

LYING AND VAGUENESS

PAUL ÉGRÉ BENJAMIN ICARD

ABSTRACT. Vagueness is a double-edged sword in relation to lying and truthfulness. In situations of in which a cooperative speaker is uncertain about the world, vagueness offers a resource for truthfulness: it avoids committing oneself to more precise utterances that would be either false or unjustifiably true, and it is arguably an optimal solution to satisfy the Gricean maxims of Quality and Quantity. In situations in which a non-cooperative speaker is well-informed about the world, on the other hand, vagueness can be a deception mechanism. We distinguish two cases of that sort: cases in which the speaker is deliberately imprecise in order to hide information from the hearer; and cases in which the speaker exploits the semantic indeterminacy of vague predicates to produce utterances that are true in one sense, but false in another. Should such utterances, which we call half-truths, be considered lies? The answer, we argue, depends on the context: the lack of unequivocal truth is not always sufficient to declare falsity.

Keywords: lying; deception; vagueness; imprecision; indeterminacy; generality; approximation; open-texture; supervaluations; half-truth; informativeness; subjectivity; equivocation

Lying may be defined as the deliberate utterance of a false sentence (or thought to be false), generally with the aim of misleading the hearer into thinking that the sentence is true. Paradigmatic examples of lies involve sentences expressing *incontrovertibly false* propositions. For example, when former French Minister of Budget Jérôme Cahuzac solemnly declared on December 5, 2012: “*I do not have, Mr Deputee, I never had, any account in a foreign country, neither now nor before*”, he made an outright false assertion whose falsity he could no longer deny after investigations found evidence that he had held bank accounts in Switzerland, Singapore, and the Isle of Man. Those investigations quickly led to Cahuzac admitting his lie and resignating.

The Cahuzac scandal is the example of a *blatant lie*: an utterance whose falsity is clear and eventually beyond doubt, including to the speaker. For many of our utterances, however, it is not so clear-cut whether they should be considered a lie or not, even after all the evidence has been collected, and even to the speaker. This happens when utterances are vague. The point is that vague sentences have unclear truth-conditions, and cannot easily be nailed down as false for that matter. Consider horoscopes, and the question of whether they are truthful or not. Suppose my horoscope tells me (an example found on the internet):

- (1) Overall, you will feel rather good, physically and morally. You won't be too inhibited this time and...

Is the sentence true or false? The answer is unclear. Characteristic of horoscopes is the exploitation of vagueness. In an expression like *overall*, it is left underspecified exactly what proportion of the time. Similarly, *feel good* is a qualitative predicate for which there is no absolute criterion of application, and likewise for *inhibited*. Further exploitation of vagueness can be found in the use of degree modifiers like *rather* or *too*, whose interpretation is characteristically speaker- and listener-dependent (see Lakoff 1973; Wright 1995). Moreover, the indexical *this time* leaves its temporal reference open, making simply unclear which context is targeted by the sentence to be counted as true or false.

To philosophers of science, horoscopes are deceitful precisely because they make exploitation of vagueness on such a large scale (Popper 1963). To casual readers, on the other hand, they are often pleasant to read, because it is easy to find confirming instances of their truth (horoscopes can be thus argued to exploit a well-documented psychological phenomenon by means of semantic vagueness, the phenomenon of *confirmation bias*, see Wason 1966). Such ambivalence suggests that vagueness can be a convenient way of calibrating the truth of an assertion. There is a tradeoff between informativeness and truth, or dually, between vagueness and falsity. That is, the more vague an utterance, the more likely it is to be true relative to some contexts of interpretation, and the less likely it is to be false as a result. The more precise an utterance, on the other hand, the narrower the range of contexts relative to which it can be true. By decreasing informativeness, a vague sentence thus increases its chances of being true (Russell 1923).¹

This inverse relationship between informativeness and truth is exploited not just by horoscopes, it is a pervasive feature of everyday conversations and exchanges, and it concerns commercial, moral and legal transactions. Consider sales and advertising: like horoscopes, ads generally use vague vocabulary to sell their products. A famous case concerns the firm Ferrero, who used to advertise for its star product as *see healthy*. The firm was sued by a Californian customer on the grounds of making a false claim, considering the high rate of sugar in its product, but the company retorted that “there are health benefits associated with eating chocolate” (more on such moves below). In ordinary exchanges, however, vagueness is not necessarily used to deceive, but simply to avoid making claims that are too committal. Vagueness in that sense is not confined to horoscopes, but concerns predictive utterances quite generally (as in medical communication, see van Deemter 2009 and below). Vagueness is a feature of language that is used to avoid flouting Grice's first Maxim of Quality (“Do not say what you believe to be false / that for which you lack adequate evidence”) while exploiting Grice's second Maxim of Quantity (“Don't make your contribution more informative than is required” see Grice 1975, 45-46).

The goal of this chapter is to clarify the ways in which the use of vague language relates to both of those maxims. Vagueness is a multifaceted notion, however. In the first part of this chapter, we start out by distinguish two main manifestations of vagueness in

¹“A vague belief has a much better chance of being true than a precise one, because there are more possible facts that would verify it” (Russell 1923, 91).

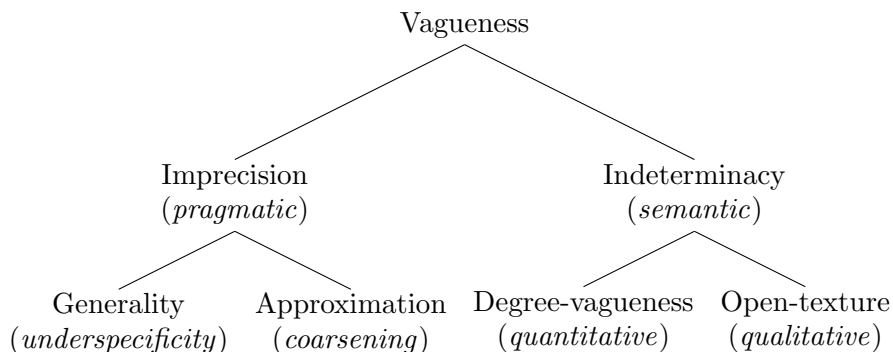


FIGURE 1. Varieties of linguistic vagueness

language: pragmatic *imprecision*, and semantic *indeterminacy*, each of which with more specific varieties. We then go on to explain in what sense vague language is a double-edged sword in relation to lying and truthfulness. First, we show that in situations in which a cooperative speaker wishes to inform about a state of affairs about which she is uncertain, vagueness offers a resource for truthfulness: it avoids making more precise utterances which may be either false or unjustifiably true (section 2). In situations in which a non-cooperative speaker is perfectly informed about the world, on the other hand, vagueness can be a deception mechanism. We distinguish two cases of that sort: cases in which the speaker is deliberately imprecise in order to hide information from the hearer, but remains literally truthful (section 3); and cases in which the speaker exploits the semantic indeterminacy of vague predicates to make utterances that are true in one sense, but false in another, what we call half-truths (section 4). The question is whether such half-truths should be counted as lies. The answer, we suggest, depends on the context: the lack of unequivocal truth is not always sufficient to declare falsity (section 5).

1. VARIETIES OF VAGUENESS

Russell (1923) offered as a general definition that “a representation is vague when the relation of the representing system to the represented system is not one-one, but one-many”. In the linguistic case, an expression is vague according to him if “there is not only one object that a word means, and not only one possible fact that will verify a proposition” [89-90]. That is, the same utterance is compatible with several distinct meanings. This one-many relationship can be realized in several ways, and a merit of Russell’s definition is that it covers a range of phenomena associated to linguistic vagueness. In what follows we distinguish four main manifestations: *generality*, *approximation*, *degree-vagueness*, and *open-texture*. Following several authors (see Pinkal 1995; Kennedy 2007; Solt 2015), we argue that generality and approximation are fundamentally cases of pragmatic imprecision, whereas degree-vagueness and open-texture are semantic phenomena, directly affecting the truth-conditions of expressions (see Figure 1).

1.1. **Generality.** A first instance of Russell’s definition concerns the phenomenon of *generality* in language (being underspecific). Consider the following dialog between a father and his son:

- (2) Q. Who did you see at the party?
A. Some friends.

Let us assume that the son has ten friends, all known to the father. The son’s answer is general in the sense that it is compatible with several more specific answers being true. The father can’t infer from the answer which exact number of friends was seen. An answer like *two friends* would be more informative in that respect, but it would still be general, leaving the father uncertain as to which two-membered subset of the relevant set includes the friends seen by his son.

Important to note is that in this context a sentence like *I saw some friends* has completely clear truth-conditions, it simply means that the number of friends seen by the speaker is greater than zero. Vagueness in that case does not mean any indeterminacy in the statement of the truth-conditions of the sentence, but simply refers to the fact that the response to the question fails to be maximally informative.

Theorists of vagueness often dismiss generality as a central aspect for that matter (see Fine 1975; Keefe 2000). We find important to keep it into consideration here, for the underspecificity of answers, although relative to the question under discussion, is a very common aspect of language use of particular relevance in relation to lying.

1.2. **Approximation.** A second illustration of Russell’s definition of vagueness pertains to approximation. In cases of approximation an expression with precise truth conditions is used to convey a meaning that differs from its literal meaning, but that is close enough. As a result, the same expression is used with a coarser meaning (larger range of interpretations than its literal meaning). Consider the following dialogues:

- (3) Q. What do you make a month?
A. 3,000 euros.
- (4) Q. What time did John arrive?
A. He arrived at 3 o’clock.
- (5) Q. How old is she?
A. She is 40.

In (3), the answer may be asserted by someone who knows the precise answer to actually be 3287,23 euros. This is a case in which the speaker rounds off the actual number to a lower number, relying on the fact that it is more relevant to set the standard of precision to multiples of 1000 euros than to a multiples of a single euro, let alone cents (see Krifka 2007). The same often happens with the other two examples: *3 o’clock* can be used when John in fact arrived at five past or five to, and *she is 40* could be used to refer to someone whose age is within a few months or even a few years around 40, depending on the context.

Approximation is not limited to numbering, but is also found in other domains, as exemplified in Austin's geometrical example (Austin 1962; Lewis 1970):

(6) France is hexagonal.

The latter sentence would be false if taken to mean that France has precisely the shape of a hexagon, but we understand it to mean that it can be circumscribed to a reasonable approximation by a hexagon.

Cases of approximation are cases in which a semantically precise expression is used with slack (Laserson 1999). Importantly, there may not be an absolutely precise convention as to the range of meanings that are compatible with the use of an expression. When is it no longer fine to say *John is 40 years old*? What if John is 35 years old? Approximation is always relative to explicit or implicit standards of precision and to rounding rules, and how close a value needs to be to the literal meaning will often be at the speaker's discretion.

1.3. Degree-vagueness. The third aspect of vagueness we isolate concerns the quantitative indeterminacy attached to gradable expressions in particular (which we call degree-vagueness, following Alston 1964; Burks 1946 talks of linear vagueness). Consider the following variation on the dialogue between a father and his son:

(7) Q. How many people were at the party?
A. Many people.

Here again, the answer is imprecise because compatible with a multiplicity of states of affairs obtaining (maybe 25 people were at the party, 50, or 100). Unlike for *some*, however, *many* is not an expression for which we can state determinate truth conditions relative to a fixed countable domain. One way of viewing that phenomenon is as a form of context-dependence (see Sapir 1944; Partee 1989; Lappin 2000; Greer 2014; Egré & Cova 2015): whereas *some As are Bs* is true exactly if the number of As that are Bs is nonzero, *many As are Bs* would be true if the number of As that are Bs exceeds a context-sensitive number n (cardinal reading), or possibly if the number of As that are Bs exceeds a context-sensitive proportion α of the As or of some other comparison class (proportional reading). The setting of such parameters is problematic: assuming such threshold values, did the son intend *many* to mean *more than 5*, *more than a third*, or some other number? A remarkable fact about vague expressions such as *many* is that the speaker himself or herself need not have a precise idea of the values of such thresholds in order to apply the expression and to convey meaning.

Beside *many*, paradigmatic examples of vague expressions in that sense include gradable adjectives like *tall*, *long*, *expensive*, *healthy*, etc., all of which accept degree-modification (as in *taller*) or modification by intensifiers (*very tall*) (see Kennedy 2007). Gradable adjectives give rise to familiar symptoms, in particular the admission of borderline cases of application and the susceptibility to sorites-reasoning (see Keefe 2000; Egré & Klinedinst 2011; Burnett 2016) for a more specific typology of gradable expressions). Borderline cases of application are cases for which it is unclear to the speaker whether the expression should

apply or not: for example, it may be unclear whether a man of 178cm should be counted as *tall* or not. An important fact about borderline cases is moreover that they give rise to inconsistent verdicts both between- and within-subjects (see McCloskey & Glucksberg 1978). Cases of between-subject inconsistencies are often viewed as manifestations of the subjectivity and evaluativity of vague expressions: *many*, *tall*, *healthy*, *beautiful*, could mean different things without error depending on the speaker (Parikh 1994; Wright 1995; Kölbel 2004; Fara 2000; Raffman 2013; Kennedy 2013; Egré 2016; Verheyen *et al.* 2017). This subjectivity is important for an assessment of the falsity of vague sentences: the same vague sentence could be used truly relative to one speaker, but be viewed as false by another, depending on their context, interests and evaluative standards (see in particular Kölbel 2004; McNally & Stojanovic 2017 on predicates of personal taste).

1.4. Open-texture. The fourth illustration of Russell’s definition we single out concerns the openness of the respects constitutive of the meaning of an expression, what we call open-texture (following Waismann 1945; Burks 1946 talks of multidimensional vagueness, and Alston 1964 of combinatorial vagueness). This openness is found at different levels, and it has to do with polysemy and multidimensionality.

Already in the case of dimensional adjectives (like *tall*), the selection of a comparison class is fundamental for the application of the adjective, but it can vary without limit, and it will impact the setting of a boundary between tall and not tall objects (*tall* for a building, for a basketball player, or for a fifth-grader, will mean different things, see Kamp 1975, Klein 1980).

For a number of gradable adjectives, moreover, several dimensions of comparison interact, and their number and structure is generally indeterminate, even when a comparison class has been fixed. Consider the adjective *healthy*. An indication that *healthy* is multidimensional is the occurrence of adjuncts such as *healthy in some respect*, *healthy in all respects* (Sassoon 2012). For example, *healthy* as applied to a meal could be predicated based on whether it provides vitamins, or based on whether it has a particular effect on blood pressure, or based on some way of integrating of those respects, and no definitive list of respects appears to be forthcoming.

The phenomenon of open texture is not limited to gradable adjectives, but it concerns the difficulty of providing necessary and sufficient conditions of applications for a vast number of expressions, including nominal expressions (Wittgenstein 1953 famously used the example of the word *game* to show the difficulty in providing a consistent and exhaustive list of defining criteria for that notion).

1.5. Representing vagueness. Degree-vagueness and open-texture can be thought of as a forms of “referential multiplicity” (Raffman 2013). A convenient way of representing the meaning of a vague expression, following the supervaluationist tradition, is thus in terms of a set of admissible sharpenings or precisifications (Mehlberg 1958; Lewis 1970; Fine 1975; Kamp 1975). For an expression like *tall*, for example, given a comparison class, the meaning can be represented by a set of precise intervals above a variable threshold; for an expression like *healthy*, given a comparison class again, it may be thought of as a set of tuples consisting of variable respects and intervals along a common dimension set by those

respects. Similarly for approximation: the meaning of *hundred* as used approximately can be represented by a set of numbers around 100 (Laserson 1999). Depending on the speaker, however, the range of such admissible sharpenings may differ.² Different speakers may also assign different weights to different sharpenings depending on the context (see Lassiter 2011; Lassiter & Goodman 2015 on probabilistic representations of vague meaning).

In this regard, the main difference between expressions like *hundred* or *some students* on the one hand, and *tall* or *game* on the other, is that the former have determinate truth conditions. Because of that, generality and approximation are cases of *pragmatic vagueness*: by being general rather than more specific a speaker chooses to be less informative than she could be, and by being approximate she gives less information than what the expression literally means. Degree-vagueness and open-texture on the other hand are cases of *semantic vagueness*: the meaning of expressions like *many*, *healthy* or *game* is “intrinsically uncertain” (in the words of Peirce 1902, 748), that is those expressions do not have constant truth conditions across contexts and speakers.

With these distinctions in mind, we are now in a position to examine the ways in which vagueness interacts with the Gricean maxims. The Gricean maxims assume that conversation fundamentally rests on cooperation. As we know from game theory, however, speaker and hearer need not have their interests perfectly aligned, and sometimes they can diverge dramatically. It may be costly to reveal the truth, or to reveal the *whole* truth. Most of the time, however, making an assertion that the listener would recognize as false can be even more costly: if a false claim is exposed, the speaker incurs the risk of losing credibility, or greater costs (see Asher & Lascarides 2013). In the rest of this chapter, we distinguish two main classes of situations that motivate the use of vague language. On the one hand, there are situations where the speaker is *imperfectly informed* about the facts, and may simply wish to avoid speaking falsely by speaking too precisely. On the other hand, there are situations where the speaker is *perfectly informed* about the facts, but has an interest to hide information from the hearer, and potentially to take advantage of the indeterminacy of vague expressions to bias or mislead.

2. AVOIDING ERROR

Grice’s Maxim of Quality enjoins one not to speak falsely, but also not to say things for which one lacks adequate evidence. One aspect in which the Maxim of Quality justifies the use of vague language concerns cases where the speaker is uncertain about which precise state of affairs obtains (see Frazee & Beaver 2010) or will obtain in the future (van Deemter 2009).

Consider a situation in which you return from a party and are a fully cooperative speaker trying to convey maximum information. The party was attended by a group of people, but

²This makes semantic vagueness close to lexical ambiguity, except that in the case of lexical ambiguity the meanings are supposed to be mentally far apart or disjoint (Keefe 2000; Pinkal 1995). Logically speaking, however, it is relevant to compare vagueness with ambiguity, since precisifications play the same role as disambiguations (see Lewis 1982). In the next section, we will see that vagueness, like ambiguity, can give rise to pragmatic equivocation.

you do not know exactly how many there were, because you could not count them. Upon returning from the party, you're asked how many people were there. In this case, there is no number n for which you can truly and justifiably say: *there were exactly n people*. In order to respond truly and justifiably, the next option would be to specify an exact interval. Suppose you are sure that there were more than 20 people, and fewer than 200 hundreds, but are uncertain in between. Then you may say:

(8) There were between 20 and 200 people.

The response is general in this case, but little informative. It would be more informative to give your best estimate of a lower bound:

(9) At least 100 people.

But suppose there were in fact 93 persons attending. The answer would be literally false, despite coming close to your assessment. On the other hand, you would not be wrong if you said:

(10) a. About 100 people.
b. Many people.

Semantic expressions like *about* and *many* allow you to convey information truly in this case, compatibly with an indeterminate range of states of affairs obtaining. They allow you to avoid error, but also, somewhat surprisingly, to be more informative than you would if you tried to specify exact intervals without error.

Importantly, the hearer may have a different understanding of what to count as *many* than you. Suppose you understand *many* to denote a range of sharp intervals (using the supervaluationist picture), with a probability distribution on them (some precisifications are more likely to you than other; (see Lassiter 2011; Lassiter & Goodman 2015)). The hearer may have a different probability distribution that rules out some of the intervals you consider possible, but you would still communicate successfully if the hearer ends up with a posterior distribution that includes the value you actually observed, and if it makes that value more likely than before you answered (see Parikh 1994; Lassiter 2011; Lassiter & Goodman 2015).

The point of the previous example is that vague language, in situations of uncertainty, may accomplish an optimal tradeoff between the need to be truthful and the need to be informative (see Frazee & Beaver 2010). Use of vague language in situations of uncertainty is also modulated by the cost of speaking falsely, compared to the benefits of speaking accurately. An example discussed by van Deemter 2009 concerns cases of medical communication. Van Deemter points out that “a doctor who says “These symptoms will disappear fairly soon” is less likely to get complaints, and to be sued, than one who says “These symptoms will have disappeared by midnight”” [8].

Vague language, in summary, is a way of speaking truly and informatively in situations of uncertainty. This does not mean that vagueness is immune to falsity: suppose the symptoms disappear only after a month, then the patient may charge the doctor of

incompetence, or even of having lied. The patient could complain that *fairly soon* was, in her perspective, incompatible with a time interval of a month. The doctor could deny having spoken falsely, on the other hand, by defending her own perspective. The relativity of vague interpretations to speakers makes charges of lies, as we will see, a delicate matter (see section 5).

3. HIDING INFORMATION

Let us now turn to cases where the speaker has no uncertainty about the world, but has an incentive to be noncooperative. Grice's Maxim of Quantity is twofold: it asks one to be as informative as required for the purpose of the conversation, but also to not be more informative than required. What counts as "required for the purpose of the conversation" is itself vague and heavily depends on the interests that speaker and hearer have in sharing information (Asher & Lascarides 2013). For a range of situations, a well-informed speaker can legitimately wish to retain information from the hearer, and so to be vague in order to limit cooperation. Cases of what we called *generality* in the previous section are very common in that regard. Consider the dialogue in (2), repeated here.

- (11) Q. Who did you see at the party?
A. Some friends.

Let us assume that the father is actually interested in knowing whether his son saw a particular person, say Ann, whom he suspects his son of dating. The son, on the other hand, wishes to keep his privacy. Assume the son saw Ann indeed, but also Don and Eli, two other friends known to the father. In this case, the son is giving a perfectly true answer, but he is not allowing the father to identify whom exactly he saw.

Compare with the example of the previous section. Assume you know this time that exactly 63 people attended the party, but have an interest not to reveal the exact number. You may choose to be underinformative by responding:

- (12) Q. How many people were at the party?
A. Fewer than a hundred.

The answer is literally true, but *partial* in the sense of Groenendijk & Stokhof (1982): it leaves possibilities open and fails to completely settle the question. Potentially, it is also misleading: for it triggers the implicature that it is compatible with your knowledge that there could have been 90 people or more attending (see Spector 2013). Such cases, in which a speaker is literally truthful but uses misleading implicatures are called cases of misdirection by Asher & Lascarides (2013), who characterize them as instances of *rhetorical* as opposed to genuine Gricean cooperativity.³

³See in particular their discussion of *Bronston vs. United States* as an exploitation of literal truth to refute perjury, as well as the presentation of the case in Tiersma (2004). See Ransom *et al.* (2017) for a recent study comparing cases in which a truthful speaker may have an incentive to be completely uninformative as opposed to partially informative depending on the level of trust in the hearer.

Neither of the previous examples relies on utterances that are vague semantically, but we can find similar cases where a semantically vague expression is used to withhold information. Imagine nosy neighbours asking how much you paid for your apartment. Assume you know the exact price you paid, but don't want to reveal it:

- (13) Q. How much did you buy your apartment?
A. It was not too expensive.

An incentive to avoid being precise in this case is that you may want to avoid appearing either lucky (in case you paid less than your neighbors for the same size) or stupid (in case you paid more), or you may just want to give no indication of your assets. Use of a qualitative expression like *expensive* is advantageous here because it avoids specifying a definite number, and it remains compatible with the preservation of truthfulness: we may assume that you are sincere in thinking that the price you paid was not expensive, even ahead of the dialogue (that assumption is not always warranted, see the next section).

Consider for comparison the following alternative answers, assuming the exact price you paid for your apartment is 220,000 euros:

- (14) a. I paid 200,000 euros.
b. I paid around 200,000 euros.
c. I paid between 50,000 euros and 300,000 euros.

Answer (14a) is approximate in this case, but it does not signal that it is approximate. As pointed out by Meibauer (2014), it may be truthfully asserted if the standard of precision in the context of the conversation is such that a difference of 20,000 euros would not be relevant. But the answer could be misleading, instead of just imprecise, if uttered with the intention of making your neighbors believe that you paid less than you actually did. For instance it would count as false in a context in which the standard of precision needs to be maximal (say in the context of declaring taxes).

Answer (14b) makes the approximation explicit, and it is also semantically vague, due to the use of the vague modifier *around*. Despite that, the answer remains more informative than the one in (13), for it lets your neighbors infer the actual price with less uncertainty than based on hearing *not too expensive*.

Answer (14c), finally, is neither approximate nor semantically vague: it states an exact interval but to create uncertainty. Like (13), it signals either that you do not know the price you paid, or that you don't want to answer the question precisely; however, the interval specified is so large here that the hearers would be better-founded to think you do not want to answer the question. Also, the answer in (13) may end up being more informative than the one in (14c) despite relying on semantic vagueness, because upon hearing "not too expensive" the hearer is likely to narrow down the range of prices you potentially paid to a smaller interval than the one specified in (14c).⁴

⁴This is because *I paid between 50,000 and 300,000 euros* scalarly implicates that it is possible you paid 51,000. With *not too expensive* this inference is not mandated at all. On the mechanism of such implicatures, see Fox (2014).

4. MAKING HALF-TRUTHS

Beside cases in which a speaker is imprecise to hide information, there is a class of cases where the speaker can exploit the semantic indeterminacy of vague expressions to produce utterances whose truth status is unclear: they are true under some way of resolving their vagueness, but that way can be tendentious or biased.⁵

Consider the following example (from C. List and L. Valentini, p.c.) where you receive an invitation for dinner. As a matter of fact, you would be free to go to that dinner, but have no inclination for it. Imagine the following dialogue:

- (15) Q. Are you free to come for supper tomorrow?
A. Sorry, I have an engagement.
- (16) Q. Are you free to come for supper tomorrow?
A. Sorry, I am busy.

In (15), your response ought to qualify as a lie. In the case of (16), the answer does not obviously count as a lie, but it does not clearly count as true either. One way of explaining the contrast is in terms of supervaluations (Fine 1975; Kamp 1975). On all admissible ways of sharpening the meaning of *I have an engagement*, the sentence would come out false (i.e. super-false). On the other hand, there are admissible ways of sharpening the meaning of *busy* for the sentence to count as true. *I am busy* may even be deemed super-true, that is true literally on all admissible ways of sharpening the meaning of *busy*, but this is moot: it depends on what to count as an admissible precisification (see below). If you end up watching TV, you would obviously be *busy watching TV*, but at the time of utterance *busy* appears to convey that you have some obligation.

In our view the answer in (16) is a half-truth, precisely because it is not clearly false, but not clearly true either. Concretely, *I am busy* offers a polite way of declining the invitation. A more informative alternative about the speaker's motives would be to say: *I am not very inclined*, but it would be clearly offending. The intent of *I am busy* is partly to mislead, therefore, but consistently with satisfying a norm of politeness.⁶

A more extreme case of exploitation of semantic vagueness concerns President Bill Clinton's declarations about the nature of his relationship with Monica Lewinsky:

- (17) I have never had sexual relations with Monica Lewinsky.

This case, importantly, is one where all parties had been fully informed of the relevant facts. To justify his claim without perjury, Bill Clinton took advantage of the open texture of the expression *sexual relations*, that is of the lack of a clear definition. However, he did it not by making up a definition, but by exploiting an attempt made by his opponents to

⁵The term *half-truth* is used in a number of different senses in the literature. Our use is broadly compatible with Carson 2010's, who defines a half-truth to be a true statement that "selectively emphasize[s] facts that tend to support a particular interpretation or assessment of an issue" [57-58]. We use *half-true* in the sense of *borderline true*.

⁶Thanks to C. List and L. Valentini for discussion of that aspect.

provide an explicit definition of the term “[engaging in] sexual relations” (see Tiersma 2004 for details).⁷ Pressed to explain himself, Clinton’s defense was:

- (18) “I thought the definition [of sexual relations, as read by Judge Wright] included any activity by the person being deposed, where the person was the actor and came in contact with those parts of the bodies with the purpose or intent of gratification, and excluded any other activity”.

The way Bill Clinton defended himself can be put in supervaluationist terms again: it is not the case that on all ways of further precisifying the explicit definition proposed by his opponents, receiving oral sex counts as engaging in a sexual relation. Interestingly, in an earlier statement Bill Clinton commented about whether Monica Lewinsky had had “a sexual affair” with him as follows:

- (19) Q. If she told someone that she had a sexual affair with you beginning in November of 1995, would that be a lie?
A. It’s certainly not the truth. It would not be the truth.

In this occurrence, Clinton appeared to concede that the allegation would not necessarily be false, but without counting as true. In supervaluationist terms again, there are some admissible ways of precisifying *sexual affair* that would make Lewinsky’s supposed statement true, yet *not all* ways of precisifying *sexual affair* would make it true. Overall, Bill Clinton was able to exploit the semantic indeterminacy of those expressions in order to avoid the charge of perjury. He would have been convicted if, from the jury’s perspective, all admissible ways of precisifying the meaning had led to the sentence being false, but the jury in that case failed to rule out Clinton’s way from being admissible.

5. ARE HALF-TRUTHS LIES?

Let us take stock. In section 3 we saw that in response to a question, a speaker can be underspecific without committing any lie. In section 4, however, we saw that semantic indeterminacy can be used to produce sentences whose truth status is unclear, what we called half-truths. Shouldn’t half-truths be considered lies, however, given that those utterances fail to be clearly true?

First of all, utterances like (16) or (17) may typically be uttered insincerely. In the case of (16), I may think to myself *in petto* “well, I am not really busy...” or “well, I am busy watching TV”, and Clinton may have silently thought to himself “well, except for an oral sexual relation”. Those utterances then may be viewed as cases of *amphiboly* or *mental reservation* (Bok 1979; Mullaney 1980; Adler 1997), whereby the actual meaning that the speaker has in mind is in fact different from the meaning the hearer can reasonably infer.

⁷The explicit definition in question is: “a person engages in “sexual relations” when the person knowingly engages in or causes contact with the genitalia, anus, groin, breast, inner thigh, or buttocks of any person with an intent to gratify or arouse the sexual desire of any person “Contact” means intentional touching, either directly or through clothing.”

To avoid that complication, let us assume that each utterance is made sincerely at the time it is uttered, and without mental reservation (without the speaker making any silent addition). In supervaluationist terms, the question we are asking is whether an utterance that fails to be super-true (true on all admissible precisifications) ought to be considered false on normative grounds. We think the answer to this question is nonobvious, for it depends on two parameters: the definition of what to count as an admissible precisification, and the choice of a standard for truth.

Regarding the first issue, most people would agree that Clinton's utterance is false *simpliciter*, despite being true under some very specific sharpening of the meaning of *sexual relation*, for they would deem that particular precisification to be inadmissible in an ordinary conversational context. In the legal context, however, Clinton was successful in making that sharpening relevant, and since it was incumbent on the jury to show that his statement was unequivocally false, it allowed for his sentence not to qualify as a lie, despite the sentence not qualifying as a clear truth either.

This brings us to the second issue. Theories of vagueness differ on the standards whereby a sentence can be truthfully uttered. Supervaluationism treats sentences as true *simpliciter* if they are true on all admissible precisifications, but there is a dual theory, subvaluationism, which treats sentences as true *simpliciter* when true under some precisification (Hyde 1997). Subvaluationism is very liberal in that it predicts that a sentence and its negation can both be true then.⁸

In practice, the standards for truth and falsity appear to depend on the context. In the Clinton lawsuit, it was sufficient for the sentence to be true under some sharpening to not be considered a lie by the jury. In the class-action lawsuit that opposed Athena Hohenberg to the Ferrero company, on the other hand, the complaint was that *healthy* was used misleadingly for a product containing too much fat and sugar. Ferrero's defense was based on the fact that *healthy* is multidimensional, and that their product was at least healthy in the respects of bringing chocolate, containing low sodium, and so on.⁹ Despite that, the court eventually forbade Ferrero from advertising the product as *healthy*. The court agreed that it is not enough for a sentence like *this product is healthy* to be true on just some ways of precisifying *healthy* in order for the sentence to avoid being misleading or to count as a lie, presumably in this case because the ways in which the sentence is false outweigh those in which it is true (Ferrero's use would in fact violate the Gricean Maxim of Relevance).

In general, however, the Ferrero example may be more emblematic of the ways in which vague language is interpreted. Grice's Maxim of Manner recommends avoiding ambiguity (see Grice 1975, 46). There is evidence, however, that in cases in which a vague predicate is used without qualification, and where two interpretations are available for the predicate, a weak one and a strong one, the stronger interpretation will be the default (see Dalrymple

⁸This implies that *I am busy* and *I am not busy* would both be true in a context in which either is true under some admissible sharpening. But each of them would also be false, since false under some sharpening. The upshot would be that the sentence both is a lie, and fails to be a lie.

⁹See <http://www.scpr.org/news/2011/02/10/23912/a-mom-sues-nutella-maker-for-deceptive-advertising/>

et al. 1998). Upon hearing *this person is tall*, the default is to get that the person is clearly tall, rather than borderline tall (Alxatib & Pelletier 2011; Cobreros *et al.* 2012, 2015). Likewise, when saying *this product is healthy*, the default is likely to hear *this product is healthy in most respects*, rather than just *some respects*. As a result, to say of a product that it is *healthy* without qualification would suggest that the product is more healthy than unhealthy: in the Ferrero case, this pragmatic enrichment is deceptive, and can legitimately be considered a lie.

We see, in summary, that often an utterance will be deemed a lie if it fails to be unambiguously true. But sometimes, as the Clinton case shows us, it might fail to be deemed a lie if it is not unambiguously false. Whichever of those two will prevail appears to depend not just on the *existence* of ways for a sentence to be true, but also on how *relevant* those ways are to the parties involved in the conversation.

6. CONCLUSIONS

Let us recapitulate the main lessons of our discussion of the relation between lying and vagueness. To begin with, we have seen that vagueness provides a way for a cooperative speaker to remain truthful in situations in which she is trying to communicate information about which she is uncertain. Vagueness may then be described as a way of avoiding error and therefore lies. This concerns all cases in which the use of qualitative but vague vocabulary (as in *many, long, expensive*) avoids committing oneself to precise quantitative expressions for which one fails to have adequate evidence. As opposed to that, we have highlighted two kinds of cases in which vagueness can be used deceptively. The first are cases in which a well-informed speaker has motives to hide or retain information. In such cases the speaker is deliberately imprecise and partial, but need not commit lies in the strict sense of the term. She may however be misleading if the partial information given triggers false implicatures. The second kind of cases concern what we have called half-truths, utterances whose status is borderline between true and false, depending on how vague expressions in them are interpreted. Such cases are more problematic. An utterance will be misleading if it is true only under some very peculiar precisification. On the other hand, the indeterminacy of vague expressions can make it difficult to prove that a vague utterance is a lie, as opposed to an expression whose intended meaning was misunderstood.

Acknowledgments: Thanks to Nick Asher, Sam Alxatib, Rae Langton, Christian List, Yael Sharvit, Stephanie Solt, Benjamin Spector, Laura Valentini, Steven Verheyen, for helpful conversations on the topics of this paper, and to Jörg Meibauer for helpful comments. We also thank Emar Maier for valuable advice regarding how to convert our paper from Latex to MSWord. PE thanks the ANR Program TrilLogMean ANR-14-CE30-0010-01 for funding. BI thanks the Direction Générale de l'Armement for doctoral funding. Both researchers acknowledge grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL* for research carried out at the Department of Cognitive Studies of ENS.

REFERENCES

- Adler, Jonathan. 1997. Lying, deceiving, or falsely implicating. *Journal of Philosophy*, **94**, 435–452.
- Alston, William P. 1964. *Philosophy of Language*. Prentice Hall.
- Alxatib, Sam, & Pelletier, Francis Jeffry. 2011. The psychology of vagueness: Borderline cases and contradictions. *Mind & Language*, **26**(3), 287–326.
- Asher, Nicholas, & Lascarides, Alex. 2013. Strategic conversation. *Semantics and Pragmatics*, **6**(2), 1–62.
- Austin, John L. 1962. *Sense and Sensibilia*. Oxford University Press.
- Bok, Sissela. 1979. *Lying: Moral Choice in Public and Private Life*. Vintage paperback editions.
- Burks, Arthur W. 1946. Empiricism and vagueness. *The Journal of Philosophy*, **43**(18), 477–486.
- Burnett, Heather. 2016. *Gradability in Natural Language: Logical and Grammatical Foundations*. Oxford University Press.
- Carson, Thomas L. 2010. *Lying and deception: Theory and practice*. Oxford University Press.
- Cobrerros, Pablo, Egré, Paul, Ripley, David, & van Rooij, Robert. 2012. Tolerant, Classical, Strict. *The Journal of Philosophical Logic*, **41**(2), 347–385.
- Cobrerros, Pablo, Egré, Paul, Ripley, David, & van Rooij, Robert. 2015. Pragmatic interpretations of vague expressions: Strongest meaning and nonmonotonic consequence. *Journal of Philosophical Logic*, **44**(4), 375–393.
- Dalrymple, Mary, Kanazawa, Makoto, Kim, Yookyung, Mchombo, Sam, & Peters, Stanley. 1998. Reciprocal expressions and the concept of reciprocity. *Linguistics and Philosophy*, **21**(2), 159–210.
- Egré, Paul. 2016. Vague judgment: a probabilistic account. *Synthese*. DOI: <https://doi.org/10.1007/s11229-016-1092-2>.
- Egré, Paul, & Cova, Florian. 2015. Moral asymmetries and the semantics of many. *Semantics and Pragmatics*, **8** (13), 1–45.
- Egré, Paul, & Klinedinst, Nathan. 2011. *Vagueness and Language Use*. Palgrave Macmillan.
- Fara, Delia. 2000. Shifting Sands: an Interest-Relative Theory of Vagueness. *Philosophical Topics*, **28**(1), 45–81. Originally published under the name “Delia Graff”.
- Fine, Kit. 1975. Vagueness, Truth, and Logic. *Synthese*, **30**, 265–300.
- Fox, Danny. 2014. Cancelling the Maxim of Quantity: Another challenge for a Gricean theory of scalar implicatures. *Semantics and Pragmatics*, **7**(5), 1–20.
- Fraee, Joey, & Beaver, David. 2010. Vagueness Is Rational under Uncertainty. In: Aloni, M., Bastiaanse, H., de Jager, T., & Schulz, K. (eds), *Logic, Language and Meaning. Lecture Notes in Computer Science*. Springer.
- Greer, Kristen A. 2014. Extensionality in natural language quantification: the case of many and few. *Linguistics and Philosophy*, **37**(4), 315–351.
- Grice, Paul H. 1975. Logic and Conversation. *Pages 41–58 of: Cole, Peter, & Morgan, Jerry L. (eds), Syntax and Semantics: Speech Acts*. New York: Academic Press.
- Groenendijk, Joroen, & Stokhof, Martin. 1982. Semantic analysis of wh-complements. *Linguistics and Philosophy*, **5**(2), 175–233.
- Hyde, Dominic. 1997. From heaps and gaps to heaps of gluts. *Mind*, **106**(424), 641–660.
- Kamp, Hans. 1975. Two theories about adjectives. In: Keenan, E. (ed), *Formal Semantics of Natural Language*. Cambridge University Press.
- Keefe, Rosanna. 2000. *Theories of Vagueness*. Cambridge University Press.

- Kennedy, Christopher. 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and philosophy*, **30(1)**, 1–45.
- Kennedy, Christopher. 2013. Two sources of subjectivity: Qualitative assessment and dimensional uncertainty. *Inquiry*, **56(2-3)**, 258–277.
- Klein, Ewan. 1980. A semantics for positive and comparative adjectives. *Linguistics and philosophy*, **4(1)**, 1–45.
- Kölbel, Max. 2004. Faultless disagreement. **104(1)**, 53–73.
- Krifka, Manfred. 2007. Approximate interpretation of number words. *Pages 111–126 of*: Bouma, G., Krämer, I., & Zwarts, J. (eds), *Cognitive Foundations of Interpretation*.
- Lakoff, George. 1973. Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of philosophical logic*, **2(4)**, 458–508.
- Lappin, Shalom. 2000. An intensional parametric semantics for vague quantifiers. *Linguistics and philosophy*, **23(6)**, 599–620.
- Laserson, Peter. 1999. Pragmatic halos. *Language*, 522–551.
- Lassiter, Daniel. 2011. Vagueness as Probabilistic Linguistic Knowledge. *Pages 127–150 of: Vagueness in Communication*. Springer.
- Lassiter, Daniel, & Goodman, Noah D. 2015. Adjectival vagueness in a Bayesian model of interpretation. *Synthese*, 1–36.
- Lewis, David. 1970. General semantics. *Synthese*, **22(1)**, 18–67.
- Lewis, David. 1982. Logic for equivocators. *Noûs*, **16**, 431–441.
- McCloskey, Michael E., & Glucksberg, Sam. 1978. Natural categories: Well defined or fuzzy sets? *Memory & Cognition*, **6(4)**, 462–472.
- McNally, Louise, & Stojanovic, Isidora. 2017. Aesthetic adjectives. *Pages 17–37 of*: Young, James O. (ed), *The Semantics of Aesthetic Judgment*. Oxford University Press.
- Mehlberg, Henry. 1958. *The reach of science*. University of Toronto Press.
- Meibauer, Jörg. 2014. *Lying at the semantics-pragmatics interface*. Walter de Gruyter GmbH & Co KG.
- Mullaney, Steven. 1980. Lying like truth: riddle, representation and treason in Renaissance England. *ELH*, **47(1)**, 32–47.
- Parikh, Rohit. 1994. Vagueness and utility: The semantics of common nouns. *Linguistics and Philosophy*, **17(6)**, 521–535.
- Partee, Barbara. 1989. Binding Implicit Variables in Quantified Contexts. *In*: Wiltshire, C., Graczyk, R., & Music, B. (eds), *Papers from the 25th Annual Regional Meeting of the Chicago Linguistic Society*. Chicago Linguistic Society.
- Peirce, Charles S. 1902. Vague. *In*: Baldwin, J. M. (ed), *Dictionary of Philosophy and Psychology*. Macmillan.
- Pinkal, Manfred. 1995. *Logic and lexicon: the semantics of the indefinite*. Kluwer Academic Publishers.
- Popper, Karl R. 1963. *Conjectures and refutations: the growth of scientific knowledge*. New York: Basic Books.
- Raffman, Diana. 2013. *Unruly words: a study of vague language*. Oxford University Press.
- Ransom, Keith, Voorspoels, Wouter, Perfors, Amy, & Dani, Navarro. 2017. A cognitive analysis of deception without lying. *Pages 992–997 of: Proceedings of the 39th Annual Conference of the Cognitive Science Society*. Cognitive Science Society.
- Russell, Bertrand. 1923. Vagueness. *The Australasian Journal of Psychology and Philosophy*, **1(2)**, 84–92.

- Sapir, Edward. 1944. Grading, a study in semantics. *Philosophy of science*, **11(2)**, 93–116.
- Sassoon, Galit W. 2012. A Typology of Multidimensional Adjectives. *Journal of Semantics*, **30**, 335–380.
- Solt, Stephanie. 2015. Vagueness and imprecision: Empirical foundations. *Annu. Rev. Linguist.*, **1(1)**, 107–127.
- Spector, Benjamin. 2013. Bare numerals and scalar implicatures. *Language and Linguistics Compass*, **7(5)**, 273–294.
- Tiersma, Peter. 2004. Did Clinton Lie: Defining “Sexual Relations”. *Chicago-Kent Law Review*, **79(3)**.
- van Deemter, Kees. 2009. Utility and language generation: The case of vagueness. *Journal of Philosophical Logic*, **38(6)**, 607–632.
- Verheyen, Steven, Dewil, Sabrina, & Egré, Paul. 2017. *Subjective meaning in gradable adjectives: The case of tall and heavy*. Manuscript, IJN, under review.
- Waismann, F. 1945. Verifiability. *Proceedings of the Aristotelian Society, Supplementary Volumes*, **19**, 119–150.
- Wason, Peter C. 1966. Reasoning. *New Horizons in Psychology*, **1**, 135–151.
- Wittgenstein, Ludwig. 1953. *Philosophical investigations*. Oxford: Blackwell.
- Wright, Crispin. 1995. The epistemic conception of vagueness. *The Southern journal of philosophy*, **33(S1)**, 133–160.