your utterance is false (condition (b) in the definitions above). But

Footnote 5 continued

and *expressives*, because I do not take our intuitions about them to be straightforward enough to establish whether a given definition should count them as lies or not.

⁶ For experimental evidence that ordinary speakers overwhelmingly classify insincere promises like (2) as lies, and a more general defence of the view that you can lie by promising, see Marsili (2016). Relatedly, authors like Ross (1930), Fried (1978) and Carson (2006, 2010) take all lying to involve the breach of an implicit promise to tell the truth; on this view, “every lie is a broken promise” (Fried 1978, p. 67).

⁷ Descriptivism is advocated by Hedenius (1963); Lewis (1970); Bach (1975); Ginet (1979); Bach and Harnish (1979).

whenever you proffer (3), you know that it is true *that you are swearing* that you saw the defendant at the crime scene (i.e. that (3) is true), because your saying so *amounts to* swearing it. Therefore, whenever you say (3) you know that (3) is true. If descriptivism is true, and the content of (3) is just (3), it follows that whenever you utter (3) you are sincere. The same diagnosis applies to any other explicit performative utterance, including (1) and (2). On a descriptivist reading, performative utterances can never be lies.⁸

It is far from obvious that descriptivism is an adequate account of performative utterances; as a matter of fact, this view is subject to a number of compelling objections (see e.g. Harris 1978; Searle 1989; Reimer 1995; Jary 2007). If descriptivism is an inadequate account of performative utterances, then there is no strong reason to accept the No-Performatives View, nor its counterintuitive consequence that (1–3) cannot be lies. But even if we leave aside the shortcomings of descriptivism, there are compelling reasons to reject the No-Performatives View: its predictions are hard to square with our most basic intuitions about lying, with our moral judgements, and with our legal practices.

To illustrate, consider the following. Every existing definition of lying converges (and rightly so) on the prediction that, uttered alone, *the starred statements* (1*), (2*), and (3*) can be lies (as long as they are uttered insincerely). This is intuitive, but it exposes some counterintuitive implications of the No-Performatives View. A speaker who disbelieves (3*) lies if she *plainly asserts* that she saw the defendant at the crime scene with (3*); but if the same speaker chooses instead to *swear* that she saw the defendant (uttering (3) instead) she is *sincere* and is *telling the truth* according to the No-Performative View. While there may be a trivial, ‘technical’ sense in which these remarks are correct (i.e. a descriptive, overly literal interpretation of what the speaker is saying), these assessments clearly do not reflect our real-world communicative practices. Clearly, choosing (3) over (3*) in court will not render you immune from a charge of perjury. By swearing, you are assuming *more* responsibility for what you say than by plainly making the same claim. Rather than freeing you from the accusation of having lied, choosing (3) over (3*) renders you liable to stronger criticisms if it turns out that (3*) is false. If lying is a concept designed to track a distinctively severe form of communicative dishonesty (Adler 1997; Williams 2002, p. 197; Krauss 2017), then it is just not clear how we can plausibly maintain that the speaker of (3*) is lying and the speaker of (3), who undertakes even *more* responsibility for the same claim, is not.

Similar considerations apply to promises. Both by *promising* that you will wear a blue dress at the wedding (2) and by merely announcing that you will do it (2*), you create an expectation that you will show up at the wedding with a blue dress. The only difference is that when you promise you take on a *stronger* and more explicit responsibility to make it happen. Oddly, the No-Performatives View predicts that only when you assume *less* responsibility you are lying. *Mutatis mutandis*, the same point

⁸ According to descriptivism, performative utterances can at most be ‘misleading’. Descriptivists will concede that with (3) the speaker can perform an *indirect* speech act with content (3*) (Bach and Harnish 1979, p. 208). On this view, (3) can be used to *imply* that the speaker saw the defendant at the crime scene, but not to *directly* claim it – so that (3) is at most deceptive or misleading. I discuss at length the implications of descriptivism for the lying/misleading distinction in Marsili (2016, pp. 275–278). For more on the distinction and its importance, see e.g. Adler (1997), Saul (2012), Stokke (2013b), Berstler (2019).

applies to the difference between simply stating that you have expressed consent from a patient (1*) and explicitly asserting it (1).

These counterintuitive predictions extend to many other performatives that are barely distinguishable from direct assertions: *warning*, *admitting*, *insisting*, *agreeing*, *denying*, *guaranteeing*, *assuring*, etc. For example, the No-Performative View predicts that under no circumstances (1a), (1b), and (1c) can be lies. And yet, these utterances are not significantly (practically, legally, morally, etc.) different from the plain assertion (1*):

- (1a) I warn you that (1*) I received expressed consent from the patient
 (1b) I admit that (1*) I received expressed consent from the patient
 (1c) I guarantee that (1*) I received expressed consent from the patient

Recapitulating, there are strong motivations to reject the No-Performatives View: it clashes with our pre-theoretical intuitions about performative utterances, and its predictions are difficult to reconcile with our moral judgments, our legal practices, and with our reactive attitudes to performative utterances in real-life situations. On the other hand, the positive case supporting the No-Performatives View is weak: the only theoretical motivation to accept it is that it is entailed by descriptivism, a view that is not exempt from objections. In what follows, I will therefore proceed on the assumption that the No-Performatives View is incorrect, and that a good definition of lying should accommodate the intuition that (1–2–3), (1a–1b–1c), and cognate utterances can be lies.

2.2 Explicit performatives that cannot be lies

Although *some* explicit performative utterances can be lies under the right conditions, not *all* performative utterances can be. Consider the following examples:

- (4) I conjecture that (4*) the blood on the blade is Reza's
 (5) I advise that (5*) you try that quiche
 (6) I command that (6*) you steal that chicken

In the previous section, we saw that (1, 2, 3) are lies whenever the speaker believes their respective content [(1*), (2*), (3*)] to be false. By contrast, it is not clear under which conditions (4), (5) or (6) could be lies. While they can surely be *deceptive* or *misleading*, it is not possible, strictly speaking, to *lie* by uttering them. For example, if I conjecture that the blood on the blade is Reza's even though I know it is not (I disbelieve (4*)), it would be appropriate to criticise me for having been deceptive, but not for having lied, since I have merely *conjectured* that (4*) is true, and conjecturing something is not yet claiming that it is true. The advice (5) can be misleading in several ways: it may falsely imply that the quiche is delicious, or falsely suggest that the hearer can (and will) eat the quiche. Similarly, the command (6) may falsely imply that it is possible to steal the chicken (even though it is well guarded), or that the speaker has the authority to command its theft (even though she is merely impersonating someone with such authority). But even though (4), (5) and (6) can be deceptive in several different

ways, it seems that under no circumstances could they be appropriately classified as lies.⁹

It should now be clear that some speech acts can be lies, while some others cannot. This is important, because it has crucial implications for theorising about lying. It establishes two key desiderata for a definition of lying to which theorists have paid little attention so far: a good definition should be able to acknowledge (a) that some performative utterances (explicit assertions, promises, sworn statements, warnings, etc.) can be lies, but also (b) that some other performative utterances (like conjectures, advices, and orders) cannot be lies. To understand the importance of these considerations for our theorising about lying, let us consider their implications for what is perhaps the most influential philosophical view about lying: deceptionism.

2.3 The puzzle applied: deceptionism

Are deceptionist accounts able to draw all the desired distinctions? The answer can only be negative, since all deceptionist definitions classify (4, 5, 6) as lies. These sentences are all in the declarative mood, so that they all meet condition (a). Furthermore, we have just seen that it is possible to imagine scenarios in which the speaker believes that the content of any of these sentences is false, and intends to make the audience believe that it is true, so that conditions (b) and (c) can also be met. Against the desiderata, deceptionist definitions classify deceptive uses of (4, 5, 6) as lies. If this is correct, deceptionist definitions are not intensionally accurate.

Appealing to a descriptivist interpretation will not help the deceptionist, for reasons that were given above (§2.1). Admittedly, a descriptivist reading of deceptionist definitions would exclude (4, 5, 6), because (so interpreted) these sentences are true in virtue of the speaker's saying so. But a descriptivist reading would also rule out *every other performative lie*. This is not a good trade-off for deceptionism, because it prevents it from counting explicit assertions, warnings, sworn statements, and other lie-apt speech acts as lies. Whichever semantics of performative utterances we favour,¹⁰ deceptionist definitions will be able to accommodate one of the required sets of intuitions, but not both.

⁹ In a recent paper, Viebahn (2019) has argued that one can lie by presuppositions. If this is right, insofar as any speech act can trigger a presupposition, any speech act can be used for lying: e.g. (5) could be a lie if the speaker knows that there is no quiche that the hearer can try. Viebahn's view can be disputed, but I do not wish to enter the debate on presuppositional lying here. If one is moved by Viebahn's arguments, my claim should be read as follows: that (4), (5), (6) cannot be used to lie *about their content* (4*), (5*) and (6*), and that a good definition of lying should predict so. For the sake of simplicity, I will assume this conditional qualification to be implicit throughout the paper.

¹⁰ Another 'semantic' strategy would be to argue that (4, 5, 6) cannot meet condition (b) because they do not possess truth-evaluable content. However, parallel problems apply. While some linguists have in fact challenged (in one way or another) the idea that every speech act possesses truth-evaluable content, what is needed here is a theory that both excludes (4, 5, 6) and includes (1, 2, 3). Proving that such a theory of content *cannot* be developed goes beyond the ambitions of this paper, but there are at least two reasons to suspect that this solution is not viable. First, despite the vast literature on explicit performatives, no theory that draws these distinctions has been defended before (see Recanati 2013 for an overview). Second, a plausible theory should employ either *syntactic features* or *direction of fit* to set apart performative sentences that have truth-evaluable content from those who don't, but neither of these features can be used to set apart the two groups of sentences under consideration (1, 2, 3 and 4, 5, 6) (see fn 20 for an example).

We will see that the challenge faced by deceptionist accounts applies to every other definition of lying. A good definition should be able to classify explicit performatives like (1, 2, 3) as lies, but also exclude performatives like (4, 5, 6), which under no circumstances can be correctly classified as lies. In the next sections, I will show that also assertion-based definitions are unable to meet these desiderata. While I will not discuss hybrid accounts, it should be noted that for any given assertion-based definition that is unable to rule out (4, 5, 6), so is the hybrid account built on that definition (because these accounts would only differ in their endorsement of the intention to deceive condition (c) which, we have seen, is unable to discard these cases). In other words: whenever an assertion-based account is proved to be too narrow, so is the hybrid account that it is built on it.¹¹

3 Testing extant definitions

Since assertion-based definitions differ primarily in how the ‘assertion condition’ (d) is fleshed out, in what follows I will only discuss how this condition is formulated by different proponents of assertion-based definitions, keeping the rest (condition (a) and (b)) fixed.¹² I will first discuss Fallis’ work.

3.1 Intentionally communicating something false

In a series of recent papers, (2009, 2012, 2013) Fallis delineates a number of ways to develop an assertion-based definition of lying. In Fallis (2012),¹³ lying is defined as the intentional, explicit communication of something that the speaker believes to be false. The following assertion-condition (d) is adopted¹⁴:

(ACF1) *S intends to communicate that p*

Fallis acknowledges that the notion of ‘communication’ plays a key role in this proposal: “what counts as communication makes a difference for what counts as a lie [according to ACF1]”. Nonetheless, he controversially adds that no particular notion of communication is needed for his account to work: “for purposes of this paper, it will

¹¹ Matters are slightly more complex for ‘Gricean’ hybrid views, according to which a speaker S asserts that *p* iff S intends her audience A to accept that *p* at least partly on the basis of the fact that A recognises S’s intention to make A accept that *p* (endorsed, slightly amended, by Meibauer 2005, 2014; Faulkner 2007, 2013). Here the deception condition (c) and the assertion condition (d) impose virtually the same constraint. I will not discuss these views here because they have already been criticised at length elsewhere (e.g. Fallis 2010, 2018), but it is worth noting that (beyond known counterexamples) they will have trouble accommodating the examples discussed in §3.2 (bets, conjectures and suppositions) and in §3.4 (provisos-lies).

¹² The recurring acronym “AC” will be meant to remind the reader that, for each view, I am reporting the ‘assertion condition’ (d) rather than the whole definition, which includes also (a) and (b).

¹³ I will not discuss Fallis’ (2009) proposal: it has been shown to be incorrect, because it counts most ironical utterances as lies (Stokke 2013a, b), and was rejected by Fallis himself (2012).

¹⁴ Fallis (2012) never presents conditions (a)–(b)–(d) separately, but rather packs them together in a single sentence. Nonetheless, he is committed to ACF1 being a necessary condition for lying in addition to (a) and (b). For ease of exposition, I will ignore this complication.

not be necessary to settle on one specific account of communication” (2012, p. 572). It is hard to agree with this claim. Absent a clear criterion to determine whether an utterance is ‘intended to be communicated’, ACF1 is underdetermined: it does not provide a clear and univocal criterion to determine whether a given utterance is a lie—in other words, it fails to *define* what lying is (cf. Keiser 2016, p. 476fn).

It could be argued, however, that failing to specify what is meant by ‘communication’ need not lead to this sort of indeterminacy. Fallis might not have specified what he means by ‘communication’ simply because he has in mind a rather ordinary notion.¹⁵ Accordingly, we may assume that ACF1 is satisfied iff an ideal English speaker would agree that the speaker intended to communicate that *p*, in the ordinary sense of the term.

However, as the predictions of ACF1 become clearer, its structural problems become clearer too. Specifically, ACF1 is unable to rule out many performative utterances that are not lie-apt. This is because virtually any speech act (and not only the ones that are lie-apt) can be accompanied by the intention to communicate that their content is true. To illustrate, consider (4) once again:

(4) I conjecture that (4*) the blood on the blade is Reza’s

Imagine a speaker (call her Luisa) who utters (4) with the intention to insinuate that the blood on the blade is indeed Reza’s. There is clearly a sense in which Luisa *intends to communicate* that the blood is Reza’s: if she believes that (4*) is false, Fallis’s definition would classify her conjecture as a lie.¹⁶ But this verdict is incorrect. If Luisa were to be accused of lying, it would be perfectly appropriate for her to object that she has merely conjectured, but never affirmed, that the blood was Reza’s. Even in a court of law, (4) could not plausibly be regarded as a lie, precisely because it is flagged as a mere conjecture (cf. S. Green 2001, pp. 176–82; Saul 2012, pp. 95–97). This is not to deny that, by uttering (4) maliciously, Luisa can *insinuate* or *imply* that the blood was Reza’s: this is exactly what happens when Luisa intends to communicate that (4*) is true, satisfying ACF1. The point here is rather that insinuating or implying something falls short of lying—it falls on the ‘misleading’ side of the lying/misleading distinction. This objection to ACF1 is not limited to conjectures: similar considerations would apply if Luisa had suggested, hypothesised, bet or guessed that (4*) is the case.

It is also possible to imagine circumstances in which ACF1 would classify directive speech acts as lies. Imagine a conversation between two individuals, A and B; A has complete authority over B. A says “What shall I do next?”; B replies with (6):

¹⁵ Although it would be a natural move, note that we cannot interpret ACF1 as appealing to *Gricean* communicative intentions. Gricean communication requires (broadly) that the speaker intends to make the audience believe what they say; pairing this requirement with the insincerity condition (b) amounts to reintroducing an intention to deceive condition (c). Since Fallis’ project is to provide an alternative to deceptionism, this interpretation is not available. Furthermore, since Gricean definitions have been defended elsewhere (see fn 12), interpreted in this way ACF1 would no longer represent an *original* proposal. To be sure: another, more modestly ‘Gricean’ reading (according to which ‘communicating’ means ‘expressing a belief’) could work for ACF1; I discuss it in §2.3.

¹⁶ Remember that what is at stake here is whether the speaker would be lying about (4*), not about (4). As we saw in (§2.1), accepting the opposite view, according to which the proposition to be evaluated is rather (4) (descriptivism), would force us to conclude that *no* performative utterance can be a lie. This is incorrect: a good definition must acknowledge that (among others) explicit assertions, sworn statements and promises can be lies.

(6) I command that (6*) you [will] steal that chicken

In this context, surely B's communicative intention is to issue a command—to tell A what she must do. But given that A has asked what to do next, in uttering (6) B may conceivably intend not only to issue a command, but also to convey an answer to A's question: to inform A of what she is doing next, namely (6*). If we postulate that B believes (6*) to be false (for instance, if B knows that the envisaged poultry theft is impossible), ACF1 would incorrectly classify this case as a lie.

Let me emphasise that the claim here is not that (4) or (6) conventionally or typically communicate contents like (4*) and (6*), but rather that there can be contexts in which it would not be blatantly irrational for the speaker to have the intention to communicate such propositions. Since both (4) and (6) can clearly meet this latter, weaker requirement, there are circumstances in which ACF1 incorrectly classifies them as lies, against our desiderata.

3.2 Representing yourself as believing

In a more recent paper, Fallis develops a different proposal; possibly, one that could be read as a refinement of ACF1. Drawing on some observations by Davidson (1985, 2001), Fallis (2013) identifies the following assertion-condition for defining lying:

(ACF2) *The speaker intends to represent herself (to her audience) as believing that p is true*

To 'represent yourself as believing something' is to present yourself as having a particular property, namely the property of believing a proposition. Fallis correctly points out that we have an intuitive grasp of the notion of 'representing yourself as having a certain property', and this becomes evident when we think about familiar cases: when you sign a cheque, you *represent yourself* as having enough money in the bank to honour the cheque (Black 1952, p. 31); by wearing a cross necklace, you *represent yourself* as being Christian, and so forth.

Even though ACF2 offers a more determinate criterion than ACF1, it is similarly unable to draw the right distinctions concerning which speech acts can be lies. This is evident when we consider conjectures. By uttering (4), Luisa can intend to represent herself as believing its literal content (4*) (that the blood is Reza's): if she believes that the blood is someone else's, ACF2 incorrectly predicts that her conjecture is a lie. To be sure: I am not claiming that whoever says (4) will *ipso facto* represent themselves as believing (4*), which is blatantly incorrect. I am merely claiming that there *can be* circumstances in which a speaker utters (4) with the *intention*¹⁷ to represent themselves as believing that (4*), which is all that ACF2 requires.

Furthermore, as for ACF1, the problem is not limited to conjectures: there are several speech acts (like guessing, supposing, hypothesising) that one can use to represent oneself as believing something (Searle 1976, p. 10), but not to lie. In sum, both ACF1 and ACF2 fail to draw the right distinctions between explicit performatives that can and cannot be lies. If lying is to be defined in terms of an insincere assertion, we need to identify an alternative account that avoids their difficulties.

¹⁷ Note, further, that whether this intention is successful is irrelevant to whether ACF2 is satisfied.

3.3 Proposing to add to the official common ground

Stokke's (2013a, b, 2018) assertion-based definition is based on the accounts of *assertion* and *conversational common ground* developed by Stalnaker (1978, 2002). According to Stalnaker (2002, p. 716), "it is common ground that p in a conversation if all members *accept* (for the purpose of the conversation) that p , and all *believe* that all accept that p , and all *believe* that all *believe* that all accept that p , etc.". Assertion is understood by Stokke as a proposal to add a proposition (specifically, the content of the sentence one utters) to the 'official' common ground:

(ACS) S proposes that p become part of the official common ground

The notion of 'official' common ground is meant to exclude speech acts that are not assertions. Consider the following cases:

- (7) Pushkin's beard never grew
- (8) Assume that (8*) I can lift weights with my mind [...]
- (9) Let us suppose that (9*) there is a demon that systematically deceives us

Although (8) and (9) are invitations to add a proposition ((8*) and (9*) respectively) to the common ground (what is accepted as true *for the purpose of the conversation*), they are not assertions. The distinction between *official* and *unofficial* common grounds (Stokke 2013a, b, 2018) handles these cases effectively. *Unofficial* common grounds are 'provisional' common grounds that open up in order to store information that is used for some temporary conversational purpose; by contrast, *official* ones are, so to say, 'permanent' common grounds. ACS only captures proposals to add a proposition to the *official*, permanent common ground. This means that it correctly rules in assertions like (7) (since (7*) is meant to be stored in the official common ground) and correctly discards assumptions like (8) and suppositions like (9) (since (8*) and (9*) are stored in the *unofficial*, temporary common ground).

Although this distinction helps with assumptions and hypotheses, it seems unable to draw all the desired distinctions. Consider commands:

- (6) I command that (6*) you steal that chicken

Here the distinction between official and unofficial common grounds is less helpful, because it is not clear how it applies to (6): without a systematic account of what qualifies as a contribution to the *official* common ground, the predictions of ACS in this sort of case are unclear. And if we attempt to extrapolate from ACS a criterion for dealing with these examples, it emerges that ACS struggles to make the required distinctions.

There are various ways to extrapolate a criterion from ACS. For the purpose of this paper, I will limit my discussion to a criterion that is explicitly defended by Stokke in his book (2018) (I pursue a more thorough analysis in Marsili 2020b). Here he suggests that we can test whether a proposition has been added to the common ground (and therefore captured by ACS) by attending to whether it can be felicitously presupposed.¹⁸ To verify whether uttering (6) adds (6*) to the common ground, for

¹⁸ A felicitous presupposition is one that does not elicit "the kinds of repair strategy that are typically prompted by unfamiliar presuppositions". Stokke (2018, p. 66), identifies two repair strategies: accom-

instance, one needs to verify whether (6*) can be felicitously presupposed after the speaker has uttered (6). To test this, imagine a conversation between three individuals: Adriano, Beppe, and Carmen. Adriano orders Beppe to steal a chicken by uttering (6), and then Carmen utters (10), which presupposes (6*):

(10) When you steal the chicken, you can use my cutters

For ACS to pass the test, there must be *no circumstances* in which (6*) can be felicitously presupposed as a result of Adriano’s command, because the possibility of felicitous presupposition would indicate that (6*) can enter the common ground as a result of Adriano’s utterance. Clearly, such circumstances are possible: whenever Beppe and Carmen take Adriano to have the authority to command (6), it is possible for Carmen to presuppose (6*) (that Beppe will steal the chicken) felicitously via (10).¹⁹ This is a problem for ACS, because it means that Stokke’s assertion-based definition counts (6) as a lie whenever Adriano successfully commands (6) and believes (6*) to be false. Perhaps there is a way to revise ACS so that it avoids these predictions. Absent major revisions, however, Stokke’s current proposal is unable to acknowledge that commands cannot be lies.²⁰ For a definition that draws the right kinds of distinctions, it is better to look elsewhere.

3.4 Warranting as true

Carson (1988, 2006, 2010, followed by Saul 2012) takes a different approach: he defines a lie as an insincere statement that you intend to warrant as true. In other words, he adopts the following assertion-condition:

(ACC) *S intends to warrant the truth of p*

Carson defines ‘warrant’ as follows: “if one warrants the truth of a statement, then one promises or guarantees, either explicitly or implicitly, that what one says is true” (2006, p. 294). According to this view, every time a speaker asserts something, they also implicitly promise that what they say is true (cf. Hawley 2019).

As I will argue in the next section, drawing the right distinctions between speech acts that can and cannot be considered lies requires adopting a view along these lines—one

Footnote 18 continued

modation (as defined by Lewis 1979), and ‘questions and rejections’—that is, (appropriate) replies of the form: “What are you talking about?”; “What makes you think *p*?” or “I never said *p*”.

¹⁹ To be sure, further conditions have to obtain for (6*) to be felicitously presupposed; for instance, it should be common knowledge that stealing the chicken is physically possible. Listing them would lead us astray and is unnecessary. As long as it is *possible* for these further conditions to obtain, the point stands: there are situations in which (6*) can be felicitously presupposed.

²⁰ A referee points out that, since the embedded that-clause (6*) could be rewritten as an infinitive to-clause (I command you *to steal that chicken*), it could be argued that (6) has no truth-evaluable content: “to steal that chicken” is not truth-apt, and therefore cannot be believed to be false. If this is right, (6) is ruled out by *every* definition. I offer a response to this sort of worries in Marsili (2020a). Simply put, as anticipated in footnote 9, this manoeuvre would prove too much: also “I promise/swear/guarantee THAT *f*” can be translated into “I promise/swear/guarantee TO *f*”, but we want to be able to count these utterances as lies. Appealing to accounts à la Portner (2004), which differentiate between the speech acts that update the *common ground* and those that update *to-do-lists* (cf. Roberts 2012), will not help for similar reasons: both promises and commands, on this view, update to-do-lists rather than the common ground.

that links the act of asserting to the acceptance of a distinctive kind of responsibility. Nonetheless, ACC is known to be vulnerable to counterexamples, such as *proviso-lies* (Fallis 2009; Arico and Fallis 2013): lies in which the speaker makes it explicit that they are not promising that what they say is true. Here is a (slightly revised) example from Arico and Fallis (2013):

Last night, after a particularly wild party, Chris found her swimming trophy broken. Today Chris is trying to figure out who broke her trophy. Chris says to Jamie, “So, somebody was in my room last night and broke my trophy. Did you see anything?”. Jamie clearly remembers that she was the one who broke Chris’s trophy. Since everyone knows that Mel is always breaking stuff, Jamie responds to Chris:

- (11) Yeah, um, Mel broke your trophy.
 (11’) But I was kinda drunk, and there were lots of people in there, so don’t take my word for it.

In this example, Jamie’s statement (11) is followed by a ‘proviso’, (11’). The proviso is meant to rectify the previous statement, and to clarify that Jamie does not intend to warrant that (11) is true. As a result, Jamie does not warrant that (11) is true, and Carson’s assertion-condition ACC is not met. Nevertheless, Jamie is clearly lying: this scenario is a counterexample to Carson’s definition.

Carson has since replied that, given that “warranting comes in degrees of strength, a moderately strong assurance of truth is all that is required for lying” (2010, pp. 36–39): the proviso (11’) reduces the assurance of truth that comes with (11), but does not eliminate it. If this is right, (11–11’) does satisfy ACC. However, the problem with this reply is that it is inconsistent with Carson’s account of warrant (Fallis 2013, pp. 347–348). Warrant is analysed as an implicit promise, and promises cannot be mitigated or downgraded. There is no sense in which they can give a “moderately strong” assurance of truth: either they guarantee that the speaker will do something, or they do not. To see this, consider the difference between adding a proviso to an assertion and adding a proviso to a promise:

- (12a) I will wake up at 7AM tomorrow, but you know that I am really unreliable in the morning, so don’t take my word for it
 (12b) # I promise that I will wake up at 7AM tomorrow, but you know that I am really unreliable in the morning, so don’t take my word for it

While (12a) is a mitigated assertion, (12b) is not a mitigated promise: it is not a promise at all. More generally, it seems that promising that p requires an outright (as opposed to “moderately strong”) assurance that p is true.²¹ Pace to Carson, ACC fails to capture *proviso-lies*.

²¹ Here’s a more precise way to put the same point: the *force* of promises cannot be mitigated. Content-mitigation (‘bushes’, in Caffi’s 1999 terminology), by contrast, is possible in promises: the content of “I promise that [I will p]” can be mitigated into “I promise that [if q , I will p]” and (for some but not all ps) [I promise that I will p a little]; cf. Holton (2008). But the possibility of content-mitigation is irrelevant to our discussion: proviso-lies are puzzling precisely because they involve the mitigation of the *force* of the utterance, not its *content*.

These difficulties could be resolved by amending the notion of warrant in a way that avoids the parallel with promises. But it should be clarified from the outset that avoiding the parallel with promises would represent more than an amendment of ACC, because Carson's original contribution to the literature resides precisely in having constructed an analogy between the breach of trust involved in unfulfilled promises and the one involved in lying (elaborating on Ross 1930; Fried 1978). Without such an analogy, ACC would no longer draw the moral parallelism that motivates Carson's overarching philosophical project. In the next section, I will present an alternative way to formulate the assertion condition, which also links assertion to a distinctive kind of responsibility, while avoiding the problematic analogy with promises.²²

4 Assertoric commitment

Before Carson, several authors have argued that asserting involves accepting some kind of responsibility for the truth of a proposition (Peirce CP 2.315, 5.29–31, 543–547, MS 280.25–26, 517.42–44, 36.104–5; Searle 1969, 1975; Brandom 1983, 1994; Searle and Vanderveken 1985; Green 1999, 2000, 2007, 2017; Alston 2000; MacFarlane 2003, 2005, 2011, Rescorla 2009a, Krifka 2014; Tanesini 2016, 2019). I have elsewhere developed (Marsili 2020b) an account of assertion in terms of commitment that falls within this tradition. Simply put, my proposal is to define assertion in terms of the acquisition of this specific kind of commitment, and lying as an insincere assertion:

Definition of Assertion

A speaker S asserts that p iff:

- (a) *S utters a sentence with content p*
- (b) *S thereby commits herself to p being the case*

Definition of Lying

S lies iff S asserts that p insincerely

Some preliminary qualifications are needed. The first is that all conditions are taken to be satisfied *intentionally* by the speaker. This is common in speech act theoretic analyses (Searle 1969; Alston 2000; but cf. Alston 2000, pp. 137–141), and it is especially uncontroversial for defining lying, as virtually every author agrees that there can be no such thing as *unintentional* lying.²³ The second is that the notion of

²² To be sure: accounts in terms of commitment like the one that I am about propose are in a very important sense in agreement with Carson's view. Crucially, they share the idea that lying requires the assumption of a distinctive kind of responsibility. But it is equally important that they take a different stance on *which* kind of responsibility is involved. Note, further, that it would be incorrect to regard commitment-based proposals as mere *refinements* of Carson's view: commitment-based analyses of assertion represent a rich, independent tradition, whose roots go back Peirce's writings, penned at the beginning of the XXth century, long before Carson proposed his alternative view in terms of warrant and promises.

²³ This requirement has the advantage of ruling out cases of misspeaking (Sorensen 2011) and may help to deal with some other puzzling cases (cf. Pepp 2018). Note that if philosophers are wrong, and there can be as unintentional lying, it does not follow that my definition is wrong: it just follows that some lies and assertions fall out of my envisaged *explanandum*. For theoretical and empirical support for the claim that unintentional lies are not lies, cf. discussion of the confused politician example in Carson (2006, p. 296) and Arico and Fallis (2013).

‘insincerity’ at play in the definition of lying is meant to be the one I advocated for in earlier work (Marsili 2014; 2018a, b, 2019): in standard cases,²⁴ I take a speaker to be insincere iff they take themselves to believe that what they are saying is more likely to be false than true.²⁵

Condition (b) does the lion’s share in the definition, and calls for some substantive elaboration. The notion of commitment is meant to capture the normative *consequences* of asserting something: it refers to a change in the speaker’s normative status that happens *in virtue of* the speaker’s act of asserting. While it has been pointed out in previous work that the notion of commitment could be helpfully put to work to define lying (Marsili 2014, pp. 165–170, 2018a, b, pp. 178–179; Leland 2015; Viebahn 2019), I am not aware of any attempt to provide a systematic proposal in this sense. Building on previous work on assertion, I will try to fill this gap by providing a fine-grained characterisation of what assertoric commitment is, and then proceed to show how this account of commitment can be put to work to draw the right distinctions about lying.

I take assertoric commitment to involve two distinct normative dimensions. The first dimension is what I call ‘accountability’. In making an assertion, the speaker becomes reproachable if the proposition turns out to be false (a point also highlighted in Carson’s analysis). An early formulation of this idea is found in Pierce: “an act of assertion [...] renders [the speaker] liable to the penalties of the social law (or, at any rate, those of the moral law) in case [the asserted proposition] should not be true, unless he has a definite and sufficient excuse” (CP 2.315). Alston (2000, p. 55) offers a more accurate definition of this distinctive kind of responsibility: a speaker accepts responsibility for p [being the case] iff the speaker “knowingly takes on the liability to (lay herself open to) blame (censure, reproach, being taken to task, being called to account), in case of *not-p*”.²⁶ Arguably, accountability plays an important role in

²⁴ By ‘non-standard’ cases I mean promises like (2), and more generally assertoric speech acts about one’s future actions. In Marsili (2016) I argued (on both theoretical and empirical grounds) that a promisor can be insincere (and lie) if she intends not to fulfil her promise, even if she believes that she will end up fulfilling it against her will (for instance: S promises *not* to f , intends to f at all costs, but believes that she will almost surely fail to f). We need not dwell on these complications here, but the interested reader can find a definition of insincerity that makes justice to both standard and non-standard cases in Marsili (2016, 2017, pp. 148–151).

²⁵ A final and perhaps less urgent qualification is that in this paper I will leave aside the issue of whether (a) needs to be expanded. While the formulation that I adopt is quite standard, it rules out presuppositional lies (Viebahn 2019) and non-literal lies (Viebahn 2017), and it may rule out non-declarative lies (Viebahn et al. 2018), depending on how the notion of ‘content’ is construed. If one is moved by some (or all) the examples presented by Viebahn, condition (a) can be expanded as required. For some further qualifications about (a), see my footnote 1.

²⁶ Alston reviews different accounts of taking responsibility for the truth of a proposition (in his terminology, “R’ing”), eventually landing on a different view that, unlike the one quoted in the main text, entails that it is only permissible to assert p if p is true (cf. Alston 2000, pp. 54–64). This requirement, also endorsed by “truth-norms” of assertion (Weiner 2005; Whiting 2012) and, indirectly, by “knowledge-norms” of assertion (Williamson 1996), is one that my notion of ‘accountability’ carefully avoids (for reasons discussed in Marsili 2018a). Accountability, as I define it here, only has to do with *downstream* normativity (the normative *effects* of asserting p), which is to be distinguished from (the related, but distinct notion of) *upstream* normativity (whether you are entitled to assert that p —*i.e.* the kind of normativity invoked by ‘norms of assertion’). For more on the irreducibility of these notions to one another, cf. Rescorla (2009a) and MacFarlane (2011).

motivating communicators not to make false claims, ensuring that assertion maintains its role as a valuable tool for sharing and acquiring information (cf. Green 2007, 2009).

In what follows, I will use the term ‘accountability’ to refer, more specifically, to the speaker’s *prima facie*²⁷ liability to be criticised if what they said turns out to be false. To verify if a given speaker is accountable for the propositional content of a given utterance, we need to ask ourselves: if that proposition turns out to be false, would the speaker be *prima facie* criticisable for the falsity of what they have said?

However, the deontic effects of assertions are not exhausted by the speaker’s liability to sanctions. By making an assertion, a speaker also becomes committed to *act* in certain ways, if the relevant conditions arise. More specifically, asserting something commits the speaker to make certain conversational steps, such as making statements that do not contradict their previous ones, or justifying their claims with adequate evidence, when they are challenged to do so (cf. Brandom 1983, 1994, pp. 172–175, MacFarlane 2003, 2005b, pp. 227–229, 2011).

Let us call this second normative component *discursive responsibility*, since it has to do with the conversational moves that a speaker is expected to make in the context of a rational discourse. Discursive responsibility has been modelled in different ways and within different theoretical frameworks (Toulmin 1958; Hamblin 1970a, b, chap. 8; Brandom 1983, 1994, pp. 172–175; MacFarlane 2003, 2005b, pp. 227–229, 2011). Within this literature, authors tend to agree that you are responsible to defend your claims (e.g. by providing evidence in their support) if *appropriately* challenged (or else take it back). To ‘challenge’ an assertion, in this sense, is to perform a speech act (typically a question²⁸) that disputes the veracity of the speaker’s claim, such as ‘How do you know that?’, or ‘Is that true?’. In turn, a challenge to *p* is ‘appropriate’ only if it is not already a settled issue in the conversation that *p* is true.²⁹ I will come back on these notions and distinctions in the next section, as I discuss some examples of conversational challenges.

Since making an assertion inevitably involves undertaking both accountability and discursive responsibility, assertoric commitment is best characterised as the conjunc-

²⁷ The “*prima facie*” qualification is meant to specify that falsity only determines a *defeasible* right to criticise the speaker. As noted by Peirce (see above), a speaker can be *excusable* for asserting something false: for instance, if their false claim was uttered under coercion, or if they had excellent reasons to think that what they said was true. Of course, if I excuse someone for not *f*-ing, I am still presupposing that that person was responsible for *f*-ing in the first place. This complicates matters: when excuses apply, responsibility for *f*-ing and criticisability for not *f*-ing can come apart, so that we cannot determine if the speaker is responsible for *f*-ing just by considering whether she is criticisable for not *f*-ing. The notion of *prima facie* accountability allows us to overcome this difficulty by pushing excuses out of the picture: this helpful notion captures both the cases in which the speaker is actually criticisable for saying something false, and the cases in which such criticism would be warranted, if it hadn’t been defeated by extenuating circumstances. Thanks to this qualification, we can define assertoric accountability in terms of one’s (actual or counterfactual) criticisability for the falsity of a proposition.

²⁸ Authors like Brandom adopt a narrower view: challenges can only be assertions that are incompatible with what the speaker said (1994, p. 178, 238, Wanderer 2010). I take Brandom’s view to be unduly restrictive (cf. Toulmin 1958; Rescher 1977, pp. 9–11; Rescorla 2009a), as it seems to me that *questions* are a paradigmatic example of challenges to the veracity of someone else’s assertion.

²⁹ Or, at least, if the speaker hasn’t already done all that she could to prove that *p* is true. In argumentation theory there is considerable disagreement as to what makes a challenge legitimate, and it would be over-ambitious for this paper to attempt to settle the issue once and for all; for further refinements, I defer to the relevant literature (see e.g. Rescorla 2009b).

tion of both normative effects. You are committed to a proposition if you are *prima facie* liable to be criticised in case the proposition is false, and *prima facie* expected to back up your claim in response to appropriate challenges (or else take it back). In sum:

Assertoric commitment

S is (assertorically) committed to *p* being the case iff

- (i) *S* is ‘accountable’ for *p*
- (ii) *S* is ‘discursively responsible’ for *p*.

In light of this characterisation, the commitment-based definition of lying presented at the beginning of this chapter can now be expounded, to display more clearly which conditions need to be satisfied for a speech act to count as a lie:

Commitment-based Definition of Lying

S lies iff

- (a) *S* utters a sentence with content *p*
- (b) In virtue of doing (a), *S* is accountable and discursively responsible for *p*
- (c) *S*’s utterance is insincere

5 Drawing the right distinctions

The commitment-based definition of lying meets the desiderata that have been identified so far. First, it differentiates between lies and other statements whose content is believed to be false but that are not lies, such as ironic and metaphoric utterances. This is because ‘accountability’ clearly does not obtain in these cases: it would be patently inappropriate, for instance, to criticise an ironic or metaphoric utterance on the grounds that its literal content is false.

Second, unlike Carson’s ACC, the proposed definition correctly identifies proviso-lies as genuine lies. While the notion of warrant cannot admit of degrees (because warranting is understood as an implicit promise), the notion of commitment can. The possibility of strengthening or diminishing the speaker’s degree of commitment to a proposition is widely acknowledged and discussed in the speech act theoretic literature (Searle 1976, p. 5; Holmes 1984; Searle and Vanderveken 1985, pp. 98–99; Coates 1987, p. 112; Sbisà 2001, pp. 1805–1806; Simons 2007; Thaler 2012; Marsili 2014, pp. 165–170),³⁰ and plays an important role in explaining the relations of ‘illocutionary entailment’ between different speech acts (Searle and Vanderveken 1985, pp. 130–131). For instance, most authors who employ the notion of commitment agree that by choosing to use the performative ‘swear’ in (13a) (instead of plainly asserting (13)) the speaker (call her Peppa) reinforces her commitment to the proposition (13*), whereas in choosing the performative ‘conjecture’ in (13b) she removes such commitment.

(13) Emma was drunk last night

(13a) I swear that (13*) Emma was drunk last night

³⁰ To be sure, there are many accounts of commitment on the market, and some authors (like Geurts 2019) adopt a different, binary conception that does not admit of degrees. Clearly, this alternative conception will not do for our purposes.

(13b) I conjecture that (13*) Emma was drunk last night

Since swearing (as in 13a) involves a stronger commitment than asserting (as in 13), its utterance is said to ‘illocutionarily entail’ the performance of an assertion, meaning that it cannot be performed without also asserting that (13*) is true. By contrast, the speaker of (13b) is merely making a conjecture, which does not commit her to the truth of (13*): (13b) is not an assertion (Searle and Vanderveken 1985, pp. 129–130; cf. Marsili 2015, pp. 124–125, 2016, pp. 277–278).

The test for discursive responsibility draws the right distinctions here. If we were to challenge (13b) with questions like “How do you know?” or “Is that true?”, Peppa would not be expected to provide evidence that (13*) is actually true. She could appropriately reply: “I don’t know, I just made a conjecture”.³¹ Contrast this with Peppa’s sworn statement (13a): the same questions (“How do you know?”, etc.), when raised in response to (13a), would indeed generate an expectation that Peppa defend her claim (e.g. “I saw her stumbling around and slurring her words”). In this case, unlike with her conjecture, Peppa is discursively responsible for the truth of (13*).

This shows that commitment can be *reinforced* (as in the sworn statement (13a)) or *removed* (as in the conjecture (13b)), but not yet that it can be *mitigated* while still asserting, which is what we need to show in order to prove that the definition can capture proviso lies. Cases of this sort are not uncommon, and typically emerge from the use of some modifiers, such as *evidentials* or *epistemic modals*. For example, suppose Peppa says:

(13c) Apparently (13*) Emma was drunk last night

With (13c), Peppa undertakes responsibility for the truth of what she has said—although less responsibility than she would have undertaken, had she uttered the unguarded assertion (13*) instead (see e.g. Caffi 1999; Sbisà 2001, 2014). This is intuitive, but we can be more precise. In which sense is Peppa accepting ‘less responsibility’ in making the mitigated assertion (13c) *in lieu* of (13)? To answer this question, let us consider each component of commitment in turn.

Accountability has to do with the social sanctions faced by the speaker if the proposition turns out to be false. Clearly, these sanctions can be more or less severe; the claim here is that mitigated assertions warrant *less severe* sanctions. This much is uncontroversial: any competent speaker knows that, *ceteris paribus*, an unguarded statement like (13) warrants more severe criticisms than a guarded statement like (13c), if (13*) turns out to be false. In fact, it is often to diminish their liability to criticisms that speakers prefer using a mitigated assertion over an unguarded one (cf. Holmes 1984; Fraser 2010).

A similar point applies to discursive responsibility. Speakers can be required to substantiate their claims with adequate evidence, but mitigation devices can affect which

³¹ At most, we may expect Peppa to explain *why* she made the conjecture, but this clearly falls short of expecting her to provide evidence that (13) is true, which is what discursive responsibility requires. After all, questions like “*Why* did you [performative verb] that *p*?” can be appropriately asked in response to virtually *any* speech act. Their availability is irrelevant to determining whether the speaker is committed (assertorically) to *p*: only the availability of challenges *to the veracity of p* reliably indicates that the speaker is discursively responsible for *p*. For more on the appropriateness of challenges to assertions, conjectures, and other assertive speech acts, see Green (2017, §2).

kind (and amount) of evidence counts as *adequate*. Evidentials such as ‘apparently’ can set the epistemic bar of adequacy to a lower standard of evidence (Sbisà 2014). In fact, it is natural to use a guarded assertion like (13c) instead more direct ones like (13) when one has *some* evidence in support of what they say, but not quite enough to license a direct assertion.

This should clarify in which sense accountability and discursive responsibility are mitigated in (13c): (13c) licenses less severe sanctions than (13), and binds the speaker to a less demanding standard of evidence. The same is not true of the conjecture (13b), where neither condition is satisfied: it would be unfair to criticise Peppa for saying (13b) in case (13*) turns out to be false, or to demand her to provide evidence in support of the truth of her conjecture.

Back to proviso-lies, the reason why they do not pose a threat to the commitment-based definition is that they behave like mitigated assertions (and unlike conjectures). In (11b) both accountability and discursive responsibility are met, although to a lesser extent:

(11b) Mel broke your trophy. But I was kinda drunk, and there were lots of people in there, so don’t take my word for it

By uttering (11b), Jamie signals that he is not willing to accept full responsibility for the proposition being true. Like the mitigated assertion (13c), and unlike the conjecture (13b), it is appropriate to inquire about the epistemic grounds for Jamie’s assertion (What evidence does he have to support the claim that Mel broke the trophy? Does he remember *seeing* him?). However, given the qualification added by Jaime, we will be satisfied with non-conclusive evidence in favour of the claim (e.g. he remembers *seeing him*, but cannot be sure). That said, the expectation that Jaime defend his claim is nonetheless clearly present: it would be inappropriate for Jaime to simply reply: “I don’t see why you’re asking these questions, I never claimed that Mel broke the trophy”. A reply of this kind would be appropriate, by contrast, if Jaime had simply made a conjecture, as in (13b). Similarly, it would be appropriate to reproach Jamie if the assertion turns out to be false (we may say: ‘You shouldn’t have accused Mel!’), although we would not be entitled to the same sort of reactive attitudes than an unguarded assertion would have warranted (after all, he invited us not to take his word for it). Like for (13c), both ‘accountability’ and ‘discursive responsibility’ are mitigated, but satisfied. This shows that, unlike Carson’s ACC, the proposed definition counts proviso-lies as mitigated assertions (and therefore as lies).³²

³² A referee points out that proviso-lies like (11) do not invite belief in their unmitigated content (Mel broke your trophy), and asks whether this is compatible with generating a commitment towards that content. My answer is positive. Simply put, the proviso at most prevents the realisation of a *perlocutionary* effect (making the hearer believe that *p*), which is logically (and pragmatically) compatible with bringing about an *illocutionary* one (committing yourself to *p*). Assertors typically intend to achieve the perlocutionary goal of convincing the hearer (usually, we aim to convince our interlocutors), but they can make assertions even if they do not have this intention (Davis 1999; Alston 2000; Green 2007; Sorensen 2007; MacFarlane 2011). If this is right, explicitly denying that you have a perlocutionary intention (“you don’t have to believe me”, “don’t take my word for it”) does not prevent you from bringing about your assertion’s illocutionary effect (committing yourself to *p*). For a discussion of some other species of provisos that threaten my view more directly, in particular in response to Rudy Hiller’s examples (2016, pp. 38–51), see Marsili (2020b).

Lastly, the commitment-based definition seems able to draw the right distinctions about explicit performatives. Since betting and swearing were discussed above (13a, 13b), we only need to consider the following cases:

- | | |
|-------------------------|---|
| (1) I assert that | (1*) I received expressed consent from the patient |
| (2) I promise that | (2*) I will wear a blue dress at the wedding |
| (5) I advise that | (5*) you try that quiche |
| (6) I command that | (6*) you steal that chicken |
| (8) Assume that | (8*) I can lift weights with my mind [...] |
| (9) Let us suppose that | (9*) there is a demon that systematically deceives us |

The predictions of the commitment-based definition are rather straightforward here. By asserting or promising that p in (1–2), the speaker becomes accountable and discursively responsible for their content, namely (1*–2*), so that these utterances are counted as lies when they are uttered insincerely. On the other hand, by uttering (5), (6), (8) and (9) the speaker does not become committed to the corresponding propositions (5*), (6*), (8*) and (9*), so that these utterances cannot be classified as lies by the definition. For instance, in response to (8) it would be inappropriate to reproach the speaker if it turns out that she has not telekinetic powers, or to challenge the speaker by asking “How do you know that you have these powers?”. It is apparent that the same tests are passed by all the other explicit performatives that cannot be lies (namely (4), (6), and (9)).

It could be objected that it is not clear that in promising (2) the speaker becomes assertorically committed to (2*), as I have claimed above. Promissory commitment and assertoric commitment differ in important respects: promising involves being responsible for *making* something true, while asserting involves being responsible for something *being* true (Watson 2004). Perhaps (2) commits the speaker to (2*) ‘promissorially’, but not ‘assertorically’. The test for discursive responsibility seems to corroborate this hypothesis: asking “How do you know?” or “What makes you think that?” in response to (2) is simply inappropriate, and it does not seem that one would be expected to support their claim with evidence in response to this sort of challenges.

Although I agree that there is more to promissory commitment than just assertoric commitment, this does not mean that the former is incompatible with the latter. Within the speech-act theoretic framework that I am adopting (Searle and Vanderveken 1985, p. 184), the relation between promissory and assertoric responsibility can be explained in terms of the notion of ‘illocutionary entailment’ introduced earlier. The underlying idea is that, if I promise that (2*) (“I will wear a blue dress at the wedding”), I am also thereby claiming that it will be true, at time of the wedding, that I will wear a blue dress: whenever promissory responsibilities arise, assertoric ones have to arise too.³³ At closer inspection, this objection is rather based on a misunderstanding of what constitutes discursive responsibility in (2).

Recall (§4) that discursive responsibility only requires the speaker to answer *appropriate* challenges (cf. MacFarlane 2005b). Challenges are *not* appropriate (in the relevant sense) if they are infelicitous for reasons that have obviously nothing to do with the force of the original utterance. A typical example is when a challenge is

³³ I defend this claim in more detail in Marsili (2016, pp. 277–278).

infelicitous because the answer is already common knowledge in the conversation. If I claim “My tooth hurts”, it would be inappropriate to challenge my claim by asking me “How do you know?”, because it is already obvious how I know that my tooth hurts—but this clearly should not be taken as evidence that my utterance is not an assertion. Similarly, since whether I wear a blue dress at the wedding will depend primarily on my decisions, asking “How do you know?” in response to (2) would not be an appropriate challenge. In both cases, the challenge is *inappropriate*, because it is obvious that the challenger already knows the answer to the question, so that considering its availability is irrelevant to determining whether the speaker is committed to the proposition.³⁴

How should we test for discursive responsibility in these cases? Since in these contexts the speaker’s *reasons for believing* (2*) are already common ground, we should consider challenges that put into question the veracity of the utterance more directly: for example, “Is that true?”, “Does it really [hurt]?”, or “Will you really [bring a blue dress]?”. Just like ‘How do you know’ challenges, these questions are appropriate only when the speaker is assertorically committed to the relevant proposition, so that they still constitute a reliable test for discursive responsibility. And these questions are clearly available in response to (2), showing that also in this case the speaker is bound by the relevant discursive obligations. In addition to this, in (2) ‘accountability’ clearly obtains: if I eventually wear a red dress to the wedding, I can be criticised for (2*) being false, and appropriately so. The right verdict is thus given also in the case of insincere promises.

It seems that the proposed account avoids all the counterexamples that affect other views. Unlike the other definitions considered so far, it deals correctly with a wide range of performative utterances, distinguishing speech acts that can be used to lie from speech acts that cannot. It captures not only standard assertions, but also assertions uttered by means of explicit performatives (e.g. ‘I hereby assert that *p*’) and explicit performatives that *illocutionary entail* an assertion, such as acts of promising or swearing. It is able to rule out illocutionary acts that are not assertions, including speech acts belonging to the class of assertives (like bets, conjectures, and suppositions), and directives (like commands, advice, and suppositions). The proposed definition brings together two philosophical traditions that analyse (respectively) assertion in terms of accountability and discursive responsibility, to deliver a fine-grained account of the distinctive responsibilities that emerge in virtue of asserting a given proposition, improving on previous attempts to characterise the distinctive responsibilities that all liars undertake. Due to its intensional accuracy, it provides a potentially insightful analysis of two concepts (assertion and lying) that are central to many contemporary philosophical inquiries in ethics, epistemology, and philosophy of language.

³⁴ Although this explanation will do for our present purposes, a further clarification may be of interest. In Marsili (2018b) I consider these issues in more depth, and distinguish between a challenge being *inappropriate* (which depends on whether the answer to the challenge is already in the common ground) and *illegitimate* (which depends on whether the speaker was committed to *p* in the first place). Only when a challenge is ‘illegitimate’ we have evidence that the speaker is not discursively responsible for *p*. Of course, challenges to promises like (2) are only ‘inappropriate’ in this sense, whereas challenges to non-assertoric acts like (6) or (8) are genuinely ‘illegitimate’.

Funding Open access funding provided by Alma Mater Studiorum - Università di Bologna within the CRUI-CARE Agreement. Research funding was provided by Ministerio de Ciencia, Innovación y Universidades (Grant No. FFI2016-80636-P, AEI/FEDER, UE).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adler, J. E. (1997). Lying, deceiving, or falsely implicating. *Journal of Philosophy*, 94(9), 435–452.
- Allcott, H., & Gentzkow, M. (2017). social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2), 211–236.
- Alston, W. P. (2000). *Illocutionary acts and sentence meaning*. Ithaca: Cornell University Press.
- Arico, A. J., & Fallis, D. (2013). Lies, damned lies, and statistics: An empirical investigation of the concept of lying. *Philosophical Psychology*, 26(6), 790–816.
- Bach, K. (1975). Performatives are statements too. *Philosophical Studies*, 28(4), 229–236.
- Bach, K., & Harnish, R. M. (1979). *Linguistic communication and speech acts*. Cambridge: MIT Press.
- Berstler, S. (2019). What's the good of language? On the moral distinction between lying and misleading. *Ethics*, 130(1), 5–31.
- Black, M. (1952). Saying and disbelieving. *Analysis*, 13(2), 25–33.
- Brandom, R. (1983). Asserting. *Noûs*, 17(4), 637–650.
- Brandom, R. (1994). *Making It explicit: Reasoning, representing, and discursive commitment*. Cambridge: Harvard University Press.
- Caffi, C. (1999). On mitigation*. *Journal of Pragmatics*, 31, 881–909.
- Carson, T. L. (1988). On the definition of lying: A reply to Jones and revisions. *Journal of Business Ethics*, 7(7), 509–514.
- Carson, T. L. (2006). The definition of lying. *Noûs*, 2, 284–306.
- Carson, T. L. (2010). *Lying and deception*. Oxford: Oxford University Press.
- Carson, T. L., Wokutch, R. E., & Murrmann, K. F. (1982). Bluffing in labor negotiations: Issues legal and ethical. *Journal of Business Ethics*, 1(1), 13–22.
- Chappell, S. G. (2016). Political deliberation under conditions of deception: The case of brexit. *Think*, 15(44), 7–13.
- Chisholm, R. M., & Feehan, T. D. (1977). The intent to deceive. *Journal of Philosophy*, 74(3), 143–159.
- Coates, J. (1987). Epistemic modality and spoken discourse. *Transactions of the Philological Society*, 85, 110–131.
- Davidson, D. (1985). Deception and division. In *The multiple self*. Cambridge: Cambridge University Press.
- Davidson, D. (2001). *Inquiries into truth and interpretation: philosophical essays* (Vol. 2). Oxford: Oxford University Press.
- Davis, W. (1999). Communicating, telling and informing. *Philosophical Inquiry*, 21(1), 21–43.
- Fallis, D. (2009). What Is lying? *Journal of Philosophy*, 106(1), 29–56.
- Fallis, D. (2010). Lying and deception. *Philosophers' Imprint* 10(11).
- Fallis, D. (2012). Lying as a violation of Grice's First maxim of quality. *Dialectica*, 66(4), 563–581.
- Fallis, D. (2013). Davidson was almost right about lying. *Australasian Journal of Philosophy*, 91(2), 337–353.
- Fallis, D. (2015). Are bald-faced lies deceptive after all? *Ratio*, 28(1), 81–96.
- Fallis, D. (2018). What is deceptive lying? In *Lying: language, knowledge, ethics, politics*. Oxford: Oxford University Press.
- Faulkner, P. (2007). What is wrong with lying? *Philosophy and Phenomenological Research* LXXV, 3, 535–557.

- Faulkner, P. (2013). Lying and deceit. *The International Encyclopedia of Ethics*, 3101–3109.
- Ferrara, E., Chang, H., Chen, E., Muric, G., Patel, J. (2020). Characterizing social media manipulation in the 2020 U.S. presidential election. First Monday.
- Fraser, B. (2010). Pragmatic competence: The case of hedging. In G. Kaltenbo, W. Mihatsch, & S. Stefan (Eds.), *New approaches to hedging* (pp. 15–34). Bingley: Emerald Group Publishing Limited.
- Fried, C. (1978). *Right and wrong*. Cambridge: Harvard University Press.
- Geurts, B. (2019). Communication as commitment making: Speech acts, implicatures. *Common Ground. Theoretical Linguistics*, 45(1–2), 1–30.
- Ginet, C. (1979). Performativity. *Linguistics and Philosophy*, 3, 245–265.
- Green, M. (1999). Illocutions, implicata, and what a conversation requires. *Pragmatics & Cognition*, 7(1), 65–91.
- Green, M. (2000). Illocutionary force and semantic content. *Linguistics and Philosophy*, 23, 435–473.
- Green, S. P. (2001). Lying, misleading, and falsely denying: How moral concepts inform the law of perjury, fraud, and false statements. *Hastings LJ*, 109(1984), 157–212.
- Green, M. (2007). *Self-expression*. Oxford: Oxford University Press.
- Green, M. (2009). Speech acts, the handicap principle and the expression of psychological states. *Mind and Language*, 24(2), 139–163.
- Green, M. (2017). *Assertion*. Oxford: Oxford Handbook Online.
- Hamblin, C. L. (1970a). *Fallacies*. Methuen, London.
- Hamblin, C. L. (1970b). The effect of when it's said. *Theoria*, 17(3), 185–185.
- Harris, R. (1978). The descriptive interpretation of performative utterances. *Journal of Linguistics*, 14(2), 309–310.
- Hawley, K. (2019). *How to be trustworthy*. Oxford University Press.
- Hedenius, I. (1963). Performatives. *Theoria*, 29, 115–136.
- Holmes, J. (1984). Modifying illocutionary force. *Journal of Pragmatics*, 8, 345–365.
- Holton, R. (2008). Partial belief, partial intention. *Mind*, 117(465), 27–58.
- Isenberg, A. (1964). Deontology and the ethics of lying. *Philosophy and Phenomenological Research*, 24(4), 463–480.
- Jary, M. (2007). Are explicit performatives assertions? *Linguistics and Philosophy*, 30(2), 207.
- Keiser, J. (2016). Bald-faced lies: How to make a move in a language game without making a move in a conversation. *Philosophical Studies*, 173(2), 461–477.
- Krauss, S. F. (2017). Lying, risk and accuracy. *Analysis*, 73, 651–659.
- Krifka, M. (2014). Embedding illocutionary acts. *Recursion: Complexity in Cognition*, no. 1995: 59–87.
- Krstić, V. (2018). Knowledge-lies re-examined. *Ratio*, 31(3), 312–320.
- Krstić, V. (2019). Can you lie without intending to deceive? *Pacific Philosophical Quarterly*, 100(2), 642–660.
- Kupfer, J. (1982). The moral presumption against lying. *The Review of Metaphysics*, 36(1), 103–126.
- Lackey, J. (2013). Lies and deception: An unhappy divorce. *Analysis*, 73(2), 236–248.
- Leland, P. R. (2015). Rational responsibility and the assertoric character of bald-faced lies. *Analysis*, 75(4), 550–554.
- Lewis, D. (1970). General semantics. *Synthese*, 22(1–2), 18–67.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1), 563–582.
- MacFarlane, J. (2003). Epistemic modalities and relative truth (Unpublished Manuscript).
- MacFarlane, J. (2005a). Making sense of relative truth. *Proceedings of the Aristotelian Society*, 105(1), 305–323.
- MacFarlane, J. (2005b). The assessment sensitivity of knowledge attributions. In T. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 4, pp. 197–233). Oxford: Oxford University Press.
- MacFarlane, J. (2011). What is assertion? In J. Brown & H. Cappelen (Eds.), *Assertion: New philosophical essays*. Oxford: Oxford University Press.
- Mahon, J. E. (2008). Two definitions of lying. *International Journal of Applied Philosophy*, 22(2), 211–230.
- Mahon, J. E. (2015). The definition of lying and deception. *Stanford Encyclopedia of Philosophy*.
- Mannison, D. S. (1969). Lying and lies. *Australasian Journal of Philosophy*, 2(47), 132–144.
- Marques, T. (2020). Disagreement with a bald-faced liar. *Ratio*, no. December 2019: 1–14.
- Marsili, N. (2014). Lying as a scalar phenomenon. In S. Cantarini, W. Abraham, & E. Leiss (Eds.), *Certainty-uncertainty—and the attitudinal space in between* (pp. 153–173). Amsterdam: John Benjamins Publishing Company.

- Marsili, N. (2015). Normative accounts of assertion: From peirce to Williamson, and Back Again. *Rivista Italiana Di Filosofia Del Linguaggio*, 112–30.
- Marsili, N. (2016). Lying by Promising. *International Review of Pragmatics*, 8(2), 271–313.
- Marsili, N. (2017). You don't say! Lying, asserting and insincerity. Dissertation, University of Sheffield. Ethos ID: uk.bl.ethos.731535. Available at: <http://etheses.whiterose.ac.uk/19068/>.
- Marsili, N. (2018a). Truth and assertion: Rules versus aims. *Analysis*, 78(4), 638–648.
- Marsili, N. (2018b). Lying and certainty. In J. Meibauer (Ed.), *The oxford handbook of lying* (pp. 169–182). Oxford: Oxford University Press.
- Marsili, N. (2019). Immoral lies and partial beliefs. *Inquiry*, 0(0), 1–11.
- Marsili, N. (2020a). Lies, Common ground and performative utterances. (Under Review).
- Marsili, N. (2020b). The definition of assertion. <https://doi.org/10.2139/ssrn.3711804>.
- Meibauer, J. (2005). Lying and falsely implicating. *Journal of Pragmatics*, 37(9), 1373–1399.
- Meibauer, J. (2014). *Lying at the semantics-pragmatics interface. Lying at the semantics-pragmatics interface*. Berlin: De Gruyter.
- Newey, G. (1997). Political lying: A defense. *Public Affairs Quarterly*, 11(2), 93–116.
- Peirce, C. S. (1960) (*CP*) *Collected papers of Charles sanders peirce*. In Hartshorne, C., Weiss, P., Burks, A. W. Cambridge: Harvard University Press.
- Peirce, C. S. (1963) (*MS*) *The Charles S. Peirce papers*. Harvard: Harvard University Library.
- Pepp, J. (2018). Truth serum, liar serum, and some problems about saying what you think is false. In *Lying: language, knowledge, ethics, politics*. Oxford University Press.
- Portner, P. (2004). The semantics of imperatives within a theory of clause types. *Proceedings of SALT, XIV*(2), 235–252.
- Primoratz, I. (1984). Lying and the “methods of ethics”. *International Studies in Philosophy*, 16, 35–57.
- Recanati, F. (2013). Content, mood, and force. *Philosophy Compass*, 8(7), 622–632.
- Reimer, M. (1995). Performative utterances: A reply to bach and harnish. *Linguistics and Philosophy*, 18, 655–675.
- Rescher, N. (1977). *Dialectics: A controversy-oriented approach to the theory of knowledge*. Albany: State University of New York Press.
- Rescorla, M. (2009a). Assertion and its constitutive norms. *Philosophy and Phenomenological Research*, LXXIX(1), 98–130.
- Rescorla, M. (2009b). Shifting the burden of proof? *Philosophical Quarterly*, 59(234), 31–36.
- Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5(6), 1–69.
- Ross, W. D. (1930). *The right and the good*. Oxford: Oxford University Press.
- Rudy Hiller, F. (2016). Assertion, justificatory commitment, and trust. *Análisis Filosófico*, XXXVI(1), 29–53.
- Saul, J. (2012). *Lying, misleading, and the role of what is said*. Oxford: Oxford University Press.
- Sbisà, M. (2001). Illocutionary force and degrees of strength in language use. *Journal of Pragmatics*, 33(September 1996), 1791–1814.
- Sbisà, M. (2014). Evidentiality and illocution. *Intercultural Pragmatics*, 11(3), 463–483.
- Searle, J. R. (1969). *Speech acts: An essay in the philosophy of language* (Vol. 626). Cambridge: Cambridge University Press.
- Searle, J. R. (1975). The logical status of fictional discourse. *New Literary History*, 6(2), 319–332.
- Searle, J. R. (1976). A classification of illocutionary acts. *Language in Society*, 5(1), 1–23.
- Searle, J. R. (1989). *How performatives work. Linguistics and philosophy*. Berlin: Springer.
- Searle, J. R., & Vanderveken, D. (1985). *Foundations of illocutionary logic*. Cambridge: CUP Archive.
- Silverman, C. (2016). This analysis shows how fake election news stories outperformed real news on facebook. *BuzzFeed News*.
- Simpson, D. (1992). Lying, liars and language. *Philosophy and Phenomenological Research*, 52(3), 623–639.
- Simons, M. (2007). Observations on embedding verbs, evidentiality, and presupposition. *Lingua*, 117(6), 1034–1056.
- Sorensen, R. (2007). Bald-faced lies! lying without the intent to deceive. *Pacific Philosophical Quarterly*, 88, 251–264.
- Sorensen, R. (2010). Knowledge-lies. *Analysis*, 70(4), 608–615.
- Sorensen, R. (2011). What lies behind misspeaking. *American Philosophical Quarterly*, 48(4), 399–410.
- Stalnaker, R. C. (1978). Assertion. In P. Cole (Ed.), *Pragmatics*. Cambridge: Academic Press.
- Stalnaker, R. C. (2002). Common ground. *Linguistics and Philosophy*, 25(5), 701–721.

- Stokke, A. (2013a). Lying and asserting. *Journal of Philosophy*, 110(1), 33–60.
- Stokke, A. (2013b). Lying, deceiving, and misleading. *Philosophy Compass*, 8(4), 348–359.
- Stokke, A. (2018). *Lying and insincerity*. Oxford: Oxford University Press.
- Tanesini, A. (2016). “Calm down, Dear”: Intellectual arrogance, silencing and ignorance. *Proceedings of the Aristotelian Society, Supplementary*, 90(1), 71–92.
- Tanesini, A. (2019). Silencing and assertion. In S. Goldberg (Ed.), *The oxford handbook of assertion*. Oxford: Oxford University Press.
- Thaler, V. (2012). Mitigation as modification of illocutionary force. *Journal of Pragmatics*, 44(6–7), 907–919.
- Toulmin, S. E. (1958). *The uses of argument*. Cambridge: Cambridge University Press.
- Viebahn, E. (2017). *Non-literal lies*. v. Dordrecht: Springer.
- Viebahn, E. (2019). *Lying with presuppositions*. *Noûs*. New York: Wiley.
- Viebahn, E., Alex, W., Neele, E., & Pascale, W. (2018). Can a question be a lie? An empirical investigation, no. July.
- Wanderer, J. (2010). Brandom’s challenges. In *Reading brandom: On making it explicit*. Abingdon: Routledge.
- Watson, G. (2004). Asserting and promising. *Philosophical Studies*, 117(1), 57–77.
- Watson, L. (2018). Systematic epistemic rights violations in the media: A brexit case study. *Social Epistemology*, 32(2), 88–102.
- Weiner, M. (2005). Must we know what we say? *The Philosophical Review*, 114(2), 227–251.
- Whiting, D. (2012). Stick to the facts: On the norms of assertion. *Erkenntnis*, 78(4), 847–867.
- Wiegmann, A., Samland, J., & Waldmann, M. R. (2016). Lying despite telling the truth. *Cognition*, 150, 37–42.
- Williams, B. A. O. (2002). *Truth and truthfulness an essay in genealogy*. Princeton: Princeton University Press.
- Williamson, T. (1996). Knowing and asserting. *The Philosophical Review*, 105(4), 489–523.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.