

Lying as a scalar phenomenon

Insincerity along the certainty-uncertainty continuum*

Neri Marsili

University of Sheffield

Every species is vague, every term goes cloudy at its edges, and so in my way of thinking, relentless logic is only another name for stupidity – for a sort of intellectual pigheadedness.

H.G. Wells, *First and Last Things* (1908)

In the philosophical debate on lying, there has generally been agreement that either the speaker believes that his statement is *false*, or he believes that his statement is *true*. This article challenges this assumption, and argues that lying is a scalar phenomenon that allows for a number of intermediate cases – the most obvious being cases of uncertainty. The first section shows that lying can involve beliefs about graded truth values (*fuzzy lies*) and graded beliefs (*graded-belief lies*). It puts forward a new definition to deal with these scalar parameters, that requires that the speaker asserts what he believes *more likely to be false than true*. The second section shows that statements are scalar in the same way beliefs are, and accounts for a further element of scalarity, illocutionary force.

o. The definition of lying and the insincerity condition

A number of philosophers have tried to define lying – *i.e.* to find the necessary and sufficient conditions for a statement to be a lie. According to the “classic” definition, first proposed in St. Augustine’s *De Mendacio* (DM, 3.3) and later developed in several works (e.g. Mannison 1969:132; Kupfer 1982:134; Williams 2002:96), “lying is to make a believed-false statement¹ to another

* The author thanks Alberto Voltolini, Matteo Grasso, Carla Bazzanella, Jennifer Saul, Rosanna Keefe, Guido Bonino and Jörg Meibauer for their helpful comments, Igor Ž. Žagar for his support, Nina Black Simone for proofreading, and three anonymous reviewers.

1. In the philosophical literature on lying, “statement” is broadly used to refer to the utterance of a declarative sentence that expresses a propositional content.

person with the intention that that other person believe that statement to be true” (Mahon 2008: 1.0). More formally:

Classic definition:

S lies to A about p iff:

- (i) S states p
- (ii) S believes p to be false
- (iii) S intends A to believe p

Even if the classic definition is the most shared view on lying, it is not universally accepted, and has been challenged in various ways. Recently, many instances of lying that involve no intention to deceive have been brought to the attention of the philosophical debate (e.g. Carson 2006; Sorensen 2007; Fallis 2009, 2010). Consequently, a growing number of authors have disputed the necessity of condition (iii), suggesting that to define lying we should instead narrow condition (i) by requiring that the speaker utter a *genuine assertion*.²

Alternative definition:

S lies to A about p iff:

- (i) S asserts p
- (ii) S believes p to be false

Virtually any philosopher endorses either the classic or the alternative definition. In this paper, I will not argue in favour of one of these definitions, and my inquiry applies to both of them. The reason is that I will discuss and attempt to revise the only condition that both definitions naively take for granted: condition (ii), the *insincerity condition*.

The insincerity condition requires that the speaker believe his statement to be *false*. The aim of the insincerity condition is to rule out believed-true statements uttered with the intention to deceive, *i.e.* misleading statements.³ Misleading

2. In this debate it is commonly accepted that assertions are a particular kind of statements, but there is disagreement about what characterises an assertion. To solve this problem, some authors expand the definition by adding a condition that excludes non-assertive statements. Carson (2006, 2010) requires that the speaker *warrant* the truth of p ; Fallis (2009, 2012) that the speaker intend to violate the norm against communicating something false; Stokke (2013b) that the speaker propose p to become common ground. Rather than one of these disputed accounts, this paper will adopt a more shared notion of assertion, grounded on speech act theory. In speech act theory (and more generally in pragmatics) it is commonly accepted that for a speaker to assert that p is for the speaker to commit himself to the truth of p (e.g. Searle 1969; Brandom 1983, 1994).

3. For a more detailed definition of the lying/misleading distinction, see Saul (2012) and Stokke (2013).

statements can be as deceiving as lies, but there is a clear difference between misleading and lying. Consider the well-known excerpt of the life of Saint Athanasius (MacIntyre 1994: 336):

Persecutors, dispatched by the emperor Julian, were pursuing him up the Nile. They came on him traveling downstream, failed to recognize him, and enquired of him: "Is Athanasius close at hand?" He replied: "He is not far from here." The persecutors hurried on and Athanasius thus successfully evaded them without telling a lie.

Athanasius tried to deceive his pursuers into thinking that he was not the man they were searching for. However, it seems wrong to say that he was lying, since he said something that he believed to be true – namely, that Athanasius was not far. The insincerity condition accounts consistently for these cases, deeming them not lies, and thus recognising the distinction between lying and misleading.⁴

However, it is not clear that the insincerity condition (ii) is a correct account of insincerity (Mahon 2008: 1.5). Most proponents of this condition (call it the “traditional insincerity condition”, or TIC) assume that either a statement is believed to be *true*, or it is believed to be *false*.⁵ From this perspective (call it the *dichotomic view*), the definition of lying correctly rules out only statements that are believed to be true.

Some philosophers contend that the dichotomic view is not an adequate assumption for the definition of lying, because a number of intermediate credal states⁶ exist between believing *p* to be true and believing *p* to be false. First, it is possible for a speaker to believe that a statement is only partly false (rather than utterly false): in this case, the speaker believes that *p* has a “graded truth value”. Second, it is possible for a speaker not to be certain; in other words, to have a graded degree of confidence (rather than a flat-out belief) in the falsity of a statement: intermediate beliefs of this kind are called “graded beliefs”. The difference between these two layers of scalarity can be difficult to grasp: in Section 1.1. and 1.2 I will explain in detail this subtle distinction.

The main reason to reject the dichotomic view on lying is that it does not allow intermediate credal states of this kind to count as lies. For instance, if A states

4. Even if there is strong philosophical consensus that the difference between lying and misleading is well grounded, some authors regard misleading as a form of lying (e.g. Bok 1978: 14; Barnes 1994: 11; Meibauer 2005, 2011, 2014: 102–155; D’Agostini 2012).

5. Another dichotomy that is taken for granted in the debate is that either the speaker has a belief about the truth value of his statement, or he does not have a belief about it. I will come back to this possibility in Section 1.2.

6. In epistemology, “credal state” indicates a specific kind of mental state: *i.e.* the mental state of having a belief.

p and he believes p to be probably false (or, alternatively, partly false) it seems clear that A is lying,⁷ but the traditional insincerity condition (TIC) does not count this statement as lying. The TIC is therefore too narrow, because it incorrectly rules out all the statements that are neither outrightly believed to be true nor outrightly believed to be false.

This article attempts to outline a non-dichotomic alternative to the TIC that counts this kind of statements as lies and allows for various degrees of insincerity in lying. Since insincerity is a relation between what is believed and what is asserted, the article is divided in two parts, the former focusing on beliefs and the latter on statements.

The first section thoroughly analyses two counterexamples to the dichotomic view: namely, lies that involve beliefs about graded truth values (1.1) and lies that involve graded beliefs (1.2). It presents a revised insincerity condition to deal with these cases, according to which the speaker must believe his statement to be *more likely to be false than true*.

The second section focuses on the speaker's statements and specifically on the linguistic devices that the speaker can use to lie about his degree of certainty or uncertainty. The inquiry shows that statements reflect the scalar structure of beliefs (2.1–2.2). Moreover, it introduces a further element of scalarity: the degree of illocutionary force of the speaker's utterance (2.3). The resulting account of lying is much more complex than the traditional ones, because it involves a complex interplay of pragmatic scalar parameters.

1. Alternatives to the insincerity condition

The only two non-dichotomic alternatives to the TIC have been proposed by Chisholm & Feehan (1977) and Carson (2006, 2010). Before examining these alternatives thoroughly, I would like to point out that neither Chisholm & Feehan nor Carson developed their alternative conditions in detail. In what follows, I will

7. Since I am focusing on the insincerity condition, I will always assume that in my examples the other conditions for lying obtain (i.e. that p is asserted with the intention to deceive). One might object that in this example (and in some of the following) condition (iii) does not obtain, because the speaker does not believe that his statement is utterly false, and thus does not believe that the statement is utterly deceiving. However, several philosophers (e.g. Chisholm & Feehan 1977:145; Fallis 2011:45; Staffel 2011:301) argue that intending to alter someone's degree of belief counts as intending to deceive them. Moreover, no intention to deceive is required by the *alternative* definition of lying.

examine the consequences yielded by the “insincerity condition” component⁸ of their definitions of lying. My discussion is thus not to be taken as an exegesis of their articles; rather, I will develop from their proposals an alternative to the TIC that allows for scalarity in lying.

1.1 Chisholm & Feehan: Graded truth values

According to the alternative insincerity condition proposed by Chisholm and Feehan, a speaker lies only if he states what he believes to be *false* or *not true* (call this the *fuzzy insincerity condition*). This alternative condition relies on the subtle difference between believing that *p* is *false* and believing that *p* is *not true*. Chisholm and Feehan (1977: 152) explicitly state that “it is logically possible to believe one of these things [e.g. *p* is not true] and not the other [e.g. *p* is false]”. One way to interpret this claim is to consider the difference between false and not true in a many-valued logic theoretical framework.

According to traditional logics, the truth value of a proposition is either true or false, and there is no point in distinguishing between a statement that is *false* and a statement that is *not true*. By contrast, many-valued logics allow for a larger set of truth degrees, thus accounting for a number of intermediate truth values between true and false.

In what follows, I will consider Chisholm and Feehan’s claim in the theoretical framework of *fuzzy logic*.⁹ Fuzzy logic is a many-valued logic conceived especially for predicates that are intrinsically vague (like being bald, or old, or happy) and that, being graded in nature, allow for a number of truth values. Fuzzy logic takes as truth values all real numbers between 0 and 1, where 0 is false and 1 is true. From this perspective, to say that the speaker believes that a proposition *p* is *not true* is to say that the speaker believes that the truth value of *p* is *x*, where x is $0 \leq x < 1$. By contrast, to say that the speaker believes that a proposition *p* is *false* is to say that the speaker believes that the truth value of *p* is 0. Against the dichotomic view, stating what is *believed to be not true* is thus not the same as stating what is *believed to be false*. Believed-false statements are a subset of insincere statements.

8. Both authors develop a definition of lying that differs from both the classic and the alternative definition also with respect to condition (i) and (iii). They also tend to focus their attention on the innovations they introduce for these latter conditions, rather than for the insincerity condition (condition (ii)).

9. Other interpretations of the claim are possible, but they will not be discussed here, since the aim of this paper is to outline the scalar dimensions of lying. For a broader discussion of many-valued and fuzzy logics, see Hajek (1998) and Gottwald (2001:423–492).

I will call lies that involve graded truth values *fuzzy lies*. Let us consider an example of *fuzzy lie* to grasp the difference from traditional lies. Suppose that Mickey utters (1) to persuade Daisy to date Donald:

(1) Donald is not bald

If Mickey thinks that Donald is *nearly* bald (e.g. he believes that (1) has a truth value of 0.3), he does not say something that he believes to be utterly false, and therefore the TIC does not count (1) as lying. However, intuitively Mickey is lying. The fuzzy insincerity condition (henceforth, FIC) accommodates our intuitions in counting (1) as a *fuzzy lie*, since (1) is believed to be *not true* (0.3 is less than 1 and more than 0).

The fuzzy insincerity condition (henceforth, FIC) is broader than the TIC: it allows all statements whose truth value is believed to be less than 1 to count as lies. Moreover, it correctly rules out misleading statements (given that they are believed to be true), while ruling in all standard lies (given that they are believed to be false). Nevertheless, it seems patent that the FIC is too broad, for the set of statements it allows to count as lies is too large: e.g. a statement whose truth value is believed to be 0.99 counts as a lie. To avoid this problem, one could modify the FIC in order to require that the believed truth value of the statement is closer to falseness than truthfulness- make it x , where x is $0 \leq x \leq 0,5$ (call this the *revised fuzzy insincerity condition*, henceforth FIC*).

This solution is nevertheless problematic, since there seems not to be a clear theoretical basis to set the limit at a precise value. If one accepts that a speaker who believes that the truth value of his statement is 0.5 is lying, then it also seems reasonable to accept that a speaker who believes that the truth value of his statement is 0.51 is lying. But the same line of reasoning would work for the successive values (0.52; 0.53; [...]; 1), so that, in the end, all statements would count as lies.

A further problem with this view concerns the very existence of such credal states: the proposed representation of the speaker's beliefs is so fine-grained that it may appear to overrate reality. It seems quite clear that in real life we do not experience the threshold between believing a statement to have a truth value of 0.49 rather than 0.51. If such a threshold exists, it is not consciously perceived; and since lying is a conscious choice, no such threshold can be taken as a necessary condition for lying.

Rather than positing a sharp "numerical" threshold that separates insincerity from sincerity, one could limit the FIC to require that the believed truth value of the statement is *perceptibly* closer to falseness than truthfulness. This condition (call it the FIC**) may seem rough compared to its "numerical" translation, but it acknowledges that our beliefs only roughly (and rarely) correspond to the subtle differences that fuzzy logic outlines. Whether or not one finds this revised

definition convincing, eventually we will be forced to abandon it: as I will show in the next section, it incorrectly rules out lies that involve graded beliefs.

1.2 Carson: Graded beliefs

Sometimes, our lies involve graded beliefs. Graded beliefs are credal states that we experience in our daily life, especially in contexts of uncertainty. In formal representations of belief, degrees of belief are represented as real numbers from 0 to 1, where 0 indicates certainty in the falsity of a proposition, 1 indicates certainty in its truth, and 0.5 indicates that the subject is uncertain about that proposition being true – in other words, that one regards the proposition just as likely to be true as false.¹⁰ Graded beliefs are hence part of a continuum whose poles are believing p to be certainly true and believing p to be certainly false – a spectrum of beliefs that Rubin (2010) calls the “certainty-uncertainty continuum”.¹¹

The FIC* accounts for lies that involve *beliefs about graded truth values* (*fuzzy lies*), but it does not account for lies that involve *graded beliefs* (call them *graded-belief lies*) about plain truth values. Such a subtle difference is worth an explanation.

Let us represent the general structure of beliefs as “B(p)”, where the variable “B” takes beliefs as values, and the variable “ p ” takes the propositional content of beliefs as values. The dichotomic view assumes that both “ p ” and “B” can assume as values only 1 or 0: either a subject believes p , or he does not believe p , and either he believes p to be true, or he believes p to be false.

Non-dichotomic accounts, by contrast, assume that “ p ” and/or “B” can take as values all the real numbers from 0 to 1. *Fuzzy lies* involve non-whole “ p ” values, while *graded-belief lies* involve non-whole “B” values. In 1.1 I provided an example of a *fuzzy lie*; let us consider a *graded-belief lie* now. Suppose Silvio utters:

(2) Ruby is Mubarak’s niece

If Silvio has a slight belief (e.g. a degree of confidence of 0.3) that Ruby is *not* Mubarak’s niece, it seems clear that Silvio is lying: his lie is a *graded belief-lie*. In this case, 0.3 expresses the value of “B”, so that 0.3 indicates Silvio’s subjective degree of confidence in the probability of (2). This case is different from the *fuzzy lie* example discussed in 1.1: in that case, 0.3 indicated the truth value of p . In

10. Graded beliefs are conscious beliefs. For a discussion of unconscious beliefs and their role in lying, see Stokke (2012).

11. There is no consensus whether outright beliefs (i.e. what we ascribe to a subject by saying that he believes that a propositional content is true) are part of this continuum or not (see e.g. Foley 1992; Frankish 2009). The present discussion of lying applies to both accounts of outright beliefs.

the fuzzy lie example, Mickey had an outright belief ($B=1$) that the truth value of p was 0.3.

All versions of the FIC allow for *fuzzy lies* but rule out *graded-belief lies*, and therefore are too narrow. Carson (2010, 17) defends a definition that encompasses this problem and allows for both kinds of lies: his alternative insincerity condition requires that the speaker *does not believe that his statement is true* (call this the *weakened insincerity condition*; henceforth WIC).

The WIC correctly rules out misleading statements (since they are believed-true statements), while ruling in traditional lies (since not-believed-true does not exclude believed-false). Moreover, it allows for garden-variety lies, *fuzzy lies* and *graded-belief lies*. Nevertheless, some philosophers (e.g. Saul 2012: 5) have argued that this condition is too broad, since it rules in what Harry Frankfurt (1986) calls ‘bullshit’.¹²

Broadly, bullshit refers to statements whose truthfulness the speaker has not assessed, so that bullshitting is generally distinct from lying. Consider the following example of deceptive bullshitting. Nick is a politician who does not know what the acronym LGBT refers to. When asked by a journalist about his opinion on LGBT rights, Nick answers:

- (3) LGBT rights are of central importance for our party

In uttering (3), Nick does not have the slightest idea whether what he said is true or false. His only concern is to trick the journalist into thinking that he knows what he is talking about. It seems clear that he is not lying, but the WIC counts incorrectly his statement as a lie, since Nick does not believe that his statement is true.

A way to respond to this objection is to narrow the WIC. We can distinguish two sets of not-believed-true statements that fall under the WIC. The first set of statements is what Frankfurt calls bullshit. These statements are not lies: they are not believed to be true because the speaker has not assessed the truth value of the asserted propositional content p , and therefore he lacks a defined belief about p .

The second set includes statements that are lies: they are not believed to be true because the speaker *does have* a belief about the truth value of p , but this belief is different from believing that p is true. This set includes cases in which the speaker believes p to be false, believes p to be not true, or has a graded belief about p .

12. Carson (2010: 17) acknowledges that his definition rules bullshit in, and he suggests the WIC to be a mere alternative to his “strong” insincerity condition: “Such statements are characteristic of bullshit. According to strong condition, this statement cannot possibly be a lie [...]. I don’t know of any decisive reason for preferring the stronger condition to the weaker condition or vice versa”.

The WIC can be narrowed to exclude the first set and include the second. It is sufficient to add an additional condition: “S has a belief about p ’s truth value, based on an assessment of p ”. Call this condition, jointly with the WIC, the revised WIC (henceforth, WIC*). The WIC* correctly rules out misleading statements and bullshit, while correctly ruling in traditional lies, *fuzzy lies*, and *graded-belief lies*.

The ability to correctly account for *graded-belief lies* is a major advantage for the WIC*. In real life, we often have partial beliefs rather than outright beliefs, and this circumstance does not seem to undermine our ability to lie. Consider again the statement:

(2) Ruby is Mubarak’s niece

If Silvio states (2) and believes that (2) is *probably* false, it seems clear that Silvio is lying: intuitively, *graded-belief lies* are lies. However, according to the TIC Silvio is not lying, since he does not have an outright belief that (2) is false. The TIC incorrectly rules out such instances of lying, while the WIC* correctly rules them in.

However, the WIC* is too broad, given that the set of degrees of belief it allows to count as lies is too large. For instance, it counts as lies statements that are believed to be *probably true*, since they are not outrightly believed to be true.

One way to settle the problem is to adhere to the second of Carson’s proposals (2010, 17), narrowing the WIC* and requiring that the speaker believe his statement to be at least *probably false*¹³ (call it the *probably-false insincerity condition*, henceforth PIC). This solution is appealing, since it rules out statements that are believed *more likely to be true than false*, while ruling in statements that are believed to be *false* or *probably false*. However, it is not clear why the assessment of probability indicated by “probably” should constitute the limit between what is a lie and what is not. It is easy to find counterexamples that suggest that such a limit is arbitrary.

For instance, consider the following case: Erwin asks Niels “Is my cat dead?” and Niels (knowing that Erwin loves his cat more than anyone else) answers “No”. Niels is uncertain about the cat being dead: he believes nonetheless that the cat is more likely to be dead than not (suppose he has a degree of belief of 0.6 in the fact that the cat is dead). Since Niels is inclined to think the opposite of what he said, there is ground to claim that Niels is lying. However, the PIC does not count these cases as

13. Generally, “probably” indicates an intermediate degree of belief between certainty and uncertainty (e.g. see Holmes 1984; Rubin 2010): I interpret Carson’s proposal in this sense. To be exact, his definition requires that the speaker believe his statement to be “false or probably false”. I preferred to use the wording “believed to be *at least* probably false” to clarify that all the intermediate degrees of belief between “believed to be false” and “believed to be probably false” are included.

instances of lying, since Niels's degree of confidence is not such that he believes that the cat is *probably* dead (if we assume that "probably" indicates a higher degree of confidence, like 0.75). The limit set by the PIC is therefore too narrow.

1.3 An intermediate position

A more promising way to settle the question is to require that the speaker believes his statement *more likely to be false than true*¹⁴ (call this the *comparative insincerity condition*, henceforth CIC). The CIC, like the PIC and the WIC*, correctly rules out misleading statements and bullshit, while correctly ruling in traditional lies, *fuzzy lies*, and *graded-belief lies*. Moreover, unlike the WIC*, it is not too broad, and unlike the PIC, it is not too narrow.

It seems that the CIC, (an intermediate position between the PIC and the WIC, both proposed by Carson) is not vulnerable to any other counterexample or objection. Therefore, the CIC represents the best non-dichotomic alternative to the TIC, and it can be embedded in both the classic and the alternative definition of lying to obtain a new definition.

In what follows, I will consider the CIC as embedded in the alternative definition.¹⁵ Let us call resultant definition the *scalar definition of lying*, because it stands against the dichotomic-discrete view, and it acknowledges that lying has many scalar dimensions of intensity.

Scalar definition

S lies to A about *p* iff:

- (i) S genuinely asserts *p*
- (ii) S believes *p* more likely to be false¹⁶ than true

14. One may want to quantify the required degree of belief by saying that the speaker must have a degree of belief in the falsity of his statement greater than 0.5. This view is compatible with my account of insincerity, but see Section 1.2 for some pessimistic remarks on the possibility to quantify numerically the precise threshold between sincerity and insincerity.

15. The CIC could as well be embedded in the *classic definition*. My choice is mainly due to the fact that in the next section assertoric force will be considered as a parameter for defining lying. While the alternative definition requires that a lie have the illocutionary force of an assertion, the classic definition vaguely requires that the speaker state a meaningful declarative sentence. The analysis that follows can be applied to both definitions, but the alternative definition allows us to avoid ambiguities and focus on the variations in the degree of assertoric force.

16. Using the expression "*more likely to be false*", I mean that the speaker must be more confident in the truth of *p* than in its falsity: this judgement must not necessarily be based on an assessment of probability.

At the beginning of this inquiry, I challenged the view that a liar must believe his statement to be outrightly false. The existence of *fuzzy lies* and *graded-belief lies* proved this view to be wrong, and revealed that lying is a scalar phenomenon. Yet our inquiry about the scalar nature of lying is not complete.

This section investigated how *beliefs* can cause insincerity (and therefore lying) to be graded. But insincerity also depends on how beliefs are communicated by means of a statement. The next section will investigate some linguistic devices that can influence the degree of certainty and the graded truth conditions that are communicated by a *statement* – thus providing further evidence to support the claim that lying is a scalar phenomenon, that allows for different degrees of intensity.

2. Layers of scalarity in insincere assertions

The philosophical literature on lying mostly endorses the dichotomic view, and thus focuses on statements that the speaker presents as outrightly believed to be true (and which are outrightly believed to be false). However, speakers often communicate from positions of subjective uncertainty, and employ several linguistic devices to express graded beliefs, or belief about graded truth values.

We know from our experience as ordinary language speakers that it is possible to modulate the intensity of a statement, either mitigating or reinforcing it. For instance, instead of simply uttering (4), a speaker can alternatively downgrade his assertion by uttering (4*) or emphasise it by uttering (4**):

- (4) Bobby is gorgeous
- (4*) Bobby is *kind of* gorgeous *somehow*
- (4**) *Believe me*, Bobby is *absolutely* gorgeous

In pragmatics, the two opposite phenomena of *mitigation* (4*) and *reinforcement* (4**) have often been studied separately, and labelled with different names: for the former, “attenuation”, “weakening” and “downgrading”; for the latter, “strengthening” and “emphasising”. The label of intensity (Holmes 1984; Labov 1984) unifies these two opposite directions of modulation. Intensity markers can operate on a number of components of the utterance: I will focus on how intensity markers affect the *propositional content* and the *illocutionary force*¹⁷ of

17. In traditional speech act theory, the illocutionary force is what characterises the utterance of *p* as being the occurrence of a specific kind of speech act (e.g. a question, an assertion, etc.). However, “along the same dimension of illocutionary point there may be varying degrees of strength or commitment” (Searle 1976, 5). Intensity markers can modify these degrees of

the speech act of asserting, and on the role they have in expressing graded truth values and graded beliefs.

2.1 Intensity and propositional content

The *propositional content* of a statement can be modulated both on the axis of quality (precision) and of quantity (augmentation or diminution). Expressions like “a little”, “very”, “much”, or “quite” are used to modify intensity on the axis of quantity. These linguistic devices determine a further scalar component of lying, because they allow the speaker to slightly alter the truth conditions of his statements (Lakoff 1975: 478–488). Thus, if Bruce utters (5*) rather than (5), he quantifies Robin’s gladness to a lower degree, thus altering the truth-conditions of his statement:

- (5) Robin is glad
 (5*) Robin is *pretty* glad

In Section 1.1, I considered *fuzzy lies*, *i.e.* lies that involve *beliefs* about graded truth values. Utterances like (5*) can express *beliefs* about graded truth values. This suggests that, with respect to *fuzzy lies*, we have to consider two scalar layers of insincerity: the layer of the speaker’s beliefs and the layer of the beliefs expressed by his statements. For instance, Bruce can tell a *fuzzy lie* either by “plainly” stating (5) while believing that (5) is partly false (e.g. 0.3-true) or by stating (5*) while believing that (5) is utterly false (that is, believing that (5*) is partly false).

2.2 Intensity and illocutionary force

Among the modulating devices that affect the *illocutionary force* of an utterance, *epistemic modality markers* are pivotal for the analysis of lying. Epistemic modality markers both “indicate the speaker’s confidence or lack of confidence in the truth of the proposition expressed” and “qualify [his] commitment to the truth of the proposition expressed in [his] utterance” (Coates 1987: 112). These two aspects are entangled, but they are two distinct phenomena, that I will present and discuss separately.

In the former sense, epistemic modality markers (henceforth, EMM) indicate the speaker’s confidence in the truth of the proposition expressed; in other words, they indicate a degree of belief. Thus, if George utters (6*) or (6**) instead of (6),

strength (Holmes 1984; Bazzanella, Caffi & Sbisà 1991; Sbisà 2001; Searle & Vanderveken 1985:99). For a discussion on the distinction between illocutionary and propositional mitigation, see Caffi (1999, 2007) and Fraser (2010: 16–17).

he communicates a degree of belief different from outright belief – respectively, probability in (6*) and certainty in (6**).

- (6) Saddam has weapons of mass destruction
- (6*) *Probably*, Saddam has weapons of mass destruction
- (6**) *Certainly*, Saddam has weapons of mass destruction

In Section 1.2, I considered *graded-belief lies*, *i.e.* lies that involve *graded beliefs* in unmarked statements. Statements like (6*) and (6**), similarly, are insincere, marked statements that *indicate* a graded belief. This suggests that, also with respect to *graded-belief lies*, we have to consider two scalar layers of insincerity: the layer of the speaker's beliefs and the layer of the beliefs communicated by his statement. For instance, George can tell a *graded-belief lie* both by plainly stating (6) while believing that (6) is probably false (e.g. holding a degree of confidence of 0.3 in its truth) and by stating (6*) while having an outright belief that (6) is false (that is, having a graded belief that (6*) is false).

Not only do EMM convey information about the speaker's degree of belief in what is said; they also alter the speaker's *commitment* to the truth of the proposition (e.g. Lyons 1977: 793–809; Holmes 1984: 349; Bazzanella 2004: 184). Modulating intensity by means of EMM, the speaker can diminish or augment his commitment to the truth of the proposition. In this sense, EMM alter the *illocutionary force* of the utterance, in addition to the degree of belief it conveys.

These two functions of EMM are entangled: if a speaker S mitigates (or reinforces) the *degree of belief* conveyed by his utterance (as in (6*) and (6**)), then S automatically mitigates (or reinforces) the illocutionary force of his utterance (that is, the degree of his *commitment* to the truth of the proposition).¹⁸

2.3 The role of illocutionary force in defining lying

On this account of EMM, insincerity is determined by the *degree of belief* expressed by the EMM (compared with the speaker's credal state), and it is unrelated to the *degree of commitment* determined by the EMM. The degree of commitment is an independent, further scalar component of lying, that is unrelated to insincerity.

18. On a standard account of assertion, the illocutionary effect of an assertion is to commit the speaker to the truth of what he says (Searle 1969; Brandom 1983, 1994). Therefore, variations in the degree of illocutionary force correspond to variation in the speaker's commitment to what is said (Holmes 1984: 349; Sbisà 2001: 1805–06). On this view, EMM are not part of the propositional content – *i.e.* they are used to modulate the speaker's commitment to the *unmarked* sentence (Holmes 1984: 352). For instance, uttering (6*) instead of (6), George strengthens his commitment to (6), rather than simply committing himself to (6*). For a different account of EMM, cfr. Footnote 21.

Even if variations in the degree of commitment do not influence the degree of insincerity, one might ask: can variations in the degree of commitment affect our evaluation of an insincere utterance as a lie? In other words, are there any *insincere* utterances that do not count as lies because their degree of assertoric force is too weak?

A positive answer to this question might be supported by appealing to condition (i) of the alternative definition of lying: in order to lie, the speaker must genuinely assert that p – that is, he must genuinely *commit* himself to the truth of p . And one might argue that when the speaker uses EMM to mitigate the illocutionary force of his assertion up to a certain extent, his utterance does not count as a genuine assertion any more.¹⁹ For instance, suppose that Pablo is certain that it is not raining in Madrid and tells Carmen:

(7) *Maybe* it is raining in Madrid

Intuitively, (7) is not a lie: Pablo is not committing himself to defend the claim that it is raining in Madrid – he has only put forward a hypothesis. To test this intuition, imagine that eventually Carmen discovers that Pablo believed that it was not raining in Madrid. Intuitively, if one lies, the addressee should be able to blame the liar (e.g. Faulkner 2007, 538). But if Carmen blames Pablo for his utterance, it would be easy for Pablo to reply that he has only suggested, or hypothesised, that it *may have been* so. This line of reply is not available to a liar – on the contrary, it resembles the line of defence that is available to a speaker that uttered a deceptive truth or a misleading statement. And in fact, on this account (7) counts as a misleading statement rather than a lie. Thus, the answer to the initial question is positive: there are indeed *insincere* utterances that do not count as lies because their degree of illocutionary force is too weak – even if these utterances are likely to be, nonetheless, deceiving.

Another issue about illocutionary force and insincerity is that EMM can express a high degree of *insincerity* and at the same time a low degree of *commitment*, or vice versa. This is because the degree of insincerity depends on the graded beliefs of the speaker, while illocutionary force does not. One might be worried that this disparity might lead to problematic cases, in which a downgraded commitment is compensated by a high degree of insincerity – or in which a low degree of insincerity is compensated by a strong degree of commitment. For

19. Carson (2010:33) seems to endorse this view: “There are weaker and stronger ways of warranting the truth of a statement. To count as a lie, a statement must be warranted to a certain minimum degree”. More generally, Searle and Vanderveken (1985:15) assume that any type of speech act has a characteristic degree of illocutionary force that needs to be achieved in order for that type of speech act to be performed.

instance, suppose that Pablo is uncertain about whether it is raining or not in Madrid, and he utters:

(8) It is *definitely* raining in Madrid

In this case, the degree of insincerity of Pablo's utterance is less than it is required by the scalar definition (to believe p more likely to be false than true) to count (8) as a lie. However, one might argue that, in committing himself so strongly and explicitly to an utterance he does not utterly believe, Pablo "compensates" for the low degree of insincerity, and therefore is lying.

Yet, this interpretation is hardly convincing. In uttering (8), Pablo is surely violating some norm (like the Gricean norm not to say that for which you lack adequate evidence) (Grice 1989:27). Moreover, in uttering (8) Pablo is behaving as an unreliable epistemic source: he is inviting us to trust him even if he has no ground to assert what he says. Nevertheless, it seems that Pablo is bullshitting, or guessing, rather than lying. His utterance might be misleading, but it is not insincere and does not count as a lie.

However, it is possible to exploit the polysemic and polypragmatic nature of EMM (cfr. Coates 1987; Fraser 2010) to argue that, in a sense, (8) is a lie. For instance, if one takes the meaning of "definitely" to be something like "there is accessible, undefeatable evidence" (that it is raining in Madrid), then (8) is a lie, because Pablo does not believe that there is such evidence. It is not clear whether one can interpret utterances like (8) in this way;²⁰ but if one accepts

20. There is disagreement about whether or not EMM are part of the propositional content. On my (call it *pragmatic*) account, EMM are not part of the propositional content: they indicate, rather than refer to, a degree of belief. In interpreting (8), I considered the opposite hypothesis: that EMM alter the truth-conditions of the utterance, and are part of the very meaning of statements that involve EMM (call this the *semantic account*). According to this view, (1) is true iff (1*) is true.

- (1) Certainly it is raining in England
 (1*) I am certain that it is raining in England

The advantage of the semantic account is that it orthodoxly assumes that the credal state relevant for lying has for its object the propositional content of the statement. By contrast, the pragmatic account unorthodoxly assumes that the credal state relevant for lying has for object the degree of belief expressed by the statement. The semantic account does not undermine the basic idea that the degree of insincerity is related to the degree of belief expressed by EMM, but it offers a different interpretation of the beliefs expressed by EMM (thus offering different responses to the examples analysed in Sections 2.1–2.3).

This paper refuses to adopt the semantic account of EMM on the ground of three reasons. First, and most importantly, the semantic account counts all misuses of EMM as

this interpretation, then (8) is a lie because it is an insincere utterance *tout court*, rather than because the speaker has strengthened an assertion he is uncertain about. Therefore, even if this interpretation is proven to be right, it does not support the idea that a stronger assertoric force can compensate for a lack of insincerity.

The hypothesis that insincerity and commitment can compensate for each other is thus to be rejected. While (8) shows that it is not possible to compensate a lack of insincerity by strengthening the commitment, (7) shows that it is not possible to compensate a lack of commitment by strengthening insincerity.

To sum up, there are three entwined scalar parameters that we have to consider when we analyse an utterance: the *graded truth value* expressed by the utterance, the *graded belief* it communicates (these two parameters determine the degree of insincerity), and the *degree of assertoric force*, or degree of commitment (this parameter determines if the speaker is asserting, and how strongly). For an utterance to count as a lie, both a certain degree of insincerity *and* a certain degree of commitment must obtain: a surplus of insincerity cannot compensate for a lack of commitment, and vice versa.

2.4 Further dimensions of scalarity

EMM are not the only way to modulate illocutionary force. This section explores which devices other than EMM can modify illocutionary force.

lies – so that, at best, it is not apt to discuss lying. For instance, according to the semantic account, if S utters (1) and he believes that (1) is probably true, S is lying (because (1) means (1*)). Second, it is not clear how EMM should be interpreted, and countless incompatible semantic accounts of EMM have been presented in the philosophical literature (e.g. Hamblin 1959; Kratzer 1981; DeRose 1991; Egan, Hawthorne & Weatherson 2005). But it is implausible to contend that EMM fix the truth conditions of statement like (1), if there is no agreement on what are the truth conditions of statement like (1). Third, there are several reasons to contend that no ordinary truth-conditions can be provided for EMM (e.g. Frege 1879, 5; Price 1983; Yalcin 2010; Swanson 2011 – but see Papafragou 2006). The pragmatic interpretation of EMM that I endorse acknowledges that EMM are polysemic and polypragmatic in nature: they are strongly context-dependent (see Kratzer 1981) and can be used to achieve simultaneously many semantic and pragmatic communicative goals (Holmes 1984; Coates 1987; Bazzanella, Caffi & Sbisà 1991; Fraser 2010). This account does not exclude the possibility that, in some cases, we can interpret EMM as altering the truth-conditions of the utterance in which they occur. But this interpretation is not to be intended as the standard interpretation; rather, it amounts to one of the possible ways to interpret an ambiguous term that can serve many communicative goals.

So far I have considered aspects of intensity that are properly linguistic. Several paralinguistic devices can intervene to modulate the illocutionary force of an assertion. A very significant factor is *prosody*, and specifically *intonation*: Gussenhoven (2002) remarks that variations in volume and tonal height are a powerful device to communicate the speakers' confidence in the truth of his assertion. Among other influential factors that can be used to influence the strength of an utterance are pauses, rhythm of speech, repetitions and proxemic signals.

Context is also an important factor to be considered. Bazzanella (2009:78) points out that, among the parameters that determine the scalar graduation of lying, also "aspects of the global context [...], of the cotext [...] and of the local context" are of central importance. For instance, an utterance like (9) can express a different degree of certainty whether (9) is uttered in (context A) a restaurant, where a person addresses it to a commensal that is known not to like cheese, or (context B) in a hospital, where a doctor addresses it to a person that he knows to have a deadly allergy to cheese. In both cases, given that (9) is insincere, (9) is a lie, but it is more intense in the first case, since the commitment to the truth of the assertion is stronger in context A than in context B.

(9) This meal contains no cheese

The very mitigating and reinforcing devices considered so far are influenced by contexts – in most cases, their very meaning is determined by contextual factors (Kratzer 1981). For instance, the EMM "definitely" in (10) expresses a different degree of certainty and commitment depending on whether (10) is uttered in context A or in context B.

(10) Definitely, this meal contains no cheese

Linguistic, paralinguistic and contextual elements determine a complex interplay of factors that influences the degree of illocutionary force of an assertion, and thus the degree of intensity of a lie. However, it seems to me that only linguistic devices (like EMM) are strong enough to affect the status of an utterance as an assertion (and therefore as a lie). Paralinguistic and contextual elements can reinforce or mitigate an assertion, but they alone cannot determine if an utterance counts as an assertion or not.

In Table 1, I summarise my taxonomy of the layers of scalarity displayed by the phenomenon of lying. This taxonomy is not meant to be exhaustive, but it succeeds in underlining a fact that has been ignored by the philosophical literature on lying for a long time: the multi-layered scalarity of lying.

Table 1. A taxonomy of the dimensions of scalarity involved in lying

Scalarity of <i>beliefs</i>	Beliefs about graded truth values Graded beliefs about truth values
Scalarity of <i>statements</i>	Stating “beliefs about graded truth values” Stating “graded beliefs about truth values” Modifying illocutionary force
Scalarity of <i>paralinguistic components</i>	Modifying intensity of illocutionary force
Scalarity of <i>contexts</i>	Modifying intensity of illocutionary force

3. Conclusions

The scalarity of lying results from several parameters: graded beliefs and beliefs about graded truth values, as much as the numerous ways to convey information about them in statements. Both the classic and the alternative definition of lying completely ignore these scalar features in assuming the dichotomic view as a starting point – a view that I proved to be inconsistent with our intuitions.

I have provided several reasons to believe that the traditional insincerity condition is also wrong, because it rules out *fuzzy lies* and *graded-belief lies*. Moreover, I have shown that the traditional insincerity condition yields a wrong *description* of lying, as it blinds us to its scalar nature. My proposed definition rules in *fuzzy lies* and *graded-belief lies*, thus allowing for many degrees of intensity in lying. Moreover, it acknowledges that *illocutionary force* is scalar, and that it can be modified by a number of linguistic and paralinguistic devices.

In this paper, I have only considered the scalar dimensions of lying that define the boundaries of the concept. Also other components of lying, nevertheless, have scalar features. The picture is complex: as Bazzanella (2009: 78) points out, “the different degrees of intensity in lying result from the complex interplay of various layers and parameters”. Some philosophers (Chisholm & Feehan 1977; Fallis 2011; Staffel 2011) contend that the intention to deceive, as well as the effects of deception, can be graded. Several pragmatic parameters determinant for lying also have scalar features, such as relevance (Wilson & Sperber 2002; Van der Henst et al. 2002), felicity, and the relations between the interactants (like social relations, and respective trust) (Bazzanella 2009). Finally, moral evaluations of lying are clearly scalar, depending (amongst other things) on the moral reprehensibility of the effects of the lie, and of the intentions of the liar.

Now that I have identified and systematised the scalar parameters that define lying, further investigation is needed to understand how they interact with these

other scalar components. A better understanding of how these parameters work together and determine various degrees of intensity in lying may be very useful not only for the philosophical understanding of lying, but also for a number of other fields of inquiry – as epistemology, ethics, forensics, and other applied disciplines.

References

- Augustine. 1887 [DM]. De mendacio (On lying). In *Nicene and Post-Nicene Fathers* [First Series, Vol. 3], Philip Schaff (ed.). Buffalo NY: Christian Literature Publishing.
- Barnes, John A. 1994. *A Pack of Lies: Towards A Sociology of Lying*. Cambridge: CUP.
DOI: 10.1017/CBO9780511520983
- Bazzanella, Carla, Caffi, Claudia & Sbisà, Marina. 1991. Scalar dimensions of illocutionary force. In *Speech Acts: Fiction or Reality*, Igor Ž. Žagar (ed.), 63–76. Antwerp & Ljubljana: IPRA Distribution Center for Yugoslavia and Institute for Social Sciences.
- Bazzanella, Carla. 2004. Atténuation et intensification en Italien: Dimensions et configuration pragmatique (Mitigation and intensification in Italian: Dimensions and pragmatic configuration). In *Plus ou moins!? L'atténuation et l'intensification dans les langues romanes*, Maria H. Araujo Carreira (ed.), 173–200. Paris: Université Paris 8, Vincennes Saint-Denis.
- Bazzanella, Carla. 2009. Approssimazioni pragmatiche e testuali alla menzogna. (Pragmatic and textual approaches to the concept of lying). In *Tra Pragmatica e Linguistica testuale. Ricordando Maria-Elisabeth Conte*, Venier Federica (ed.), 67–90. Dell'Orso: Alessandria. (Translation of the excerpts is mine).
- Bok, Sissela. 1978. *Lying: Moral Choice in Public and Private Life*. New York NY: Random House.
- Brandom, Robert. 1983. Asserting. *Nous* 17(4): 637–650. DOI: 10.2307/2215086
- Brandom, Robert. 1994. *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge MA: Harvard University Press.
- Caffi, Claudia. 1999. On mitigation. *Journal of Pragmatics* 31: 881–909.
DOI: 10.1016/S0378-2166(98)00098-8
- Caffi, Claudia. 2007. *Mitigation*. London: Elsevier.
- Carson, Thomas L. 2006. The definition of lying. *Nous* 2: 284–306.
DOI: 10.1111/j.0029-4624.2006.00610.x
- Carson, Thomas L. 2010. *Lying and Deception*. Oxford: OUP.
DOI: 10.1093/acprof:oso/9780199577415.001.0001
- Chisholm, Roderick M. & Feehan, Thomas D. 1977. The intent to deceive. *Journal of Philosophy* 74(3): 143–159. DOI: 10.2307/2025605
- Coates, Jennifer. 1987. Epistemic modality and spoken discourse. *Transactions of the Philological Society* 85(1): 110–131. DOI: 10.1111/j.1467-968X.1987.tb00714.x
- D'Agostini, Franca. 2012. *Menzogna* (lying). Torino: Bollati Boringhieri.
- DeRose, Keith. 1991. Epistemic possibilities. *The Philosophical Review* 100(4): 581–605.
DOI: 10.2307/2185175
- Egan, Andy, Hawthorne, John & Weatherson, Brian. 2005. Epistemic modals in context. In *Contextualism in Philosophy: Knowledge, Meaning and Truth*, Gerhard Preyer & Georg Peter (eds), 1–48. Oxford: Clarendon Press.
- Fallis, Don. 2009. What is lying? *The Journal of Philosophy* 106(1): 29–56.

- Fallis, Don. 2010. Lying and deception. *Philosophers' Imprint* 10: 1–22.
- Fallis, Don. 2011. What liars can tell us about the knowledge norm of practical reasoning. *The Southern Journal of Philosophy* 49(4): 347–367. DOI: 10.1111/j.2041-6962.2011.00078.x
- Fallis, Don. 2012. Lying as a violation of Grice's first maxim of quality. *Dialectica* 66(4): 563–581. DOI: 10.1111/1746-8361.12007
- Faulkner, Paul. 2007. What is wrong with lying? *Philosophy and Phenomenological Research* LXXV(3): 535–557. DOI: 10.1111/j.1933-1592.2007.00092.x
- Foley, Richard. 1992. The epistemology of belief and the epistemology of degrees of belief. *American Philosophical Quarterly* 29(2): 111–124.
- Frankfurt, Harry G. 1986. On Bullshit. *Raritan Quarterly Review* 6:2.
- Fraser, Bruce. 2010. Pragmatic competence: The case of hedging. In *New Approaches to Hedging*, Gunther Kaltenböck, Wiltrud Mihatsch & Stefan Schneider (eds), 15–34. Bingley: Emerald. DOI: 10.1163/9789004253247_003
- Frege, Gottlob. 1879. Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens (Concept Script, a formal language of pure thought modelled upon that of arithmetic). Halle: L. Nebert. In *From Frege to Gödel: A Source Book in Mathematical Logic*, Jan van Heijenoort (eds). Cambridge MA: Harvard University Press.
- Frankish, Keith. 2009. Partial belief and flat-out belief. In *Degrees of Belief*, Franz Huber & Christoph Schmidt-Petri (eds), 75–93. Berlin: Springer. DOI: 10.1007/978-1-4020-9198-8_4
- Gottwald, Siegfried. 2001. *A treatise on Many-valued Logics*. Baldock: Research Studies Press.
- Grice, Herbert P. 1989. *Studies in the Way of Words*. Cambridge MA: Harvard University Press.
- Gussenhoven, Carlos. 2002. Intonation and interpretation: Phonetics and phonology. In *Proceedings of the Speech Prosody International Conference*, 47–57. Aix en Provence: Laboratoire Parole et Langage.
- Hájek, Petr. 1998. *Metamathematics of Fuzzy Logic*. Dordrecht: Kluwer. DOI: 10.1007/978-94-011-5300-3
- Holmes, Janet. 1984. Modifying illocutionary force. *Journal of Pragmatics* 8: 345–365. DOI: 10.1016/0378-2166(84)90028-6
- Hamblin, Charles L. 1959. The modal 'probably'. *Mind* 68(270): 234–240. DOI: 10.1093/mind/LXVIII.270.234
- Kratzer, Angelika. 1981. The notional category of modality. In *Worlds, Words, and Contexts*, Hans-Jürgen Eikmeyer & Hannes Rieser (eds), 38–74. Berlin: De Gruyter.
- Kupfer, Joseph. 1982. The moral presumption against lying. *The Review of Metaphysics* 36(1): 103–126.
- Labov, William. 1984. Intensity. In *Georgetown University Round Table on Language and Linguistics*, Deborah Schiffrin (ed.), 43–70. Washington DC: Georgetown University Press.
- Lakoff, George. 1975. Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic* 2(4): 458–508.
- Lyons, John. 1977. *Semantics*. Cambridge: CUP.
- Macintyre, Alasdair. 1994. Truthfulness, lies, and moral philosophers: What can we learn from Mill and Kant? *The Tanner Lectures on Human Values* 16: 307–61.
- Mahon, James E. 2008. The definition of lying and deception. In *The Stanford Encyclopedia of Philosophy*, Fall 2008 edn, Edward N. Zalta (ed.), (<http://plato.stanford.edu/archives/fall2008/entries/lying-definition/>)
- Mannison, Don S. 1969. Lying and lies. *Australasian Journal of Philosophy* 2(47): 132–144. DOI: 10.1080/00048406912341141
- Meibauer, Jörg. 2005. Lying and falsely implicating. *Journal of Pragmatics* 37(9): 1373–1399. DOI: 10.1016/j.pragma.2004.12.007

- Meibauer, Jörg. 2011. On lying: Intentionality, implicature, and imprecision. *Intercultural Pragmatics* 2(8): 277–292.
- Meibauer, Jörg. 2014. *Lying at the semantic-pragmatic interface*. Berlin: De Gruyter.
- Papafragou, Anna. 2006. Epistemic modality and truth conditions. *Lingua* 116(10): 1688–1702. DOI: 10.1016/j.lingua.2005.05.009
- Price, Huw. 1983. Does ‘probably’ modify sense? *Australasian Journal of Philosophy* 61: 37–41. DOI: 10.1080/00048408312341151
- Rubin, Victoria L. 2010. Epistemic modality: From uncertainty to certainty in the context of information seeking as interactions with texts. *Information Processing & Management* 46(5): 533–540. DOI: 10.1016/j.ipm.2010.02.006
- Saul, Jennifer M. 2012. *Lying, Misleading, and what is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford: OUP. DOI: 10.1093/acprof:oso/9780199603688.001.0001
- Sbisà, Marina. 2001. Illocutionary force and degrees of strength in language use. *Journal of Pragmatics* 33: 1791–1814. DOI: 10.1016/S0378-2166(00)00060-6
- Searle, John R. 1969. *Speech Acts. An Essay in the Philosophy of Language*. Cambridge: CUP. DOI: 10.1017/CBO9781139173438
- Searle, John R. 1976. A classification of illocutionary acts. *Language in Society* 5(1): 1–23. DOI: 10.1017/S0047404500006837
- Searle, John R. & Vanderveken, Daniel. 1985. *Foundations of Illocutionary Logic*. Cambridge: CUP.
- Sorensen, Roy. 2007. Bald faced lies! Lying without the intent to deceive. *Pacific Philosophical Quarterly* 88: 251–264. DOI: 10.1111/j.1468-0114.2007.00290.x
- Staffel, Julia. 2011. Reply to Sorensen, ‘knowledge-lies’. *Analysis* 71: 300–303. DOI: 10.1093/analys/anr021
- Stokke, Andreas. 2012. Insincerity. *Noûs*. DOI: 10.1111/nous.12001
- Stokke, Andreas. 2013. Lying, deceiving, and misleading. *Philosophy Compass* 8: 348–359. DOI: 10.1111/phc3.12022
- Stokke, Andreas. 2013b. Lying and asserting. *Journal of Philosophy* 110(1): 33–60.
- Swanson, Eric. 2011. How not to theorize about the language of subjective uncertainty. In *Epistemic Modality*, Andy Egan & Brian Weatherson (eds), 249–269. Oxford: OUP. DOI: 10.1093/acprof:oso/9780199591596.003.0009
- Van der Henst, Jean-Baptiste, Carles, Laure & Sperber, Dan. 2002. Truthfulness and relevance in telling the time. *Mind and Language* 17(5): 457–466. DOI: 10.1111/1468-0017.00207
- Williams, Bernard. 2002. *Truth and Truthfulness: An Essay in Genealogy*. Princeton NJ: Princeton University Press.
- Wilson, Deirdre & Sperber, Dan. 2002. Truthfulness and relevance. *Mind* 25: 1–41.
- Yalcin, Seth. 2010. Probability operators. *Philosophy Compass* 5(11): 916–937. DOI: 10.1111/j.1747-9991.2010.00360.x