

Saying, presupposing and implicating: How pragmatics modulates commitment

Diana Mazzarella, Robert Reinecke, Ira Noveck, & Hugo Mercier

This is the peer-reviewed version of the following article: Mazzarella, D., Reinecke, R, Noveck, I. & Mercier, H. (2018) Saying, presupposing and implicating: How pragmatics modulates commitment. *Journal of Pragmatics*, 133, 15-27, which has been published in final form at <https://doi.org/10.1016/j.pragma.2018.05.009>.

Saying, presupposing and implicating: How pragmatics modulates commitment

Abstract Commitment plays a crucial role in the stabilization of communication. While commitment increases the acceptance of the message communicated, it comes with a price: the greater the commitment, the greater the cost (direct or reputational) the speakers incur if the message is found unreliable (Vullioud, Clément, Scott-Phillips & Mercier, 2017). This opens up the question of which linguistic cues hearers deploy in order to infer speaker commitment in communication. We present a series of empirical studies to test the hypothesis that distinct meaning-relations – *saying*, *presupposing* and *implicating* – act as pragmatic cues of speaker commitment. Our results demonstrate that, after a message *p* is found to be false, speakers incur different reputational costs as a function of whether *p* had been explicitly stated, presupposed, or implicated. All else being equal, participants are significantly more likely to selectively trust the speaker who implicated *p* than the speaker who asserted or presupposed *p*. These results provide the first empirical evidence that commitment is modulated by different meaning-relations, and shed a new light on the strategic advantages of implicit communication. Speakers can decrease the reputational damages they incur by conveying unreliable messages when these are implicitly communicated.

Keywords Commitment, Implicature, Presupposition, Experimental pragmatics

1. Introduction

On June 1st, 2016, during a rally in Sacramento, Donald Trump accused his Democratic rival, Hilary Clinton, of telling “such lies about his foreign policy.” The issue at stake concerned his position on nuclear weapons and Japan, and he forcefully denied having claimed that Japan should obtain nuclear power. In fact, Trump had expressed his opinion in two distinct interviews. In a first interview at a CNN town hall in March, he had said that “At some point we have to say, you know what, we’re better off if Japan protects itself against this maniac in North Korea”. Clearly, while he had not explicitly stated it, Trump had (strongly) suggested the idea of Japan getting nuclear power. However, in a second interview with Fox News Sunday a few days later, he claimed: “Maybe they would in fact be better off if they defend themselves from North Korea [...] Maybe they would be better off - including with nukes, yes, including with nukes.” The difference between Trump’s statements relies on the fact that while the former merely implicates that Japan should get nuclear power, the latter explicitly states it (Grice, 1989). The intuition here is that Trump’s *commitment* to what he communicated - and his accountability for it – increased from the first to the second interview. Consequently, his denial appeared not only implausible but also illegitimate.

The notion of commitment, widely employed in the linguistic literature, aims at capturing the fact that senders can endorse or distance themselves to differing degrees from what they communicate (for a review see Boulat & Maillat, 2017; Brabanter & Dendale, 2008). The study of commitment has traditionally focused on a variety of linguistic devices - evidentials, epistemic modals, verbal expressions of confidence, as well as reported speech - whose semantics constrains the attribution of

speaker commitment (e.g., see Lyons, 1977; Ifantidou, 2001; Noveck, Ho & Sera, 1996; Palmer, 1986; Papafragou, 2000, 2006; among many others). More recently, however, linguists have started investigating the ways in which commitment can be pragmatically modulated (Moeschler, 2013; Morency, Oswald & Saussure, 2008; Saussure & Oswald, 2009). This research has opened up the question of whether a communicated assumption carries distinct degrees of speaker's commitment depending on its relation to the semantic content of the utterance as well as on the role it plays in the overall interpretation of the speaker meaning.

In what follows, we address the question of whether the meaning relations of *saying*, *presupposing*, and *implicating* are pragmatic cues of the degree to which a speaker is committed to the proposition conveyed. Our goal is two-fold: on the one hand, we aim to provide a theoretical framework in which one can investigate the pragmatic modulation of commitment; on the other hand, we explore this modulation in a series of empirical studies. We approach this linguistic endeavour by borrowing theoretical and methodological tools from evolutionary and cognitive psychology (Vullioud, Clément, Scott-Phillips & Mercier, 2017, see also, e.g., Anderson, Brion, Moore, & Kennedy, 2012; Fusaroli *et al.*, 2012; Tenney, MacCoun, Spellman, & Hastie, 2007; Tenney, Spellman, & MacCoun, 2008)).

2. Commitment across meaning-relations

In this section, we introduce the notion of *saying*, *implicating*, and *presupposing* and discuss different theoretical proposals regarding their relative degree of commitment. The obvious starting point of this investigation is the Gricean distinction between 'what is said' and 'what is implicated' (Grice, 1989). Grice argued that speakers typically communicate more than they linguistically encode. That is, a speaker can *say* something while *implicating* further propositional contents. Crucially, implicatures are by definition cancellable, either explicitly or implicitly:

[...] a putative conversational implicature that *p* is explicitly cancelable if, to the form of words the utterance of which putatively implicates that *p*, it is admissible to add *but not p*, or *I do not mean to imply that p*, and it is contextually cancelable if one can find situations in which the utterance of the form of words would simply not carry the implicature. (Grice, 1989, p. 44)

Several authors maintain that, because of the cancellability of 'what is implicated,' *implicating* is non-committal, or at least less committal than *saying*. Cancellability is described as closely intertwined with the following notions: (i) non truth-conditionality, (ii) deniability, and (iii) certainty about the intended interpretation. All these are relevant to the study of commitment. Implicatures are non truth-conditional content, that is, their truth-value has no bearing on the truth of the utterance that carries them: if they are false, the utterance is odd, but not necessarily false (Carston, 2004). According to Moeschler (2013),

this makes *implicating* a weak meaning-relation, “which means that the commitment of the speaker is not as strong as with the other relations” (Moeschler, 2013, p. 96), such as *saying*.

Furthermore, the cancellability of ‘what is implicated’ opens the door to its deniability (Pinker, 2007; Lee & Pinker, 2010). A content is deniable if the speaker can deny (when openly challenged) to have had the intention to communicate it in the first place. Deniability and cancellability do not overlap: what is deniable is cancellable, but what is cancellable is not necessarily deniable (at least not plausibly). Consider the following example adapted from Grice (1989). Mary is worried that her husband might be having an extra-marital affair and asks her friend Susy what she thinks about it. Susy replies:

(1) I saw your husband with a woman the other day at the cinema.

Mary interprets Susy’s utterance as implicating that her husband is indeed likely to be having an affair. It turns out that Mary’s husband was at the cinema, but in the company of his sister, whom Susy knows very well. When confronted with this, Susy defends herself by claiming: “I didn’t mean to suggest that he had an affair. In fact, the woman he was with was his sister.” In this case, the implicature is cancellable but hardly deniable, as suggested by the fact that Susy’s defense is very much unlikely to convince Mary. While cancellability is a binary category (cancelable/non-cancelable), deniability is a matter of degree and it is a function of the discourse-related properties of the context of utterance (see, e.g., Sternau, Ariel, Giora & Fein, 2015, 2016, 2017). Importantly, Pinker (2007) suggests that the deniability of ‘what is implicated’ allows a speaker to convey some content (sexual innuendos, bribes, threats, etc.), without running into the risk of paying its potential cost (direct and/or reputational). For instance, by implicating a bribe to a police officer (“So maybe the best thing would be to take care of it here”), a speaker can avoid the risk of being arrested for bribery by an honest cop or a speaker can avoid the embarrassment of having a sexual advance turned down, if it had been merely implicated. That is, in social interactions, *implicating* is reputationally less costly for the speaker than *saying*.

Finally, the cancellability of ‘what is implicated’ depends on its context-dependency: ‘what is implicated’ is cancellable because it is possible to find contexts in which the speaker could use the same utterance without conveying the same set of implicatures. Implicatures are the result of an inferential process that takes ‘what is said’ as a premise, and together with available contextual assumptions, leads to an implicated conclusion or implicature (Sperber & Wilson, 1995). The selection (or construction) of the intended contextual assumption introduces an additional risk of misunderstanding. Because of this, Morency *et al.* (2008) maintain that *implicating* involves a more risky interpretative process than *saying*, which typically generates a lower degree of confidence in the hearer. This in turn leads to a weaker attribution of commitment towards the implicated content.

The picture that emerges is clear: *implicating* is taken to be less committal than *saying*. This conclusion is further echoed by work in the philosophy of language and epistemology of testimony. For instance, the philosopher Elizabeth Fricker has argued that insinuation and hints allow a message to be

conveyed while being disavowable by its sender. The distinction between *saying* and *implicating* is conceived as a “*socially entrenched distinction*” to the extent that social norms make speakers accountable for the former, but not for the latter. She writes:

It is only what a speaker explicitly states that she incurs overt, full and undeniable responsibility for the truth of. Things she leaves it to her audience to figure out, even if she fully expects and intends the audience to figure them out, and this is part of the perlocutionary point of her utterance, are not committed to by her – not, at any rate, to the same full and undeniable extent. (Fricker 2012, p. 85)

In spite of the wide consensus that the *saying/implicating* distinction is relevant to the domain of speaker’s commitment, some scholars have recently challenged this conclusion. For instance, Meibauer (2014) argues that speakers are committed to the truth of an implicature in the same way as they are committed to the truth of what they explicitly say. This is because – according to Meibauer – a deliberately false implicature qualifies as a genuine *lie*, that is, a content that the speaker believes to be false and puts forth with the intention of creating a false belief in the audience.

Let us now develop our discussion beyond the Gricean categories of *saying* and *implicating*, and include the meaning-relation of *presupposing* (Ducrot, 1984; Stalnaker, 1974). In fact, while it is uncontroversially assumed that a speaker is committed to the truth of what her utterance presupposes, less attention has been devoted to the relative degree of commitment with regard to other meaning-relations. The question of whether *presupposing* is more or less committal than *saying* or *implicating* has not received a univocal answer. For instance, Moeschler (2013) has suggested that the implicit nature of presupposed content – typically what is presupposed is not explicitly articulated in discourse – should lower the degree of speaker commitment (as it does with *implicating*). On the other hand, he argues that the strength of the inference drawn to recover the presupposed content (linguistically triggered and semantic in nature) should make *presupposing* more committal than *saying*. Unfortunately, no principled way to establish the relative weight of these conflicting determinants of speaker commitment is provided, and the issue remains open.

Intuitions about the relative degree of commitment of *presupposing* with respect to *saying* and *implicating* are not clear-cut. ‘What is presupposed’ is typically background information, that is, information which is old, previous, or given (van der Sandt, 1992) or presented as such (Saussure, 2013). It is information that “the hearer will not want to dispute” (von Stechow, 2000), and that the speaker puts forth as uncontroversial. As Simons (2005) points out, presuppositions tend to be non-controversial, independently of whether they are assumed to be shared prior the time of the utterance. This is due to the role they play in the interpretation of the utterance. Presuppositions are meant to ‘establish’ the relevance of the utterance. As a result, if presuppositions were not to be accepted by the addressee, the utterance would not provide any relevant contribution to the conversation. This suggests that the degree

of commitment associated to ‘what is presupposed’ might be stronger than that associated with foreground contents, like ‘what is said’ and ‘what is implicated.’

However, the backgrounded nature of presupposed content – it being typically presented as part of the common ground (Stalnaker, 1974, 2002) or as likely to be accepted uncontroversially by the interlocutors (Simons, 2005) – may also indicate that interlocutors share the responsibility for ‘what is presupposed’ in the conversation. In line with this intuition, Lombardi Vallauri (2016a, p. 3) argues that “messages containing presuppositions [...] conceal the very act of proposing it [*their content*] as true, as if the speaker has *no commitment* to transferring it” (my emphasis). In the same vein, Ducrot (1984) conceives of presuppositions as echoing a “collective voice” and, as a result, the speaker is assumed not to take direct responsibility for the presupposed content.

To sum up, the linguistic literature on commitment modulation via meaning-relations is theoretically and experimentally lacking. On the theory side, the apparent consensus that commitment is a graded notion is not grounded on any well-established criteria for classification. That is, there is no unanimously accepted criterion to determine the degree of speaker commitment that pertains to each meaning-relation. The notions of ‘truth-conditionality,’ ‘deniability,’ ‘hearer’s certainty,’ ‘explicitness,’ ‘accessibility,’ ‘inferential strength’ have all been employed to investigate the gradedness of commitment, generating orderings which are often in conflict with one another. Critically, we find no studies that address this issue experimentally. These two shortcomings are clearly linked to one another: the lack of a systematic treatment of commitment in the theoretical literature has arguably prevented any sound empirical investigation.

This state of affairs calls for new approaches to the study of commitment. To move the field forward we suggest turning one’s attention to the social and interactional dimension of commitment, as well as the sender’s liability to criticism and blame, which goes hand in hand with any commitment violation (Harnish, 2005; Haugh, 2013). As already discussed in the literature, this approach has important consequences for linguistic theorizing. For instance, it allows one to overcome some problems with defining commitment in terms of the speaker’s mental states or propositional attitudes (e.g., as the strength of the speaker’s belief towards a proposition). Among these we find the issue of the ‘inscrutability’ of commitment, as well as the possibility of a mismatch between the speaker’s mental states (her beliefs) and the degree of commitment that she intends to communicate in interaction (see, e.g. Saussure & Oswald, 2009). In the following section, we examine how recent insights from psychology can provide us with methodological tools to help investigate commitment experimentally.

3. The role of commitment in the evolution of communication

Reputation plays a crucial role in our daily interactions, as it shapes the beliefs that others hold about one’s worth as a potential cooperator (Heintz, Karabegovic & Molnar, 2016; Sperber & Baumard, 2012). This is also true for communicative interactions. Communication creates a complex variety of cooperative opportunities, but it amplifies the risk of deception too. For this reason, it requires

interlocutors to select vigilantly which conversational partners are worth their trust, and to punish – at least reputationally – those who are not (Sperber *et al.*, 2010). This calls for a better understanding of the linguistic devices that allow to manage one's reputation as a trustworthy conversational partner.

In a recent contribution, Vulloud, Clément, Scott-Phillips and Mercier (2017) have suggested that commitment plays a crucial role in reputation management, and have focused their attention on strategic concerns in the understanding of the dynamics of commitment. They maintain that one of the effects of commitment is to moderate the acceptance of the information communicated (see also Anderson *et al.*, 2012). Specifically, they argue that people tend to be more accepting of statements to which the speaker is more committed (e.g. statements that are expressed more confidently). However, as Vulloud *et al.* (2017) emphasize, there has to be a potential downside to commitment, otherwise the equilibrium would be for everyone to commit maximally to any statement they want the interlocutor to accept. This would render commitment a useless cue, and interlocutors would stop paying attention to cues to commitment. In line with this, a review of the literature suggests that commitment comes at a price. Several experiments have shown that among speakers whose statements have been found to be false, speakers who had expressed more commitment suffered more reputational damage than those who had expressed less commitment (Tenney *et al.*, 2007; Tenney, Small, Kondrad, Jaswal, & Spellman, 2011; Tenney *et al.*, 2008; Vulloud, Clément, Scott-Phillips, & Mercier, 2017).

By focusing on the use of confidence expressions as commitment signals, Vulloud *et al.* (2017) ran a series of studies in which they showed that a confident speaker incurs a greater reputational damage than an unconfident speaker when her testimony turns out to be unreliable. Crucially, they also show that this damage is more severe when the message is accepted based on the speaker's commitment than for reasons other than that (such as the speaker's competence). Their studies share the following template: participants are presented the testimonies of two senders, which differ from each other with respect to their degree of confidence. It is then revealed that they are both wrong and participants are asked to decide which of the two senders they would like to punish and which one they would like to trust in the future, in a different domain from that in which he was found to be wrong. That is, their dependent measures are represented by participants' answers to a 'punishment question' and a 'trust question.' While the first aims at capturing the direct cost of commitment violation, the second is a measure of reputational damage. Both the punishment and the trust question provide an indication of the blameworthiness and the liability of a committed sender.

Drawing on this research, we can operationalize commitment as a function of the direct and reputational costs the senders incurs when her message is found to be false. This provides us with an empirically tractable notion of commitment, one that capitalizes on its interpersonal nature and that does not require any metalinguistic judgement from the participant's side. In the following sections we present a series of studies that aims at investigating the relative degree of speaker commitment towards a message when this is explicitly communicated, implicated or presupposed.

4. Experiment 1

By adopting Vullioud *et al.*'s (2017) paradigm, we aim at empirically investigating the way in which meaning-relations modulate speaker commitment. With regard to this, our investigation is genuinely explorative, as it does not rely on widely shared predictions. As noted above, the strongest prediction which we can draw from the existing literature concerns the relationship between what is said and what is implicated, and goes in the direction of attributing a stronger commitment to the former than to the latter (but see Meibauer, 2014). The picture is far less clear when it comes to presuppositions, and no clear prediction is available with regard to whether *presupposing* would be more or less committal than *saying* and *implicating*. By subjecting the literature to empirical testing, we aim at moving this debate forward and at providing data on which to build new, potentially finer-grained, theoretical contributions.

4.1 Methods

Participants

We recruited 291 participants through Amazon Mechanical Turk (163 men, 127 women, 1 others, $M_{age} = 34.94$, $SD = 10.08$).

Materials and procedure

We created four stories each describing a professional context in which the participant receives the testimonies of two speakers. Both speakers convey the same piece of information, varying in their degree of confidence (condition 1) or with respect to the level of meaning at which the information is conveyed (condition 2-4). In light of the information received, the participant is asked to type in a message to a third story character. The participant then receives additional information which falsifies the two testimonies. Finally, the participant answers two forced-choice test questions: a 'punishment question' and a 'trust question'. In Conditions 3 and 4, the participant is also asked a third 'implicature question'. Question order for the punishment question and the trust question as well as order of presentation of the speakers were counterbalanced across participants. The implicature question – when present – was always displayed at the end. The implicature question served as a criterion to ensure that we would only retain the data of those participants that had indeed derived the relevant implicature on the basis of the available contextual information.

The experiment was comprised of the following four conditions:

- Condition 1: confident speaker vs. unconfident speaker
- Condition 2: *presupposing* speaker vs. *saying* speaker
- Condition 3: *saying* speaker vs. *implicating* speaker
- Condition 4: *implicating* speaker vs. *presupposing* speaker

Each story contained a different presupposition trigger belonging to the following four categories: it-clefts, iteratives (*too*), change of state verbs (*repair*) and emotive factives (*be relieved that*). This selection involves both syntactic constructions and lexical items and it crosscuts alternative

Analysis

We removed from our analysis participants who answered either “No” or “I don’t know” to the implicature question in Conditions 3 and 4 (130 participants). Furthermore, we excluded 1 non-native English speaker. Our final sample included 160 participants. Participants were equally distributed across the four conditions. In order to have comparable samples for each condition, we had anticipated the exclusion rate for Condition 3 and Condition 4 (based on a pilot), and recruited proportionately more participants for these conditions. The writing task, i.e. typing a message, was meant to increase the engagement of the participant in the fictional story. As it was not directly relevant to our concerns, we did not perform any qualitative analysis on the participants’ answers.

4.2 Results

Condition 1 (confident vs. unconfident speaker). The results replicated Vullioud *et al.* (2017). Participants were significantly more likely to punish the confident speaker (73%, 27/37, binomial $p = .008$) and to trust the unconfident one (84%, 31/37, binomial $p < .001$).²

Condition 2 (presupposing vs. saying speaker). Participants were as likely to punish the *presupposing* speaker as they were to punish the *saying* speaker (49%, 21/43, binomial $p = 1$). Similarly, there was no significant difference with regard to the trust question: 42% of the participants trusted the *saying* speaker and 58% trusted the *presupposing* speaker (25/43, binomial $p = .36$)

Condition 3 (saying vs. implicating speaker). The punishment question did not show any significant difference, with 63% of participants punishing the *saying* speaker and 37% the *implicating* speaker (26/41, binomial $p = .12$). By contrast, participants were significantly more likely to trust the *implicating* speaker than the *saying* speaker (73%, 30/41, binomial $p = .004$).

Condition 4 (implicating vs. presupposing speaker). Participants were as likely to punish the *implicating* speaker as they were to punish the *presupposing* speaker (51%, 20/39, binomial $p = 1$). However, a clear difference emerged with regard to the trust question: only 23% of the participants trusted the *presupposing* speaker, while 77% preferred to trust the *implicating* one (30/39, binomial $p = .001$).³

4.3 Discussion

² All the binomial tests here report two-sided p -values.

³ The pattern of results for Conditions 3 and 4 does not change if we include in the analysis all participants (independently of their answer to the implicature question). Participants were significantly more likely to trust the *implicating* speaker than the *saying* speaker (Condition 3: 86/104, binomial $p < .001$) and the *presupposing* speaker (Condition 4: 88/106, binomial $p < .001$). Furthermore, a statistical significance difference also emerges with respect to the Punishment question: participants were significantly less likely to punish the *implicating* speaker than the *saying* speaker (Condition 3: 27/104, binomial $p < .001$) and the *presupposing* speaker (Condition 4: 38/106, binomial $p = .005$).

A first important result of our experiment is the replication of Vullioud *et al.* (2017). It not only established the reliability of our adapted paradigm, but it also provided us with a way to compare the behaviour of explicit commitment signals (confidence expressions) with tacit cues of commitment (meaning-relations). The comparison between explicit and implicit cues suggests that the former leads to a stronger modulation of speaker commitment (as evidenced by the fact that both the punishment and the trust questions showed a significance difference in Condition 1). While the participants are inclined to impose direct costs (punishment question) as well as reputational ones (trust question) to a confident but mistaken speaker, when it comes to meaning-relations the speaker's damage is only reputational. Crucially, our results show that interlocutors are sensitive to the way in which content is communicated – whether it is presupposed, said or implicated. This provides the first empirical evidence that participants modulate the degree of speaker commitment as a function of a provided meaning-relation.

Our data show that *implicating* is taken to be less committal than *saying* and *presupposing*. This result suggests that the relative degree of commitment attributed to a message does not entirely depend on its degree of explicitness. While implicatures and presuppositions are contents typically left implicit (with the exception of some presupposition triggers), they clearly commit the speaker to different degrees. *Implicating* allows speaker to get their message across without incurring the same reputational damage as *saying* and *presupposing*, and it thus represents a powerful way to minimize drops of trust that result from unreliable testimony.

5. Experiment 2 (a and b)

Experiment 1 did not provide evidence of a difference in terms of the reputational damage incurred by a speaker who presupposes false information compared to a speaker who explicitly communicates it. To assess the robustness of this result, in Experiment 2 (a and b) we include new stimuli covering a broader range of presupposition triggers (Karttunen, 1969; Levinson, 1983). Along with the stimuli used in Experiment 1, we use items from the following categories: definite descriptions, temporal clauses, counterfactual conditionals and *only* (Experiment 2a), as well as new iteratives and focus-sensitive particles (Experiment 2b).

5.1 Methods (Experiment 2a)

Participants

We recruited 151 participants through Amazon Mechanical Turk (79 men, 72 women, $M_{age} = 35.11$, $SD = 9.98$).

Materials and procedure

Our stimuli comprised the set of stories used in Experiment 1 (Condition 2) and 4 additional new stories which included the following presupposition triggers: the definite description *the*, the temporal clause *after*, a counterfactual conditional and *only*. See Appendix B for a detailed list. The procedure was the

same as Experiment 1's. Participants were presented with the testimonies of a *presupposing* speaker and a *saying* speaker and they were asked to select which of the two they wished to punish or trust after receiving negative feedback about the information transmitted.

5.2 Results (Experiment 2a)

The results confirmed the null finding of Experiment 1 (Condition 2). Participants were equally likely to punish the *presupposing* speaker and the *saying* speaker (50%, 76/151, binomial $p = 1$), and equally likely to trust them (50%, 75/151, binomial $p = 1$).

We then looked at each presupposition trigger individually in order to detect any potential difference in behaviour. The only presupposition trigger that stood out was the iterative *again*: participants were more likely to punish the *saying* speaker than the *presupposing* one. However, this result is only significant if we do not correct for multiple comparisons (76%, 16/21, binomial $p = .03$).

5.3 Discussion (Experiment 2a)

Overall, data from Experiments 1 and 2 are consistent with each other. Participants do not seem to be more prone to attribute different reputational costs to a speaker that has presupposed a message which turns out to be false than to a speaker who has explicitly stated it. This in turn suggests that to the extent that *presupposing* and *saying* are taken as cues to attribute a certain degree of commitment towards the message conveyed, they behave very similarly to each other.

However, given the suggesting that the presupposition trigger *again* might differ from others, it leaves open the question of whether different presupposition triggers might yield different degrees of speaker commitment. With regard to this, it is worth noting that *again* belongs to a class of presupposition triggers that Glanzberg (2003, 2005) has called 'weak presupposition triggers.' This class includes iteratives and focus-sensitive particles and it is characterised by the fact that in case of presupposition failure, repair is typically optional. That is, even if the presupposition is false or not part of the common ground, the utterance explicitly communicates a proposition whose truth-value can be independently evaluated. Take the following example from our set of stimuli: "He is late again". Even if the presupposition fails, the proposition that *He is late (today)* can be coherently assessed as true or false. Because of this feature of weak presupposition triggers, it is plausible to hypothesize that, in these linguistic contexts, the speaker might commit more to what is said than to what is presupposed. This hypothesis sets the ground for our Experiment 2b.

5.4 Methods (Experiment 2b)

Participants We recruited 77 participants through Amazon Mechanical Turk (48 men, 29 women, $M_{age} = 33.02$, $SD = 10.21$). We excluded 3 participants as they were not English native speakers.

Materials and procedure Our stimuli comprised 4 stories which included the following presupposition triggers: the iteratives *again, also, too* and the focus-sensitive particle *even*. See Appendix B for a detailed list. The procedure was the same of Experiment 2a.

5.5 Results (Experiment 2b)

The results did not confirm our hypothesis. Overall, participants were equally likely to punish the *presupposing* speaker and the *saying* speaker (57%, 32/74, binomial $p = .30$), and trusted them to the same extent (51%, 38/74, binomial $p = .91$).

5.6 Discussion (Experiment 2b)

The results of Experiment 2b confirm that participants do not treat *presupposing* and *saying* as fundamentally different with respect to degrees of speaker commitment. This holds even when looking at more subtle distinctions within the heterogeneous class of presuppositions triggers, such as the distinction between weak and strong triggers (Glanzberg, 2003, 2005). This suggests that *presupposing* and *saying* expose the speaker to the same reputational damage if the message conveyed is found to be unreliable.

6. Experiment 3

In line with the predictions of most of the linguistic and philosophical literature on the distinction between what is said and what is implicated, Experiment 1 shows that *implicating* is less committal than *saying*. Specifically, it reveals that senders can safeguard (at least partially) the trust receivers grant them by communicating unreliable messages by means of implicatures. To confirm the robustness of our results, we conducted a replication study.

6.1 Methods

We recruited 85 participants through Amazon Mechanical Turk (50 men, 34 women, 1 other, $M_{age} = 34.48$, $SD = 10.60$). The material, procedure and criteria for data cleaning were the same as Experiment 1 (Condition 3). Our final sample included 50 participants.

6.2 Results

The results confirmed our data from Experiment 1. While participants were as likely to punish the *saying* speaker as the *implicating* one (50%, 25/50, binomial $p = 1$), they clearly prefer to trust the *implicating* speaker over the *saying* one (70%, 35/50, binomial $p = .007$).

7. General discussion

In this article, we provide the first - to the best of our knowledge - empirical evidence that the meaning-relation through which the speaker conveys her message affects the degree of speaker commitment to

its truth. Specifically, we show that, holding everything else equal, the meaning-relation of implicating is taken to be less committal than both saying and presupposing (which appear to be on a par). We compared speakers *implicating*, *saying*, or *presupposing* the same piece of information, and assessed the extent to which receivers adjusted their trust towards the sender after her testimony was revealed to have been misguided. Speakers are judged as less blameworthy when they implicate rather than explicitly communicate or presuppose a false piece of information. This result cannot be attributed to receivers believing that the message was only accidentally conveyed, that is, that speakers did not *intend* to implicate it. Thanks to the presence of the ‘implicature question,’ our studies assessed what participants took to be the speaker’s intended meaning. By restricting our analysis to participants who recognized the implicature as part of what the speaker intended to communicate with her utterance, we show that trustworthiness is less damaged even when false information is taken to be intentionally implicated (and not the result of accidental misunderstanding). Meaning-relations act as cues for the addressee to infer the speaker’s degree of commitment towards the information communicated (Moeschler, 2013), and are thus relevant to establishing the social implications of an act of communication. These implications include speakers’ liability to criticism and blame. As a result, speakers can strategically deploy these cues with the aim of managing their reputation.

To conclude, we would like to discuss our results in light of recent proposals concerning the evolution of implicit communication. Implicit communication raises a *prima facie* puzzle, which is the following: why is implicit communication pervasive despite it being costly and open to misunderstandings. It has been well documented that implicit communication is costly, as implicature derivation typically imposes extra processing costs to the receiver (see, e.g. Bott & Noveck, 2004; Tomlinson, Bailey, & Bott, 2013). Furthermore, because the recovery of implicatures exclusively relies on contextual cues, implicit communication is more prone to misunderstandings and, as a result, it reduces the likelihood that the sender will get her message across. To solve this puzzle, Reboul (2017a, 2017b) proposes what she calls the “manipulation hypothesis,” that is, the hypothesis that implicit communication would have emerged in order to facilitate manipulation. Specifically, Reboul suggests an explicit link between commitment and manipulation. The argument goes as follows. First, *implicating* does not commit the speaker to the truth of the message conveyed. Second, a message on the truth of which the speaker does not commit is less likely to be critically evaluated by the receiver. It follows that *implicating* increases the chance of the receiver accepting the speaker’s message without further scrutiny.

Our results offer a different perspective on the relation between commitment and manipulation, which enriches and complements the discussion above. Meaning-relations are linguistic tools that can modulate the degree of speaker commitment and thus be used to influence interlocutors. By lowering her commitment, a speaker can attempt to get her message across while reducing social sanctions. Our results show that implicating is taken to be less committal than saying and presupposing. As a result, the advantage of implicating relies on the fact that the speaker can reduce the costs associated with the

transmission of false information, and preserve some of her reputation as a reliable source of information. This requires calibrating one's commitment towards the unreliable pieces of information that are part of the overall communicated message. It is worth noting, though, that one's reputation is affected by, not only the outcomes of his or her actions but also by his or her underlying intentions, preferences and task specific capabilities (Heintz *et al.*, 2016). While our studies do not explicitly distinguish between deceptive and honest but mistaken communicators, future research should look at the effect of intentional or accidental misinformation on the relation between commitment and reputation.

Reputation managing is an essential component of our social life and helps us navigate the intricate dynamics of social interactions. This work ultimately contributes to advance our understanding of how reputation is managed in a market of potential communicative partners and opens up an interesting new line of research at the interface between pragmatics and social psychology.

References

- Abusch, D. (2010). Presupposition triggering from alternatives. *Journal of Semantics*, 27(1), 37-80.
- Anderson, C., Brion, S., Moore, D. A., & Kennedy, J. A. (2012). A status-enhancement account of overconfidence. *Journal of Personality and Social Psychology*, 103(4), 718.
- Boulat, K. & Maillat, D. (2017). She said you said I saw it with my own eyes: A pragmatic account of commitment. In J Blochowiak., C Grisot., S Durrleman., C Laenzlinger (eds), *Formal Models in the Study of Language* (pp. 261-279). Springer, Cham.
- Brabanter, P. de, & Dendale, P. (2008). Commitment: the term and the notions. In P. de Brabanter & P. Dendale (eds.), *Commitment* (pp. 1-14), John Benjamins.
- Bott, L., & Noveck, I.A. (2004). Some utterances are underinformative: The onset and time course of scalar inferences. *Cognition*, 51, 437-457.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Carston, R. (2004). Truth-conditional content and conversational implicature. In C. Bianchi (ed.), *The Semantics/Pragmatics Distinction* (pp. 65-100). CSLI Publications.
- Ducrot, O. (1984). *Le dire et le dit*. Paris: Minuit.
- Fricker, E. (2012). Stating and insinuating. *Proceedings of the Aristotelian Society Supplementary*, LXXXVI, 61-94.
- Fusaroli, R., Bahrami, B., Olsen, K., Roepstorff, A., Rees, G., Frith, C., & Tylén, K. (2012). Coming to terms quantifying the benefits of linguistic coordination. *Psychological Science*, 0956797612436816.

- Glanzberg, M. (2003). Felicity and presupposition triggers. *Workshop in Philosophy and Linguistics*, University of Michigan, Michigan.
- Glanzberg, M. (2005). Presuppositions, truth values and expressing pro-positions. In G. Preyer & G. Peter (eds.), *Contextualism in philosophy: Knowledge, meaning, and truth*. Oxford, UK: Oxford University Press.
- Grice, P.H. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Harnish, R. M. (2005). Commitment and speech act, *Philosophica*, 75, 11-41.
- Haugh, M. (2013). Speaker meaning and accountability in interaction. *Journal of Pragmatics*, 48(1), 41-56.
- Heintz, C., Karabegovic, M., & Molnar, A. (2016). The co-evolution of honesty and strategic vigilance. *Frontiers in Psychology*, 7, 1503.
- Ifantidou, E. (2001). *Evidentials and Relevance*. Amsterdam, Philadelphia: John Benjamins.
- Karttunen, L. (1969). Pronouns and variables. In R. Binnick, A. Davison, G. Green & J. Morgan (eds.), *Syntax and semantics, 11: Presupposition*; New York, NY: Academic Press.
- Lee, J. J., & Pinker, S. (2010). Rationales for indirect speech: The theory of the strategic speaker. *Psychological Review*, 117(3), 785-807.
- Levinson, S.C. (1983). *Pragmatics*. Ambridge: Cambridge University Press.
- Lombardi Vallauri, E. (2016a). The “exaptation” of linguistic implicit strategies, *SpringerPlus*, 5, 1106.
- Lombardi Vallauri, E. (2016b) Implicits as Evolved Persuaders. In K. Allan, A. Capone & I. Kecskés (eds.), *Pragmemes and Theories of Language Use* (725-748). Cham, Switzerland: Springer International Publishing.
- Lyon, J. (1977). *Semantics*. London, New York : Cambridge University Press.
- Meibauer, J. (2014). A truth that’s told with bad intent. Lying and implicit content. *Belgian Journal of Linguistics*, 28, 97-118.
- Moeschler, J. (2013). Is a speaker-based pragmatics possible? Or how can a hearer infer a speaker’s commitment? *Journal of Pragmatics*, 48(1), 84-97.
- Morency, P., Oswald, S., & Saussure, L. de (2008). Explicitness, implicitness and commitment attribution: A cognitive pragmatic approach. *Belgian Journal of Linguistics*, 22, 197-219.
- Noveck, I., Ho, S., & Sera, M. (1996). Children’s understanding of epistemic modals. *Journal of Child Language*, 23(3), 621-643.
- Palmer, F.R. (1986). *Mood and Modality*. Cambridge: Cambridge University Press.
- Papafragou, A. (2000). On speech-act modality. *Journal of Pragmatics*, 32, 519-538.
- Papafragou, A. (2006). Epistemic modality and truth conditions. *Lingua*, 116, 1688-1702.
- Pinker, S. (2007). The evolutionary social psychology of off-record indirect speech acts. *Intercultural Pragmatics*, 4(4), 437-461.
- Reboul, A. (2017a). Is implicit communication a way to escape epistemic vigilance? In: S. Assimakopoulos (Ed.), *Pragmatics and Its Interfaces* (pp. 91-112), Berlin: Mouton de Gruyter.

- Reboul, A. (2017b). *Cognition and communication in the evolution of language*. Oxford: Oxford University Press.
- Saul, J. (2012). *Lying, Misleading, and What is Said*. Oxford: Oxford University Press.
- Saussure, L. de (2013). Background relevance. *Journal of Pragmatics*, 59, 178-189.
- Saussure, L., & Oswald, S. (2009). Argumentation et engagement du locuteur : Pour un point de vue subjectiviste. *Nouveaux Cahier de Linguistique Française*, 29, 215-243.
- Simons, M. (2005) Presuppositions and relevance. In Z. Gendler Szabo (ed.), *Semantics vs. Pragmatics* (pp. 329-355), New York: Oxford University Press.
- Sperber, D., & Baumard, N. (2012). Moral reputation: an evolutionary and cognitive perspective. *Mind & Language*, 27, 495-518.
- Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origg, G., Wilson, D. (2010). Epistemic vigilance. *Mind and Language*, 25(4), 359-393.
- Sperber, D., & Wilson, D. (1995). *Relevance: Communication and Cognition*. Oxford: Blackwell Publishing.
- Stalnaker, R. (1974). *Inquiry*. Cambridge, MA: MIT Press.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25(5-6), 701-721.
- Sternau, M., Ariel, M., Giora, R. & Fein, O. (2015). Levels of interpretation: New tools for characterizing intended meanings. *Journal of Pragmatics*, 84, 86-110.
- Sternau, M., Ariel, M., Giora, R. & Fein, O. (2016). A graded strength for privileged interactional interpretations. In K. Allan, A. Capone & I. Kecskés (eds.), *Pragmemes and Theories of Language Use* (703-724). Cham, Switzerland: Springer International Publishing.
- Sternau, M., Ariel, M., Giora, R. & Fein, O. (2017). Deniability and explicatures. In R. Giora & M. Haugh (eds.), *Doing Pragmatics Interculturally* (pp. 97-120). Berlin/Boston: Walter de Gruyter GmbH.
- Tenney, E. R., MacCoun, R. J., Spellman, B. A., & Hastie, R. (2007). Calibration trumps confidence as a basis for witness credibility. *Psychological Science*, 18(1), 46–50.
- Tenney, E. R., Small, J. E., Kondrad, R. L., Jaswal, V. K., & Spellman, B. A. (2011). Accuracy, confidence, and calibration: how young children and adults assess credibility. *Developmental Psychology*, 47(4), 1065.
- Tenney, E. R., Spellman, B. A., & MacCoun, R. J. (2008). The benefits of knowing what you know (and what you don't): How calibration affects credibility. *Journal of Experimental Social Psychology*, 44(5), 1368–1375.
- Tomlinson, J., Bailey, T. M. & Bott, L. (2013). Possibly all of that and then some: Scalar implicatures are understood in two steps. *Journal of Memory and Language*, 69, pp. 18-35.
- van der Sandt, R. (1992). Presupposition projection as anaphora resolution. *Journal of Semantics*, 9, 332-377.

- von Fintel, K. (2000). What is Presupposition Accommodation? ms. MIT, Cambridge, <http://web.mit.edu/fintel/www/accomm.pdf>
- Vullioud, C., Clément, F., Scott-Phillips, T. & Mercier, H. (2017). Confidence as an expression of commitment: why misplaced expressions of confidence backfire. *Evolution & Human behavior*, 38(1):9-17.

Appendix A

Full list of stimuli for Experiment 1. The punishment and the trust questions were the same for all stories (see body text).

Story 1

You are the new manager of a small advertising company. The company policy states that employees are not allowed to receive a weekly bonus if they are late more than once. After having been on a two-day business trip, you are back at work. You see that one of your employees, Richard, is not at his desk. So, you send an email to his two office mates and ask whether Richard is coming in. The replies come back as follows:

- Peter [*confident speaker*]: He is late, and I'm sure he was late yesterday too.
- Adam [*unconfident speaker*]: He is late, and I think he was late yesterday as well, but I'm not sure.
- Peter [*presupposing speaker*]: He is late again.
- Adam [*saying speaker*]: He is late today, like yesterday.
- Joe [*implicating speaker*]: Joe: He is late. Yesterday morning his car broke down, I don't know what happened today.

In light of this, you write the following message to Richard: [*type text*]

Later on, you discover that Richard presented a project at a customer's office yesterday morning and that he was late just today.

Implicature question: Remember what happened. Your employee Richard was not at his desk this morning. You asked his office mates whether he was coming in, and Joe said:

Joe: He is late. Yesterday morning his car broke down, I don't know what happened today.

Does Joe mean that Richard was late yesterday?

Story 2

See body text.

Story 3

You are the new supervisor of a team of shoe designers. Your team has a deadline for submitting a prototype to an important customer by tomorrow. You have not received the prototype yet, so you send an email to your team members and ask whether there is any progress on the project. The replies come back as follows:

- Peter [*confident speaker*]: We got some delays but it is almost ready! I'm sure the 3D printer hasn't been working until yesterday.

- Adam [*unconfident speaker*]: Nearly done - We were a bit delayed! I think the 3D printer hasn't been working until yesterday, but I'm not sure.
- Peter [*presupposing speaker*]: Almost ready! The 3D-printer was repaired yesterday.
- Adam [*saying speaker*]: Almost ready! The 3D-printer hasn't been working but as of yesterday it is okay.
- Joe [*implicating speaker*]: Almost ready! We could work with the 3D-printer as of yesterday.

In light of this, you write the following message to the IT support: [*type text*]

Later on, you talk to the IT support and you discover that the 3D printer has been working fine for the whole time.

Implicature question: Remember what happened. Your team had a deadline for submitting a project to an important customer by tomorrow, but you had not received the prototype yet. You asked your team members about it, and Joe said:

Joe: Almost ready! We could work with the 3D-printer as of yesterday.

Does Joe mean that the 3D-printer had not been working before yesterday?

Story 4

You are the new supervisor of the IT department at a large university. Last week the main server broke down. Researchers cannot store their data until the new server is in place. Your team received many complaints from several academic units and the atmosphere in the office was very tense. Today, you send an email to your two assistants and ask whether they feel less stressed about the server issue. The replies come back as follows:

- Peter [*confident speaker*]: Yes, and I'm sure the server arrived this morning.
- Adam [*unconfident speaker*]: Yes, and I think the server arrived this morning, but I'm not sure.
- Peter [*presupposing speaker*]: Yes, I'm relieved that the server arrived this morning.
- Adam [*saying speaker*]: Yes, the server arrived this morning. I'm relieved!
- Joe [*implicating speaker*]: Yes, we got a big delivery this morning. I'm relieved!

In light of this, you write the following message to the academic staff: [*type text*]

Later that day, you discover that the new server has not arrived yet and that the delivery concerned the printers you ordered last week.

Implicature question: Remember what happened. The main server broke down and your team received many complaints from the academic staff. This morning you asked your team members whether they felt less stressed about the server issue, and Joe said:

Joe: Yes, we got a big delivery this morning. I'm relieved!

Does Joe mean that the new server arrived in the morning?

Appendix B

The material of Experiment 2a was comprised of eight stories: the four story from Experiment 1 and four additional stories (see below).

Story 5

You are the new supervisor of the IT department at a large university. Last week the main server broke down. Researchers cannot store their data until the server is either repaired or replaced with a new one. Your team received many complaints from several academic units. Today, you send an email to your two assistants and ask whether there is any progress on the server issue. The replies come back as follows:

- Peter [*presupposing speaker*]: Yes, the new server arrived this morning.
- Adam [*saying speaker*]: Yes, there's a new server, it arrived this morning.

In light of this, you write the following message to the academic staff: [*type text*]

Later that day, you discover that the server is going to be repaired rather than replaced and that the department received a delivery in the morning concerning new printers.

Story 6

You are the new supervisor of the IT department at a large university. Last week the main server broke down. Researchers cannot store their data until the new server is in place. Your team received many complaints from several academic units. Today, you send an email to your two assistants and ask whether there is any progress on the server issue. The replies come back as follows:

- Peter [*presupposing speaker*]: Yes, after the server arrived this morning, I have informed the academic staff about this.
- Adam [*saying speaker*]: Yes, the server arrived this morning and I have informed the academic staff about this.

In light of this, you write the following message to the university mailing list: [*type text*]

Later that day, you discover that the new server has not arrived yet and that the department received a delivery in the morning concerning some new printers.

Story 7

You are the new supervisor of the IT department at a large university. Last week the main server broke down. Researchers cannot store their data until the new server is in place. Your team received many

complaints from several academic units. Today, you send an email to your two assistants and ask whether there is any progress on the server issue. The replies come back as follows:

- Peter [*presupposing speaker*]: No, if the new server had arrived, I would have informed the academic staff about this.
- Adam [*saying speaker*]: No, the new server hasn't arrived, so I didn't write to the academic staff.

In light of this, you write the following message to the academic staff: [*type text*]

Later that day, you discover that the new server had arrived yesterday.

Story 8

You are the new manager of an advertising company. Tomorrow you have a meeting with two important and well-known costumers, Ms. Leroy and Ms. Thomas. You are planning to give each of them a gift on behalf of the company: an ink pen to Ms. Leroy and a highly rated bottle of wine to Ms. Thomas. You send an email to two colleagues telling them about your plan and asking them whether Ms. Thomas drinks wine. The replies come back as follows:

- Peter [*presupposing speaker*]: Only Ms. Leroy drinks wine.
- Adam [*saying speaker*]: Ms. Leroy drinks wine but Ms. Thomas doesn't.

In light of this, you write the following message to your assistant: [*type text*]

The day after you decide to give the ink pen to Ms. Thomas and the bottle of wine to Ms. Leroy. After receiving your gift, Ms. Leroy tells you that she has always been allergic to wine.

The material of Experiment 2b was comprised of four stories: Story 1 from Experiment 1 (*again*) and three additional stories (see below).

Story 9

You are the new manager of a small advertising company. The company policy states that employees are not allowed to receive a weekly bonus if they are late more than once. After having been on a two-day business trip, you are back at work. You see that one of your employees, Richard, is not at his desk. So, you send an email to his two office mates and ask whether Richard is coming in. The replies come back as follows:

- Peter [*presupposing speaker*]: He is also late today.
- Adam [*saying speaker*]: He is late today, like yesterday.

In light of this, you write the following message to Richard: [*type text*]

Later on, you discover that Richard had taken yesterday morning off and that he was late just today.

Story 10

You are the new supervisor of the HR team at a large university. Last week the online application platform went down. Prospective applicants cannot send their application until the problem is fixed. Your team received many complaints and you had asked your two assistants to seek a solution. Today, you send an email to them and ask whether there is any progress about it. The replies come back as follows:

- Peter [*presupposing speaker*]: Even John couldn't solve the problem.
- Adam [*saying speaker*]: John was the most likely to solve the problem and he couldn't do it.

In light of this, you write the following message to applicants who reached out for you: [*type text*]

Later that day, you discover that John, the technical assistant of the HR group, is not an expert on maintenance of online platforms and that the head of the IT department had not been informed about the problem.

Story 11

You are the new supervisor of the HR team of a small company. There is an open position for an IT job. If you receive more than one application, you need to organize an interview panel to select among the candidates. A few days ago you have been informally contacted by a potential applicant, Mr. Smith. Today is the deadline and you send an email to your two assistants to ask whether Mr. Smith has indeed applied. The replies come back as follows:

- Peter [*presupposing speaker*]: Yes, Mr. Smith sent an application too.
- Adam [*saying speaker*]: Yes, two people have sent applications, including Mr. Smith.

In light of this, you write the following email to the colleagues who had given their availability to be part of the interview panel: [*type text*]

Later on, you discover that Mr. Smith is the only candidate. Given this, the interview panel is not needed.