

Empirically Investigating the Concept of Lying

Alex Wiegmann¹ · Ronja Rutschmann^{1,2} · Pascale Willemsen³

Received: 9 March 2017 / Revised: 16 June 2017 / Accepted: 20 June 2017 /
Published online: 28 August 2017
© ICPR 2017

Abstract Lying is an everyday moral phenomenon about which philosophers have written a lot. Not only the moral status of lying has been intensively discussed but also what it means to lie in the first place. Perhaps the most important criterion for an adequate definition of lying is that it fits with people’s understanding and use of this concept. In this light, it comes as a surprise that researchers only recently started to empirically investigate the folk concept of lying. In this paper, we describe three experimental studies which address the following questions: Does a statement need to be objectively false in order to constitute lying? Does lying necessarily include the intention to deceive? Can one lie by omitting relevant facts?

Keywords Concept of lying · Lying · Definition of lying · Experimental philosophy

Introduction

“We’ve got to keep our country safe. You look at what’s happening in Germany. You look at what’s happening last night in Sweden. Sweden, who would believe this? Sweden. They took in large numbers. They’re having problems like they never thought possible. You look at what’s happening in

Alex Wiegmann, Ronja Rutschmann, and Pascale Willemsen have contributed equally to this work.

✉ Alex Wiegmann
awiegma@uni-goettingen.de

¹ University of Goettingen, Goettingen, Germany

² University of Potsdam, Potsdam, Germany

³ Ruhr-University Bochum, Bochum, Germany

Brussels. You look at what’s happening all over the world. Take a look at Nice. Take a look at Paris.”

Shortly after Trump made these statements, especially the “you look at what’s happening last night in Sweden”—part, there was a heated debate over whether Trump lied. The verdicts ranged from “flat out lie” to “no lie at all.” But how come that some people believed Trump to have clearly lied, while other believed his statement to be no lie at all? In general, it might seem that whether something is considered lying should not be open for strong disagreement—at least not if it is clear what actually did happen.

Such bafflement can be reduced once we noticed the fact that whereas the features of prototypical lying are clear, it is less clear where the lines restricting this concept should be drawn. For instance, many people argued that Trump could not have lied because nothing he said was false. But can you only lie by saying something that is objectively false? Many people also believed that no matter what he explicitly said, Trump clearly had the intention to deceive people. But how important is such an intention for lying? A further suspicion that has been raised in another context is that Trump might not have told the whole truth concerning his relationship with the Russian government. But can you actually lie by omitting facts? It seems that addressing these and similar issues is a necessary precondition for having a meaningful discussion about whether certain statements constitute lying.

In this paper, we address the just described three questions:

1. Can an agent only lie by saying something false?
2. What is the role of the agent’s intention to deceive?
3. Can an agent lie by omitting relevant facts?

To do so, we will present empirical experiments that systematically investigate people’s intuitions on the concept of lying.

Standard Definition of Lying

Philosophers and linguists have presented a variety of definitions of lying. Typically, philosophical definitions are rather technical and designed for theoretical purposes, and they sometimes deviate quite strongly from how the term is used in everyday communication. In contrast, when looking at the concept of lying, a primary goal is to find a definition that best captures the ordinary concept (see Carson 2006; Fallis 2009). That is, philosophers try to capture lay people’s intuitions about what situations constitute cases of lying. However, there has been a lot of disagreement in the debate about what intuitions lay people actually have. For this reason, it is surprising that only recently some philosophers become more and more convinced that the concept of lying cannot be defined by armchair philosophy alone, but that it needs to be built on systematic empirical investigations of how the folk use the term (Carson 2010).

The standard philosophical definition of lying holds that an agent lies if she makes “a believed-false statement to another person with the intention that that

other person believes that statement to be true” (Mahon 2016). Such a definition of lying entails four necessary and jointly sufficient conditions, namely

1. the *statement condition* requires that a person makes a statement,
2. the *untruthfulness condition* requires that a person makes a statement that she believes to be false,
3. the *addressee condition* requires that a person makes an untruthful statement to another person,
4. and the *intention to deceive the addressee condition* requires that the speaker intends to make the addressee believe the untruthful statement to be true.

According to the statement condition, every lie requires the use of linguistic symbols (Mahon 2016). Typically, these linguistic symbols are words and their grammatical combinations, but they also include shaking one’s head to answer a question or pointing towards a certain direction when asked for the way. In contrast, philosophers typically hold that we cannot lie to someone by wearing a wig or taking off the wedding ring when going to a bar if it is granted that a person does not make a statement by showing these behaviours. One might mislead others by doing so, but such behaviour would not be considered a lie. Second, for a statement to count as a lie, it needs to be uttered untruthfully. This condition does not require the uttered statement to be objectively false but merely that the speaker *believes* his statement to be false. As a consequence, an agent might be lying even if what she believes to be false turns out to be true. According to the addressee condition, the speaker needs to lie *to someone*. Finally, the intention to deceive the addressee condition states that lies aim to deceive and make the addressee believe what the speaker believes to be false. Thus, irony, jokes, or science fiction is ruled out as cases of lying.

In the following, we will present a couple of studies investigating whether the standard definition of lying fits with people’s intuitions.

Aim and Overview of this Paper

In this paper, we discuss a series of experiments that have been conducted on the folk concept of lying and thereby draw attention to the usefulness of an empirical approach to systematically inform the philosophical debate.

In “[Lying despite telling the truth?](#)” section, we will discuss Turri and Turri’s (2015) work on the untruthfulness condition. Against the standard definition, they argued that for something to be a lie, the statement’s falsity is required. More specifically, they argued that if a speaker believes to have said something false which was, however, true, people would not consider this statement a lie. Wiegmann et al. (2016) raised some objections against the experimental design that Turri and Turri (2015) used and tested these objections empirically. Based on their experimental findings, Wiegmann, Samland, and Waldmann argue, in line with the standard definition, that objective falsity is not required for something to count as a case of lying. Instead, what is important is the speaker’s subjective believe that her statement is false. We will further use this debate to show how supposedly minor differences in the wording of test questions can dramatically affect the results.

“[Lying without intention to deceive?](#)” section focuses on the question of whether lies necessarily require the agent’s intention to deceive the addressee. Rutschmann and Wiegmann (2017) have presented initial evidence that in some contexts, an agent is judged to have lied despite his lack of intention to deceive. While these results pose an interesting challenge to the standard definition, they also call for additional research on the exact role of intentions to deceive.

In “[Lies of omission](#)” section, we will discuss some new empirical evidence that is very much a work in progress. We will demonstrate that against the standard view in philosophy the folk believe that an agent can lie by omitting the truth. We will discuss potential interpretations of the experimental results and outline a possible avenue for follow-up studies.

Lying Despite Telling the Truth?¹

Jacob’s friend Mary recently posted information on the internet that will alert the public to serious government corruption. Soon some federal agents visit Jacob and ask where Mary is in order to detain her. Jacob thinks that Mary is at her brother’s house, so he tells the agents, “She is at the grocery store.” In fact, Mary is at the grocery store (Turri and Turri 2015, p. 167).

Did Jacob lie to the agents about Mary’s location? On the one hand, Jacob intended to deceive the agents by uttering a statement which he believed to be false. However, what he believed to be false turned out to be true. So, did Jacob lie? The answer to this question seems to depend on whether is it more important what Jacob believed to be the true or what actually was true.

Turri and Turri (2015) recently engaged in empirical research to answer this question. Most philosophers assume that what is relevant is only that the speaker utters something he believes to be false (Mahon 2016). Accordingly, even if a statement turns out to be true, it would count as a lie if the speaker believed it to be false. Let us call this view the *subjective view* on lying. In contrast, Turri and Turri hypothesized that the concept of lying presupposes that the statement was objectively false. Call this the *objective view* on lying.

To test which of the two views lay people endorse, Turri and Turri (2015) presented participants with the above-mentioned story of Jacob and the federal agents. When asked “Did Jacob lie about Mary’s location?”, most participants (78%) answered in line with the subjective view and said “yes.” However, Turri and Turri did not believe that participants’ responses reflected their real concept of lying. They speculated that the responses might have been an artefact of the question mode. More specifically, they reasoned that participants might have answered with “yes” in order to express blame for the agent’s intent to deceive the federal agents, or they might have taken the agent’s perspective who did not know that he was telling the (objective) truth. Hence, Turri and Turri conducted two

¹ This section is based on the article by Wiegmann et al. (2016) in the journal *Cognition*. Some parts of this section are identical to parts of the article and reprinted with permission by the journal (*Cognition*).

further experiments in which the same vignette was used but the answer options were modified. In experiment 2, they asked:

Choose the option that better describes Jacob in the story. When Jacob spoke to the agents about Mary's location:

- (a) He tried to tell a lie but failed to tell a lie.
- (b) He tried to tell a lie and succeeded in telling a lie.

Turri and Turri believed that if participants tried to express their blame judgment by agreeing that the agent lied, giving them an option to say that the agent “tried to lie” (expressing blame), but saying that, in fact, this attempt failed, should help clarify whether people considered Jacob's utterance a lie.

In Experiment 3, Turri and Turri (2015) used the following wording of the response options that was supposed to deal with their worry that participants' responses in their first experiment (in which the majority of participants answered in line with the subjective view) were “biased” by basing their response solely on Jacob's perspective.²

- (c) He tried to lie but only thinks he lied.
- (d) He tried to lie and actually did lie.

When the answer options were phrased in this way, the overwhelming majority chose options (a) and (c), respectively, suggesting that objective falsity is necessary for a statement to constitute lying, which is the upshot of Turri and Turri's paper.

However, Wiegmann et al. (2016) were not convinced by their conclusion. They suspected that conversational and experimental pragmatics (Grice 1991; Noveck and Reboul 2008; Royzman and Hagan 2017) could account for participants' seemingly objective responses. More specifically, they believed that the answer options Turri and Turri provided had led participants to misinterpret the test question. In Turri and Turri's second and third experiments, the two answers options each consisted of two parts, namely a “trying” and a “result” part. In experiment 2, for instance, option (a) is taken to express participants' objective view on lying; the subjective view, in contrast, is represented by option (b).

- (a) He tried to tell a lie but failed to tell a lie.
- (b) He tried to tell a lie and succeeded in telling a lie.

According to Wiegmann et al. (2016), a crucial difference between the subjective and the objective view on the untruthfulness condition is that lying is a more “difficult” act under the objective view compared to the subjective view. For something to be a lie under the subjective view, it suffices that the agent says something which he believes to be false. However, the objective view adds an additional criterion: In addition to all the conditions required by the subjective view, the objective view also requires the statement to objectively false. As a consequence of this additional criterion, trying to lie can result in not lying if you fail to state something objectively false. However, it is rather hard to imagine how a lie can fail according to the subjective view—here, the agent has already lied the moment he

² We put “bias” in quotation marks because it is only a bias if you assume that the objective view is right.

has uttered something he believes to be false. As Marsili (2016) puts it “If you think you lied, you lied.” No further checks are necessary.

Based on these observations, Wiegmann et al. (2016) pointed out that if it is so easy to lie according to the subjective view, splitting the response option into a trying and a result part seems rather unnatural. If trying is succeeding, then why mentioning the obvious success separately? To get a better grasp on the problem, imagine another case. Billy had ice cream for lunch, no specific incidents. A friend asks what Billy had for lunch. Does it sound natural to say “Billy tried to eat ice cream, and he succeeded in eating ice cream”? Not really. Usually, we only split up the description into a trying and a result part when the action was hard to perform (“he tried to break the world record and succeeded in breaking the world record”). Accordingly, if the question is split up for an act that is normally easy to accomplish, an additional unusual complication seems to be pragmatically implied. For instance, one might think that Billy’s hands might have been tied behind his back, the portion was very big for a normal person, or Billy has never eaten ice cream before without making a mess.

As a consequence, Wiegmann et al. (2016) speculated that the unnaturalness of splitting up the response might trigger pragmatic reinterpretations of the story. In the eating example, the two-part format of the question may imply that there was something special about eating ice cream, something that goes beyond opening the mouth and swallowing food.

If subjects have a subjective view on lying, presenting them with the split-up response options used by Turri and Turri (2015) might have led them to interpret the test question not as merely asking whether Jacob lied. If the question had been only about that, less complex response options would have been adequate (such as “he did lie” and “he did not lie”). Wiegmann et al. (2016) hypothesized that participants reinterpreted the response options as referring not merely about the question of whether Jacob lied but about some additional relevant component of the story, namely the fact that what Jacob said turned out to be true although he intended the statement to be objectively false. Since the trying part is identical in both response options and only the objective option states that he failed, subjects might have chosen the objective option not because they think that lying requires falsity but because of the fact that this response option expresses the failure of the agent to say something that is objectively false.

To test this hypothesis, Wiegmann et al. (2016) rerun the two experiments the findings of which Turri and Turri (2015) had used to argue for an objective view of the untruthfulness condition. Participants were presented with the vignette and answer options used by Turri and Turri, but in each experiment, a new condition was added. In these new conditions, response options were identical to the ones used by Turri and Turri, but they were extended by stating that what Jacob said turned out to be true. The aim of this addition was to prevent subjects from misinterpreting the response options as asking for the protagonist’s success or failure to state something objectively false. Moreover, if the objective view was true, this addition should rather increase than decrease the number of responses indicating that objective falsity is necessary for lying because it highlights the fact that what Jacob has said is the objective truth. Accordingly, in their replication of Turri and Turri’s Experiment 2, participants’ task was:

Choose the option that better describes Jacob in the story. When Jacob spoke to the agents about Mary’s location:

In one condition (participants were randomly assigned to only one condition), participants could choose between the two options provided in Turri and Turri’s Experiment 2, namely:

- (a) He tried to tell a lie but failed to tell a lie.
- (b) He tried to tell a lie and succeeded in telling a lie.

In the other condition, Wiegmann et al. (2016) added the fact that what Jacob said turned out to be true. Hence, these participants could choose between:

- (a) He tried to tell a lie but failed to tell a lie because what he said turned out to be true.
- (b) He tried to tell a lie and succeeded in telling a lie although what he said turned out to be true.

The results are summarized in Fig. 1. The original finding by Turri and Turri could be replicated. However, the pattern was reversed in the new augmented condition in which a clause was added saying that what the protagonist said turned out to be true. Here, the majority of subjects (72%) chose the option stating that Jacob did lie, although what he said turned out to be true.

Wiegmann and colleagues applied the same rationale to Turri and Turri’s third experiment (2015, Experiment 3) and manipulated an additional factor. In Turri and Turri’s third experiment, participants were asked the following question before they answered the lie question: “What Jacob said is objectively ____.” The options were “true” and “false.” The aim of this question, according to Turri and Turri (2015), was to encourage participants to think about the objective truth or falsity of Jacob’s statement. Wiegmann et al. (2016), however, speculated that the prime question may have alerted subjects that, in this experiment, objective truth or falsity is particularly relevant for answering the test question. Therefore, two conditions were added without the question about the objective truth or falsity of Jacob’s statement.

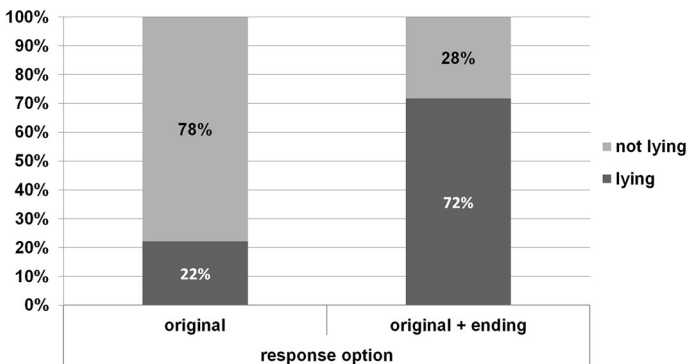


Fig. 1 Percentage of participants’ response choices

In all conditions, the test question was: “Which better describes Jacob?”. In two conditions (Prime/Original vs. No Prime/Original), the original response options used by Turri and Turri were presented, namely:

- (a) He tried to lie and actually did lie.
- (b) He tried to lie but only thinks he lied.

In the other two conditions (Prime/Original+Ending vs. No Prime/Original+Ending), participants could choose between

- (c) He tried to lie and actually did lie although what he said turned out to be true.
- (d) He tried to lie but only thinks he lied because what he said turned out to be true.

The results are summarized in Fig. 2.

Wiegmann et al. (2016) obtained the same result pattern as in the previous experiment just described. Hence, their alternative explanation was supported by the data. They were, therefore, justified to conclude that people’s responses were in line with a subjective view according to which a statement to constitute a lie even if it turns out to be true—the crucial point is that the agent *believes* it to be false.

One might question whether Wiegmann et al. (2016) initial line of reasoning actually explains the new results. They argued that the shift towards the objective options in Turri and Turri’s (2015) Experiments 2 and 3 is due to a pragmatic reinterpretation of the test question induced by the splitting of the response options in two parts. While the results of the new experiments support this interpretation, Wiegmann et al. (2016) also provided converging evidence from a domain in which the meaning of the target concept is not under dispute. The act of making a promise is such a case. Linguists consider promising a performative act which just depends on the context of utterance and which can neither be true nor false (Austin 1962). Most importantly, making a promise is not invalidated by breaking the promise

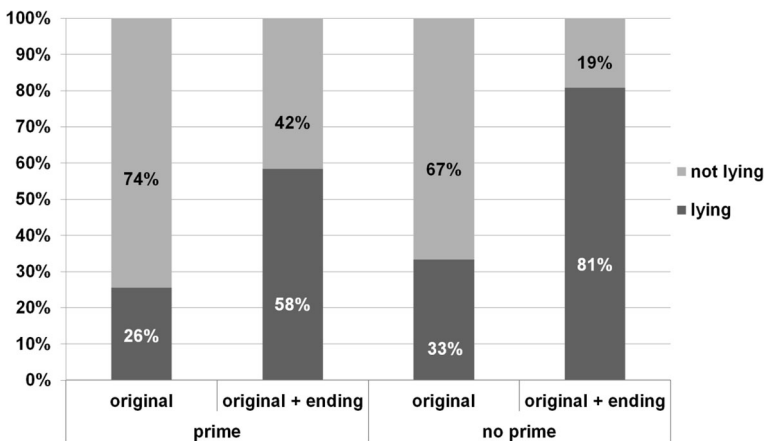


Fig. 2 Percentage of participants’ response choices

later. Wiegmann and colleagues' goal was to investigate whether split response options in cases of promises also lead to responses that seem inconsistent with the universally accepted meaning of the concept of making a promise. They, therefore, conducted a further experiment to show how strongly pragmatic considerations influence participants' responses. Subjects were presented with the following story:

Jacob and Anna have been a couple for the last five years. At the moment, however, Anna is thinking about splitting up because Jacob has been drinking a lot of alcohol lately. In order to save the relationship, Jacob promises Anna to never ever touch any alcohol again. To keep his promise and thereby saving the relationship, he throws away all the alcohol that he used to keep in his flat. However, two weeks later, an old friend is visiting Jacob and brought, as a gift, a quite old and expensive whiskey. Due to his promise, Jacob first denies the offer to drink the whiskey but is not able to overcome the temptation after his friend has asked him several times. The two get really drunk. Because of Jacob having been drunk once again, Anna breaks up with Jacob (Wiegmann et al. 2016, p. 41).

Participants were randomly assigned to one of three conditions that merely differed in the offered response options (Wiegmann et al. 2016). In the Control condition, participants were asked to choose between two simple options:

- (a) He made a promise.
- (b) He did not make a promise.

This condition resembles the options offered in Experiment 1 of Turri and Turri (2015). In the Two-Option condition, two different response options were offered that were inspired by the two options in Turri and Turri's Experiment 2:

- (c) He tried to make a promise but failed to make a promise.
- (d) He tried to make a promise and succeeded in making a promise.

In the Three-Option condition, participants were presented with the two response options of the two-option condition but in addition a third option was offered:

- (e) He tried to make a promise but failed to make a promise.
- (f) He tried to make a promise and succeeded in making a promise.
- (g) He made a promise but failed to keep it.

The results are shown in Fig. 3. The vast majority of participants indicated that the agent in the story made a promise in both the Control condition and the Three-Option condition. Most importantly, the dominant answer chosen in the experimental Two-Option condition was that the agent tried but failed to make a promise.

Thus, as in the case of "lying," participants could be pushed towards answering that the agent failed to make a promise although in the other conditions they clearly expressed that he did make a promise. Apparently, wording the response options in an unnatural way led to a reinterpretation of the meaning of the test question.




Fig. 3 Percentage of participants' response choices

Lying Without Intention to Deceive?³

Imagine the following story:

Charlie has a gambling addiction, so he often sneaks away from work to bet on horses. He has just come home after spending another thousand dollars at the racetrack. As he enters the house, his wife Natalie says to him, “I tried calling you at work today, but no one could find you. You didn’t skip out of work to go gambling again, did you?” Charlie knows that his wife doesn’t approve and will likely leave him if she finds out he’s been betting again. Before Charlie can say anything, Natalie reaches into his coat pockets and pulls out betting tickets from that afternoon’s races. Despite the evidence, Charlie automatically responds, “No, Honey, I wasn’t gambling.” Natalie doesn’t believe him (Arico and Fallis 2013, as used in Rutschmann and Wiegmann 2017, p. 17–18)

Did Charlie lie? Charlie made a statement which he believed to be false, and he addressed another person, namely his wife. Consequently, it seems that Charlie lied. However, the evidence clearly indicates him being guilty of gambling. Therefore, Natalie knows that Charlie is not telling the truth and Charlie knows that she knows. According to some philosophers, Charlie cannot have had the intention to deceive Nathalie. As a consequence, not all the supposedly necessary conditions of the standard definition of lying are fulfilled, and, thus, we would have to say that Charlie did not lie to wife.

Cases like the one just described are often referred to as *bald-faced lies*. Bald-faced lies are cases in which speaker and hearer/s share the belief that the speaker’s statement is false (Dyner 2015). Many philosophers have argued that because of the shared belief the speaker cannot have had an intention to deceive in these cases, yet they still qualify as lies. In consequence, bald-faced lies are regarded as strong evidence for the claim that the intention to deceive is not a necessary condition for lying (Carson 2006, 2010, Fallis 2009, 2010; Saul 2012; Sorensen 2007; Stokke

³ This section is based on Rutschmann and Wiegmann (2017).

2013). Positions arguing for such a view are called non-deceptionist positions. So-called deceptionists, on the other hand, who dominate the standard definition, reject bald-faced “lies” as actual lies when assuming that they do not match the intention to deceive condition (Dynel 2015; Meibauer 2011, 2014, 2015, 2016a), or argue that some intention to deceive might still be involved in these cases (Dynel 2015; Mahon 2008; Meibauer 2016a).

Does the intention to deceive constitute a necessary condition for lying? Since the best concept of lying is, as philosophers themselves claim, one that covers the folk’s intuition, the question of who is right cannot be answered a priori. For this reason, one might engage in actual empirical studies as to how the term is actually used. Surprisingly, very little research has addressed this question yet. The only empirical studies on the matter were conducted by Coleman and Kay (1981), Arico and Fallis (2013), Meibauer (2016b) and Rutschmann and Wiegmann (2017).

Coleman and Kay (1981) used a social lie vignette where presumably no intention to deceive was involved. They asked participants to judge on a seven-point scale whether the presented scenario was a lie or not. Subjects rated the scenario above the midpoint ($M = 4.70$). However, the authors failed to check whether participants also ascribed an intention to deceive. Arico and Fallis (2013) presented participants with bald-faced lies (such as the Charlie scenario) and asked them to judge on a seven-point scale whether they would classify these cases as lies. 98% of the participants rated these cases as lies, but again the authors failed to check whether they ascribed an intention to deceive to the liars.

Meibauer (2016b) and Rutschmann and Wiegmann (2017) were the first to actually ask participants whether they believed that the speaker had an intention to deceive. Meibauer (2016b) showed participants eight vignettes of bald-faced lies and asked them to answer on a seven-point scale whether the speaker lied, whether he acted deceptively, and whether he acted brazenly. He found that all utterances were judged as lies and that all utterances were considered deceptive.

Here is another, often cited example from the literature that both Meibauer (2016b) and Rutschmann and Wiegmann (2017) used in their studies:

Marc has witnessed a murder and clearly seen that Bob, the leader of a local Mafia, committed the crime. A CCTV camera has recorded everything. Later, Bob is accused of the homicide and Marc, as a witness in court, is asked whether or not he has seen Bob committing the crime. Marc makes the false statement that he did not see Bob committing the crime, for fear of being killed by his gang members. Marc knows that nobody will believe him, since the videotape clearly shows him witnessing the murder. In addition, Marc hopes that the jury won’t believe his testimony and convict Bob in spite of it. The jury doesn’t believe his testimony (Carson 2006, as used in Rutschmann and Wiegmann 2017, p. 18).

As mentioned before, all bald-faced lies share one important feature, namely the speaker and hearer/s shared knowledge that the statement is false. On closer examination, though, bald-faced lies differ dramatically in the possible consequences that would arise in the unlikely case that the hearer/s believes the speaker. In some scenarios (as in Charlie’s case), the speaker would profit from being

believed. Rutschmann and Wiegmann (2017), therefore, called this category *consistent bald-faced lies*. In other scenarios (as in Marc's case), the speaker explicitly hopes not to be believed. According to Carson (2010), Marc makes the false statement in order to reach another goal (saving his life); thus, deceiving others is only an unintended side effect and not the intention of the speaker. Rutschmann and Wiegmann (2017) called cases like these, where the agent hopes not to be believed, *conflicting bald-faced lies*.

Rutschmann and Wiegmann (2017) proposed a third category of potential lies that might lack an intention to deceive. In these cases, the speaker is indifferent as to whether he is believed or not:

Charlotte bets with her friend Grace, that she is able to answer all questions with “yes” for a whole day. If Charlotte wins, Grace has to pay her 10 lb. Later that day, Charlotte's mother asks Charlotte, if her brother was secretly smoking. Charlotte tells her, “yes”, though she knows that her brother is actually not smoking. Charlotte doesn't want to deceive her mother, but she is obeying the rules of the bet. Her mother doesn't believe her (Rutschmann and Wiegmann 2017, p. 18).

Rutschmann and Wiegmann (2017) called cases like these *indifferent lies* (not bald-faced because the common knowledge that what the speaker is saying is wrong does not hold in the cases we tested) and suggested that indifferent lies are in fact cases that lay people will consider lies without simultaneously ascribing an intention to deceive to the speaker. If people actually believe that the agent lied without any intention to deceive, this would raise some questions for the standard definition of lying.

Rutschmann and Wiegmann tested these hypotheses in two experiments (2017). Participants were confronted with one out of five scenarios: one consistent bald-faced lie, two conflicting bald-faced lies, and two indifferent lie scenarios. Then they answered the questions: “Did x lie?”, “Did x *intend to deceive* somebody?”, “Did x *actually deceive* somebody?”. In the first experiment, participants answered all questions on a seven-point scale, ranging from 1 (*definitely did not*) to 7 (*definitely did*). We used this scale to be able to directly compare our results with previous findings by Arico and Fallis (2013).

The results of the first experiment are displayed in Fig. 4. Participants significantly rated all scenarios as lies. The results of the intention to deceive rating on the other hand differed highly between the scenarios. However, in line with the theoretical assumptions, the intentionality response patterns were most similar within the three categories of lying, namely consistent bald-faced, conflicting bald-faced, and indifferent lies. For indifferent lies, participants significantly ascribed no intention to deceive. For conflicting bald-faced lies, the rating did not differ significantly from the midpoint. In the case of the consistent bald-faced lie, though, participants clearly ascribed an intention to deceive.

In the second experiment, participants answered the three questions in a forced-choice format (“yes” or “no”) for two reasons. First, asking in this format provides direct percentages of agreement or disagreement which are easier to interpret. Second, different formats of asking for responses might lead to different response




Fig. 4 Mean ratings of participant's response choices. *Error bars* represent standard error of means

patterns. If, however, previous results are robust, it should be possible to replicate the overall response patterns using a forced-choice format.

The consistent bald-faced lie scenario was not used again because the previous experiment showed that participants clearly ascribed an intention to deceive to the agent in this scenario. The results, which are shown in Fig. 5, were similar to the first experiment. A significant majority of participants (89%) judged both indifferent lie scenarios to be lies. The same is true for the conflicting bald-faced lies. Again, a significant majority of subjects (69%) ascribed no intention to deceive at indifferent lie scenarios. For conflicting bald-faced lie scenarios, though, the pattern of results was reversed, with a significant majority of participants (64%) ascribing an intention to deceive.

Two important findings result from this study. First, the results challenge the line of argumentation of many non-deceptionists who use bald-faced lies as their main argument to prove that intention to deceive is not a necessary condition for lying. Participants clearly ascribed an intention to deceive to consistent bald-faced liars and did the same in conflicting bald-faced lie scenarios, when confronted with a forced-choice format (“yes” or “no”). Thus, bald-faced lies might not be suitable counterexamples against the intention to deceive condition. Secondly, and most important, Rutschmann and Wiegmann (2017) found two scenarios that lay people judged as lies without ascribing an intention to deceive to the speaker. It




Fig. 5 Percentage of participant's response choices

seems plausible to assume that lay people will judge similar scenarios in which an agent lies to achieve a goal that is independent from being believed, similarly. Thus, intention to deceive does not seem to be necessary for the folk concept of lying.

But what are implications of these findings for the philosophical debate about the concept of lying? One possibility seems to be to simply drop the fourth condition. But philosophers have argued that the traditional definition of lying, including the intention to deceive condition, has some great advantages: It matches the prototypical cases of lying and excludes many cases that might intuitively not constitute lies, such as irony, sarcasm, or playing (Fallis 2010). Therefore, it would only make sense to drop the fourth condition if lay people actually judge sarcasm or irony as lies. That is not very likely the case, but should be empirically investigated.

Another possible solution is to replace the intention to deceive condition by another condition, instead of simply leaving it out. Some non-deceptionists proposed to focus on the context in which the statement is made. According to Fallis, for example, a believed-false statement is a lie if “you believe that you make this statement in a context where the following norm of conversation is in effect: *Do not make statements that you believe to be false*” (2009, p. 34). Carson suggests, among others, the following two conditions: “3. S states X in a context in which S thereby warrants the truth of X to S1, and 4. S does not take herself to be not warranting the truth of what she says to S1” (2010, p. 39). Similarly, Saul states that a believed-false statement is a lie if “[the speaker] take herself [*sic*] to be in a warranting context” (2012, p. 12).

Fallis (2009), Carson (2010), and Saul (2012) based their demand for a context condition to replace the intention to deceive condition on intuitive thought experiments about bald-faced lie scenarios. Rutschmann and Wiegmann’s (2017) approach is similar in the sense that they also used scenarios to argue that an intention to deceive is not a necessary condition for lying. However, they have empirically shown that lay people do not follow this intuition regarding bald-faced lies but found instead other scenarios (indifferent lies) that lay people regarded as lies without ascribing an intention to deceive. Nevertheless, the kind of definitions that highlight the role of context might actually be able to capture lay people’s intuition about these cases. Hence, further empirical research on such approaches seems a promising direction to find a definition of lying that fits with most people intuitions.

Lies of Omission?

In *Star Trek: Next Generation*, Cadet Wesley Crusher and four of his comrades performed a dangerous manoeuvre with their spaceships. However, the manoeuvre did not work as planned and there was an accident. In an inquiry, Wesley explains: “The accident occurred after the loop.” While everything Wesley said was true, the captain confronted him with being a liar: “What you neglect to mention is that following the loop, your team attempted a manoeuvre that was the direct cause of the crash. You told the truth up to a point, but a lie of omission is still a lie.”

Did Wesley really lie? Since Wesley omitted the crucial part but did not say anything he believed to be false, his statement does not meet the untruthfulness condition. So according to the standard definition of lying, Wesley did not lie and the captain is wrong in accusing Wesley of lying (Dyrel 2011). However, one might feel that it does not really matter if what he said was true. What Wesley did is to say something which he knew would lead the others to believe that it was the whole story—which it was not.

In this experiment, we aim to pump the folk's intuition about lies of omission. What we are interested in is whether people believe that a case like Wesley's should be considered a lie, even though it does not meet the standard definition of lying. We tested two different cases of omissions. On the one hand, a speaker might completely refrain from making a declarative statement. For instance, he might simply not answer the question by asking to change the topic, leaving the room, or declaring his unwillingness to provide an answer. On the other hand, lies of omissions might also be told by using half-truths (Vincent and Castelfranchi 1981). This is exactly the case of Wesley: He does give an answer, the answer is true, but it omits pieces of information Wesley knows to be relevant. We therefore created two stories, one in which the agent omits the answers entirely, and another one in which he only tells half the truth.

Methods

The two vignettes we used in our experiment are presented in Table 1.

Participants for our experiment were recruited on Prolific Academics, an online-based recruiting platform. Each participant was paid £0.20 for doing the study. A total of 210 participants started the experiment. However, we had to exclude 23 participants who did not finish the study, failed an attention check, or finished the experiment in less than 40 s. Thus, results are reported on 187 participants.

Participants first read general instructions. They were then randomly assigned to one of the two conditions, Change Topic or Half-Truth. Immediately after reading the vignette, we asked participants: Do you rather agree or disagree with the following statement: "Tom lied to Jane". Participants then chose between two options, namely "I (rather) agree" and "I (rather) disagree".

Results

The results are summarized in Fig. 6. When Tom changes the topic, 46% of participants (45 out of 94) participants considered Tom's behaviour a case of lying. In contrast, 78% did so in half-truth (74 out of 93).

For Change Topic, agreement is not statistically significantly different from chance (50%) level, (binomial test, 45 out of 94, test value = .5, $p = .76$). In contrast, agreement for half-truth was significantly above chance level (binomial test, 74 out of 93, test value = .5, $p < .0001$). This difference between Change Topic and half-truth was significantly different, $\chi^2 (df = 1, N = 187) = 21.29$, $p < .0001$.

Table 1 Scenarios used in the experiment

Change Topic	Half-truth
Jane recently found out that her son Tom got a couple of bad marks at school and might have to repeat a year if his marks do not improve. When Jane talks to Tom’s teacher, she learns that Tom often does not do his homework.	
The next day, Tom gets a lot of homework. He knows that if he does all of it, he will not be able to meet his friend in the afternoon.	
When he comes home, Jane asks: “What’s your homework?”	
In order to avoid telling his mother that he got a lot of homework to do, Tom says: “Let’s talk about something other than school. What’s for lunch? It smells delicious!”	In order to avoid telling his mother that he got a lot of homework to do, Tom only tells his mother about the English homework but does not mention the homework in maths and biology: “I have to read a chapter of this novel for my English lessons.”
Just as Tom intended, his mother believes that Tom has no homework to do, and she never asks about it again.	Just as Tom intended, his mother believes that Tom’s only homework is to read one chapter, and she never asks about it again.

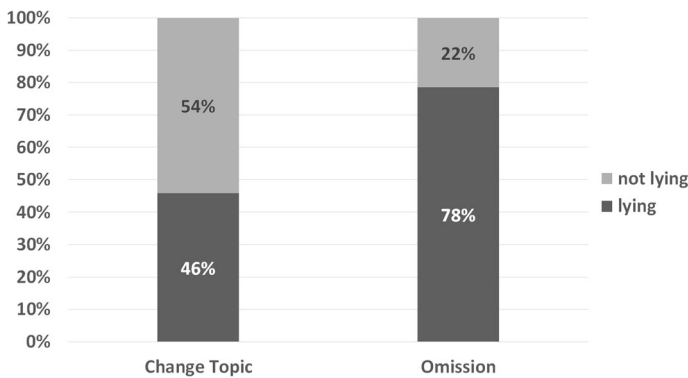


Fig. 6 Percentage of participant’s response choices

Discussion

The results of this experiment provide some first good indication whether people can lie by omission. The majority of participants believe that we can lie by omitting the truth. How can we make sense of this result?

Some authors have suggested that lying was not only about what is said, but also about what is implied (Meibauer 2005; Wiegmann and Willemsen 2017). Such a view builds on the philosophical and linguistic position that we do not only communicate by what we actually say, but also by what is implied by saying it. Grice (1975) emphasized the role of conversational implicatures and argued that every conversational context rests on the assumption that both speakers share a conversational goal. In the light of this shared goal, both speakers can expect the other to be cooperative, that is to make their “conversational contribution such as is

required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which [they] are engaged” (Grice 1975, p. 46). In order to reach this shared goal, both speakers will typically conform to four conversational maxims:

1. Quantity,
2. Quality,
3. Relation, and
4. Manner.

Can conversational pragmatics account for our findings? In cases of half-truths, the speaker violates the maxim of quantity as she does not say everything that would be required given the conversational situation. In doing so, she makes the addressee believe something which is, in fact, not the case. It seems plausible to assume that in our half-truth condition, mentioning only the English homework while omitting the home work in maths or biology makes the mother believe that Tom only has homework in English. Thus, it seems that Tom can make his mother believe that he only had homework in English both by implying and by explicitly saying that this is the only homework that he has. As a consequence, understanding lying at the pragmatic level can account for why people believe that the agent lied. While the agent did not say anything he believed to be false, he did imply something which he knew was false.

Can conversational pragmatics also explain our results for changing the topic? Other than when the agent only said half the truth, participants’ responses were not significantly different from chance. We believe that conversational pragmatics can account for this fact as well. When the speaker changes the topic, the conversational implicature that he does not have homework is less strong. In cases of half-truths, it seems rather unlikely that the mother becomes suspicious and asks whether this is really all homework he has. In cases of changing the topic, it seems much more likely that the mother insists on an answer to the question and does not form the belief that Tom has no homework. This difference indicates that something about Tom’s half-truth is more informative compared to when he changes the topic.

In related studies, Wiegmann and Willemsen (2017) demonstrated that the pragmatic level seems to play an important role for the folk concept of lying. They presented subjects with different cases in which the speaker violated either the maxim of quantity, relation, or manner and asked participants whether the agent lied. Even though everything the agent said was true, participants still believed the agent to have lied. For this reason, Wiegmann and Willemsen (2017) argue for a pragmatically grounded concept of that includes, in addition to what is said, also what is implied.

General Discussion

In this paper, we demonstrated how empirical investigations can inform the debate about philosophical questions. Given that it is taken for granted that an adequate definition of lying should capture lay people’s use and understanding of this

concept, it seems quite surprising that only very little empirical research about people's concept of lying has been conducted. The studies reported in this paper constitute a first step to fill this empirical gap.

In “[Lying despite telling the truth?](#)” section, we have shown that people exhibit a subjective view on lying. Even if a statement turns out to be true, making this statement can be considered lying. What seems crucial for something to count as a lie is that the speaker *believes* her statement (or what is implied by it) to be false. Hence, the mismatch between what a speaker believes and what she says seems to be more important for lying than the match between what she says and what is actually the case. The reported experiments in this section are not only relevant to identify people's concept of lying. They also nicely illustrate potential perils of empirical research on people's concept (Grice 1991; Noveck and Reboul 2008; Royzman and Hagan 2017). Researchers should be aware that participants tend to make all kinds of pragmatic assumptions when answering test questions.

In “[Lying without intention to deceive?](#)” section, we reported experiments that suggest that in some cases the intention to deceive the addressee condition is not necessary for lying. If the protagonist did not really care about deceiving the addressee because his primary concern was winning a bet, people judged the protagonist to be lying while not ascribing him an intention to deceive the addressee. Hence, at least in such cases the philosophical standard definition of lying does not capture people's intuitions. More research is needed to figure out the general importance of the intention to deceive in people's concept of lying.

In “[Lies of omission](#)” section, we examined whether lay people consider so-called lies of omission to be actual lies. While they do think so when the agent tells only half the truth, significantly fewer people agree when the agent changed the topic. These results point at an interesting new direction for future research. One might think that if the agent lied when telling only half the truth, he lied at the level of what is implied. What he said was not false, but it conversationally implied false information (see Willemsen and Wiegmann 2017). Some researchers have already argued that you can not only lie by means of what is said but also by what is implied (Meibauer 2005, Viebahn 2017). More systematic empirical research will have to follow to shed light on the linguistic level at which lying works.

We believe to have we demonstrated that empirical studies can provide valuable contributions to philosophical discussions about the concept of lying. However, the utility of the experimental philosophy approach is certainly not limited to this topic. Rather experimental tools can be applied to a wide range of important philosophical topics, more specifically to all debates in which philosopher aim to cover the folk's intuitions.

References

- Arico, A. J., & Fallis, D. (2013). Lies, damned lies, and statistics: An empirical investigation of the concept of lying. *Philosophical Psychology*, 26(6), 790–816.
- Austin, J. L. (1962). *How to do things with words*. Oxford: Oxford University Press.

- Carson, T. L. (2006). The definition of lying. *Nous*, 40(2), 284–306. doi:[10.1111/j.0029-4624.2006.00610.x](https://doi.org/10.1111/j.0029-4624.2006.00610.x).
- Carson, T. L. (2010). *Lying and deception: Theory and practice (1st publ.)*. Oxford: Oxford University Press.
- Coleman, L., & Kay, P. (1981). Prototype semantics: The English word lie. *Language*, 57(1), 26. doi:[10.2307/414285](https://doi.org/10.2307/414285).
- Dynel, M. (2011). A web of deceit: A neo-Gricean view on types of verbal deception. *International Review of Pragmatics*, 3(2), 139–167.
- Dynel, M. (2015). Intention to deceive, bald-faced lies, and deceptive implicature: Insights into lying at the semantics–pragmatics interface. *Intercultural Pragmatics*, 12(3), 309–332. doi:[10.1515/ip-2015-0016](https://doi.org/10.1515/ip-2015-0016).
- Fallis, D. (2009). What is lying? *The Journal of Philosophy*, 106(1), 29–56.
- Fallis, D. (2010). *Is a lie a lie if everyone knows it's a lie?* Rochester, NY: Social Science Research Network. Retrieved from <https://papers.ssrn.com/abstract=1668919>.
- Grice, H. P. (1975). Logic and Conversation. In R. J. Stainton (Ed.), *Perspectives in the Philosophy of Language: A Concise Anthology* (pp. 271–288). Toronto: Broadview Press.
- Grice, P. (1991). *Studies in the way of words*. Cambridge: Harvard University Press.
- Mahon, J. E. (2008). Two definitions of lying. *International Journal of Applied Philosophy*, 22(2), 211–230.
- Mahon, J. E. (2016). The definition of lying and deception. In E. N. Zalta (Ed.) *The Stanford encyclopedia of philosophy* (Spring 2016). Retrieved from <http://plato.stanford.edu/archives/spr2016/entries/lying-definition/>.
- Marsili, N. (2016). Lying by promising. *International Review of Pragmatics*, 8, 271–313.
- Meibauer, J. (2005). Lying and falsely implicating. *Journal of Pragmatics*, 37, 1373–1399.
- Meibauer, J. (2011). On lying: intentionality, implicature, and imprecision. *Intercultural Pragmatics*, 8(2), 277–292. doi:[10.1515/iprg.2011.013](https://doi.org/10.1515/iprg.2011.013).
- Meibauer, J. (2014). *Lying at the semantics–pragmatics interface*. Berlin: De Gruyter Mouton.
- Meibauer, J. (2015). Konzepte des Lügens. *Zeitschrift Für Sprachwissenschaft*, 34(2), 175–212. doi:[10.1515/zfs-2015-0010](https://doi.org/10.1515/zfs-2015-0010).
- Meibauer, J. (2016a). Topics in the linguistics of lying: A reply to Marta Dynel. *Intercultural Pragmatics*, 13(1), 107–123. doi:[10.1515/ip-2016-0004](https://doi.org/10.1515/ip-2016-0004).
- Meibauer, J. (2016b). Understanding bald-faced lies. *International Review of Pragmatics*, 8(2), 247–270.
- Noveck, I. A., & Reboul, A. (2008). Experimental Pragmatics: a Gricean turn in the study of language. *Trends in Cognitive Sciences*, 12(11), 425–431. doi:[10.1016/j.tics.2008.07.009](https://doi.org/10.1016/j.tics.2008.07.009).
- Rozyman, E., & Hagan, J. P. (2017). The shadow and the tree. In J.-F. Bonnefort & B. Trémolière (Eds.), *Moral inferences* (pp. 56–74). London: Routledge.
- Rutschmann, R., & Wiegmann, A. (2017). No need for an intention to deceive? Challenging the traditional definition of lying. *Philosophical Psychology*, 30(4), 434–453. doi:[10.1080/09515089.2016.1277382](https://doi.org/10.1080/09515089.2016.1277382).
- Saul, J. (2012). *Lying, misleading, and what is said: an exploration in philosophy of language and in ethics (1. publ. in paperback.)*. Oxford: Oxford Univ. Press.
- Sorensen, R. (2007). Bald-faced lies! Lying without the intent to deceive. *Pacific Philosophical Quarterly*, 88(2), 251–264. doi:[10.1111/j.1468-0114.2007.00290.x](https://doi.org/10.1111/j.1468-0114.2007.00290.x).
- Stokke, A. (2013). Lying, deceiving, and misleading. *Philosophy Compass*, 8(4), 348–359. doi:[10.1111/phc3.12022](https://doi.org/10.1111/phc3.12022).
- Turri, A., & Turri, J. (2015). The truth about lying. *Cognition*, 138, 161–168. doi:[10.1016/j.cognition.2015.01.007](https://doi.org/10.1016/j.cognition.2015.01.007).
- Viebahn, E. (2017). Non-literal Lies. *Erkenntnis*, 1–14. doi:[10.1007/s10670-017-9880-8](https://doi.org/10.1007/s10670-017-9880-8).
- Vincent, J. M., & Castelfranchi, C. (1981). On the art of deception: how to lie while saying the truth. In H. Parret, M. Sbisà & J. Verschueren (Eds.), *Possibilities and Limitations of Pragmatics* (pp. 749–777). Amsterdam: John Benjamins.
- Wiegmann, A., Samland, J., & Waldmann, M. R. (2016). Lying despite telling the truth. *Cognition*, 150, 37–42. doi:[10.1016/j.cognition.2016.01.017](https://doi.org/10.1016/j.cognition.2016.01.017).
- Wiegmann, A., & Willemsen, P. (2017). How the truth can make a great lie: An empirical investigation of lying by falsely implicating. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 3516–3621). Austin, TX: Cognitive Science Society.