

University of Groningen

Editors' Review and Introduction

van Ditmarsch, Hans; Hendriks, Petra; Verbrugge, Rineke

Published in:
Topics in Cognitive Science

DOI:
[10.1111/tops.12492](https://doi.org/10.1111/tops.12492)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Final author's version (accepted by publisher, after peer review)

Publication date:
2020

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
van Ditmarsch, H., Hendriks, P., & Verbrugge, R. (2020). Editors' Review and Introduction: Lying in Logic, Language, and Cognition. *Topics in Cognitive Science*, 12(2 SI), 466–484.
<https://doi.org/10.1111/tops.12492>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Editors' Review and Introductions: Lying in Logic, Language and Cognition

Hans van Ditmarsch, Petra Hendriks, Rineke Verbrugge

Abstract

We describe some recent trends in research on lying from a multidisciplinary perspective, including logic, philosophy, linguistics, psychology, cognitive science, behavioral economics, and artificial intelligence. Furthermore, we outline the seven contributions to this special issue of *TopiCS*.

Keywords: Lying, deception, lie detection, cognitive load, cognitive resources, theory of mind.

1. Introduction

Lying is an everyday phenomenon that is central to social interaction. For most people, lying and being lied to are common experiences that may involve white lies to spare someone else's feelings, small dishonesties (for example, on a social network site) to improve one's self-presentation, and blatant falsities for financial or other profit.

A comprehensive account of lying requires a multidisciplinary approach. A speaker wishing to tell a lie must not only determine how to phrase the lie (which requires linguistic knowledge), but must also estimate how the listener will understand the utterance (requiring pragmatic skills), compute the change in beliefs brought about by the lie (studied in dynamic epistemic logic), and possess the cognitive abilities to carry out these computations (studied in cognitive science and developmental and cognitive psychology).

Likewise, a listener trying to determine whether the speaker is lying must use a variety of cues (e.g., amount of detail provided, pitch and internal consistency, see the meta-analysis of DePaulo et al., 2003) and will need to estimate how likely it is for the speaker to lie in the specific situation, in order to detect the lie. All of these aspects involved in the act of

lying are most likely connected. However, due to the different questions being asked in the various research fields, the different assumptions entertained, and the different methods and formalisms used (ranging from logical analysis, to computational modeling, to behavioral and psychophysiological experiments, to discourse analysis), revealing the connections between the different approaches is challenging.

Developing a comprehensive, multidisciplinary account of lying would not only extend the scope of the separate disciplines, but would also contribute to a more complete picture of human cognition, its capabilities and its limitations. According to Talwar, Gordon and Lee (2007), lying is theory of mind in action. If true, we expect to see this action reflected at all levels of analysis of the phenomenon of lying. For example, we then expect the linguistic choices made by the speaker to depend on his or her estimation of the interpretation by the listener, we expect epistemic reasoning to be bounded by cognitive limitations (see Flobbe, Verbrugge, Hendriks, & Krämer, 2008, for limitations on children's application of theory of mind), and we also expect to see a relation between the features of linguistic style that have been argued to be indicative of lying (e.g., a different use of pronouns and causal terms, see Hancock, Curry, Goorha, & Woodworth, 2008) and the aspects of language that have been found to be affected in groups with impaired theory of mind abilities, such as individuals with autism. Developing a comprehensive account of lying would also benefit technology and society at large, as understanding lying and deception in humans is essential for human-computer interaction, for example by making a virtual training environment more realistic in areas such as safety and security, education, digital games and health care (e.g., Nijholt et al., 2012).

In recent years, the phenomenon of lying and the human capacity and tendency to lie have led to the publication of several volumes and special issues. A recent addition is the broadly multidisciplinary Oxford Handbook on Lying (Meibauer 2018). Compared to the handbook, this topic more strongly focuses on the logical, philosophical, linguistic and psychological aspects of lying.

This issue grew out of an interdisciplinary workshop called "The Invention of Lying: Language, Logic & Cognition", which took place at the Lorentz Center in Leiden, The Netherlands on January 9-13, 2017. The seven articles in this issue have been written by workshop participants, together with colleagues. In the remainder of this

introduction, we will first give a short overview of different definitions of lying, followed by sections on the logic of lying, on linguistic aspects of lie telling and lie understanding, and on psychological approaches to lying and lie detection. Finally, we will briefly introduce the seven articles that comprise this topic.

2. What is lying?

Let us present a number of well-known definitions of lying, that have all been discussed at length in the literature (see e.g. Mahon, 2008; Fallis, 2009); we also review some in the next section.

You lie if you say something that is false.

That is what happens if Johnny ate the chocolate and his mother asks him if he did, and he says no.

You lie if you say something that you believe to be false.

However, sometimes you say something that you believe to be false but that is actually true. Priscilla's irate husband confronts you: "Where the hell is she?" You think that she is already at the railway station with her packed bags and you want to protect her. You therefore say: "She is at home." So the husband goes home and finds her, and bad things happen. In fact you said the truth, but you believed what you said to be false. So, were you lying?

You lie if you say something that you believe to be false with the intention to deceive.

The intentional aspect of lying is important but difficult to capture. In the previous two examples there was a clear intention to deceive: we may assume that Johnny wanted to be believed by his mother, otherwise, why bother trying to lie? You wanted to shield Priscilla from her irate husband. The intentional aspect can be seen as execution of a partial plan in

order to achieve a belief or ignorance objective. Now consider sarcasm: “Everybody is going to be rich.” You believe this to be false. But you do not intend to deceive. You know that nobody will believe you. Or consider statements about fictional characters: “Wonder Woman has great spirit.” You believe this to be false (or that it has no truth value), as Wonder Woman does not exist. There is no make-believe either, you just went to this movie with a friend and you discuss it with her.

You lie if you say something that you believe to be false with the intention that the addressee believes that it is true.

If the addressee believes what is said to be true, she is deceived. However, she may also be deceived if she does not believe what is said to be false, without believing what is said to be true. So this definition is a further restriction. However, every further precision presents a new trapdoor to fall through:

You lie if you say something that you believe to be false with the intention that the addressee believes that you believe that it is true.

Alice and Carol both know that Bob is not a spy, but they do not know this of each other. Alice now lies to Carol that Bob is a spy. Carol now believes that Alice believes that Bob is a spy. But she is skeptical: what Alice said does not affect her (strong) belief that Bob is not a spy. She merely believes that Alice has an incorrect belief. Her own belief about Alice is also incorrect.

You lie if you say something that you believe to be false with the intention that the addressee believes that you and the addressee commonly believe that it is true.

This definition is the outcome of further and further complicating the previous scenario. It is not merely important that Alice knows that Carol knows this, but also that Carol knows that Alice knows that Carol knows this, and that Alice knows that Carol knows that Alice knows that Carol knows this, and so on. The result of the ‘and so on’ is what is

called shared knowledge or common knowledge (van Ditmarsch, van Eijck, & Verbrugge, 2009). Instead of such arbitrarily long iterations of what people *know* about each other, we can also have arbitrarily long iterations of what people *believe* about each other: common belief plays an important role in lying phenomena as well.

Lying should be distinguished from other actions and other, similar, speech acts. You can *deceive* without lying, for example by non-verbal behavior, such as acting as if you do not see your friend walking out of the bar. Another form of verbal deception would be if you truthfully say something, but you believe that this will make the addressee believe a falsehood, such as: you believe and say p , you also believe that q is false, and that the addressee believes that p implies q ; what you say will make the addressee incorrectly believe q . Lying is different from *bluffing*. You are bluffing if you say that something is true but in fact you are uncertain whether it is true. "John is married". But you have no idea. If John turns out to be actually married, you will not be surprised (and also not when he is not married). This is different from lying, when you say what you believe to be false. In the case of a lie, if John turns out to be actually married, this is contrary to your beliefs, and you are surprised. A *white lie* has a moral connotation: you say it with the intention to deceive. But this intention is good and not bad. An *omission* is not a lie. For example: Jane says that she is looking forward to meet John - in the next 5 minutes, at a party - and that he is married. You know that he is not married but do not say anything to correct Jane. However, an omission under circumstances where a response was expected or obligatory can certainly be called *deception*. It is not always clear when a response is obligatory in everyday communication, as this is not always strictly procedural, unlike, for example, in games: If you hold the Queen of Hearts, and another player asks "Does somebody have the Queen of Hearts", then not responding is an omission that is deception.

3. Logic and philosophy: Can lying be precisely defined?

Few topics in logic or philosophy have roots as ancient as lying. Epimenides of Knossos (on Crete), active around 600 BC, is credited with the so-called Liar Paradox (see, e.g., Visser, 1984). Let us assume that what a liar says is always false. Epimenides, a Cretan,

now says that “Cretans are liars”. Either this is true, or this is false. If it is true, then Epimenides is a liar. And if Epimenides is a liar, “Cretans are liars” must be false. But if it is false, then (according to the classical analysis) Cretans tell the truth. So Epimenides tells the truth, namely that Cretans are liars, so, as he is Cretan, he is lying, so he is not telling the truth. Either way we end up with a contradiction.

Hang on: the negation of “All Cretans are liars” is not “All Cretans tell the truth” but “Some Cretan tells the truth”. This truth teller does not have to be Epimenides, so there is no paradox. However, it is easy to change the sentence into a real paradox. A more modern version is “This sentence is false”. With a similar analysis, this really is a paradox. Without delving further into lying paradoxes: these are more paradoxes of *meaning* than paradoxes about lying.

Respecting chronology, but not the topic, almost any work on lying will quote Augustine (4th century) on lying – this is not about paradoxes but is purely descriptive:

*That man lies, who has one thing in his mind and utters another in words.
The fault of him who lies, is the desire of deceiving in the uttering of his mind.*
(Augustine, 1956)

Liar paradoxes emerged again from the 12th century onward in the Scholastic tradition of the Middle Ages, for example in the work of Adam of Balsham. This later also included more complex versions that one could call “multi-agent lying” *avant la lettre*, for example:

*Socrates says 'What Plato says is false', and Plato says 'What Socrates says is true',
and neither says anything else. Is what Socrates says true or false?*
(Kneale & Kneale, 1962, p. 227-228).

Around 1900, Russell’s set-theoretical paradoxes, that laid the foundations for mathematical logic, led to a renewed interest in liar paradoxes. One might even see them as precursors of (unparadoxical) sentences like “This sentence cannot be proved” as employed by Gödel in incompleteness results in the foundations of mathematics (Gödel, 1931).

Descriptive investigations of lying in philosophy, and now mainly in formal logic,

inspired by Augustine's analysis, have been a thriving topic in the philosophical community for a long, long time since Siegler (1966), Bok (1978), and Mahon (2006). They focus on the truthfulness of statements, their intention, the various roles of speaker and addressee as mentioned above, and also higher-order interactions. The philosophical literature also clearly distinguishes between false propositions and propositions believed to be false by the speaker but in fact true, so that when you lie about them, you actually tell the truth. Whether such utterances can then still be called lies is questionable. Consider the following scenario involving delayed (moral) justification:

"Suppose that you believe that (not p) and that you lie that p . Later, you find out that your belief was mistaken because p was really true. You can then with some justification say "Ah, so I was not really lying." (Rott, 2003, translated from German).

What a relief! You do not feel bad now. It all depends on the setting. Consider Sartre's short story *The Wall* (Sartre 1937): a Republican prisoner believing to be lying to his Fascist jailers about the whereabouts of a comrade erroneously gives away his real hiding place (Mahon, 2008). Not much of a relief. The lying prisoner feels guilty now. Much is written on the morality of lying (Bok, 1978) and on its intentional aspect (Mahon, 2008). Frankfurt (2005) distinguished what he called bullshit from lying. A modal logical analysis of lying, bluffing, and their interaction with intention is found in Sakama, Caminada, and Herzig (2015), and a survey of such modal logics combining knowledge and intention is found in Sakama (2011). A survey of philosophical approaches to lying is Fallis (2009).

Let us now continue with analyses of lying in logics of *action*, where lying is such an action. The basic idea is that we cannot statically determine whether an utterance is a lie, but that its meaning lies in the transition from one state of information to another state of information. The problems then already start with the truth, not only with lies. Suppose that Petra likes cheese but that you do not know whether Petra likes cheese. Given that, you are now being told: "You don't know that Petra likes cheese." You may infer from this that:

"Petra likes cheese and you do not know whether Petra likes cheese."

So this is indeed true. However, after you have been informed of it, you now know that Petra likes cheese. The statement “*Petra likes cheese and you do not know whether Petra likes cheese*” is therefore no longer true. It was only true in the previous information state. It seems paradoxical to be informed of something that is true at the moment of utterance but that, as in this example, as a consequence of being uttered becomes false. Distinguishing states of information before the utterance and after the utterance resolves this seeming paradox. Such phenomena are called Moorean, after G. E. Moore (1942).

We can also investigate lying in this setting: The same utterance suffices as an example but now on the factual condition that Petra does not like cheese and that (as in the previous case) you do not know whether she likes cheese. The utterance is false, as the first conjunct “Petra likes cheese” is false and thus so is the entire conjunction. However, if you believe the speaker, after the utterance you will now have a false belief that Petra likes cheese, so again you do not believe the utterance itself anymore (you no longer believe that you don’t know whether Petra likes cheese), you merely (incorrectly) believe that it was true before the utterance. Such analysis of lying as an action that transforms states of information into new states of information have been investigated in, among others Baltag (2002), Steiner (2006), Baltag and Smets (2008), Kooi and Renne (2011), van Ditmarsch et al. (2012), van Ditmarsch (2014), and Ågotnes, van Ditmarsch, and Wang (2017). Baltag and collaborators have discussed lying by an external observer from the inception of dynamic epistemic logic onward (Baltag, 2002). Baltag and Smets (2008) and van Ditmarsch (2014) also discuss lying in logics with both knowledge and plausible belief.

In van Ditmarsch, van Eijck, Sietsma, and Wang (2012), the conscious update of Gerbrandy and Groeneveld (1997) is applied to model lying by an external observer. Kooi and Renne (2011) investigate the *cautious update*, which is called *lying to skeptical agents* by van Ditmarsch (2014). The issue here is consistency preservation for belief: if you hear a lie but you are already convinced of the opposite, i.e., if it is an *unbelievable update*, then you simply ignore it. Ågotnes, van Ditmarsch, and Wang (2017) carry the unbelievable update into the realm of higher-order aspects of lying, and analyze the so-called *true lie*: the lie that becomes the truth after being announced. The typical example is being informed that “ p or you believe that p ”, given that p is false and that you are uncertain about p . So it is a lie. Afterwards, you believe that p . So the lie has now become the truth. Liu and Wang (2013)

propose a dynamic epistemic logic of *agent types*, such that in the given logical language one can distinguish truthful, lying and bluffing agents, thus enabling some form of self-reference, so that a kind of liar paradox can be modelled. However, a fully-fledged logic for self-reference involving knowledge and actions seems still beyond the horizon.

4. Linguistics and developmental psychology: How do people lie?

Lying is a form of verbal behavior and as such – at least partially – belongs to the domain of linguistic pragmatics. A dominant view in pragmatics, originally put forward by Grice in 1967 (Grice, 1989), is that what is communicated often goes beyond what is literally said. For example, by saying “This suitcase weighs a ton!” the speaker probably does not mean that the suitcase literally weighs a ton, but rather that the suitcase feels very heavy. To account for the distinction between what is said and what is communicated, Grice proposes that speakers and listeners are generally cooperative and assume that participants in a conversation observe the so-called maxims of conversation. When the speaker observes these maxims, what is communicated is essentially identical to what is said. However, when the speaker flouts – i.e., overtly violates – these maxims, the listener must make a reasoning step (called an *implicature*) to infer what the speaker intends to communicate. One of Grice’s maxims of conversation is the first maxim of Quality, which says: “Do not say what you believe to be false” (Grice, 1989: 27). By uttering a statement that is false, the speaker does not behave in accordance with the maxim of Quality. According to the Gricean view, when speakers overtly violate the maxim of Quality by uttering a falsehood, this is because they wish to communicate something else and trust that the listener, by drawing the implicature, will be able to infer what they intend to communicate.

4.1 Linguistic falsehoods

In pragmatics, two types of falsehoods can be distinguished. On the one hand, there are cases such as the hyperbole above, where the speaker exaggerates the weight of the suitcase but wants the listener to recognize this as a falsehood. To give another example, a speaker who describes her lawyer by uttering the metaphoric statement “My lawyer is a shark” does not intend the listener to believe that her lawyer actually is a shark and has sharp teeth and can

swim well. Rather, her intention is for the listener to recognize that the utterance must be understood metaphorically and that the lawyer shares certain characteristics with a shark, for example being aggressive. Thus, the listener should recognize that the utterance is not literally true and therefore flouts Grice's first maxim of Quality. This should lead the listener to infer that the meaning the speaker intended to communicate was that her lawyer is rather aggressive. The listener should engage in a similar type of reasoning when hearing an ironic statement such as "What great weather!" on a rainy day. Understanding falsehoods such as hyperbole, metaphor, irony and other non-literal language thus depends on the listener's ability to recognize the speaker's intention to communicate something else than what was literally said.

A second type of falsehood involves cases where the speaker does *not* want the listener to recognize the falsehood. This holds for lying. If Johnny tells his mother "I did not eat the chocolate" when actually he did eat the chocolate, he does not intend his mother to believe anything else than what he just said. That is, he intends his mother to believe that his utterance is true and that he did not eat the chocolate. So the distinction between metaphoric and ironic statements, on the one hand, and lies, on the other hand, is in the intention of the speaker and whether the speaker does or does not intend the listener to recognize the falsehood.

Because of the focus of linguistic research in the past decades on the conventional rules of discourse and cooperative linguistic behavior, only few linguistic analyses exist of lying; notable exceptions are Meibauer (2011, 2014) and Dynel (2011). Meibauer (2014) emphasizes that, in order to produce a lie, a liar must make various linguistic decisions. The liar must, for example, decide on how explicit to be in expressing the lie (either as an assertion or as an implicature, the latter heavily reducing the risks involved in lying), what sentence type to use (since not only declarative sentences can convey lies, but also exclamative sentences and sentences containing a relative clause), and what syntactic construction, words and intonation to use. These linguistic decisions all influence the 'lying potential' of an utterance, that is, its capacity to convey a lie.

4.2 Learning to lie

Although lying is often seen as antisocial and negative behavior, it is also quite common in our society. Imagine that you are a parent and just watched your child playing in a school play. Even if your child's performance was terrible, you still say "Well done!" to protect your child's feelings and perhaps even boost their self-esteem. Falsehoods like these, which are told about unimportant matters in order to avoid hurting another person's feelings, are generally called white lies. White lies are told without malicious intent and serve prosocial purposes such as maintaining social relationships. Besides being frequent liars ourselves, most of us also occasionally appreciate lies by others. For example, we generally prefer others to be polite rather than truthful when commenting on the present we gave them, or the painting we did our very best on, or our new hairstyle.

Since lying is highly pervasive in daily life, it may seem reasonable to assume that children are able to lie as soon as they start talking. But is this indeed the case? The answer is no. Whereas adults are able to tell lies and frequently do so, young children cannot yet convincingly tell lies. Most 2-year-olds are still unable to tell lies in laboratory experiments (Evans & Lee, 2013), although 2-year-olds have occasionally been reported by their parents to produce lies, mostly as attempts to avoid punishment or for other self-interested reasons (Talwar & Lee, 2008). For example, Charles Darwin reports that his son, at that time 2 years and 8 months old, attempted to deceive him about a pinafore that he stained (Darwin, 1877, p. 292). While such parental reports provide important information, several studies have attempted to more systematically chart children's development of lying using experimental methods. A common experimental paradigm to investigate children's lie-telling behavior is the temptation-resistance paradigm. In this paradigm, children play a guessing game with the experimenter. The experimenter has a toy and the child has to guess what the toy is, based on its sound, but is instructed not to peek at the toy. During the experiment, the experimenter leaves the room, while the child's behavior is videotaped. When the experimenter returns, the child is asked whether or not (s)he peeked at the toy.

According to Talwar and Lee (2008), children's development of lie telling proceeds through various stages. Between 2 and 3 years of age, children start to deliberately make factually untrue statements, which however are still infrequent and may not yet be made with the intention to instill a false belief in the mind of the listener. In the next stage, from age 4, most children are able to tell lies, but often fail to maintain the lie in subsequent statements. Finally, from age 7 or 8, children are able to prevent 'semantic leakage' of the lie and can

maintain the lie in subsequent statements. Children's development of the identification and understanding of lies told by others seems to proceed in a similar gradual fashion (Bussey, 1999; Williams et al., 2016).

The development of lying seems to reflect increased cognitive development. Hence, the ability to tell lies is generally viewed as an important developmental milestone (e.g., De Villiers & De Villiers, 1978; Evans & Lee, 2013). Children's lie telling is directly related to their inhibitory control and working memory (Talwar & Lee, 2008): telling a lie requires that the child remembers the content of the lie, while simultaneously suppressing the truth. Furthermore, several studies have found that the ability to tell lies requires theory of mind. This is the cognitive ability to understand that other people have mental states (beliefs, desires, intentions) that may differ from one's own (Wimmer & Perner, 1983). It has been argued that first-order theory of mind (i.e., the ability to infer what another person thinks, which is attested from around 4 years of age in explicit false-belief tasks) is needed to be able to tell lies, and second-order theory of mind (i.e., the ability to infer what one person thinks about another person's thoughts, which begins to emerge around 6 years of age) is needed to be able to maintain a lie over multiple statements (e.g., Talwar & Lee, 2008). Thus, there seems to be a close connection between children's development of lie telling and their development of theory of mind.

Regarding the development of theory of mind, a puzzling observation is that 15-month-old infants already appear to possess some implicit theory-of-mind abilities, suggesting that theory of mind starts to develop quite early in life. To account for the discrepancy between these early theory-of-mind abilities and children's later difficulties with explicit false-belief tasks, a dual-system account has been proposed for theory of mind (Apperly & Butterfill, 2009). The early-available system is argued to be fast and automatic but limited and inflexible, whereas the later-developing system is argued to be highly flexible but slower and more effortful. Similar dual-system/two-stage accounts have also been proposed for reasoning and social cognition (e.g., Kahneman & Frederick, 2002), grammar (Hendriks, 2014) and referential communication (e.g., Epley, Morewedge & Keysar, 2004). A relevant question is whether these accounts can be connected in such a way that this could explain the various stages in children's development of lie telling and lie understanding, a phenomenon that lies at the interface between these domains.

4.3 Verbal cues to lying

Because lying is a verbal act, one might wonder whether there are any verbal cues that distinguish lying from truth telling. Due to the increasing availability of digital language data and specialized software to analyze these data, this question has received increasing attention in recent years.

One prominent idea is that, because liars are fabricating facts, they do not know the details of what they are talking about and hence can be expected to use fewer words and provide fewer details than truth tellers. Indeed, this expectation was borne out by some studies (e.g., DePaulo et al. 2003), but not by all studies. For example, Hancock et al. (2008) and Braun, Van Swol, and Vang (2015) found that liars instead use more words than truth tellers, a phenomenon that Van Swol, Braun and Malhotra (2012) have called the *Pinocchio effect*. This difference in findings between studies could be due to the fact that different communicative contexts and modes of communication offer different possibilities and pose different constraints on lying. In the study of Hancock et al. (2008), for example, liars communicated through a computer interface with a conversational partner they did not know and were instructed to lie about a personal topic. Since the liars' false statements were not verifiable by the addressee in this experimental setting, there was no risk that the lie would be detected, which could have led the liars to provide more details to appear credible to the addressee (see Hancock et al, 2008, for discussion of this matter).

Another idea that has received much attention is that liars, due to lack of personal experience with the events being lied about, or perhaps from a desire to distance themselves from the lie being told, use fewer first-person pronouns (e.g., *I, me, my*) and more third-person pronouns (e.g., *she, they, their*) than truth tellers. This has been confirmed in several studies (e.g., Newman et al., 2003; DePaulo et al. 2003; Hancock et al., 2008). Some studies also found a higher use of negative emotion words (e.g., *hate, worthless*) in lies compared to truths (e.g., Newman et al., 2003), which has been explained from the assumption that liars feel guilty about their lie. Furthermore, lies have been argued to contain fewer negations (*no, not, never*) and fewer exclusive discourse markers such as *but, except* and *without* (see Newman et al., 2003), possibly because their use would require the speaker to provide more details and not only specify what did happen, but also what did not happen. In addition to these features,

several other linguistic features have been investigated. Just as for the number of words liars use, also for these other linguistic features results are mixed: not all studies find the expected differences between lying and truth telling, and some studies even find effects in the opposite direction. This suggests that whether a particular linguistic feature signals a lie, and thus informs us of how people lie, heavily depends on its context of use (cf. Burgoon, 2018).

5. Cognitive psychology and behavioral economics: When and why do people lie?

Even though people in general want to be moral and to speak the truth, adults report telling on average one or two lies a day (De Paulo et al., 2003; DePaulo & Kashby, 1998). Why do they do that and what kinds of lies are prevalent?

5.1 Prevalence of lying by adults

Many everyday lies are not especially harmful to others. An example of such a white lie would be to tell your friend after her concert: “That was a really special rendition of the aria ‘Voi che sapete’ ”, even though you think that she was singing out of tune. The white lie is meant to show support for your friend and thereby to cement the relationship between the two of you. More in general, in game experiments, altruistic and pro-social people turn out to prefer using white lies only in cases when that would benefit the other player and thus signal cooperation (Biziou-van Pol et al., 2015).

As in the singing example above, people in close relations or in everyday situations often tell lies about feelings and opinions rather than about world facts. They claim to feel more positively or to agree more with the listener than they do in reality (DePaulo & Kashby, 1998), in order to appear nicer or more sophisticated than they are (DePaulo et al., 2003). Moral evaluations of such white lies differ in the literature. Some authors claim that all lies, including white lies, harm close friendships and relationships because when using white lies about who they are and what they feel, people are not understood or accepted by the other for who they really believe themselves to be (DePaulo & Kashby, 1998). Other studies find that pro-social lies may actually

strengthen trust between people (Levine & Schweitzer, 2015). Most people have moral qualms about their own “innocuous” lies and sometimes fool themselves into thinking that they did not lie at all (Merzel et al., 2015).

The case is quite different for more serious lies about world facts, which are meant to hide, for example, that the speaker cheated on a test or in a relationship. Such lies are quite rare and when discovered, they may have grave consequences for the liar. For example, a lawyer may incur a prison sentence after it is discovered that he lied under oath in order to hide his powerful employer’s transgressions. Still, in experiments in behavioral economics, most participants appear to lie quite easily to their opponent in cases where it benefits them and where their lie cannot be detected: percentages such as 36%-44% of all messages being lies are quite normal (Gneezy, 2005; Sutter, 2008). However, in a meta-study, it was found that the propensity to lie is far from evenly divided over the population. Over many different experiments, participants appear to be divided into a relatively small group who almost always lie versus a relatively large group who almost always speak the truth (Rosenbaum, Billinger, & Stieglitz, 2014).

5.2 Producing and maintaining a lie requires cognitive resources

Intuitively, lying takes more effort than telling the truth: It is easier for a speaker to access a true memory than to construct a false but consistent story, especially when taking the listener’s beliefs and previous knowledge into account as well. In general, it has been shown that understanding the concept of lying and being able to maintain a lie over time requires second-order theory of mind (Gordon, & Lee, 2007), which develops only after age 5 and remains effortful even for adults (Flobbe et al., 2008).

Many experiments have provided support for the added cognitive load of lying. For example, Vrij et al. (2008) showed that when asked to speak about a recent negative event in chronologically reverse order, liars had a significantly harder time to tell their story than truth tellers, including slower speech, less details, more speech errors and more eye blinks. Van ‘t Veer, Stel & Van Beest (2014) showed that participants under high cognitive load truthfully reported the outcomes of the rolls of dies, whereas participants under low cognitive load often reported higher outcomes than the true ones

in order to gain more money. Consequently, much recent literature on lie detection has focused on increasing cognitive load in order to catch liars.

Currently, several researchers present a more nuanced view, namely, that the cognitive load of lying as compared to telling the truth depends on several aspects of the task, the speaker, and the context (Sporer, 2006; Burgoon, 2015).

5.3 How to catch a lie?

Catching mundane everyday lies is not the focus of much literature on lie detection: The consequences of catching those lies are relatively small and in the case of white lies, may not even be beneficial for the parties involved. In court, however, accurately catching lies is important: Both false negatives and false positives have serious consequences for the people involved and possibly even for society. A recent overview of the topic of lie detection can be found in Granhag, Vrij and Verschuere (2015). Lie detection is surprisingly hard. It turns out that in laboratory experiments, the differences in behavior between liars and truth tellers are only barely perceptible (DePaulo et al., 2003). Therefore, it is no surprise that people, amateurs and experts alike, can detect whether somebody is lying to them at only slightly better than chance level (Hartwig & Bond, 2011).

Still, there has been a large literature about a plethora of cues to deception. Ekman (1985) argued that liars often experience different emotions than truth tellers: fear of being caught and guilt about the deception or, at the other extreme, 'duping delight'. All these emotions may 'leak' in such a way that the receiver catches them. The liar may also try to simulate emotions such as enjoyment that the receiver may recognize as artificial. The question what emotion-related cues are relevant depends very much on the context and the relationship between speaker and listener.

DePaulo et al. (2003) present a meta-analysis of cues to deception in self-presentation. They posit that speakers who present deceptive self-images do so with less conviction, more deliberateness, and less details than when being truthful. The relevant cues of lying are longer response latency, slower speech, less immediacy, more tension, and less pleasant and positive impression. These predicted cues are

corroborated by the experiments in their meta-study, with the caveat that many liars become more fluent and less easy to catch by using the strategy of basing their deceitful self-presentation on their actual experiences, with only some crucial details changed (DePaulo, 2003).

6 What lies in the future: Can there be artificial liars?

Computational systems may assist us in detecting deception, for example by checking for certain telltale elements in texts (Hancock et al., 2008). In addition to assisting in lie detection, artificially intelligent systems may also try their hand at lying themselves (Chakraborti & Kambhampati, 2019). Isaac and Bridewell (2017) present the view that robots sometimes need to lie in order to fulfill their design goals to optimally support people. Sarkadi (2018), in contrast, emphasizes the risk of an artificially intelligent agent being able to outsmart its human interrogators who want to test whether the agent is truthful; after all, artificial systems can be programmed to have even more sophisticated theory of mind than the two levels that we humans usually require to maintain lies and to detect that others lie to us. Thus, artificially intelligent systems could take the spread of false information and manipulation of people's opinions to an even more worrying level than today's human influencers.

7 Overview of the articles in the special issue

People frequently lie for prosocial reasons, for example when they want to surprise someone with a nice gift. In the article "Care to Share? Children's cognitive skills and concealing responses to a parent", authors *Jennifer Lavoie and Victoria Talwar* experimentally investigated this type of prosocial lie telling by children. In their study, children made a surprise gift for their parent while the parent was away, and were later questioned by their parent about what they had been doing. Examining children in a wide age range from 4 to 11 years old, Lavoie and Talwar found that the children gave a variety of answers in response to the questions asked by their parent: while some

children were unable to conceal the secret and disclosed information about the activity to their parent, other children concealed the secret from their parent. Children's responses highlight the fact that concealment methods occur on a spectrum, ranging from partial concealment ('leaking' verbal evidence) to passively hiding information (omission) to actively giving incorrect information (lying). Whereas the youngest children tended to disclose information about the secret, the oldest children tended to use lying to conceal the secret. Importantly, the study suggests that as children's theory of mind abilities and working memory improve, their abilities to conceal information from others develop.

In the article "Being deceived: Information asymmetry in second-order false belief tasks", the authors *Torben Braüner, Patrick Blackburn and Irina Polyanskaya* also relate children's theory of mind skills to deception, but in a quite different context. Second-order false-belief tasks are often used to test children's second-order theory of mind development. For example, in the typical 'birthday puppy story' they have to correctly infer that, although his mother has tried to deceive Frederik into thinking otherwise, Frederik *knows* that he'll get a mountain-bike for his birthday but Mum *doesn't know* that he *knows*. The article gives an interesting logical analysis of the reasoning needed to solve four types of second-order false belief tasks, distinguished on two dimensions: whether or not a story character is deceived, and whether or not the story hinges on facts in the world changing (an ice-cream van changes location) or only the characters' beliefs (a child inadvertently gets a peek preview of his birthday present). It turns out that for all four types of story, the principle of inertia, first discussed in the context of first-order false-belief tasks by Stenning and Van Lambalgen (2008), plays an important role. In an experimental study, the 'ice cream van' story, without deception but with a location change, is much harder than the other three types of stories for children with autism spectrum disorder (but not for typically developing children) – success on that task even predicts success on all three others for children with autism spectrum disorder.

In the article "Shedding light on keeping people in the dark", author *Don Fallis* analyzes from a logical and epistemological perspective what it means to hide or conceal information from others: keeping information in the dark means hiding or concealing

that information. Prior investigations incorrectly focused on possible means of keeping someone in the dark rather than on what it is to keep someone in the dark. In this article, Don Fallis argues that you keep X in the dark about a proposition P if and only if you intentionally cause X not to have a true belief that P. He also extends this analysis from a categorical belief model of epistemic states to a credence (or degree of belief) model.

Not all lies are overtly communicated in the form of an assertion that is believed to be false by the speaker. Sometimes, lies can also be covertly communicated. In the article "To say the least: Where deceptively withholding information ends and lying begins", *Marta Dynel* addresses the notion of deceptively withholding information. She argues that deceptively withholding information is a form of deception that rests on a covert violation of Grice's first maxim of Quantity: "Be as informative as is required". When violating this maxim, the speaker fails to provide information that the speaker believes the listener requires, given the question under discussion (i.e., what is relevant in the context). Because the speaker did not tell the whole truth, the listener develops a false belief about what was not said. This linguistic analysis of deceptively withholding information is an addition to Grice's framework of cooperative communication and leads to a reanalysis of some of the examples presented in the literature as withholding information, analyzing them instead as lies.

Some contexts may invite uncooperative communicative behavior such as lying or being uninformative, for example when speaker and addressee have competing goals. In the article "Strategies of deception: Under-informativity, uninformativity and lies – misleading with different kinds of implicature", the authors *Michael Franke, Giulio Dulcinati and Nausicaa Pouscoulous* investigate to what extent speakers in an uncooperative communicative context use implicatures to deceive their addressee. Playing a signaling game, the participants in their online experiment had to select a sentence completion to describe one card from a pair of cards to either help a virtual co-player to select this same card in the cooperative version of the strategic communication game, or compete against this virtual co-player and trick him into selecting the other card in the competitive versions of the game. In the competitive versions of the game, participants were found to use literally false sentences (lies) as well as implicatures to deceive their co-players. These results suggest that, while speakers themselves are

uncooperative in competitive contexts and produce lies to deceive their addressees, they expect their addressee to be cooperative and infer scalar implicatures from the speaker's utterances.

In the article "Memory-based deception detection", the authors *Linda Geven, Gershon Ben-Shakhar, Merel Kindt and Bruno Verschuere* demonstrate experimentally that lying is not restricted to explicitly instructed cheating, but that it can also be observed for self-initiated cheating. Clearly, from a cognitive perspective, lying can be regarded as a complex process requiring the interplay of several executive functions. Many investigations suggest that computerized paradigms can reliably assess the cognitive burden of lying, with large reaction time differences between lying and truth telling. These studies, however, lack a key ingredient of real-life deception, namely self-initiated behavior. Research participants have typically been instructed to commit a mock crime and conceal critical information, whereas in real life, people freely choose whether or not to engage in antisocial behavior. This article compares the former to the latter. In this article, a large number of participants engaged in a trivia quiz and were provided with a monetary incentive for high accuracy performance. Participants were randomly allocated to either a condition where they were instructed to cheat on the quiz or to a condition in which they were provided with the opportunity to cheat. Assessments of response times in a subsequent Concealed Information Test revealed that both instructed cheaters and self-initiated cheaters showed the expected slowing of the response time for concealed information.

Whereas some accounts of lying argue that dishonesty is more cognitively demanding than honesty, other accounts suggest that a person's default response depends on the salience of the action. If the impact of a person's action on others is *not* salient, as in the situation studied by *Yoella Bereby-Meyer, Sayuri Hayakawa, Shaul Shalvi, Joanna Corey, Albert Costa and Boaz Keysar* and reported on in the article "Honesty speaks a second language", according to the latter account the dominant impulse is to behave in a self-serving manner. Participants speaking different foreign languages were asked to roll a die, the outcome of which was only known to them, and were paid a reward according to the outcome they reported. They were found to inflate their outcomes less when using a foreign language compared to using their native

language. According to the authors, this can be explained by a dual-system account that assumes that lying for reasons of self-interest is the automatic, affective response in this situation, but that honesty can be promoted if processing relies less on the affective system, such as when speaking a foreign language.

Acknowledgments

We would like to thank all attendees and speakers at the interdisciplinary Lorentz workshop on *Lying: Language, Logic and Cognition* in January 2017 in Leiden for their inspiring presentations and contributions to the discussions. We especially thank Eline Pollaert of the Lorentz Center in Leiden, who supported the organization of the workshop in every possible way. Furthermore, we would like to thank the authors who submitted papers to this issue and the reviewers, who have done a thorough job of making fruitful suggestions from the perspectives of their different fields of expertise. Finally, we would like to thank Wayne Gray, Caroline Verdier and Marta Ahmed of *Topics in Cognitive Science* for their continuous support and their confidence.

References

Ågotnes, T., van Ditmarsch, H., & Wang, Y. (2018). True lies. *Synthese*, 195, 4581–4615.

Apperly, I., & Butterfill, S. (2009). Do humans have two systems to track beliefs and belief-like states? *Psychological Review*, 116, 953–970.

Augustine (1956). De Mendacio. In: P. Schaff (ed.), *A select library of the Nicene and Post-Nicene fathers of the Christian church*, volume 3. Grand Rapids, MI: Eerdmans.

Baltag, A. (2002). A logic for suspicious players: Epistemic actions and belief updates in games. *Bulletin of Economic Research*, 54(1), 1-45.

Baltag, A. and Smets, S. (2008). The logic of conditional doxastic actions. In *New perspectives on games and interaction*, Texts in Logic and Games 4 (pp. 9–31). Amsterdam: Amsterdam University Press.

Biziou-van-Pol, L., Haenen, J., Novaro, A., Occhipinti Liberman, A., & Capraro, V. (2015). Does telling white lies signal pro-social preferences? *Judgment and Decision Making*, 10(6), 538–548.

Bok, S. (1978). *Lying: Moral choice in public and private life*. New York: Random House.

Braun, M. T., Van Swol, L. M., & Vang, L. (2015). His lips are moving: Pinocchio effect and other lexical indicators of political deceptions. *Discourse Processes*, 52(1), 1-20.

Burgoon, J. K. (2015). When is deceptive message production more effortful than truth-telling? A baker's dozen of moderators. *Frontiers in Psychology*, 6.

Burgoon, J. K. (2018). Predicting veracity from linguistic indicators. *Journal of Language and Social Psychology*, 37(6), 603-631.

Bussey, K. (1999). Children's categorization and evaluation of different types of lies and truths. *Child Development*, 70(6), 1338-1347.

Chakraborti, T., & Kambhampati, S. (2019). (When) can AI bots lie? In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 53–59).

Darwin C. (1877). A biographical sketch of an infant. *Mind*, 2, 285–294.

DePaulo, B. M., & Kashy, D. A. (1998). Everyday lies in close and casual relationships. *Journal of Personality and Social Psychology*, 74, 63–79.

DePaulo, B.M., Lindsay, J.J., Malone, B.E., Muhlenbruck, L., Charlton, K., & Cooper, H. (2003).

Cues to deception. *Psychological Bulletin*, 129(1), 74–118.

De Villiers, J.G. & De Villiers. P.A. (1978). *Language acquisition*. Cambridge, MA: Harvard UP, 1978.

van Ditmarsch, H., van Eijck, J., Sietsma, F. & Wang, Y. (2012). On the logic of lying. In: J. van Eijck and R. Verbrugge (eds.), *Games, actions and social software* (pp. 41-72). Berlin: Springer LNCS 7010.

van Ditmarsch, H., van Eijck, J., & Verbrugge, R. (2009). Common knowledge and common belief. In J. van Eijck and R. Verbrugge (eds.), *Discourses on social software*, Texts in Logic and Games 5 (pp. 99–122). Amsterdam: Amsterdam University Press.

van Ditmarsch, H. (2014). Dynamics of lying. *Synthese*, 191(5), 745–777.

Dynel, M. (2011). A web of deceit: A neo-Gricean view on types of verbal deception. *International Review of Pragmatics*, 3(2), 139–167.

Epley, N., Morewedge, C. K., & Keysar, B. (2004). Perspective taking in children and adults: Equivalent egocentrism but differential correction. *Journal of Experimental Social Psychology*, 40(6), 760-768.

Evans, A. D., & Lee, K. (2013). Emergence of lying in very young children. *Developmental Psychology*, 49(10), 1958–1963.

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

Flobbe, L., Verbrugge, R., Hendriks, P., & Krämer, I. (2008). Children's application of theory of mind in reasoning and language. *Journal of Logic, Language and Information*, 17(4), 417–442.

Frankfurt, H.G. (2005). *On bullshit*. Princeton, NJ: Princeton University Press.

Gerbrandy, J.D., & Groeneveld, W. (1997). Reasoning about information change. *Journal of Logic, Language, and Information*, 6, 147–169.

Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, 95(1), 384–394.

Gödel, K. (1931). Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38(1), 173–198.

Granhag, P. A., Vrij, A., & Verschuere, B. (Eds.), (2015). *Detecting deception: Current challenges and cognitive approaches*. London: Wiley-Blackwell.

Grice, H. P. [1989]. Logic and conversation. In *Studies in the way of words* (pp. 22-40). Cambridge, MA: Harvard University Press.

Hancock, J.T., Curry, L.E., Goorha, S. & Woodworth, M. (2008). On lying and being lied to: A linguistic analysis of deception in computer-mediated communication. *Discourse Processes*, 45(1), 1–23.

Hartwig, M., & Bond Jr, C. F. (2014). Lie detection from multiple cues: A meta-analysis. *Applied Cognitive Psychology*, 28(5), 661–676.

Hartwig, M., & Bond Jr, C. F. (2011). Why do lie-catchers fail? A lens model meta-analysis of human lie judgments. *Psychological Bulletin*, 137(4), 643–659.

Hendriks, P. (2014). *Asymmetries between language production and comprehension*. Studies in Theoretical Psycholinguistics, Vol. 42. Dordrecht: Springer.

- Isaac, A. M., & Bridewell, W. (2017). Why robots need to deceive (and how). In: P. Lin, K. Abney, & R. Jenkins (Eds.): *Robot ethics 2.0: From autonomous cars to artificial intelligence* (pp. 157–172). Oxford: Oxford University Press.
- Kahneman D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgement. In T. Gilovich, D. Griffin, & D. Kahneman (eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49–81). Cambridge, UK: Cambridge University Press.
- Kneale, W., & Kneale, M. (1962). *The development of logic*. Oxford: Clarendon Press.
- Kooi, B., & Renne, B. (2011). Arrow update logic. *Review of Symbolic Logic*, 4(4), 536–559.
- Levine, E. E., & Schweitzer, M. (2015). Prosocial lies: When deception breeds trust. *Organizational Behavior and Human Decision Processes*, 26, 88–106.
- Liu, F., & Wang, Y. (2013). Reasoning about agent types and the hardest logic puzzle ever. *Minds and Machines*, 23(1), 123–161.
- Mahon, J.E. (2006). Two definitions of lying. *Journal of Applied Philosophy*, 22(2), 21–230.
- Mahon, J.E. (2008) The definition of lying and deception. *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/lying-definition/>.
- Meibauer, J. (2011). On lying: Intentionality, implicature, and imprecision. *Intercultural Pragmatics* 8(2), 277–292.
- Meibauer, J. (2014). *Lying at the semantics-pragmatics interface*. Berlin: De Gruyter Mouton.
- Meibauer, J. (ed.), (2018). *The Oxford handbook of lying*. Oxford: Oxford University Press.

- Merzel, A., Ritov, I., Kareev, Y., & Avrahami, J. (2015). Binding lies. *Frontiers in Psychology*, 6, 1566.
- Moore, G. E. (1942). A reply to my critics. In: P.A. Schilpp (ed.), *The philosophy of G.E. Moore*, Evanston, IL: Northwestern University Press.
- Newman, M. L., Pennebaker, J. W., Berry, D. S., & Richards, J. M. (2003). Lying words: Predicting deception from linguistic styles. *Personality and Social Psychology Bulletin*, 29(5), 665-675.
- Nijholt, A. et al. (2012). Trends & controversies: Computational deception and noncooperation. *Intelligent Systems, IEEE*, 27(6), 60-75.
- Rosenbaum, S. M., Billinger, S., & Stieglitz, N. (2014). Let's be honest: A review of experimental evidence of honesty and truth-telling. *Journal of Economic Psychology*, 45, 181-196.
- Rott, H. (2003). Der Wert der Wahrheit. In: M. Mayer (ed.), *Kulturen der Lüge*, (pp. 7-34). Köln und Weimar: Böhlau-Verlag.
- Sakama, C., Caminada, M., & Herzig, A. (2015). A formal account of dishonesty. *Logic Journal of the IGPL*, 23(2), 259-294.
- Sarkadi, S. (2018). Deception. In: *Proceedings of the 27th International Joint Conference on Artificial Intelligence* (pp. 5781-5782). Cambridge: AAAI Press.
- Sartre, J-P. (1937). Le mur, *La Nouvelle Revue Française*, 286, 38-62.
- Sporer, S. L. (2016). Deception and cognitive load: Expanding our horizon with a working memory model. *Frontiers in Psychology*, 7.

Sutter, M. (2008). Deception through telling the truth?! Experimental evidence from individuals and teams. *The Economic Journal*, 119 (534), 47–60.

van Swol, L. M., Braun, M. T., & Malhotra, D. (2012). Evidence for the Pinocchio effect: Linguistic differences between lies, deception by omissions, and truths. *Discourse Processes*, 49(2), 79-106.

Talwar, V., & Lee, K. (2008). Social and cognitive correlates of children's lying behavior (2008). *Child Development*, 79, (4), 866–881.

Talwar, V., Gordon, H., & Lee, K. (2007). Lying in the elementary school years: Verbal deception and its relation to second-order false-belief understanding. *Developmental Psychology*, 43(3), 804–810.

van 't Veer, A., Stel, M., & van Beest, I. (2014). Limited capacity to lie: Cognitive load interferes with being dishonest. *Judgment and Decision Making*, 9(3), 199–206.

Visser, A. (1984). Semantics and the liar paradox. *Journal of Philosophical Logic*, 13(2), 181–212.

Vrij, A., Mann, S. A., Fisher, R. P., Leal, S., Milne, R., & Bull, R. (2008). Increasing cognitive load to facilitate lie detection: The benefit of recalling an event in reverse order. *Law and Human Behavior*, 32 (3), 253–265.

Williams, S., Leduc, K., Crossman, A., & Talwar, V. (2016). Young deceivers: Executive functioning and antisocial lie-telling in preschool aged children. *Infant and Child Development*, 26(1), e1956.

Wimmer, H., & Perner, J. (1983). Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 13(1), 103–128.

Papers in this topic

Bereby-Meyer, Y., Hayakawa, S., Shalvi, S., Corey, J., Costa, A., & Keysar, B. (YEAR). Honesty speaks a second language. *Topics in Cognitive Science*, X(X), PAGE. DOI URL.

Braüner, T., Blackburn, P., & Polyanskaya, I. (YEAR). Being deceived: Information asymmetry in second-order false belief tasks. *Topics in Cognitive Science*, X(X), PAGE. DOI URL.

van Ditmarsch, H., Hendriks, P., & Verbrugge, R. (YEAR). Editors' review and introductions: Lying in logic, language and cognition. *Topics in Cognitive Science*, X(X), PAGE. DOI URL.

Dynel, M. (YEAR). To say the least: Where deceptively withholding information ends and lying begins. *Topics in Cognitive Science*, X(X), PAGE. DOI URL.

Fallis, D. (YEAR). Shedding light on keeping people in the dark. *Topics in Cognitive Science*, X(X), PAGE. DOI URL.

Franke, M., Dulcinati, G., & Pouscoulous, N. (YEAR). Strategies of deception: Under-informativity, uninformativity and lies – misleading with different kinds of implicature. *Topics in Cognitive Science*, X(X), PAGE. DOI URL.

Geven, L., Ben-Shakhar, G., Kindt, M., & Verschuere, B. (YEAR). Memory-based deception detection. *Topics in Cognitive Science*, X(X), PAGE. DOI URL.

Lavoie, J. & Talwar, V. (YEAR). Care to share? Children's cognitive skills and concealing responses to a parent. *Topics in Cognitive Science*, X(X), PAGE. DOI URL.