

Leiden University

Humanities

M.A. Linguistics

Pragmatics

Instructor: Dr. M. Westera MSc

**Project on
Pragmatics of Lying and Disinformation**

Oleg Astafyev

Student number: 3280551

18/12/2021

Table of contents:

1.0 Introduction	2
2.1 Defining a lie	2
2.2 Prototype definition of a lie	3
2.3 Revisiting Chen et al. 2013	5
2.4 The improved definition of a lie	8
3.0 Conclusion and questions for further research	10
References	11

1.0 Introduction

Lies and lying have long been of significant interest for a great many researchers from such fields as philosophy, psychology, forensics and linguistics (Meibauer 2018). This particular project was decided to have a focus on the very fundamental level, namely the definition of a lie. Thus, the central questions of this paper are:

1. What is/(are) the most effective definition(s) of a lie, i.e. what definition allows one to judge with a considerable degree of precision whether and to what extent an utterance is a lie?
2. If there is a way, how can such definition(s) be improved?

The project is structured as follows: Subsection 2.1 explains why defining a lie is an arduous task. The subsections 2.2 and 2.3 introduce the available definitions within the prototype approach and showcase how well they can be applied to a range of example-scenarios (all of the examples used in the project are invented from scratch unless indicated otherwise). The subsection 2.4 provides a side-by-side overview of the performance of the definitions in question and introduces a newly-synthesised, improved definition of a lie. Section 3.0 concludes the project and outlines a set of questions for further research.

2.1 Defining a lie

What is a lie? Something so usual and ingrained in our minds, it is understood intuitively. However, if one was asked to give a concise definition for it or describe the typical features of a lie, complications arise. The most straightforward approach would be to state that a lie is anything that is not true. Such definition is indeed concise, universal and straight to the point. The problem with this definition is that, perhaps unfortunately, truth-value is not the sole factor of a lie. In fact, a lie has an extremely convoluted nature. It exists in the domain of communication, which is regulated by society, in a sense that certain constraints are imposed on the speaker. Despite the putative freedom of speech one is only allowed to say this much without being frowned upon. Crucially, quite often blatantly telling the truth is a sure way to incur disfavour. At the same time, since the very childhood one is taught that telling lies is bad and wrong. Thus, a lie is somewhat of an intersection between falsity, moral objectionability, context, and, perhaps, a range of other different things.

A decent definition then must take into account all the nuanced aspects of our perception of what a lie is. This endeavour has proven to be challenging, since after centuries of debates there still is no universally accepted definition. Since listing and describing a vast number of theories is beyond the scope of this project, the reader is heartily referred to a more than extensive overview of the existing approaches to defining lying and the history behind them in *The Oxford Handbook of Lying* (Meibauer 2018).

The following discussion will focus on approaches which share the structural template of a prototype definition including a number of elements, i.e. certain features or conditions, which can be either present/fulfilled or not, and are inevitably representing the most crucial aspects of the subject defined.

2.2 Prototype definition of a lie

Coleman & Kay (1981) were the first to propose a *prototype* definition of a lie containing three main elements:

1. the falsity of an utterance
2. the speaker's belief in the falsity of an utterance
3. the speaker's intention to deceive.

The elements, albeit being dichotomous, are intended to allow for varying degree of membership or scalar framework of lies, i.e. utterances endowed with all the three elements are considered proper, full-fledged lies, whereas those lacking some elements are then less lie-like. It is important to note that elements do not equally contribute to the lie-likeness of an utterance, according to results of the C&K's experiment, the belief in falsity is the most important element, intent to deceive is the second most important and the actual falsity is the least important (1981:35).

The motivation for using this kind of a dynamic structural template with variable elements instead of the 'rigid' notion of lie, e.g. anything that is untrue is a lie, stems from the fact that speakers produce *technically false* utterances daily as members of a society. Almost any situation involving a polite remark may lead to the production of a *social lie* (C&K 1981:29). Consider [1]:

[1]

- a. What a lovely surprise!
- b. This is the best birthday present I've ever received!
- c. Come by anytime you want.

In day-to-day situations all these remarks may not be wholeheartedly meant, one could even say that most of the times they are not, thus the utterances are false and a speaker knows that, i.e. believes in their falsity. Two out of three elements of the definition are present, the only one lacking is the intent to deceive. It can be argued that the absence of the intent to deceive is not guaranteed, and, indeed, it is not impossible to imagine scenarios, where social lies are meant to deceive, since, ultimately, it is decided by the speaker. However, just as Grice presupposed that speakers and listeners are, in general, willing to cooperate, I dare to take the liberty of presupposing that social lies lack the intent to deceive, and are merely a part of the game of manners all members of society are obliged to play. Now consider [2]:

[2]

Josh did not do his homework. When asked to hand it in, he tells the teacher:

- My dog ate my homework.

The utterance is false, Josh believes it to be false and also intends to deceive the teacher in order to escape the punishment. All three elements are present, thus Josh's utterance is considered a full-fledged lie. Hence, the prototype definition by C&K allows one to

differentiate between [1] and [2], i.e. judge the lie-likeness in a more precise way, instead of just labelling them equally as lies due to the utterances being false.

The ability of a definition to account for a multitude of contexts and assign the degree of a lie with a considerable precision is what is argued here to be the *power* of a definition. C&K's prototype definition is quite powerful, since apart from the more or less straightforward examples above, it facilitates the judgement on the lie-likeness of the less clear-cut cases, such as jokes, metaphors and even lying with implicatures, e.g. lying by saying the truth.

Since both jokes and metaphors often involve technically false statements and are produced by speakers who are aware of that, they are inherently bound to scoring more than zero on the lie-likeness scale. The remaining element of intent to deceive, however, is, as a rule, absent, since the speaker intends the untruth to be uncovered, recognised by the hearer, for the desired communicative effect to be achieved. Thus, the existence of the third element in the definition prevents jokes and metaphors from being assigned to the same 'group' as other, far less amicably intended lies.

What about false implicatures, e.g. cases of lying by telling the truth? Consider [3]:

[3]

Suzy and Lucy are friends. Lucy has recently started dating Peter, and Suzy is jealous and upset, since her friend is not spending as much time with her as she used to. One day the girls arrange a sleepover, and Suzy is happy to spend some quality time with her friend. "It's just like in the good old times", she thinks to herself and decides that it is time to get her friend back. She tells Lucy that the day before she saw Peter with another girl in the cafe. Peter, in fact, was in the cafe with his sister, and Suzy is aware of this.

Two first elements of the definition will not detect any lie-like activity in Suzy's utterance, since Suzy is telling a truth. However, she does that with the intent to deceive Lucy, to make her think that Peter is having an affair. Thus, even though Suzy is telling a 100% truth, what she intentionally implicates by that is false, and the third element of the definition secures the increase in the lie score of the utterance.

Thus, C&K's prototype definition has proven itself robust and quite powerful. It allows for determining the extent to which an utterance is a lie. A salient advantage is the context-dependent nature of the elements 2 and 3, it facilitates the flexible application of the generalised template in a variety of situations, as illustrated by the examples above, ranging from usual conversational utterances to jokes, metaphors and more intricate cases like lying by telling the truth. However, there still remain some contexts in which the lie-likeness cannot be adequately determined by means of the C&K's definition:

[4]

a. A parent tells a kid that Santa is going to bring him/her presents if they behave well - 100% lie according to C&K definition, has all three elements:

1. falsity - Parents will be the providers of the presents, not Santa
2. belief in falsity - Parents know that they will be the providers of the presents
3. intent to deceive - Parents want the kid to believe that Santa is real and he will be the provider of the presents

b. Sergei Mavrodi tells people that they will become significantly richer if they invest into his company. - 100% lie according to C&K definition, has all three elements:

1. falsity - Mavrodi's company is a fraud, people won't get richer
2. belief in falsity - Mavrodi know for sure that the only person getting richer is himself
3. intent to deceive - Mavrodi wants people to believe they will get richer by investing

Per C&K definition both situations are full-fledged lies, however it is quite clear that [4a] is significantly different from [4b] at least in a sense that [4a] is far more harmless. Cases like [4a] are the most problematic, since it is due to the societal or cultural norms people might not necessarily perceive [4a] as a lie. Clearly, some other factors, apart from those mentioned as elements in the C&K's definition, are at play here, e.g. moral objectionability to an utterance, whether its effect has the harmful potential or not. This, in turn, complicates things even further, since coming up with a universal template for judging whether something is moral or not is, to say the least, far from an easy task.

Coleman & Kay pioneered the prototype approach with their definition of lying. They made it powerful enough to account for a wide range of contexts, however such complex cases like [4a] indicate that, perhaps, the definition is underspecified and calls for introducing new elements or reworking the present elements from a new perspective. This challenge, in a way, was taken up by Chen et al. 2013, whose approach will be introduced in the next section.

2.3 Revisiting Chen et al. 2013

The definition of a lie by Chen et al. 2013 is partially built on C&K's legacy, with a range of newly introduced and reworked old elements. Chen et al. noticed that even though C&K's definition indeed supports the idea of scalarity of the notion *lie*, the elements in question are not scalar themselves - an element is either present or absent. To compensate for this conceptual drawback they propose the version of the definition with only one bivalent element (See [5(A)]):

[5] *Pragmatic definition of lying*

Lying is a speech act that

(A) presupposes the semantic untruth of an assertion, *p*, and

(B) displays various degrees of lie-likeness depending on

i. the extent to which the speaker intends to conceal the untruth of *p*

(Other things being equal, the more the speaker intends to conceal the untruth of *p*, the more lie-like *p* is)

- ii. the extent to which p benefits self and/or hurts others
(Other things being equal, the more p benefits self, the more lie-like it is)
- iii. the extent to which p benefits other
(Other things being equal, the more p benefits other, the less lie-like it is).

(Chen et al. 2013:390)

Scalarity of the elements in [5(B)] and newly added variables of *benefit* and *intent to conceal the untruth* (which in essence are byproducts of C&K's *intent to deceive*) make a significant contribution to increasing the power of the definition, since due to these changes it allows for a much more fine-tuned assessment of the lie-likeness and is supposed to account for even more complex contexts. The only way to find out whether this is truly the case is to apply it to various scenarios.

While the benefit variables can indeed help in demarcating 'social' lies, which according to Chen et al. (2013) are often connected to politeness, from full-fledged lies that contribute to socially reprimanded acts, there are cases where it is extremely hard to think in categories of benefit, especially when no material benefit is involved. Consider [6]:

[6]

Parents tell their son that he was found in a cabbage patch.
Is it a lie?

The assertion is false. The intention to conceal the untruth of the assertion must be present to some degree, because otherwise parents could have just told the truth. Now when it comes to variables of benefit, everything gets more complicated. It is possible to claim that such an assertion benefits parents (self) in a sense that it liberates them from having a 'serious' talk with their son, which might be an issue for someone. Does such an assertion benefit their son (other)? Hardly so, as being uneducated on such topics like birth and sex could lead to disastrous consequences in his future life. Naturally, such claim is an exaggeration, however it does not take away the fact that the judgement process about the variables of benefit can become extremely complicated, as it is largely dependent on personal views and beliefs.

As was demonstrated in the previous section the element of *intent to deceive* is what allowed the definition by C&K to judge jokes and metaphors (those which involve technically false statements) adequately, i.e. not labelling them as 100% lies. The definition by Chen et al. lacks this element in principle, however it still is able to account for such contexts, due to the presence of the element B(i). Consider [7]:

[7]

She is a bomb.
Is it a lie?

The assertion is false. It is highly unlikely that the speaker intends to conceal the falsity, in fact, the untruth must be immediately recognisable for the hearer to perceive the utterance as an instance of a figurative language and thus produce the desired effect of language play.

The same line of reasoning holds for jokes. The hearer must be able to discern the playfulness of the speaker’s intention to properly engage with a joke.

In contrast with the C&K’s definition, the definition by Chen et al. 2013 can handle contexts like [4], mainly due to the added variable of benefit. See [8]:

[8] Analysis of [4] using the definition by Chen et al. in [5]

<i>Elements</i>	<i>Santa-story</i>	<i>Mavrodi-story</i>
(A) untruth	+	+
(B) i. the degree of intent to conceal the untruth	Probably lower, since nothing significantly dramatic will happen if the truth is found out	Probably higher, if the truth is found out, he will have to deal with the legal consequences
ii. the degree of benefit to self	Arguably, some degree of benefit for parents in a form of well-behaved children might be present	The highest degree possible, lots of money made fast
iii. the degree of benefit to other	Arguably, higher degree of benefit to other, the child has both the motivation to behave well, has something magical to believe in and to look forward to	X

As illustrated by [8], using the definition by Chen et al. allows for a much more adequate and accurate judgement of the given contexts - the Mavrodi case will have a significantly higher lie-score than Santa case. Hence, in these circumstances, the definition by Chen et al. clearly outperforms the definition by C&K.

Thus, the definition by Chen et al. indeed shows an improved performance in various contexts in comparison to the C&K’s definition. Not only is it able to handle [4], it also allows for a much more precise judgement of the contexts covered by C&K’s definition, by virtue of the elements being scalar, rather than bivalent. However, there still remain some problematic contexts even for the improved definition by Chen et al. Consider [9]:

[9]

An elderly woman comes home from her grocery shopping and sees that the window is broken. She sees the grandson of her friend named Jack passing by and asks whether he saw who did it. Jack smirks and tells her:

- Bobby and Pete were playing with a ball nearby.

Bobby and Pete were indeed playing with a ball that day, however they did not break the window. Did Jack lie?

This context becomes problematic for the definition, if we assume that Jack intentionally misled the elderly woman, i.e. made her think that Bobby and Pete broke the window and at the same time made sure that he cannot be accused of anything, since he merely stated the truth. Jack’s assertion indeed was true, hence Jack does not have to conceal anything, so elements A and B (i) of the definition render a zero lie-score for Jack’s assertion so far. B(ii) and B(iii) will thus be the only elements which will provide an opportunity to increase the total lie-score of his utterance, however it will still score lower than it should. This problem arises due to the fact that the definition by Chen et al. does not explicitly account for cases of false implicatures. Suggestions on how the definition can be improved will be introduced in the next section.

2.4 The improved definition of a lie

In the previous sections it was shown that while both definitions by C&K and Chen et al. are quite robust and have certain advantages over each other, there still remains a range of possible improvements. [10] provides a side-by-side overview of both definitions’ performance:

[10]

<i>Context</i>	<i>C&K</i>	<i>Chen et al.</i>
Social lies as in [1]	+ Accounted for due to the element of <i>intent to deceive</i>	+ Accounted for due to the element of <i>benefit</i>
Full-fledged lies as in [2]	+	+
Metaphors/Jokes (involving technically false statements)	+ Accounted for due to the element of <i>intent to deceive</i>	+ Accounted for due to the element of <i>intent to conceal the untruth</i> , also due to the element of <i>benefit</i>
False implicatures (e.g lying by saying the truth as in [3])	+ Accounted for due to the element of <i>intent to deceive</i>	-
Complex cases like [4]	- Unable to demarcate ‘harmless’ lies from harmful ones	+ Accounted for due to the element of <i>benefit</i>

As is evident from [10] the element of *intent to deceive* in C&K’s version plays a significant role, whereas in the version by Chen et al. variable of *benefit* has greater influence. In principle, the *intent to deceive* is indirectly embedded in the definition by Chen et al. via the interplay of the *intent to conceal the untruth* and the *benefit* variables. However, as

demonstrated by [9], when lying does not necessarily involve untruth, the lack of *intent to deceive* element is responsible for the incorrect lie-likeness prediction.

Perhaps, if both definitions were combined in a meaningful way, an improved, even more powerful, and accurately accounting for the greatest number of contexts definition could be created. The main problem of the discussed definitions is that the C&K's version is underspecified and the version by Chen et al. is, in a way, overspecified. Hence, the synthesised definition must aim at maintaining the balance between the two.

The strongest feature of C&K's definition is its simplicity. Even though scalarity of elements in the definition by Chen et al. increase its power, it also introduces complications in the judgement process, e.g. determining the extent of benefit to self/other is often far from a straightforward task. The bivalency of elements in the definition by C&K allows to avoid this at the expense of accuracy of the judgement. Finding a middle ground between bivalency and scalarity is the most difficult challenge on the way to a perfect definition, and as of now I have no solution to offer.

To construct the improved definition it is necessary to weigh the existing elements of the two definitions and decide which ones are worthy to be included. Although *the intent to deceive* element has shown to be important and reliable, the *intent to conceal the untruth* seems to be a preferable option, it facilitates better judgement on the range of contexts. The *belief in falsity* element allowed C&K's definition to accurately judge the contexts where lies are transmitted by virtue of the speaker's unawareness of the falsity and thus, it would be plausible to include it into the improved definition, however its function is already covered by the *intent to conceal the untruth*, i.e. speaker would not intend to conceal the untruth if he is unaware of it. Furthermore, since lying is an intrinsically pragmatic phenomenon, i.e. it is the speaker who lies and not the words themselves, it appears logical to reformulate the first element in a way that it is more bound to the speaker-meaning, i.e. exchange 'utterance' or an 'assertion' for 'explicature'. Moreover, it is necessary to expand the first element beyond just the explicature and add the implicature, this way the definition will handle cases of lying by telling the truth. Finally, in pursuance of decreasing the complexity of judging on the benefit to self or others, the B(ii) element of Chen et al.'s definition "the extent to which *p* benefits self and/or hurts others" needs to be decomposed, that is, the extent to which *p* hurts others should be an independent element. This would allow the improved definition to handle cases where hurting others is not necessarily linked to benefitting self.

An attempt to synthesise a definition that accounts for all the nuances mentioned above is presented in [11]. Even though the scalarity of the elements has shown to introduce complications in the judgment process, it still remains a preferred option, since it provides more accuracy and flexibility (at least in some contexts).

[11] The improved definition of lying

A lie

(A) presupposes the falsity of an explicature or an implicature, *p*, and

(B) displays various degrees of lie-likeness depending on

i. the extent to which the speaker intends to conceal the untruth of *p*

(Other things being equal, the more the speaker intends to conceal the untruth of p , the more lie-like p is)

ii. the extent to which p benefits self

(Other things being equal, the more p benefits self, the more lie-like it is)

iii. the extent to which p benefits other

(Other things being equal, the more p benefits other, the less lie-like it is).

iv. the extent to which p hurts other

(Other things being equal, the more p hurts other, the more lie-like it is).

Albeit lacking the absolute novelty in structure and most of the elements, this definition is supposed to comprise the most valuable features of the definitions by C&K and Chen et al., as well as account for the shortcomings of each. Does this definition represent a significant improvement in defining a lie? The reader is more than welcome to put it to test and find out.

An important aside must be made before proceeding to the concluding section. The entirety of the discussion above is to be understood in direct relation to the English language. The idea of a prototype definition is inherently bound to a specific language and/or culture. Unfortunately, the aspect of language specificity in lying cannot be covered in full detail in this project. Thus, the reader is referred to Hardin 2010, Mondry & Taylor 1992 and Nishimura 2018.

3.0 Conclusion and questions for further research

This project confirms the vastness of the challenge to define a lie. Forty years have passed since Coleman & Kay set the ground for the prototype approach, and it would not be an overstatement to say that the quest for the most effective definition is still not over. The modified fusion of the two most prominent, at least in the prototype tradition, definitions by C&K and Chen et al. in [11] aspires to be a humble contribution to this on-going mission. [11] is powerful and flexible enough to account for a wide range of contexts mentioned above while facilitating an arguably high degree of lie-likeness precision. Still, a number of points can be further improved. As discussed in subsection 2.4, a major step towards a better definition would be finding the golden mean between bivalency and scalarity of the elements, i.e. solving the under- vs overspecification conflict. Perhaps, fruitful observations can be made while studying other approaches to defining a lie, a prototype approach is just one out of a myriad. Another line of research which should be explored is the cross-linguistic investigation of lies. Looking at what is considered a lie by speakers of different languages should reveal whether even more generalised definitions or schemata with cross-linguistic validity can be introduced.

Bibliography

- Chen, Rong, Chunmei Hu, and Lin He. "Lying between English and Chinese: An intercultural comparative study." *Intercultural Pragmatics* 10.3: 375-401. 2013.
- Coleman, Linda, and Paul Kay. "Prototype semantics: The English word lie." *Language* 57.1: 26-44. 1981.
- Hardin, Karol J. "The Spanish notion of lie: revisiting Coleman and Kay." *Journal of pragmatics* 42.12: 3199-3213. 2010.
- Meibauer, Jörg, ed. *The Oxford handbook of lying*. Oxford Handbooks, 2018.
- Mondry, Henrietta, and John R. Taylor. "On lying in Russian." *Language & Communication* 12.2: 133-143. 1992
- Nishimura, F. "Lying in Different Cultures". *The Oxford handbook of lying*: 565-577. 2018