



**Universiteit
Leiden**

The Netherlands

A speech-act theoretic analysis of white (pro-social) lies

Terkourafi, M.; Caponetto, L.; Labinaz, P.

Citation

Terkourafi, M. (2023). A speech-act theoretic analysis of white (pro-social) lies. In L. Caponetto & P. Labinaz (Eds.), *Philosophers in Depth* (pp. 245-273). Cham: Palgrave MacMillan.
doi:10.1007/978-3-031-22528-4_11

Version: Publisher's Version

License: [Licensed under Article 25fa Copyright Act/Law \(Amendment Taverne\)](#)

Downloaded from: <https://hdl.handle.net/1887/4210014>

Note: To cite this publication please use the final published version (if applicable).



A Speech-Act Theoretic Analysis of White (Prosocial) Lies

Marina Terkourafi

1 Introduction

A friend invites you over to show you their apartment, which they spent months renovating. Room after room, you find the new style of decoration abominable but don't want to hurt your friend's feelings; after all, you won't be living there yourself. So you utter small phrases like "That looks really nice!" and your friend responds with a grateful smile, beaming happiness. Have you lied? Most people might say yes, inasmuch as you stated something you did not believe to be true. Were you wrong to lie? Most people might say no, inasmuch as telling your friend your honest opinion would achieve little under the circumstances beyond hurting their feelings (they do not have the resources to engage in further re-decorating anyway). Now consider an alternative scenario: friends are visiting you from out of town and you shower their young children with gifts. At the end of their visit, they forget to take their gifts with them and

M. Terkourafi (✉)

Centre for Linguistics, Leiden University, Leiden, The Netherlands

e-mail: m.terkourafi@hum.leidenuniv.nl

when you remind them, the mother responds: “Our kids don’t play with such nonsense anyway.” Did she lie? Most people might say no, inasmuch as she voiced her true opinion about the gifts. Was she wrong to do so? Most people might say yes, inasmuch as she gratuitously hurt your feelings and displayed her lack of gratitude to you as a host.

We can all think of similar situations where telling the truth might be less socially adept than telling a lie. Called “white lies” in English, such lies have been defined as falsehoods “not meant to injure anyone and of little moral import” (Bok 1978, 58). In what follows, I define white lies generally as lies that benefit the addressee and are uttered out of concern for them. Researchers have long recognized that white lies can be so routine as to constitute a “particular sort of communicative competence” (Camden et al. 1984, 321) and people’s ratings of the truthfulness of their own statements support this. Naïve subjects asked to rate the truthfulness of their statements in everyday conversation rated only about a third of their statements as completely honest (Turner et al. 1975), suggesting that the majority of what we casually say consists of various degrees of falsehoods. But if white lies are so common, are they still lies? This hinges on one’s definition of lying, as well as on the different “shades” of lying this allows.

2 Lying: A Locutionary Definition

Despite centuries of philosophical, and more recently experimental, engagement with lying, no widely agreed upon definition of lying is available.¹ Equally long-lived is the awareness that not all lies are morally reprehensible: according to Zembaty (1993, 29), Aristotle himself considered lies that “harm no one, involve no undeserved disrespect, and stem from excellence and self-sufficiency rather than some deficiency of character” as morally acceptable, yet he saw lies that simply benefit others as not justified (Zembaty 1988, 528 and 540). To justify the “lie” part of

¹ Mahon (2019) provides a critical overview of several modern definitions.

their name, then—that is, to provide a definition general enough to include malicious and white lies under the same umbrella—what we need is a locutionary definition of lying, such as the one provided by Saul (2012, 29; cf. Fallis 2009):

- (1) If the speaker is not the victim of linguistic error/malapropism or using metaphor, or irony, then they lie if and only if
 1. they say that *P*
 2. they believe *P* to be false
 3. they take themselves to be in a warranting context

where a warranting context is one where sincerity is expected (as opposed to, e.g., being in a play, or telling a joke). Saul's definition excludes (unintentional) errors like "She's a great flamingo dancer," where the speaker means *flamenco* dancer, as well as (intentional) floutings of Quality ("Do not say that which you believe to be false"; Grice 1975, 28), such as metaphor and irony, where correctly interpreting the speaker's utterance necessitates recognizing that the speaker is speaking untruthfully.

The definition in (1) is locutionary in the sense that what matters is the fact that the speaker has uttered the untruthful words.² It can thus be used to distinguish between lying and misleading: under a locutionary definition such as (1), lying amounts to saying_{LOC} *p* while believing *p* is false, whereas misleading involves meaning *p* (without saying_{LOC} *p*) while believing *p* is false. On this definition, both malicious and white lies count as lies—because in both cases the speaker says_{LOC} something they believe to be false—but false implicatures do not. Rather, false implicatures, where the falsehood is meant but not said_{LOC}, fall under misleading.³ While this accounts for the observation that falsehoods

² On a locutionary sense of saying (Bach 2001), "S has said_{locutionary} that *p*" does not entail meaning *p*. This is different from Grice's (1975) sense of saying, in which saying *p* entails meaning *p*; that is, for Grice, saying *p* commits the speaker to the truth of the propositional content expressed.

³ For experimental support of this claim, see Weissman and Terkourafi (2018).

benefitting the addressee are characterized as “lies” cross-linguistically,⁴ it does not help us account for the difference between malicious, morally reprehensible lies and “innocent,” “pious,” “harmless” white lies. To account for this last difference, we need to consider the different intentions with which these two kinds of lies are uttered.

3 Two Kinds of Intentions

Several researchers have raised the question whether lies must be intentional, and if so, exactly what the lying intention consists in: is it an intention to have the listener *believe something* that the speaker herself does not believe, or is it an intention to have the listener *believe that the speaker believes* something she does not in fact believe? At issue here is the content of the false belief that the speaker wants the listener to entertain: is it a false belief about the *world* or a false belief about the *speaker's* beliefs about the world?⁵ Alternatively, lying can be a kind of perlocutionary effect, assessed depending on its results: if the target of lying (who may not always be the addressee)⁶ is deceived, then the speaker has lied, whereas if they are not deceived, then she has not lied.⁷ However, given the emphasis on lying as an indicator of the speaker's morality cross-linguistically, it seems awkward to claim that lying can be judged exclusively based on its outcome, over which the speaker has only partial control. Therefore, I take the view that lying must involve some kind of

⁴ An informal polling of native speakers of different languages yielded the following terms (the original is given first, followed by an IPA transcription or transliteration in italics and English translation in parenthesis): Egyptian Arabic: ارضي ب قبدك /*kidba be.dʕal* (white lie); Chinese: *yuǎnchāng huǎng* (lie smoothing over a situation); Danish: hvid løgn (white lie); Farsi: Duruq-e-Maslahati (lie with good intentions); French: pieux mensonge (pious lie), German: weiße Lüge (white lie); Greek: κατά συνθήκην/αθώα ψέματα *kata sinthēkin/aθoa pse mata* (lies by convention/innocent lies); Hungarian: füllentés (white lie); Italian: bugie innocenti (innocent lies); Norwegian: hvite løgner (white lies); Russian: невинная ложь *nevinnaia lozh* (innocent lie); Spanish (Peninsular/Latin American): mentira (or mentirita) piadosa/bianca (pious/white (little) lie); Thai: เกรงใจ /*kreŋ.ɿcail* (to be (too) courteous); Turkish: zararsız yalan (harmless lie); Urdu: دیفس ٹوٹھج *sufaid jhoot* (white lies).

⁵ Kupfer (1982) defends the former view while Chisholm and Feehan (1977) defend the latter.

⁶ According to the *Oxford English Dictionary*, lying to overhearers is possible (Mahon 2008, 218).

⁷ Coleman and Kay (1981) offer a definition of this type.

lying intention on behalf of the speaker. Still, this leaves open what the lying intention involves.

I propose that lies (both white and malicious) come with two kinds of intentions. The first intention all lies come with is a classic Gricean *r*(eflexive)-intention (Grice 1969), such as is necessary to invest an act of verbal behavior with meaning (Grice 1957). According to Grice (1969, 92), for the speaker to mean something by means of her utterance, she must have the following tripartite intention:

(2) “U[tterer] meant something by uttering [expression] *x*” is true if [and only if], for some audience *A*, U uttered *x* intending:

1. *A* to produce a particular response *r*.
2. *A* to think (recognize) that U intends 1.
3. *A* to fulfill 1. on the basis of his fulfillment of 2.

It is the third clause of this requirement that is usually identified as the *r*-intention: the speaker wants the listener to think about what the speaker is trying to achieve through her utterance and to use that information to help him figure out what her utterance means.⁸ This first, tripartite intention must be recognized in order to be fulfilled and is fulfilled simply by being recognized (i.e., it is not necessary for the listener to act on it). *R*-intentions are necessary for linguistic communication inasmuch as communication involves a change in both the speaker’s and the addressee’s cognitive environments.⁹

Nevertheless, lies also come with a second, “lying” intention. The content of this second intention is that the speaker wants the listener to believe something that she does not in fact believe herself. Inasmuch as this lying intention does not have to be recognized in order to be fulfilled, it is not a Gricean *r*-intention. Rather, it is a type of non-communicative intention similar to those that guide many of our solitary acts, such as the intention to eat, sleep, stand, and many other acts expressed as one-place

⁸This is recast in terms of a communicative intention embedding an informative intention in Relevance Theory (Sperber and Wilson 1995).

⁹On the Gricean picture, propositional contents that the listener entertains as a result of the speaker’s utterance but were not part of what the speaker wanted the hearer to understand are not linguistically communicated.

predicates. These intentions do not have to be recognized by someone else in order to be fulfilled. In everyday life, the word “intention” is typically used in this more mundane, non-technical sense.

Malicious and white lies differ in their treatment of this second intention. In the case of malicious lies, the lying intention must remain hidden from the addressee, or the lie will be unsuccessful (the speaker’s goal to deceive the addressee will not be achieved). However, in the case of white lies, it is permissible for this second “lying” intention to be recognized without jeopardizing the goal of the speaker. That is because the goal of the speaker this time is different. Rather than deceiving the addressee, in white lies, the goal of the speaker is first and foremost to constitute their face, and lying to them is one means of achieving that—this goal may still be achieved (for instance, by recognizing it directly) even if the means fails.¹⁰

In fact, recognition of the speaker’s lying intention in the case of a white lie can result in strengthening the politeness of the speaker’s utterance. That is because, assuming the speaker takes themselves to be in a warranting context (see (1) above), recognition of the speaker’s lying intention by the addressee generates a search for reasons why the speaker transgressed expected standards of truthfulness. If as a result of this search the listener concludes that the speaker lied in order to constitute the addressee’s (and potentially also the speaker’s) face, then they end up attributing

¹⁰ A reviewer points out that in some cases recognition of the speaker’s lying intention in white lies can jeopardize the speaker’s goal not to hurt the hearer’s feelings, and offers the following as an example: suppose I have a new haircut and ask my partner if he likes it. In order not to hurt my feelings, he mumbles a “Yes.” Suppose also that I perfectly get that he’s lying: I can read it in his face. Moreover, the fact that he doesn’t tell me what he really thinks ends up hurting my feelings more than a criticism of my new look on his part would have. So in this case, the one who uttered a white lie ends up hurting the other’s feelings precisely because the other recognized that the former lied. I would argue that in this example the listener’s being hurt hinges precisely on her having first recognized the speaker’s “Yes” as a *white lie*, that is, her recognizing that the speaker’s goal was face-constituting. This upset her because, by prioritizing face-constituting, the speaker appears to have treated her as he might treat a stranger to whom politeness rather than sincerity is due. As his partner, she experiences his choice as creating distance between them and an affront to their relationship, and that is why her feelings are hurt. Rather than challenging the permissibility of the recognition of the lying intention in white lies, then, this example confirms that, when recognized, such lies are attributed to the speaker’s face-constituting goal. Whether that goal is felt to be appropriate under the circumstances or not is another matter; and it is precisely that appropriateness that the “white liar” seems to have miscalculated in this case.

a face-constituting—that is, polite¹¹—intention to the speaker.¹² If I say that your hair looks nice and you recognize that I don't really believe that, you realize that my reason for saying so is not a belief I hold about the world independently of how I feel about you (I think that your hair looks nice in the same way I think that the sun is shining), but that how I feel about you has precisely something to do with my saying that your hair looks nice: how I feel about you is my reason for saying I like your hair. In other words, politeness is in this case inferred as the reason why the speaker transgressed expected standards of truthfulness: if truthfulness is prized in conversation, then sacrificing it at the altar of politeness actively shows the lengths to which the speaker is willing to go to please (i.e., constitute the face of) the addressee.

Prima facie support for the claim that the lying intention is handled differently in the case of malicious vs. white lies comes from the fact that different responses are appropriate in case the lie is recognized. If a malicious lie is called out as such, the best response is generally to admit the lie. If the speaker keeps insisting they have not lied, they risk adding layers of lying and perjury to their initial lie, casting further doubt on their morality and the integrity of their character. However, if a white lie is called out as such, the best (and most common) response is denial. That is because admitting the lie in this second case—for which, once the lying intention has been recognized, politeness is the sole plausible motivation—has the effect of canceling that self-same politeness. That is, admitting a white lie would defeat the purpose of engaging in (white) lying in the first place.

Isaacs and Clark (1990) propose a similar, two-level account for “ostensible invitations,” which they treat as part of a larger class of ostensible speech acts. According to them, ostensible speech acts are a type of non-serious language use which involves mutually recognized pretense and collusion by interlocutors to respond appropriately to the pretense. It is precisely this sense of collusion that would be destroyed if the utterer of

¹¹ Face-constituting constitutes the core of a second-order definition of “politeness” (Terkourafi 2008).

¹² This search may alternatively lead to the attribution of a face-*aggravating* intention to the speaker, as with bald-faced lies (Meibauer 2014). Formal and situational factors (see Section 4) help determine which of the two—the speech act of white lying (resulting in politeness) or a bald-faced lie (resulting in aggression)—is taking place.

a white lie were to admit the lie. Crucially, ostensible speech acts' primary purpose (or perlocutionary effect) is face-management;¹³ yet, this remains off-record and is achieved by making evident that the perlocutionary effect canonically associated with the speech act (e.g., in the case of invitations, that the listener attend a certain event) is situationally blocked.

However, ostensible speech acts also have a further feature setting them apart from other types of nonserious language use, namely their ambivalence: if asked, the speaker is unable to say categorically whether they intended to bring about the perlocutionary effect canonically associated with the relevant type of speech act or not; simply put, they are unable to say whether they "really meant" it. This is unlike other types of nonserious language use such as irony, where the speaker is clear about what they mean (and it is not what they say). This ambivalence regarding the speaker's commitment to the meaning of their words creates the possibility that the speaker's utterance may perform the speech act of white lying only up to a degree—that is, that the speech-act-hood of the utterance as a white lie will be graded (see Sect. 5).

Isaacs and Clark's account of ostensible invitations draws on the following exchange between Ross and Cathy, two students who had a date that same evening:

- (3) Ross: Cathy, Scott just called and told me that Brad and Dave and Rich and a lot of other guys from UCLA are going to be there tonight, so I guess we're going to go a night early. [He explains the plans for the night.] Do you want to come?

Cathy: That's all right. I'll pass.

Ross: Okay.

According to their analysis of this exchange, ostensible invitations involve

¹³ This follows from Isaacs & Clark's "P[erlocution]2. B comes to feel that A likes or approves of B to an extent consistent with P[erlocution]1" (Isaacs and Clark 1990, 502).

two layers. At the top layer, Ross makes an invitation and Cathy declines it. At the bottom layer, Ross and Cathy take collusive actions toward each other with the mutual recognition that the top layer is a pretense. It is the discrepancy between the two layers that gives Ross's and Cathy's actions their functions. What gets put on record, via the pretense, is that Ross would like Cathy to go but that she can't manage it. Off record, though, he breaks their date while assuring her he still enjoys her company, and she assures him she is not offended. (Isaacs and Clark 1990, 497)

In terms of conversational structure, this example empirically supports the claim made above that in the case of white lies (similar to Ross's ostensible invitation), the preferred response format is reversed. While the preferred response to a genuine invitation is acceptance, Cathy colludes with Ross's pretense by refusing, thereby indicating that she recognized his invitation as ostensible. What is more, Cathy's polite refusal is immediately accepted by Ross, unlike in genuine invitations where an initial refusal normally prompts a repeat until the invitation is accepted (or refused for good reason). The reversal of the canonical pattern for invitations indicates that both interlocutors recognize the situation for what it is, and the ease with which they do this (both the refusal and its acceptance are formulated succinctly and without hesitation) signals that they know they are doing this.

In sum, all lies (white and malicious) come with two types of intentions. The first is a classic Gricean *r*-intention. In virtue of this intention, lies, like all utterances, communicate some meaning about the world and perform a speech act such as an assertion, an invitation, and so on. The second is a lying intention, in virtue of which the speaker wants the listener to believe something she herself does not believe. While this second intention is *not* permissible to be recognized in malicious lies, it *is* permissible for it to be recognized in white ones. Differently put, the lying intention *must not* be recognized in the case of malicious lies but *may* be recognized (although it *need not* be) in the case of white lies. Permissibility of the recognition of the lying intention is different from its actual recognition and it is the former (whether permissibility is available as an option or not) that sets white lies apart from malicious ones. This means that, unlike the first intention all lies come with, which *qua* Gricean *r*-intention




Fig. 1 The two types of intentions (top: Intention 1 = Gricean r-intention, bottom: Intention 2 = the lying intention) involved in malicious and non-malicious (gray and white) lies

must be recognized in order to be fulfilled, whether the second, lying intention is itself a Gricean r-intention varies for different types of lies: for malicious lies, the lying intention cannot be a Gricean r-intention, but for gray and white lies, it can (to different extents). Figure 1 represents this diagrammatically.¹⁴

On the left of this diagram, we have malicious lies, for which the lying intention (Intention 2) must not be recognized (that is why for these lies this intention lies below the level of awareness of the addressee and outside the domain of speech acts in the upper part of the diagram). As recognition of the lying intention becomes an option without jeopardizing the speaker's goal (to the right of the red dashed line), we move toward the domain of white lies. However, we do so gradually, through shades of "gray" lies (Sect. 5). Once we acknowledge that permissibility of recognition of the lying intention is available for non-malicious (gray and white) lies but unavailable for malicious ones, we can both maintain a clear distinction between malicious and white lies (of various shades), and define white lies as a type of speech act that different utterances perform to various degrees.

¹⁴ For ease of exposition, Fig. 1 shows only assertions as a vehicle for (malicious or white) lies (that's why Intention 1 at the top right of Fig. 1 corresponds to Assertion). That does not mean that non-assertive speech acts, such as invitations, compliments and so on cannot also be used ostensibly. Indeed, the type of two-layered account developed by Isaacs and Clark (1990) and here is meant to account for ostensible speech acts more generally, not just assertions, though that remains to be worked out in detail.

4 White Lies Are a Type of Speech Act But Malicious Lies Are Not

On Why There Is No Speech Act of Malicious Lying

The above two-level analysis of the intentions involved in lying has interesting consequences for a speech-act theoretic account of lying in general, and of white lies in particular. Specifically, the non-permissibility of the recognition of the second, lying intention makes malicious lies different from speech acts such as promising, requesting, threatening, complaining, and so on. In promising, requesting, threatening, complaining, and so on, the speaker's intention to promise, request, threaten, complain, and so on *must* be recognized by the addressee for the act to count as a promise, a request, a threat, or a complaint respectively—or, in Sbisà's (2009) terms, in order to secure “uptake,” understood as the achievement of the characteristic conventional effect associated with this illocutionary act. However, in malicious lying, the speaker's (non-Gricean) intention to lie must precisely *not* be recognized in order to be fulfilled (for the speaker to have successfully lied). Recognizing a lie as a lie makes a malicious lie unsuccessful. This means that, in the case of malicious lies, the only uptake the liar can aim for is as an assertion (or whatever other speech act the utterance is demonstrably performing)—corresponding to the recognition of the first, Gricean *r*-intention that all lies come with. Malicious lies are indistinguishable from assertions (or invitations, etc.) in this respect. To listeners, they are assertions (or invitations, etc.) *tout court*.

While several theorists have proposed that lies are insincere assertions (Meibauer 2011; cf. Searle 1969, 42–49),¹⁵ others have been quick to point out the theoretical fraughtness of this move (see esp. Vincent and Castelfranchi 1981, 754, fn. 14; Reboul 1994). For instance, Reboul (1994, 297) writes: “[i]f a speaker produces an utterance which is a lie, it

¹⁵ Metaphors, ironies etc. can also be considered insincere assertions but there the insincerity is transparent and meant to be recognized for the speaker's *r*-intention to be fulfilled. This is indicated by Grice's phrase “make as if to say” in cases of metaphor etc., which indicates that their insincerity is communicative, that is, it is meant to be recognized by the addressee.

is necessary for the success of the lie that the illocutionary act of assertion should be successful. But if the perlocutionary act of lying is successful, then the illocutionary act of assertion is not successful [read: felicitous; A/N].” That happens, according to Reboul, because the speaker who lies is by definition not in the psychological state of belief toward the propositional content expressed by her utterance. In other words, the sincerity felicity condition, which stipulates that “a speech act is sincere only if the speaker is in the psychological state that her speech act expresses” (Green 2017) can never be met in the case of lies. On this view, an insincere assertion is necessarily an infelicitous assertion. But if lies are infelicitous assertions, how can they fulfil their perlocutionary goal of deceiving the addressee?

The puzzle highlighted by Reboul arises only if we try to account for lies on a single level, as underlain by only one intention which is a communicative one. By contrast, on a two-level account such as the one proposed here, in which lies come with two types of intentions only the first of which is a Gricean *r*-intention intended to be recognized (see Fig. 1, Intention 1), while the second one is, in the case of malicious lies, a non-communicative intention that must remain hidden (see Fig. 1, Intention 2), a lie can be *both* a felicitous assertion (because the speaker sincerely wants the listener to believe something about the world) and a lie (because what the speaker sincerely wants the listener to believe is something she does not believe herself; see also Sect. 3).¹⁶ But if this second intention is not communicative, then in terms of the speech acts they perform, malicious lies perform only one type of speech act: they are assertions *tout*

¹⁶Reboul’s way out of this conundrum is to lower the bar on sincerity, so to speak (Reboul 1994, 298). In her view, all that is required for the speaker to be sincere is that she intend her utterance to be taken by the addressee as *a reason to believe* that she herself believes what she says. This is a condition on how the speaker would like the addressee to take her utterance (as evidence of belief), not on how she actually positions herself toward the propositional content of her utterance (whether she believes it or not). On this revised (compared with the classical view reiterated by Green 2017) understanding of the sincerity condition, a lie is as sincere as any other assertion. The account I sketch immediately below also considers lies to be sincere assertions but additionally, by singling out the lying intention as a separate intention that is handled differently by malicious and non-malicious lies, offers the possibility of a unified account of these two types of lies.

court.¹⁷ Lying is indeed something that is done through language but it is not done linguistically: it is not communicated.

The Speech Act of White Lying

The fact that the lying intention in malicious lies is not communicative means that malicious lying cannot be a separate speech act, which in turn suggests that malicious lies perform only one speech act, that of assertion. However, one can go further and claim that, precisely because of the permissibility of the recognition of the lying intention in white lies, white lies *can* be a separate speech act. In this case, an utterance which is a white lie performs two speech acts at the same time: an assertion and a white lie (but see Sect. 5 on why this second speech act may be performed only up to a degree). Three types of evidence support this claim: (i) the availability of conventional means for their performance, (ii) their online processing, and (iii) their acquisition.

Before proceeding, it is important to be clear about what we mean by “speech act.” In what follows, I will be adopting Sbisà’s (2009) neo-Austinian approach to speech acts, according to which a speech act is a socially constituted act that becomes meaningful against the background of existing norms and conventions in a group; it is a recognizable structured unit of communicative behavior aiming to bring about certain effects conventionally associated with this behavior. In other words, a speech act is a socially constituted tool for doing the things that a culture deems worthy of getting done linguistically.¹⁸ This neo-Austinian understanding of speech acts is especially apt in the case of white lies, since the extent to which it is permissible for white lies to be recognized is not up

¹⁷ Meibauer (2011, 280) also correctly states that lies are not a separate kind of speech act. However, he considers them to be assertions, specifically insincere assertions, which is problematic for the reasons highlighted by Reboul. Moreover, he does not consider the (second) type of intention that makes an assertion additionally a lie or the fact that this lying intention is not meant to be recognized.

¹⁸ By contrast, the fact that speakers can use these tools in unexpected ways and listeners can understand unexpected meanings out of them goes beyond conventionality and into the realm of speaker’s intentions and perlocutionary effects achieved.

to the speaker to decide but rather hinges on the societal consequences of such recognition.

These consequences, in the case of white lies, can include a range of positive outcomes. I have already commented in Sect. 3 on the possibility that, as a result of uttering a white lie, the speaker may be perceived as polite, and how this perception can be strengthened if the lie is recognized. Further to this possibility, research on trust has shown that prosocial lies (a type of white lie with high stakes) can enhance benevolence-based trust (Levine and Schweitzer 2015). Having asked experimental participants to play a trust game with a confederate who had previously either told them a prosocial lie or been selfishly honest, researchers found that participants trusted more those willing to tell a prosocial lie, despite recognizing them as deceptive, compared with those who stuck with the truth when it was selfish to do so. They explained these findings by appealing to “benevolence, demonstrating concern for others, [which] can be far more important for fostering trust than either honesty or selflessness” (Levine and Schweitzer 2015, 102). Findings such as these depend on the extent to which benevolence-based trust trumps other kinds of trust (e.g., integrity- and competence-based trust) in the situation at hand and need not generalize across cultures.

Another positive outcome related with recognizing white lies concerns knowing how to respond to them. As Isaacs and Clark (1990, 496) note, with respect to example (3) above, it can be anything from embarrassing to hurtful not to recognize a white lie for what it is:

Mutual recognition is important for several reasons. Without it, Cathy might genuinely accept the invitation, not realizing it was intended to be seen as merely a pretense. Or she might take it as obviously insincere, without realizing she was intended to see that, and she might feel insulted. Or if Ross made his invitation without expecting her to recognize the pretense, then it would be merely insincere, bearing the same relation to sincere invitations as lies bear to assertions. It would simply deceive.

The reversed preference organization in example (3), with Cathy's unhesitating and brief refusal and Ross's equally brief acceptance of it, helps them correctly identify what is going on.

In sum, having one's white lies recognized can be advantageous for several reasons including the attribution of politeness, benevolence, or simply eliciting an appropriate response. This observation provides a reason why languages/cultures should have recognizable ways of getting this done—in other words, it motivates the need for a speech act of white lying. Claiming the existence of a speech act, however, necessitates more than establishing a motivation for it. As Sbisà (2009, 46) reminds us:

There must exist an accepted conventional procedure having a certain conventional effect, that procedure to include the uttering of certain words by certain persons in certain circumstances. (Austin 1975, 14, 26–28)

The next three subsections provide evidence in this direction.

Conventionality of Means

The first type of evidence supporting the existence of a speech act of white lying concerns the availability of conventional means for its performance. We may distinguish between two types of conventionality in this respect: conventionality of form and conventionality of content. Starting with conventionality of form, Brown and Levinson (1987, 172) highlight the existence of a high pitch in Tzeltal, a Mayan language, which functions precisely to signal a white lie:

In Tzeltal, there is a highly conventionalized use of high pitch or falsetto, which marks polite or formal interchanges, operating as a kind of giant hedge on everything that is said. [...] Use of it seems to release the speaker from responsibility for believing the truth of what he utters so that the presence of this falsetto in an otherwise normal conversation may well mark the presence of a social lie.

More recently, in a lab setting, Fish et al. (2017) identified a set of acoustic cues predicting the perception of a compliment as sincere or insincere (a white lie) in Canadian English. Keeping the wording of the compliment stable, they found that listeners rated as insincere those compliments which were spoken more slowly, started lower, and tended

to get louder as the utterance unfolded. The same stimuli were rated as less sincere if they had been elicited by an uncertain question (Asker asking for reassurance) compared to a confident one (Asker just curious about Responder's opinion). In other words, not only the acoustic make-up of the compliment, but also the social requirements of the context (Asker's expectations) affected listeners' evaluation of a compliment as insincere. This finding supports the claim that white lies are expected in certain contexts, which has been independently advanced based on data about their online processing (see Sect. 4.2.2).

The successful identification of acoustic cues for low-stakes white lies such as those investigated by Fish and her colleagues contrasts with the general difficulty of identifying reliable cues to lying behavior. According to Fish et al., this is because “control behaviours [taxed by prosocial lying] sometimes produce detectable vocal indications of the speakers' insincere opinions, which are qualitatively different than when speakers are unfettered by the need to conceal their negative attitudes to the listener” (Fish et al. 2017, 157). It may well be that reliable cues are only available for white lies, whereas real lies are not consistently signposted by their utterers—precisely because, unlike white lies, they are not supposed to be recognized by their recipients.

Coming to conventionality of content, perhaps the most conventionalized means for white lying can be found in Persian *ta'arof*, a highly ritualized practice involving figurative language and extreme “self-lowering” referring expressions that to outsiders can sound like “a collection of lies” (Miller et al. 2014, 15). By contrast, native speakers consider *ta'arof* to be a token of goodwill and respect, a strategic move to achieve particular perlocutionary goals, and/or an indication of the speaker's good manners and upbringing (Miller et al. 2014, 19). In fact, *ta'arof* is thought to occur less in lower socio-economic classes and among villagers (Beeman 1986, 197). Example (4) illustrates this practice (for the original, see Izadi 2015, 86):

- (4) Ali and Reza are close friends returning home from an evening out. Reza has given Ali, who lives with his family, a ride home. It is almost midnight, not an appropriate time for receiving guests, especially if it is not prearranged.

1. Ali: Come in.
2. Reza: Thanks a lot.
3. Ali: Come in.
4. Reza: Thanks. I've got to go. I have work.
5. Ali: Well, just come in for a minute then you can go.
6. Reza:[extending hand] May I sacrifice for you. Don't you need favor?
7. Ali: [refuses to shake hands] Are you doing ta'arof? (standing on ceremony?)
8. Reza: (1.0) No (I swear) by God, convey my hello.
9. Ali: [extending hand] I'm at your service.
10. Reza: I'm your slave.

Several features highlighted earlier as characteristic of white lies can be identified in this exchange, including the use of lexical formulae ("may I sacrifice for you," line 6; "I'm your slave," line 10) and the ceremonial denial of the white lie (lines 7–8). Being part of a leave-taking sequence, this exchange functions as an ostensible invitation "solicited by context" (Eslami 2005, 464). Nevertheless, it is not pointless: as Koutlaki (1997, 119) explains, "the fact that [the] speaker takes the trouble to use a socially enjoined formula indicates [his] intention to accord respect to [his] interlocutor and takes on therefore a phatic function."

Qua widely recognized cultural practice, ta'arof involves expressions that can be conventionalized to various degrees. While example (4) includes some highly conventionalized ones, similar contents can be expressed in more creative ways. What is recognizable in such cases is the attitude (of generosity, self-lowering, etc.) expressed by the utterance rather than the specific words used. The expression of these semantic meanings at particular moments during the interaction (especially during offers/invitations, requests/orders, thanking, complimenting, greetings and leave-taking) and in ways that are appropriate to one's gender and age jointly co-constitute an exchange as an instance of ta'arof. In all of these ways, ta'arof resembles ritualized types of speech acts found in a variety of cultures (e.g., ritual refusals in Chinese; Yang 2008). Such acts are at one end of a continuum of ostensible speech acts (Isaacs and Clark 1990), the other end of which is populated by acts that are not necessarily as strictly

ritualized, yet can be conventionalized as regards their conversational make-up, sequential placement in interaction, objective felicity, and manner of delivery (Isaacs and Clark 1990, 499–502). The existence of these conventional signals suggests that speakers are aware of when their utterance is a white lie and have subtle ways of (consciously or subconsciously) signaling that to their addressees.

The Online Processing of White Lies

Experimental results from the online processing of white lies suggest that in contexts where they are expected, white lies are treated no different from true statements in non-biasing contexts. Moreno et al. (2016) presented 26 female native speakers of Spanish with a context in which the truth was unpleasant (e.g., a host asking her guests whether they liked the dinner she had cooked for them, when the dinner had been burned). They then measured their ERP responses while reading one of three possible answers to the host's question: (i) White Lies (WL) (e.g., *The meat sauce was tasty*); (ii) Blunt Truths (e.g., *The meat sauce was overcooked*); or (iii) Semantic Violations (e.g., *The meat sauce was romantic*). These answers had been previously elicited as likely continuations to *The meat sauce was ...* in a paper-and-pencil task by a separate group of participants. They were then normed to obtain equally expected target words for each condition, factors such as word frequency biasing the results.

What the researchers found was that white lies (*The meat sauce was tasty*) were not treated as semantically anomalous (did not induce an N400 response) in social situations where they were expected, whereas blunt truths (*The meat sauce was overcooked*) in the same situations were. This effect was heightened in situations strongly biasing toward white-lying behavior, that is, situations for which participants in the paper-and-pencil task agreed that a white lie was expected (34/93 scenarios in their study). Based on their findings, Moreno et al. (2016, 624) argued that

the processing of white lies lacks of any semantic (N400) or interpretative (P600) difficulty. White lies are processed neither as false nor as ironic messages. Their immersion in a social context overrules the neural conse-

quences that have previously been linked to the processing of factually untrue statements. Indeed, closer to truth statements [blunt truths] became more difficult to process, as indexed by an enhanced N400 for those sentences strongly/moderately biased toward white lying, as well as by a frontal late positivity in all cases.

In other words, taking N400 to be “not as sensitive to truth/falsity computations as [...] an online predictor for upcoming information based on world knowledge” (Moreno et al. 2016, 617), white lies are treated as expected, at least in some situations.

These results are interesting for a couple of reasons. First, both parts of the experiment (the paper-and-pencil task, and the online ERP task) confirm that there are specific moments in interaction which ratify the telling of a white lie and when participants expect a white lie to be told. This is in line with Fish et al.’s (2017) finding that exactly the same acoustic stimuli were judged as less sincere if prompted by an uncertain question of the Asker, that is, if the situation called for a white lie (irrespective of whether the propositional content of the utterance was truthful or not). The finding that white lies are what is normatively expected in some situations is also in line with Wilson and Sperber’s claim that “language use is not governed by any convention or maxim of truthfulness in what is said” but rather “expectations of truthfulness—to the extent that they exist—are a by-product of expectations of relevance” (Wilson and Sperber 2002, 583–584). The experimental evidence from the processing of white lies supports this claim, adding to our reasons for considering sincerity to be a maxim governing optimal performance of a speech act (along the lines of Sbisà 2018) rather than constitutive of its performance.

The Acquisition of White Lies

Several studies have shown that children as young as three can produce white lies, that is, lies that show consideration for another’s well-being (Talwar and Lee 2002; Talwar et al. 2007). What is more, from a young age, children are able to do so *appropriately*: that is, they discriminate between situations that call for the truth versus those that call for a white lie and favor white lying in the latter (Walper and Valtin 1992; Bussey

1999). These findings are in line with Fish et al. (2017) and Moreno et al. (2016), who showed that there are specific moments in interaction when a white lie is likely to be told. By around age seven, children learn to align with adult preferences and “read” these moments correctly (Heyman and Lee 2012, 169).

Crucially, these situations can be different in different cultures suggesting that children’s ability to read these situations correctly is not (just) part of cognitive maturation but rather of early socialization into the group norms surrounding them—that is, it is part and parcel of acquiring the repertoire of socio-culturally constituted speech acts available in their surrounding community. Specifically, although all children tend to value white lying when required by the social context (Xu et al. 2010), in cultures that place emphasis on autonomy, children justify the telling of white lies primarily with respect to their effect on the individual recipient’s *emotional* well-being, whereas in cultures prizing societal interdependence, they highlight the *social* implications for the recipient (Heyman and Lee 2012, 169).

Focusing on the latter, experimental results from 7–11-year-olds in China suggest that their sensitivity to a public/private distinction and the greater risk of face loss in the former lead them to evaluate white lies less negatively in public settings vs. private ones, while they do the opposite for blunt truths (Ma et al. 2011, 314). Moreover, “children in all of the age groups rated white lies less negatively than harmful lies and rated blunt truths less positively than helpful truths. This finding suggests that even 7-year-olds are aware that the moral value of a verbal statement depends not only on its truthfulness but also on whether it serves to help or harm its recipient” (Ma et al. 2011, 314)—a claim that resonates with Wilson and Sperber’s (2002) argument that relevance, rather than truthfulness, is what is expected in interaction.

Similar results were obtained in a study comparing classifications of lie- and truth-telling in prosocial situations by Persian and Canadian children aged 5 through 11 (Mojdehi et al. 2020). Not only did Persian children generally rate politeness and ta’arof lies more positively compared with the Canadians, as they got older, Persian children also

evaluated untruthful statements in ta'arof situations less negatively than Canadians. All in all, the evidence from acquisition shows that children are not only able to produce white lies from early on but they do so in accordance with adult norms regarding the situational appropriateness of these lies, which can be different in different cultures. They thus support the claim that white lies constitute a type of speech act which different groups have developed different conventional procedures to signpost (Sbisà 2009).

5 The Context-Dependence of (White) Lies

The three types of evidence surveyed above jointly suggest that white lies can be viewed as a type of speech act, understood as a socially constituted cultural practice aiming to bring about certain conventional effects along the lines of Sbisà (2009). However, it does not follow from this that white lies are always and unambiguously “white.” The existence of a continuum from malicious (real) to white lies has been empirically demonstrated by Bryant (2008), who combined interviews and focus-group discussion in an attempt to capture the perspective of the participants themselves.

As Table 1 shows, next to real and white lies, participant views led Bryant to identify two categories of “gray” lies, meaning lies in which the various dimensions (intention, consequences, beneficiary, etc.) are at odds with each other, making it impossible to classify them as either real or white. An example of a gray lie is that of a vegetarian asking, after a meal that her friend had cooked, if the meal had contained any meat products and her friend assuring her that it did not, despite knowing that it did (since she had cooked it). Since the question was asked after the fact, when nothing could be done to rectify the situation, a positive answer would have only saddened or angered the questioner. Hence the falsehood can be said to benefit the addressee as well as the speaker and the relationship between them. Yet the lie is more consequential than an innocuous, white lie, making that label inappropriate.

Table 1 Bryant's open-ended lie classification (from Bryant 2008, 32)

Factors						
Types of Lies		Intention	Consequences	Beneficiary	Truthfulness	Acceptability
	Real Lies	Malicious Deliberate Deceptive Deceitful	Serious Direct	Self-Serving Egotistical	Complete Fabrication Blatant Un-truth Zero truth	Unacceptable Not Justified
	White Lies	Benign Pure	Trivial Meaningless Harmless	Altruistic Other-Focused Protecting Helpful	Partial truth Half truth Bending the Truth Stretching the Truth	Acceptable Justified Expected Common
	Gray Lies					
	Ambiguous Gray Lies	Ambiguous Intention	Ambiguous Consequences	Ambiguous Beneficiary	Ambiguous Level of Truth	Open to Interpretation
	Justifiable Gray Lies	Malicious	Direct	Self-Serving	Complete Fabrication	Justified Acceptable

That what is a white lie depends on the context and is co-constructed between speaker and addressee is also supported by Fish et al.'s experimental results, who showed that there is no clear break between sincere and insincere compliments (i.e., white lies), but rather the interpretation of an utterance as a white lie also depends on the expectations of the Asker (2017, 156). As several scholars have noted, it is possible that, in the end, whether a lie is harmful or not depends on the recipient of the lie and how they feel about it: "what is a vicious, harmful lie for one person may be an act of loving concern for another. [...] Lies can only 'be' as they are perceived by specific involved people" (Knapp and Comadena 1979, 271; cf. Bok 1978, 60).

These empirical findings are especially pertinent to the claim that the hallmark of a white lie is *permissibility* of the recognition of the lying intention (Sect. 3) and bring us back to ambivalence as a defining feature of ostensible speech acts (Isaacs and Clark 1990). In Terkourafi (2014), I identified several reasons for such ambivalence, which extend beyond

face-management served by ostensible speech acts: speakers may be happy to let addressees be the ultimate arbiters of what they mean, because it saves them the effort of spelling it out or because they do not want to take responsibility for it; but they may also not have a fully formed intention that they can communicate, as with children whose early vocalizations are requests only to the “ears” of caretakers who rush to attend to them and only gradually become requests (intended to be recognized) as the child observes their perlocutionary effects on others (Huls and van Wijk 2012, 92–97).

To account for all of these cases we need a way of talking about an utterance’s performing a speech act as a graded notion. This idea is entertained in an early article by Lakoff (1992). Citing several legal cases where the courts decided differently, variably prioritizing the speaker’s (intended) meaning or the listener’s (perceived) one, Lakoff observes that “language and linguistic behaviour (and human behaviour in general) by their nature may not be reducible to yes/no decisions, *nor amenable to being assigned to only one pile among many* [...]”; to get anywhere near a deep understanding of the forms, functions, and properties of language, a theory which includes the possibility of continuous classification is necessary” (Lakoff 1992, 317; emphasis added). She then proposes such a classification, assigning utterances along a linguistic continuum that ranges from constatives to performatives corresponding to an action continuum that ranges from expression to conduct (Lakoff 1992, 320). What is useful to retain from this discussion is Lakoff’s observation that “words can have the effect of actions, be tantamount to action, *sometimes* [...] but not all utterances are equally “active” that is, reality-changing” (Lakoff 1992, 318).

A similar intuition is developed by Sbisà (2001, 1797), who writes of “a conception of illocutionary force according to which illocutionary acts have conventional effects” where such effects “may be described in terms of ‘deontic modality’, namely, as assignments to or cancellations from each one of the participants of modal predicates related to the necessity or possibility of actions with respect to norms.” In other words, speech acts, when performed felicitously, alter the rights and obligations of interactants and thereby the future courses of action they may plausibly take.

Since utterances can do this to different degrees, such a conception of illocutionary force

allows for degrees of strength. What the speaker has done (the effect of the illocutionary act) is no longer bound to mirror a discrete intention of the speaker. Rather, since there are aspects of the interpersonal relationship that are settled on the basis of intersubjective agreement, the conventional effects of speech acts may be considered as affecting them, and as playing a role in their adjustment and fine tuning. (Sbisà 2001, 1797)

Simply put, this means that an utterance may be a request, an invitation, a compliment and so on more or less; that is, it may carry the corresponding force to a greater or lesser extent. In the case of white lies, it would seem that, once it is permissible for the speaker's lying intention to be recognized, what would otherwise be a malicious lie starts to become socially justifiable. Yet this leaves an enormous range over which language users can agree or disagree about the beneficial or damaging effects of an otherwise non-malicious (white or gray) lie.

6 Summary

Potential lies come with two types of intentions: a Gricean *r*-intention, in virtue of which all lies are assertions; and a "lying" intention that must remain hidden in the case of malicious lies (i.e., is not a Gricean *r*-intention) and is variably permissible to be recognized (i.e., may be a Gricean *r*-intention) in white ones. Due to their different handling of this second intention, malicious lies are not a separate kind of speech act, while white lies can be. Three types of evidence (conventionality of means, their online processing, and their acquisition) support analyzing white lies as a separate type of speech act. However, white lies are not always as clearly signposted in terms of where they occur in the exchange and the form they take, and a continuum of cases exists from the most innocuous white lie to the most malicious real one. While speaker and addressee(s) may disagree about whether the speech act of white lying or, conversely, malicious lying has occurred, the permissibility of the

recognition of the speaker's lying intention (i.e., the extent to which her goals would be served if her lying intention were recognized) provides a principled basis for analytically distinguishing between these two phenomena.

Acknowledgments This chapter is based on presentations given at “New Directions in Linguistic Pragmatics” (NDLP2017), the International Symposium on Advances in (Im)politeness Studies 2018 at Fujian Normal University, and the workshop on Perspectives on Speech as Action in Trieste, in 2018 and its writing was made possible in part through a research leave granted by the Leiden University Center for Linguistics. I am grateful to the audiences on these occasions for their questions, as well as to the editors of this volume for their insightful comments on a previous draft. All remaining errors are my own.

References

- Austin, John L. 1975. *How to do Things with Words*. Oxford: Oxford University Press, 1962 (2nd ed.).
- Bach, Kent. 2001. You don't say? *Synthese* 128: 15–44.
- Beeman, William. 1986. *Language, Status, and Power in Iran*. Bloomington: Indiana University Press.
- Bok, Sissela. 1978. *Lying. Moral Choice in Public and Private Life*. New York: Pantheon.
- Brown, Penelope, and Stephen Levinson. 1987. *Politeness: Some Universals in Language Usage*. Cambridge: Cambridge University Press.
- Bryant, Erin. 2008. Real Lies, White Lies and Gray Lies: Towards a Typology of Deception. *Kaleidoscope* 7: 23–50.
- Bussey, Kay. 1999. Children's Categorization and Evaluation of Different Types of Lies and Truths. *Child Development* 70: 1338–1347.
- Camden, Carl, Michael Motley, and Ann Wilson. 1984. White Lies in Interpersonal Communication: A Taxonomy and Preliminary Investigation of Social Motivations. *The Western Journal of Speech Communication* 48: 309–325.
- Chisholm, Roderick, and Thomas Feehan. 1977. The Intent to Deceive. *Journal of Philosophy* 74: 143–159.
- Coleman, Linda, and Paul Kay. 1981. Prototype Semantics: The English Word Lie. *Language* 57: 26–44.

- Eslami, Zohreh. 2005. Invitations in Persian and English: Ostensible or Genuine? *Intercultural Pragmatics* 2: 453–480.
- Fallis, Don. 2009. What Is Lying? *Journal of Philosophy* 106: 29–56.
- Fish, Karyn, Kathrin Rothermich, and Marc Pell. 2017. The Sound of (In) Sincerity. *Journal of Pragmatics* 121: 147–161.
- Green, Mitchell. 2017. Speech Acts. In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/win2017/entries/speech-acts/>.
- Grice, Paul H. 1957. Meaning. *The Philosophical Review* 66, 377–88. Reprinted in *Studies in the Way of Words* (1989), 213–223. Cambridge, MA: Harvard University Press.
- . 1969. Utterer's Meaning and Intentions. *The Philosophical Review* 78. Reprinted in *Studies in the Way of Words* (1989), 86–116. Cambridge, MA: Harvard University Press.
- . 1975. Logic and Conversation. In *Syntax and Semantics, Vol. 3: Speech Acts*, edited by Paul Cole and Jerry Morgan. New York: Academic Press. Reprinted in *Studies in the Way of Words* (1989), 22–40. Cambridge, MA: Harvard University Press.
- Heyman, Gail, and Kang Lee. 2012. Moral Development: Revisiting Kohlberg's Stages. In *Developmental Psychology: Revisiting the Classic Studies*, ed. Alan Slater and Paul Quinn, 164–175. London: Sage.
- Huls, Erica, and Carel van Wijk. 2012. The Development of a Directive Repertoire in Context: A Case Study of a Dutch Speaking Young Child. *Journal of Pragmatics* 44: 83–103.
- Isaacs, Ellen, and Herbert Clark. 1990. Ostensible Invitations. *Language in Society* 19: 493–509.
- Izadi, Ahmad. 2015. Persian Honorifics and Im/Politeness as Social Practice. *Journal of Pragmatics* 85: 81–91.
- Knapp, Mark, and Mark Comadena. 1979. Telling It like It Isn't: A Review of Theory and Research on Deceptive Communication. *Human Communication Research* 5: 270–285.
- Koutlaki, Sophia. 1997. *The Persian System of Politeness and the Persian Folk Concept of Face, with Some Reference to Efl Teaching to Iranian Native Speakers*. Ph.D. Dissertation, University of Wales.
- Kupfer, Joseph. 1982. The Moral Presumption against Lying. *Review of Metaphysics* 36: 103–126.
- Lakoff, Robin. 1992. Pragmatics and the Law: Speech Act Theory Confronts the First Amendment. In *Papers from the 27th Regional Meeting of the Chicago*

- Linguistic Society, Vol. 1: The Main Session*, ed. Lise Dobrin, Lynn Nichols, and Rosa Rodriguez, 306–323. Chicago: Chicago Linguistic Society.
- Levine, Emma, and Maurice Schweitzer. 2015. Prosocial Lies: When Deception Breeds Trust. *Organizational Behavior and Human Decision Processes* 126: 88–106.
- Ma, Fengling, Xu Fen, Gail Heyman, and Kang Lee. 2011. Chinese Children's Evaluations of White Lies: Weighing the Consequences for Recipients. *Journal of Experimental Child Psychology* 108: 308–321.
- Mahon, James. 2008. Two Definitions of Lying. *International Journal of Applied Philosophy* 22: 211–230.
- . 2019. Contemporary Approaches to the Philosophy of Lying. In *The Oxford Handbook of Lying*, ed. Jörg Meibauer, 32–55. Oxford: Oxford University Press.
- Meibauer, Jörg. 2011. On Lying: Intentionality, Implicature, and Imprecision. *Intercultural Pragmatics* 8: 277–292.
- . 2014. Bald-Faced Lies as Acts of Verbal Aggression. *Journal of Language Aggression and Conflict* 2: 127–150.
- Miller, Corey, Rachel Strong, Mark Vinson, and Claudia Brugman. 2014. *Ritualized Indirectness: Explaining the Language Practice Taarof as a Reflection of Persian Speakers' Cultural Values*. Technical Report, Center for Advanced Study of Language. University of Maryland.
- Mojdehi, Atiyeh Shohoudi, Asadeh Shohoudi, and Victoria Talwar. 2020. Children's Moral Evaluations of Different Types of Lies and Parenting Practices and across Cultural Contexts. *Current Psychology*. <https://doi.org/10.1007/s12144-020-01059-7>.
- Moreno, Eva, Pilar Casado, and Manuel Martin-Loeches. 2016. Tell Me Sweet Little Lies: An Event-Related Potentials Study on the Processing of Social Lies. *Cognitive Affective and Behavioural Neuroscience* 16: 616–625.
- Reboul, Anne. 1994. The Description of Lies in Speech Acts Theory. In *Pretending to Communicate*, ed. Herman Parret, 292–298. Berlin, New York: de Gruyter.
- Saul, Jennifer. 2012. *Lying, Misleading, and What Is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford: Oxford University Press.
- Sbisà, Marina. 2001. Illocutionary Force and Degrees of Strength in Language Use. *Journal of Pragmatics* 33: 1791–1814. Reprinted in Marina Sbisà (forthcoming), *Essays on Speech Acts and Other Topics in Pragmatics*. Oxford: Oxford University Press.

- . 2009. Uptake and Conventionality in Illocution. *Lodz Papers in Pragmatics* 5: 33–52. Reprinted in Marina Sbisà (forthcoming), *Essays on Speech Acts and Other Topics in Pragmatics*. Oxford: Oxford University Press.
- . 2018. Varieties of Speech Act Norms. In *Normativity and Variety of Speech Actions*, edited by Maciej Witek and Iwona Witczak-Plisiecka, 23–50. Leiden: Brill. Reprinted in Marina Sbisà (forthcoming), *Essays on Speech Acts and Other Topics in Pragmatics*. Oxford: Oxford University Press.
- Searle, John. 1969. *Speech Acts. An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
- Sperber, Dan, and Deirdre Wilson. 1995. *Relevance. Communication and Cognition*. 2nd ed. Oxford: Blackwell.
- Talwar, Victoria, and Kang Lee. 2002. Emergence of White-Lie Telling in Children between 3 and 7 Years of Age. *Merrill-Palmer Quarterly* 48: 160–181.
- Talwar, Victoria, Sean Murphy, and Kang Lee. 2007. White Lie-Telling in Children for Politeness Purposes. *International Journal of Behavioral Development* 30: 1–11.
- Terkourafi, Marina. 2008. Toward a Unified Theory of Politeness, Impoliteness, and Rudeness. In *Impoliteness in Language: Studies on Its Interplay with Power in Theory and Practice*, ed. Derek Bousfield and Miriam Locher, 45–74. Berlin: Mouton de Gruyter.
- . 2014. The Importance of Being Indirect: A New Nomenclature for Indirect Speech. *Belgian Journal of Linguistics* 28: 45–70.
- Turner, Ronny, Charles Edgely, and Glenn Olmstead. 1975. Information Control in Conversation: Honesty Is Not Always the Best Policy. *Kansas Journal of Sociology* 11: 69–89.
- Vincent, Jocelyne, and Cristiano Castelfranchi. 1981. On the Art of Deception: How to Lie While Saying the Truth. In *Possibilities and Limitations of Pragmatics*, ed. Herman Parret, Marina Sbisà, and Jef Verschueren, 749–777. Amsterdam: Benjamins.
- Walper, Sabine, and Renate Valtin. 1992. Children's Understanding of White Lies. In *Politeness in Language: Studies in its History, Theory and Practice*, ed. Richard Watts, Sachiko Ide, and Konrad Ehlich, 231–251. Berlin: Mouton de Gruyter.
- Weissman, Benjamin, and Marina Terkourafi. 2018. Are False Implicatures Lies? An Empirical Investigation. *Mind and Language* 34 (2): 221–246.
- Wilson, Deirdre, and Dan Sperber. 2002. Truthfulness and Relevance. *Mind* 111: 583–632.

- Xu, Fen, Xuehua Bao, Fu Genyuen, Victoria Talwar, and Kang Lee. 2010. Lying and Truth-Telling in Children: From Concept to Action. *Child Development* 81: 581–596.
- Yang, Jia 2008. How to Say “No” in Chinese: A Pragmatic Study of Refusal Strategies in Five TV Series. In *Proceedings of the 20th North American Conference on Chinese Linguistics (NACCL-20) 2*, edited by Marjorie Chan and Hana Kang, 1041–1058. Columbus: The Ohio State University.
- Zembaty, Jane. 1988. Plato’s Republic and Greek Morality on Lying. *Journal of the History of Philosophy* 26: 517–545.
- . 1993. Aristotle on Lying. *Journal of the History of Philosophy* 31: 7–29.