Lying and Certainty

Neri Marsili

University of Sheffield

n.marsili@sheffield.ac.uk

ABSTRACT

In the philosophical literature on the definition of lying, the analysis is generally restricted to cases of flat-out belief. This chapter considers the complex phenomenon of lies involving partial beliefs – beliefs ranging from mere uncertainty to absolute certainty. The first section analyses lies uttered while holding a graded belief in the falsity of the assertion, and presents a revised insincerity condition, requiring that the liar believes the assertion to be more likely to be false than true. The second section analyses assertions that express graded beliefs, exploring how mitigation and reinforcement can alter the insincerity conditions for lying. The last section considers the case of lies that attack certainty (*knowledge-lies*), understood as attempts to alter the hearer's graded beliefs.

Introduction

A prominent characteristic of lies is that they come in a variety of forms and kinds, and this is part of what makes them elusive and difficult to identify. As Montaigne nicely stated, while truth is unique, "the opposite of truth has many shapes, and an indefinite field" (Montaigne *E:*1.IX). One of the ambitions of this book is to categorise these different shapes, and to provide systematic criteria to distinguish lies from other utterances.

This can be a hard challenge, and indeed for a theory of lying "the more difficult task [is] that of drawing lines" (Bok 1989:49). There is a whole grey area of deceptive utterances that are difficult to classify (for instance, the so-called "half-truths") and, quite importantly, it is in this grey zone that liars strive. To shed some light in this obscure area, this chapter will consider the problem of classifying statements that are not fully believed to be false, but that are nevertheless not believed to be true. Are these statements lies? And how much confidence in their falsity is required for them to count as lies? We will focus on such questions, exploring the thin, elusive line that distinguishes a sincere assertion from an insincere one.

On a standard view, an utterance is a lie only if (1) it is an *assertion* and (2) the speaker believes it to be false[1]. However, the expression "believe to be false" is not really helpful to deal with intermediate cases, as it does not specify which degree of confidence in the falsity of *p* counts as believing to be false. As we analyse this issue, we will see that lies are difficult to categorise because lying is often a matter of degree (Bazzanella 2009, Isenberg 1964:470). This is not because lying is a scalar predicate (one cannot say that p is more of a lie than q), but rather because being sincere is: a speaker can be more or less sincere – so that, indirectly, a lie can be a more or less severe violation of the linguistic (and moral) norm of sincerity.

1. Speaker certainty

1.1. Certainty, uncertainty and graded beliefs

In this world nothing can be said to be certain, except death and taxes
Benjamin Franklin (1789)

What is certainty? Most authors agree that it is a kind of attitude that a subject can have towards a proposition. More specifically, certainty (or what philosophers call 'psychological' certainty[2]) can be defined as the *highest degree of confidence* that a subject can have in the truth of a proposition. Thus understood, certainty is always relative to *someone's standpoint*: it does not matter if the subject has no ground (or bad grounds) for holding that belief, because certainty only requires that *the subject* be supremely convinced of its truth.

This conception of certainty, defined as the highest degree of belief a speaker can have in a proposition, presupposes that believing comes in degrees – and that there are degrees of belief lower than certainty. That beliefs can be graded is evident if one thinks about daily situations in which a subject lacks certainty in a proposition that he nonetheless, to some extent, believes. Too see this, consider some examples: suppose Groucho believes that (i) he

---

[1] Some authors require a further condition: that the speaker intends to deceive the hearer (e.g. Mannison 1969: 132; Kupfer 1982:134; Williams 2002:96). For the purpose of this chapter, I will assume that while prototypical lies are intended to deceive, this intention is not a necessary condition for lying.

[2] Philosophers often distinguish *psychological* certainty from *epistemic* certainty (Klein 1998, Reid 2008, Stainley 2008). Epistemic certainty refers to the *degree of epistemic warrant* that a proposition has, independently of the speaker's confidence in it (*i.e.* independently of psychological certainty). While psychological certainty is purely 'subjective' (it only depends on the subject's confidence), epistemic certainty is in a sense 'objective' (it depends on the *actual solidity* of the subject's reasons to believe in that proposition). The literature on lying is generally concerned with *psychological certainty*, since the strength of the speaker's grounds for disbelieving an assertion is irrelevant to assess whether he is insincere or not. Consequently, in this chapter, "certainty" (and "uncertainty") will refer to *psychological* certainty (and uncertainty).

has a pair of moustaches, (ii) Bulgaria will beat Azerbaijan in their next football match, and (iii) there is life on some other planet in the universe. Groucho regards (i) as certain, (ii) as probable, (iii) as merely more likely true than not.

Groucho neither fully believes nor fully disbelieves (ii) or (iii). These intermediate, partial beliefs (believing to be probable, believing to be unlikely, etc.) are called in the literature "graded beliefs", because they can be ordered in a graded scale[3]: Groucho is more confident in the truth of (i) than he is in (ii), and in (ii) than he is in (iii). Formal accounts of degrees of belief (namely Bayesian accounts) represent this scale with real numbers from 0 to 1, where 0 indicates certainty in the falsity of p, 1 indicates certainty in the truth of p, and 0.5 indicates uncertainty (cases in which the subject regards p as just as likely to be true as to be false). On this view, uncertainty is the middle point (0.5) of a continuum of degrees of belief whose poles are certainty in falsity (0) and in truth (1) of the proposition (see Figure 1). To provide a formal account of the previous example, one could say that Groucho has degree of belief of 1 in (i), of 0.75 in (ii), of 0.51 in (iii).
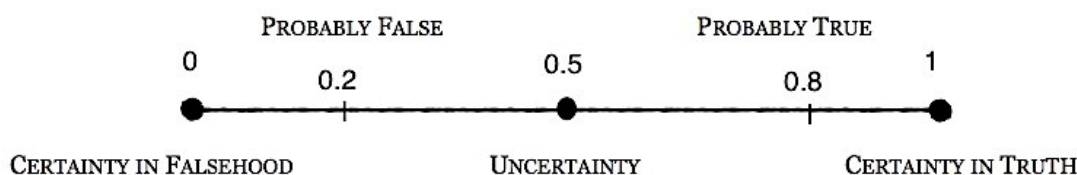

Figure 1: A visual representation of the certainty-uncertainty continuum

Any account of sincerity that takes into consideration this wide array of graded beliefs will have to explain how they relate to the boundaries of the concept of lying. For instance, suppose that Groucho states that Bulgaria will beat Azerbaijan while believing that it is probably false, or as likely to be false as true. Would these utterances be sincere or insincere? And more generally, how are we to draw the boundary between sincere and insincere utterances, and (consequently) between lies and not lies?

1.2. Insincerity and (un)certainty

What constitutes the essence of lying,
if not saying what you don't have in your heart?
Rosmini, Apologetica, 7-131

---

[3] For a discussion of the mutual relations between flat-out beliefs and graded beliefs, see Frankish (2009).

A necessary condition for lying is that the speaker utters an *insincere* statement. There are several ways to define insincerity. In a general sense, insincerity refers to a *discrepancy between the psychological state of the speaker* (e.g. believing, intending, desiring) *and the psychological state expressed by his linguistic action* (e.g. asserting, promising, requiring) (Falkenberg 1988:94, Searle & Vanderveken 2005:111). Defining 'insincerity' amounts to defining the nature of this discrepancy.

Limiting the present discussion to assertion[4], and accepting the standard view that assertions express beliefs (Searle 1969: 65), this chapter will analyse insincerity as a discrepancy *between what is asserted* and *what is believed* by the speaker. Once graded beliefs are taken into account, the main challenge is to define how large this discrepancy has to be for a statement to count as insincere – and hence as a lie.

In the philosophical literature, this challenge is generally overlooked or ignored. It is taken as an uncontroversial claim that a statement is *sincere* when the speaker believes it to be *true,* and *insincere* when the speaker believes it to be *false*[5], and that a more fine-grained analysis would be unnecessarily intricate (Saul 2012:5, fn10). The standard "insincerity condition" for lying (IC) is generally phrased as follows:

IC = S believes p to be false

This condition correctly accounts for prototypical cases of lying, but is ill-suited to consider cases of uncertainty and graded beliefs. Since graded beliefs and uncertainty are ordinary psychological states, it seems that a theory of lying should account for them (Meibauer 2014: 223, D'Agostini 2012:41, Isenberg 1964: 468).

To see that the standard account of insincerity struggles to handle graded beliefs in a satisfactory way, consider the following example, inspired by historical events (cf. Carson

---

[4] Arguably, assertion is not the only speech act that can be used to lie. However, it seems that non-assertive speech acts can be used to lie only insofar as they *entail* an assertion (see Marsili 2016), so that an analysis of assertion is sufficient for our purposes. I am here referring to *illocutionary entailment* (Searle & Vanderveken 1985:130), where F1 *illocutionary entails* F2 iff a speaker cannot perform F1 without also performing F2 (e.g. you cannot promise that you will *p* without also asserting that you will *p*).

[5] A third option is that the speaker has *no opinion about p* (he *lacks a credal state about p*); I will come back on this in the next section.

2010:212-21): George is a political leader, and tells (1) to a journalist. Propositions (a), (b), and (c) indicate George's degree of confidence in his utterance, in three possible scenarios[6].

(1) Iraq has weapons of mass destruction

    (a)   (1/¬p)   [Iraq has *certainly no* weapons of mass destruction]

    (b)   (0.75/p)   [*Probably*, Iraq has weapons of mass destruction]

    (c)   (0.75/¬p)   [*Probably*, Iraq *does not* have weapons of mass destruction]

Scenario (1a) is a clear-cut case of lying, since George believes (1) to be *certainly false:* the standard account correctly tracks the intuition that this is a lie. In (1b), by contrast, George believes the statement to be *probably true:* even if he is not completely confident that the statement is true, it seems that in this case he is not lying (Austin 1946:65). The utterance is inaccurate, and perhaps misleading, because it misrepresents George's degree of belief in (1). However, being inaccurate or misleading is clearly not the same as lying (Saul 2012, Stokke 2013): condition IC is again on the right track, since it predicts that this utterance is not a lie. Problems arise for scenario (1c), where George believes (1) to be *probably false*. It seems that condition IC does not count it as a lie, because George does not utterly believe (1) to be false[7]. However, intuitively this is a case of lying, because George is saying something he believes to be very likely false. Since it excludes this sort of cases, (IC) is too narrow, and needs some refinement.

Cases like (1b,c) suggests that a more fine-grained account of lying is needed, one that appreciates how lying can involve graded beliefs. Such a definition of lying should offer an account of what we might call '*graded-belief lies*': statements that are not outrightly believed to be false (nor true), but are nonetheless lies. The next section will review a few attempts to revise the definition of lying in this direction[8].

---

[6] Assigning a defined, numeric degree of belief to these linguistic expressions (e.g "probably", "perhaps") merely aims to indicate how these expressions can be ordered on a scale that goes from certainty to doubt (Holmes 1982, Levinson 1983:134, Hoye 1997). Only their reciprocal relation in the scale matters to the present discussion – the accuracy of the numeric values is not important.

[7] To save (IC) against this objection, a partisan of the standard view might suggest to interpret (IC) in a non-literal sense, so that (2) counts as a case of believing p to be false, and hence as lying. However, this broad interpretation would open the problem of which intermediate credal states count as believing false and which do not. Since this is exactly the problem that the sincerity condition should solve, (IC) would still be an unattractive option to settle the issue.

[8] Further complications for a definition of insincerity, such as cases of *self-deception* (in which the speaker is mistaken about his own beliefs or mental states, cf. Moran 2005, Erikssonn 2011, Chan & Kakane 2011, Stokke 2014) and *malapropism* (in which the speaker is mistaken about what he said, cf. Reimer 2004, Sorensen 2011, Saul 2012:15-9) will not be discussed here; for our purposes, to deal with these cases it is sufficient to require that the speaker satisfy (any version of) the insincerity conditions for lying *advertently*.

1.3. Degrees of (dis)believing

Carson (2006: 298) has offered a definition of lying that captures graded-belief lies. His proposal presents a strong and a weak version of the "insincerity condition" for lying. The first, "strong" version requires that the speaker believe his assertion to be "false or probably false". Let us call Carson's first condition the "strong insincerity condition" for lying:

SIC= S believes p to be at least *probably* false[9]

SIC correctly captures prototypical cases of lying like (1a). Unlike the traditional definition, it also includes lies that are not believed with certainty to be false, like (1c), that George believes to be *probably false*. This is an advantage of SIC over IC, since it seems intuitive that saying what you believe to be *probably* false counts as lying – even if it is arguably less insincere, and less deceptive, than a full-fledged lie.

However, it is not clear that the boundary between sincerity and insincerity lies exactly on the degree of confidence indicated by 'probably', and not on another. The term 'probably' indicates a degree of confidence in the truth (or falsity) of the proposition higher than uncertainty and lower than certainty: for the sake of the argument, let us assume it stands for a degree of belief of 0.75. If a degree of belief of 0.75 in the falsity of the proposition is enough for lying, there seems to be no reason to exclude lower graded beliefs like 0.7, or 0.6, that are perceivably higher than uncertainty (0.5). To see this, consider the following scenarios:

(1) Iraq has weapons of mass destruction

    (c)   (0.75/¬p) [*Probably* Iraq *does not* have weapons of mass destruction]

    (d)   (0.6/¬p) [*Presumably* Iraq *does not* have weapons of mass destruction]

In (1d), George utters what he believes to be *more likely to be false than true*, so that it seems that he is lying. However, SIC does not capture (1d), because by hypothesis George's degree of confidence is higher than uncertainty but falls short of believing (1) to be probably false. Since it fails to account for the intuition that also (1d) is a lie (even if arguably less insincere than (1c)), SIC is too restrictive.

---

[9] I rephrased Carson's condition to avoid the counterintuitive consequence that degrees of belief included between "believing false" and "believing probably false" would not be counted as lies.

Carson's second, "weak" proposal avoids this problem. The "weak insincerity condition" posits that lying requires that the speaker "does not believe [the asserted proposition] to be true" (Carson 2006, cf. also Sorensen 2007:256, 2011:407).

WIC: S *does not* believe p to be true

Since it acknowledges that utterances like (1d) are lies, WIC is preferable to SIC. However, WIC is too broad: it incorrectly captures cases in which the speaker has no idea whether what he says is true or false, but goes on saying it for some other reasons. These cases are classified in the literature as bullshit (Frankfurt 1986). The typical example of bullshitter is the politician who "never yet considered whether any proposition were true or false, but whether it were convenient for the present minute or company to affirm or deny it" (Swift 1710). Now, as long as the speaker has no opinion about the veracity of what he is saying, his utterance is better classified as a misleading utterance than as a lie (Saul 2012:20, Meibauer 2014: 103, but cf. Falkenberg 1988:93, Carson 2010:61-2) and WIC is too broad to account for this intuition.

Given that SIC is too narrow and WIC is too broad, an ideal condition has to lie somewhere in the middle. Marsili (2014:162) proposes a middle ground between these two proposals: on this view, lying requires that the speaker believes *p more likely to be false than true*. Call this the *comparative insincerity condition:*

CIC: S believes p more likely to be false than true

Unlike WIC, CIC correctly rules out bullshit and statements uttered in cases of uncertainty. Unlike IC, it counts graded-belief lies as lies. And unlike SIC, it rules in the other cases in which the speaker does not believe the statement to be true – like (1c) and (1d).

A possible worry about the CIC is that it implicitly accepts that every belief can be represented as an assessment of probability, and that the speaker would find any such difference to be significant. To avoid this worry, CIC can be revised into CIC**:**

CIC': S's degree of confidence in ¬p is stronger than his degree of confidence in p

Of course, it is possible to challenge the very assumption that there is a clear-cut boundary between insincerity and sincerity, so that we should allow for intermediate indeterminate cases (thereby treating insincerity and lying as vague predicates, see Isenberg 1964:470). But

this intuition can be accommodated by the CIC which, unlike the other accounts, allows for a progressive transition from sincerity to insincerity. At the same time, if a neat point of transition is to be individuated, the CIC is fine-grained enough to identify a boundary that meets our intuitions and avoids the counterexamples to which the alternative accounts fall victim[10].

 The proposed view correctly accounts for assertions that do not specify the speaker's degree of belief in the truth of the proposition (*i.e.,* assertions that do not express graded beliefs): the next section will extend this view, and provide a broader framework to also treat assertions that express graded beliefs.

2. Insincerely expressing a graded belief

An assertion is insincere if there is a significant discrepancy between *the speaker's degree of belief* (henceforth BΨ) and *the degree of belief expressed by the sentence* (henceforth BΛ). The previous section has shown that this discrepancy can come in degrees, because of the graded nature of certainty – i.e. the graded nature of BΨ. A complete explanation of insincerity needs to account for the other side of the coin: the different degrees of belief that an assertion can express – the graded nature of BΛ.

Assertions that express graded beliefs are generally overlooked in the literature on lying. This is because, in standard cases, statements express a flat-out belief in the truth of the proposition, rather than a graded belief. For instance, (1) expresses a flat-out belief in the asserted proposition:

(1) Iraq has weapons of mass destruction

Not all statements, however, are as simple as (1), for some express graded beliefs. For instance, in (1e) the speaker believes that (1) is probably true, and in (1f) he expresses uncertainty in the truth of the proposition:

(1e) (0.75/p) *Probably* Iraq has weapons of mass destruction
(1f) (0.5/p) *Maybe* Iraq has weapons of mass destruction

_____

[10] Note that this blocks counterexamples based on borderline cases of uncertainty (see Krauss 2017). Krauss has proposed an alternative condition to CIC, that has been proven to be untenable by Benton (2018).

Now, the previous section has considered graded-belief lies, in which BΨ is graded. Graded assertions like (1e) and (1f), by contrast, are cases in which BΛ is graded. Three kind of graded-belief lies are hence possible: (A) *BΨ-graded*: plain assertions like (1), uttered while holding a graded belief; (B) *BΛ-graded*: assertions expressing a graded belief, like (1e) and (1f), uttered while holding an outright belief; (C) *complex cases*, where both BΛ and BΨ are graded. In what follows, I will consider how the insincerity condition applies to cases (B) and (C), and discuss the difference between mitigating and reinforcing an assertion.

2.1. Two directions of belief misrepresentation

Few authors have raised the question how assertions that express graded beliefs are to be analysed within a theory of lying. Meibauer (2014: 225) suggests that there are three kinds of BΛ-graded assertions that may qualify as lies: those "(i) expressing certainty when [you] are uncertain, those (ii) expressing uncertainty when [you] are certain, and those (iii) expressing certainty or uncertainty to a higher degree than being adequate with respect with [your] knowledge base". Since the third case seems to include the previous two, to simplify this taxonomy I will simply distinguish between two "directions" in misrepresenting your degree of belief: namely, pretending to have a *higher degree of belief* or a *lower degree of belief* than the one you have (cf. Falkenberg 1988:93).

A first, tempting idea is to assume that these two directions are equivalent. This would mean that, from the point of view of the analysis of lying, "pretending to be more certain than you are" is *as insincere as* "pretending to be less certain than you are" (Schiffrin 2014:12n15). A reason to make this assumption is that the "discrepancy" between your state of mind and the state of mind expressed by the statement is the same in both cases. However, at a closer look this assumption reveals to be naïve, as the first case (overstating) is often perceived as being more insincere, or more misleading, than the second (understating). To see this, imagine two utterances:

(1g) (1/p) Certainly Iraq has weapons of mass destruction
(1h) (0.5/p) Perhaps Iraq has weapons of mass destruction

Imagine that in both cases George's mental state is in between certainty and uncertainty, so that he believes:

(0.75/p) [Probably Iraq has weapon of mass destruction]

According to the 'naïve' view, (1g) and (1h) are equally insincere, because the discrepancy between BΨ and BΛ is the same (0.25). These scenarios differ only in the direction of misrepresentation: (1g) represents the speaker as having a higher degree of belief than he has, while (1h) as having a lower degree of belief. Interestingly, however, it is natural to assess (1g) as more insincere than (1h). The reason is that we tend to judge (1h) as a prudent statement, that cooperatively avoids saying more than the speaker knows, while (1g) is perceived a misleading overstatement, that the speaker lacks sufficient knowledge to assert. In other words, *ceteris paribus*, understating your degree of belief is generally seen as a cooperative linguistic practice, while overstating it is generally regarded as uncooperative.

In line with this intuition, Falkenberg (1988: 94, 1990) proposes to distinguish between "hard lies" (overstatements, like (1g)) and "soft lies" (understatements, like (1h)). However, this taxonomy is misleading in two respects. First, not all overstatements and understatements are lies: if the CIC is a condition for lying, only statements displaying a certain level of discrepancy between BΨ and BΛ can be lies. Second, it is not clear if an *overstatement* (hard lie) is necessarily more of a lie than an *understatement* (soft lie): next section will show that the direction of misrepresentation is just one of the parameters of intensity that must be considered, another one being the discrepancy between BΨ and BΛ.

2.2. Epistemic modality markers and degrees of commitment

There is a vast literature exploring the various ways in which an assertion can be mitigated (expressing a lower degree of belief, as in (1g)) or reinforced (expressing a higher degree of belief, as in (1h)) (see Fraser 1980, Holmes 1984, Coates 1987, Bazzanella et al. 1991, Caffi 2007, Egan & Weatherson 2011). The most prominent linguistic devices used for these purposes (expression like 'certainly','probably', 'perhaps') are called *epistemic modals*. This section will explain their pragmatic function, and clarify why we generally assess *understatements* as more sincere (or more honest) than *overstatements*.

Epistemic modals 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, italic is mine). In other words, they act on two components of the assertion, altering both (1) the psychological state expressed by the speaker (the degree

of belief), and (2) his degree of commitment to the truth of the proposition (the illocutionary strength[11]) (cf. Sbisà & Labinaz 2014:52, Lyons 1977: 793-809; Holmes 1984: 349).

These two functions are distinct in nature, but are entangled: if a speaker S mitigates (or reinforces) the degree of belief conveyed by his assertion, then S automatically mitigates (or reinforces) the illocutionary force of his assertion (that is, the degree of his commitment to the truth of the proposition). For instance, if you state (2b) instead of plainly stating (2), you both mitigate the degree of belief expressed ((2b) expresses uncertainty in (2)) and the degree of your commitment to the truth of the asserted proposition (you are committed to the truth of (2) to a much lower degree if you utter (2b))[12].

(2) Plato will quit smoking tomorrow

(2b) *Perhaps* Plato will quit smoking tomorrow

The role that epistemic modals play in reinforcing/weakening the illocutionary force of assertions explains why *understatements* are perceived as more honest than *overstatements*. *Ceteris paribus* (given the same degree of insincerity, like in (1g)-(1h)) a reinforced assertion has a stronger illocutionary force than a mitigated assertion, so that the speaker has a stronger commitment to its truth. And if the commitment to sincerity is stronger in reinforced statements, then violating that commitment is more serious in that statements than in mitigated ones.

Variations in illocutionary force induced by epistemic modals can also affect whether the speaker is asserting the proposition or not – and hence whether he is lying, because lying requires asserting. This is because epistemic modals can downgrade the degree of

---

[11] In speech act theory, the illocutionary force is what characterises the utterance of *p* as being the occurrence of a specific kind of illocutionary act (e.g. a question, an assertion, etc.). The illocutionary force of an assertion can be reinforced or mitigated (Bazzanella, Caffi & Sbisà 1991; Sbisà 2000; Searle & Vanderveken 1985: 99), thus altering the speaker's degree of commitment to the truth of the proposition.

[12] On this '*espressivist*' interpretation, epistemic modals are not part of the proposition asserted (at least not of the proposition against which speaker sincerity and commitment is assessed). A '*descriptivist*' might object that we should instead take them to be part of the content of the assertion (and hence of the proposition against which sincerity is measured). However, this would often yield counterintuitive predictions for the sincerity conditions of assertions. For instance, on a descriptive interpretation of "certainly p" as true iff (q): "the speaker is certain that p", a speaker that believes that there are 9/10 chances that p is true would counterintuitively be counted as insincere (as S would be certain that q is false). It should be noted that even if this section provides sincerity conditions for marked assertions interpreted in an expressivist fashion, it is not committed to expressivism: a descriptivist can still adopt the model proposed in section 1 (CIC). I follow Coates' (1987:130) view that epistemic modals can be appropriately used and interpreted in both ways. When they are used 'literally' to assert the epistemic or psychological (un)certainty of a proposition (rather than *express* that the proposition asserted is (un)certain, the simple sincerity conditions provided by CIC will apply; in the other cases (that I take to be the prevalent uses), the expressivist explanation outlined in this section will apply instead. On the debate over the semantics of epistemic modals, cf. Kratzer (1981), DeRose (1991), Egan, Hawtorne & Weatherson (2005), Papafragou (2006), Fintel & Gillies (2008), Yalcin (2007, 2011), Swanson (2011).

illocutionary force of a declarative sentence to such an extent that it longer counts as an assertion, but rather as a supposition or an hypothesis (Sbisà & Labinaz 2014:52-3). For instance, (2b) is a supposition rather than an assertion: its insincere utterance does not amount to lying, while insincerely uttering its unmitigated version (2) does. Carson (2010: 33,38) shares this intuition: "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". This is even more evident in other speech acts. For instance, if Matteo utters (3b) instead of (3), it is clear that he has not promised that he will buy you an elephant (he is merely suggesting it), while he would be promising it if he uttered (3). It seems that an insincere utterance of the first amounts to lying, while this is not true for the second[13].

(3) Tomorrow I will buy you an elephant
(3b) *Perhaps* tomorrow I will buy you an elephant

To sum up, this section considered how to account for assertions of type (ii), expressing graded beliefs (BΛ-graded lies), and (iii), complex cases where both BΛ and BΨ are graded. Both cases could have been dealt simply by appealing to condition CIC, but such an explanation would not have been able to account for differences determined by the direction of misrepresentation (overstatements vs understatements). This difficulty dissipates once it is understood that epistemic modals influence not only whether the sincerity condition is satisfied (by altering the degree of belief expressed), but also whether the assertion condition is satisfied (by altering the speaker's degree of commitment). This clarifies why Falkenberg's characterisation of overstatements and understatements as "hard" and "soft" lies is misleading: first, assertions that represent the speaker as less certain than he is (understatements) can be "hard" lies, if they achieve a sufficient degree of illocutionary force and insincerity; second, not all understatements and overstatements are lies, because they both may fail to achieve the sincerity and/or assertion condition for lying.

---

[13] One might wonder whether uttering (2b) or (3b) while being psychologically *certain* that the mitigated proposition is false would count as lying – *i.e.* if a high degree of insincerity can compensate for a low degree of commitment. Marsili (2014: 166-8) argues against this view, claiming that these utterances are to be classified as misleading statements rather than lies.

3. **Lies that attack certainty**.

So far, this chapter has focused on speaker certainty. This last section will consider hearer certainty, and how lies can affect the hearer's degree of belief in the asserted proposition. The primary goal of lying is generally to attack a hearer's beliefs (or his grounds for believing): typically, to make the hearer believe that the asserted proposition is true, *i.e.* to *deceive* him. Traditionally, this intention to deceive was believed to be necessary for lying, but recently, on the grounds of several convincing arguments (e.g. Carson 2006, Sorensen 2007, Fallis 2010, 2014) most authors believe that this intention is merely *frequently associated* with lying.

One prominent (but, we shall see, unsuccessful) argument against the idea that lying requires intending to deceive is the argument from *knowledge-lies* (Sorensen 2010), or lies that only attack certainty. Knowledge-lies aim to *prevent the hearer from being certain* of the truth of ¬*p* without *intending him to believe* that *p*[14]. Consider an example: Pietro has won a lottery and is ready to send his ticket to the address A1, to reclaim his prize. Luca wants to steal Pietro's ticket, so he tells him that he believes that the correct address is A2 instead. He knows that this way Pietro will have to go to an Internet café to be sure that the address is indeed A1, and plans to steal the ticket in the meanwhile. This is a knowledge-lie because Luca has no intention to make Luca believe that the address is A2 (otherwise Pietro would send it to A2, and the plan would fail): he merely aims to undermine Pietro's certainty (and knowledge) that the correct address is A1.

According to the argument from knowledge-lies against the intention to deceive, (i) knowledge-lies are not intended to deceive, but since (ii) knowledge-lies are clearly lies, (iii) intending to deceive is not a necessary condition for lying. The problem with this argument is in premise (i), which holds only if one endorses a very narrow account of intended deception – one that includes attempts to make someone believe *p*, and excludes attempts to modify his degree of belief in *p*. However, it seems that intending to alter someone's degree of belief may count as intending to deceive him (Chisholm & Feehan 1977:145, Fallis 2009:45, but cf. Carson 2010:181), and consequently that (i) is false. In fact, there is an obvious sense in which knowledge-lies are intended to deceive (Staffel 2011, Fallis 2011): even if Luca does not intend to make Pietro believe that the address is A2, it aims to undermine Pietro's conviction that the address is not A2, thereby deceiving him.

---

[14] Sorensen's original definition is slightly broader: it captures any believed-false statement that aims "to prevent the addressee from *knowing* that p is untrue" without intending him to believe that p, including statements that attack the hearer's ground for believing (i.e.. epistemic certainty). However, since virtually every attack to epistemic certainty via lies is ultimately meant to attack psychological certainty (cf. Fallis 2011:360-4), I will again restrict the discussion to lies that attack psychological certainty.

The argument from knowledge-lies hence fails to prove that lying does not require intending to deceive, but it should not be concluded that intending to deceive is a necessary condition for lying (other compelling counterexamples still hold), nor that knowledge-lies are of no interest for theorising about lying. On the contrary, knowledge-lies are remarkable in several respects. They effectively illustrate that manipulating the hearer can be more important than convincing him of the truth of the utterance, and that altering his degree of belief can be a powerful way to manipulate his behaviour: the lower somebody's confidence in a belief, the lower their disposition to act on its basis.

Knowledge-lies help us understanding the essence of calumny, a common kind of political knowledge-lie. Calumnies are lies that aim to undermine someone's reputation. Often, calumnies work like knowledge lies: this happens when they merely aim to put into question someone's reputation, without aiming to convince anybody that the propositional content of the calumny is true. Here is an example from Italian politics: during the 2011 Milan mayoral elections, Letizia Moratti claimed that her opponent Giuliano Pisapia had been condemned for stealing a van used to kidnap and beat a young man. Pisapia had indeed been accused of the crime in 1985, but he was also acquitted of all charges. This calumny can be read as a knowledge-lie because it arguably did not aim to convince the electorate that Pisapia had committed the crime (a very difficult aim to achieve, given that all media quickly rectified Moratti's statement), but rather to publicly put into question Pisapia's criminal record, thus undermining the electorate's certainty that Pisapia had never committed a crime[15].

Finally, lies that alter hearer certainty are interesting as they represent a further parameter that influences the "intensity" of lies, that adds to the ones considered in the previous sections: the higher the expected epistemic damage, the more reprehensible the lie.

This concludes the analysis of the concept of lying with respect to degrees of certainty, that has identified three dimensions of certainty that are relevant to lying: the speaker's degree of certainty, the degree of certainty expressed by the assertion, and the degree to which the speaker aims to modify hearer's certainty. Only the last two are relevant to define lying, but all affect the relative 'strength' of the lie. Lying can thus be considered a scalar phenomenon (a phenomenon that comes in degrees), with three dimensions of gradability that are relevant to its evaluation from a linguistic, epistemic and ethical point of view.

---

[15] Cf. Poggi et al. (2011). Pisapia eventually won the elections.

References

Austin, J.L. 1946. Other Minds. *Aristotelian Society Supplementary Volume*. 20:1.

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. 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.

Benjamin, Franklin. 1789. "Letter to Jean-Baptiste Leroy" (13/11/1789), in *Bartlett's Familiar Quotations*, 10th ed. (1919/2002), New York: Bartleby.

Benton, Matthew Aaron. 2018. Lying, Accuracy, and Credence. *Analysis*. doi:10.1093/analys/anx132.

Bok, Sissela. 1978. *Lying: Moral Choice in Public and Private Life*. New York NY: Random House.

Caffi, Claudia. 2007. *Mitigation*. London: Elsevier.

Carson, Thomas L. 2006. The definition of lying. *Nous* 2: 284–306.

Carson, Thomas L. 2010. *Lying and Deception*. Oxford: OUP.

Chan, T., & Kahane, G. 2011. The trouble with being sincere. *Canadian Journal of Philosophy* 41(2): 215–234.

Chisholm, Roderick M. & Feehan, Thomas D. 1977. The intent to deceive. *Journal of Philosophy* 74(3): 143–159.

Coates, Jennifer. 1987. Epistemic modality and spoken discourse. *Transactions of the Philological Society* 85(1): 110–131.

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.

Egan, Andy & Weatherson, Brian (eds). 2011. *Epistemic Modality*. Oxford: OUP. DOI: 10.1093/acprof:oso/9780199591596.001.0001.

Eriksson, John. 2011. Straight Talk: Conceptions of Sincerity in Speech. *Philosophical Studies* 153 (2): 213–234. doi:10.1007/s11098-009-9487-2.

Falkenberg, Gabriel. 1988. Insincerity and disloyalty. *Argumentation* 2.1: 89-97.

Falkenberg, Gabriel. 1990. Searle on sincerity. *Speech Acts, Meaning and Intentions. Walter de Gruyter, Berlin*, 129-145.

Fallis, Don. 2009. What is lying? *The Journal of Philosophy* 106(1): 29–56.

Fallis, Don. 2010. What is Deceptive Lying?, Paper presented at the *Pacific Division Meeting of the American Philosophical Association*, Available at SSRN:http://ssrn.com/abstract=1702023 or http://dx.doi.org/10.2139/ssrn.1702023.

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. 2014. "Are Bald-Faced Lies Deceptive after All?" *Ratio* (May 3) http://doi.wiley.com/10.1111/rati.12055.

Frankfurt, Harry G. 1986. On Bullshit. *Raritan Quarterly Review* 6:2.

Fraser, Bruce. 1980. Conversational mitigation. *Journal of Pragmatics* 4: 341-350.

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.

Holmes, Janet. 1984. Modifying illocutionary force. *Journal of Pragmatics* 8: 345–365. DOI: 10.1016/0378-2166(84)90028-6.

Hoye, Leo. 1997. *Adverbs and modality in English*. London, New York: Longman.

Isenberg, Arnold. 1964. Deontology and the Ethics of Lying. *Philosophy and Phenomenological Research* 24 (4): 463–80.

Klein, Peter. 1998. Certainty. In Edward Craig, *Routledge Encyclopedia of Philosophy* 264-267.

Kratzer, Angelika. 1981. The notional category of modality. In *Worlds, Words, and Contexts*, Hans-Jürgen Eikmeyer & Hannes Rieser (eds), 38–74. Berlin: De Gruyter.

Krauss, Sam. 2017. Lying, risk and accuracy. *Analysis*. 77 (4).

Kupfer, Joseph. 1982. The moral presumption against lying. *The Review of Metaphysics* 36(1): 103–126.

Labinaz, Paolo, and Marina Sbisà. 2014. Certainty and Uncertainty in Assertive Speech Acts. In *Communicating Certainty and Uncertainty in Medical, Supportive and Scientific Contexts*, Ilaria Riccioni, Carla Canestrari, Andrzej Zuczkowski, and Ramona Bongelli (eds.). Amsterdam: John Benjamins Publishing Company.

Levinson, Stephen. C. 1983. *Pragmatics*. Cambridge: CUP.

Lyons, John. 1977. *Semantics*. Cambridge: CUP.

Marsili, Neri. 2014. "Lying as a Scalar Phenomenon: Insincerity along the Certainty-Uncertainty Continuum." In *Certainty-Uncertainty – and the Attitudinal Space in between*, Sibilla Cantarini, Werner Abraham, and Elisabeth Leiss (eds.), 165:153–173. Amsterdam: John Benjamins Publishing Company.

Marsili, Neri. 2016. Lying by Promising. *International Review of Pragmatics* 8 (2). Brill: 271–313. doi:10.1163/18773109-00802005.

Montaigne, Michel. 1595 (E). *Essais,* Verdun: P. Villey et Saulnier.

Moran, Richard. 2005. Problems with sincerity. *Proceedings of the Aristotelian Society*, 105, 325–345.

Mannison, Don S. 1969. Lying and lies. *Australasian Journal of Philosophy* 2(47): 132–144. DOI: 10.1080/00048406912341141.

Meibauer, Jörg. 2014. *Lying at the semantic-pragmatic interface*. Berlin: De Gruyter.

Papafragou, A. 2006. Epistemic modality and truth conditions. *Lingua*, *116*(10), 1688–1702. http://doi.org/10.1016/j.lingua.2005.05.009.

Poggi, I., D'Errico, F., & Vincze, L. 2011. Discrediting moves in political debates. *In Proceedings of second international workshop on user models for motivational systems: the affective and the rational routes to persuasion*. Springer LNCS: 84-99.

Reed, Baron. 2008. Certainty. In *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL=http://plato.stanford.edu/archives/win2011/entries/certainty.

Reimer, Marga. 2004. What malapropisms mean: A reply to Donald Davidson, *Erkenntnis* 60(3): 317– 334.

Saul, Jennifer M. 2012. *Lying, Misleading, and what is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford: OUP.

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. & Vanderveken, Daniel. 1985. *Foundations of Illocutionary Logic*. Cambridge: CUP.

Searle, John R., and Daniel Vanderveken. 2005. Speech acts and illocutionary logic. *Logic, Thought and Action*. Springer Netherlands. 109-132.

Shiffrin, Seana V. 2014. *Speech Matters: on lying, morality, and the law*. Princeton University Press.

Sorensen, Roy. 2010. Knowledge-Lies. *Analysis* 70 (4) (August 7): 608–615. doi:10.1093/ analys/anq072.

Sorensen, Roy. 2011. What lies behind misspeaking. *American Philosophical Quarterly* 48(4): 399-409.

Staffel, Julia. 2011. Reply to Sorensen, 'knowledge-lies'. *Analysis* 71: 300–303.

Stanley, Jason. 2008. Knowledge and certainty. *Philosophical Issues* 18.1 35-5.

Stokke, Andreas. 2013. Lying, Deceiving, and Misleading. *Philosophy Compass* 8(4).

Stokke, Andreas. 2014. Insincerity. *Noûs* 48 (3): 496–520. doi:10.1111/nous.12001. http:// doi.wiley.com/10.1111/nous.12001.

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.

Swift, Jonathan. 1710. The Art of Political Lying. *The Examiner*, 10.

Von Fintel, Kai, & Gillies, Anthony. S. 2008. CIA leaks. *Philosophical review*, 117(1), 77-98. doi: 10.1215/00318108-2007-025

Yalcin, Seth. 2007. Epistemic Modals. *Mind*, *116*(464), 983–1026. http://doi.org/10.1093/mind/fzm983

Yalcin, Seth. 2011. Nonfactualism about epistemic modality. In Andy Egan & B. Weatherson (eds.), *Epistemic Modality*. Oxford University Press.

Williams, Bernard. 2002. *Truth and Truthfulness: An Essay in Genealogy*. Princeton NJ: Princeton University Press .