

A cognitive analysis of deception without lying

Keith Ransom (keith.ransom@adelaide.edu.au)

School of Psychology, University of Adelaide

Wouter Voorspoels (wouter.voorspoels@kuleuven.be)

Brain & Cognition, University of Leuven

Amy Perfors (amy.perfors@adelaide.edu.au)

School of Psychology, University of Adelaide

Danielle J. Navarro (d.navarro@unsw.edu.au)

School of Psychology, University of New South Wales

Abstract

When the interests of interlocutors are not aligned, either party may wish to avoid truthful disclosure. A *sender* wishing to conceal the truth from a *receiver* may lie by providing false information, mislead by actively encouraging the receiver to reach a false conclusion, or simply be uninformative by providing little or no relevant information. Lying entails moral and other hazards, such as detection and its consequences, and is thus often avoided. We focus here on the latter two strategies, arguably more pernicious and prevalent, but not without their own drawbacks. We argue and show in two studies that when choosing between these options, senders consider the level of suspicion likely to be exercised on the part of the receiver and how much truth must be revealed in order to mislead. Extending Bayesian models of cooperative communication to include higher level inference regarding the helpfulness of the sender leads to insight into the strategies employed in non-cooperative contexts.

Keywords: deception; Inductive inference; communication; pragmatics

Introduction

“You can tell he’s lying because his lips are moving.”

If only detecting lies were that simple! Despite its importance, people generally perform at chance when classifying liars and truth tellers (C. Bond & DePaulo, 2006). Indeed, most verbal and nonverbal cues have only marginal diagnostic value (DePaulo et al., 2003). Instead of focusing on cues to deception, a promising new approach considers how the *cognitive processes* involved in deception may differ from telling the truth. It has been suggested, for example, that deception imposes higher cognitive demands on liars, who may find it more difficult to furnish details when interviewed (Vrij & Granhag, 2012). A good understanding of the cognitive mechanisms underlying deception, taking into account the complexities of the strategies employed, would be a tremendous asset (G. D. Bond, 2012; Blandon-Gitlin, Fenn, Masip, & Yoo, 2014).

In the present research we analyse the challenge faced by would-be deceivers seeking to conceal the truth. We begin with a brief analysis of the deceiver’s perspective, identifying the main deceptive strategies, and outline a preliminary study which illustrates people’s preferences for these strategies. We then present two studies where we ask people to conceal the

truth, manipulating the level of suspicion of the hypothetical receiver and the information content of the available message options.

How to deceive in ten steps (with pictures)

Communication relies on principles of cooperation (Grice, 1989). The intended meaning of a sender rarely coincides perfectly with the “literal” content of a message, but by making assumptions about why the sender chose that particular message, the receiver may infer the intended meaning. By assuming that a sender is cooperative and produces messages that follow the Gricean maxims (described below), a receiver can increase the speed and strength of the inferences they draw from those messages (Horn, 2004).

But what if the sender is *not* trying to be cooperative? In that case, there are three main strategies the sender can rely on, each corresponding to different violations of the Gricean maxims. Consider the following scenario: You shot your neighbour’s hamster with your shotgun while she was away for the weekend. Obviously, you’d prefer that she didn’t learn the truth. However, you were given the key to her house to take care of her pet, so you are definitely a person of interest in her investigation. How can you conceal the truth from her?

- **Lying:** You might try an outright lie: “I did not shoot your hamster”. Lying involves communicating a proposition to the receiver with the full knowledge that it is false. From a Gricean perspective, lying is a violation of the supermaxim of quality, stating that your contribution should be true.
- **Being uninformative:** It seems very sensible to be uninformative. Neighbour: “Did you shoot my hamster?” You: “Have you heard the new Justin Bieber album? It’s fantastic!”¹ With this kind of utterance, it would seem that the receiver can infer nothing beyond her prior beliefs. But this violates the Gricean maxims of relevance and quantity, and these violations can themselves be informative about the sender’s intentions even if not the actual facts of the matter.
- **Misleading:** A third option is to mislead your neighbour by implicature. “I was not at home when your hamster was shot!” You tell her a truth very relevant to the issue at hand, but from which you believe a false conclusion will be drawn (you were not at home when you shot the hamster; you took it with you to a nearby park for target practice). Misleading involves covertly violating the maxims of quantity, but may be harder to detect.

¹ Admittedly, this could well be considered a bald-faced lie.

In a preliminary study, we asked 96 first year psychology students (87 women) at the University of Leuven to imagine seven different scenarios like the following:

A man arrives home after a weekend in Vegas, during which he won \$2000 playing Poker, but lost \$500 at the slot machines and \$4000 on Blackjack. When he returns, he does not know how to tell his wife. His wife knows that he has had gambling problems in the past, but is convinced that they are resolved. Their relationship is currently strained, so the man would rather not cause any additional problems. His wife asks him directly if he has gambled. Which of these answers would you give if you were in his situation?

Participants selected a response from seven options comprised of two lies (e.g. “I didn’t gamble”), two uninformative statements (e.g. “there were a lot of people gambling”), two misleading statements (e.g. “I won \$2000 playing Poker”), and the truth. Figure 1 presents the preference of the participants (collapsed across scenarios and equivalent response options).




Figure 1: When choosing how to communicate in a variety of different scenarios with a clear motivation to deceive, participants showed a strong tendency to mislead rather than be uninformative.

Two important conclusions emerge from Figure 1. Firstly, people were uncomfortable with deception: 37% of responses involved telling the full truth and only 10% were outright lies: a surprising number perhaps given that each scenario provided a clear motivation to deceive. However, this finding is consistent with previous research showing that in general people avoid lying through concerns regarding self-image, guilt, and anxiety (Aquino & Reed, 2002). Of more relevance to our present purposes, we found that among those who chose not to tell the truth, people showed a clear preference for misleading over lying or being uninformative (37%, 10% and 15% respectively). This finding is consistent with earlier work on the topic (Montague, Navarro, Perfors, Warner, & Shafto, 2011; Rogers, Zeckhauser, Gino, Schweitzer, & Norton, 2014).

Balancing suspicion and information

So, why do people seem to prefer to actively mislead rather than be entirely uninformative? At first glance, it seems rational to be as uninformative as possible: the receiver cannot revise her beliefs on the basis of your utterances. Misleading on the other hand, involves salting your statements with a grain of truth – something which the receiver may build upon to infer the whole truth.

An important motivation for choosing a misleading utterance over a strictly uninformative one is because the latter

raises suspicion. Consider the likely response of choosing to be uninformative, as in the Las Vegas scenario:

Spouse: Did you gamble?

Gambler: Where shall we go for dinner? I’m hungry.

Spouse: You lost money didn’t you?

Gambler: Some of the guys won big.

Spouse: How much did you lose?

As Sperber et al. (2010) points out, people have a toolbox of cognitive mechanisms for *epistemic vigilance* that reduce the risk of being deceived. One such tool supports tracking the cooperation of others; as a result, obvious departures from that cooperation are noted (e.g., Mills, 2013). Responding in an uninformative way violates the principle of cooperation so blatantly that the deception is revealed.

A deceiver, sensitive to the epistemic vigilance of his counterpart may prefer instead to provide truthful but misleading utterances, but in doing so faces a delicate trade-off. Chosen well, such utterances may not only allay the receiver’s suspicion, but by virtue of the inferential boost accorded to cooperative speakers, the receiver may be led to a false conclusion, terminating the search for further information. Yet if suspicion is already raised, the receiver is unlikely to fall for the false implicature, which relies on her assumption of cooperation (Dyrel, 2011), and may use the information to get closer to the truth.

This analysis points to two opposite forces, balanced in the selection of one strategy over another. On one hand, the knowledge that the receiver may engage in inference about the helpfulness of the statement may lead the sender to opt for a misleading yet informative statement. On the other hand, if the sender considers that the receiver will be suspicious a priori, he may resort to being uninformative. We examine these factors in two experiments.

Experiment 1: The deception game

As the basis for our empirical investigation, we use a simple two person communication game in which the interests of the sender and receiver are opposed. In the game, the sender (cast in the role of a *pirate*) and the receiver (cast as an *explorer*) see four alternative maps, only one of which is genuine. Each map purports to show the true extent of a contiguous region where treasure is buried. The pirate, who knows the identity of the genuine map and must seek to protect it, is required to provide a hint to the explorer in the form of a small number of locations that lie within the region. The explorer must use this information to guess the identity of the genuine map. Both players know that lying is not allowed; so the pirate can only reveal locations where treasure is actually buried.

An example trial faced by participants, who took the role of the pirate in this experiment, is shown in Figure 2. Figure 2(a) illustrates the four maps shown to the pirate and (hypothetical) explorer for an example trial. In providing a hint to the explorer, participants were restricted in their choice to one of three kinds of hints : one uninformative, one mislead-

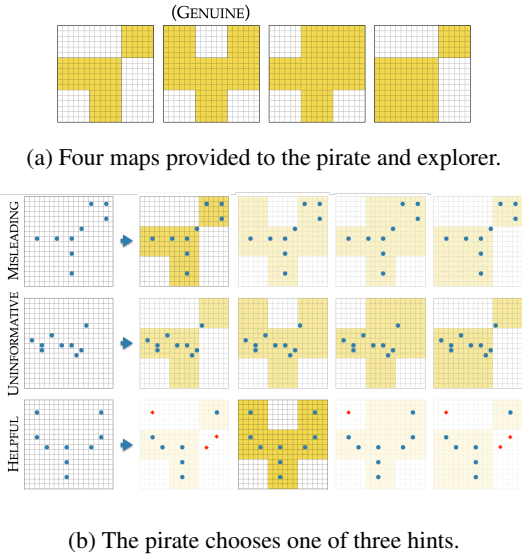


Figure 2: **Experiment 1: The deception game.** (a) Both the pirate and explorer see the same four maps; the shaded area marks the region where treasure be buried. Only one of the four maps is genuine: the pirate seeks to conceal its identity, the explorer seeks to discover it. (b) The pirate must provide a hint to the explorer, and is given three hints from which to choose: a MISLEADING, UNINFORMATIVE, or HELPFUL hint. Each hint marks a subset of locations where treasure is buried (blue dots). Since lying is not permitted all marked points lie within the shaded area when overlaid on the genuine map. The hints vary in their potential to drive inference. The MISLEADING hint appears to clearly point out the wrong map, despite being consistent with all of them. The UNINFORMATIVE hint neither points to nor excludes any of the maps. The HELPFUL hint excludes all of the maps except the genuine one.

ing, and one genuinely helpful. Figure 2(b) shows examples of the three options, and illustrates the likely effect that each hint would have on the inference of a trusting explorer.

The **informativeness** of a hint was manipulated by varying the number of treasure maps (out of four) that were excluded by the hint. When maximally informative (HELPFUL), a hint excludes all treasure maps, except the true one (bottom row of Figure 2b). When maximally UNINFORMATIVE, it excludes no treasure maps (middle row). The MISLEADING hints were designed to closely resemble one of the three false maps (top row), and to exclude none (as in the example shown), one or two of the treasure maps.

People’s beliefs about the **suspicion level** of the receiver (the hypothetical explorer) were manipulated by changing the proportion of deceivers in the population that were supposedly providing information. Participants were told that they were part of a crew of six sailors providing a hint to the explorer. Participants were also told that the explorer knew that the hint came from an unknown (but randomly selected) crew member and knew the proportion of deceptive crew members. Varying the number of deceivers in the crew from one to five was intended to raise the perceived suspicion level of the explorer. Our question was whether the pirate would track and use this information when providing hints, being more likely to mislead when the suspicion level was lower. When facing

a trusting receiver, we expect people to be more inclined to mislead, provided that the amount of information disclosed is acceptable. But if the receiver is likely to be suspicious or too much information would otherwise be revealed, we expect people to be uninformative.

Participants were 120 undergraduates from the University of Leuven (86% female, ages 18-24, median 18) participating for course credit. Participants faced 30 trials in all: six sets of four maps were presented in conjunction with each of five crew configurations.

Results and discussion

Our first question was whether people’s decision to mislead or not depended on how much information the misleading option gave away. To address this, we examined whether the proportion of participants choosing each kind of hint depended on the number of treasure maps that were excluded by the MISLEADING one. A chi-square test confirmed the dependency shown in Figure 3 ($\chi^2(4, 3600) = 93.31, p < 0.001$): as the amount of information revealed by the MISLEADING hint increased, people were less willing to select it, preferring instead to choose the UNINFORMATIVE option. Indeed, in contrast to the pilot data in Figure 1, the UNINFORMATIVE option was the most favoured in this task. This is unsurprising perhaps, since a hint that excludes even one of four maps is extremely informative.

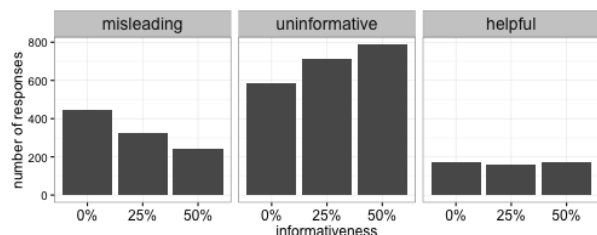


Figure 3: **Experiment 1: Informativeness.** The y-axis shows the number of times each type of hint (given by column label) was selected as a function of the number of treasure maps excluded by the misleading hint (x-axis): 0% means none were excluded, 25% means one was, and 50% means that two were. People were sensitive to informativeness: when the misleading hint was more informative, people were less likely to mislead.

Our second question was whether people’s decision was affected by their estimate of how suspicious the receiver (the hypothetical explorer) was likely to be. In light of this, we examined the relationship between participants’ choice of hints and the number of deceptive crew members providing hints for the explorer. Curiously, as Figure 4 illustrates, there was no evidence for a relation between the level of suspicion and the type of hint selected ($\chi^2(8, 3600) = 12.96, p = 0.11$). Even when people knew that the explorer thought that five out of the six possible senders was acting deceptively, they did not alter their selection of hints.




Figure 4: **Experiment 1: Suspiciousness.** The y-axis shows the number of times each type of hint (given by column label) was selected as a function of the implied suspicion level of the explorer. The numbers on the x-axis represent the number of deceptive members in the crew that provides the explorer with a hint (more members suggests that the explorer should be more suspicious). There was no dependency between choice of hint and number of deceptive crew members.

Experiment 2: Increasing suspicion

Experiment 1 found convincing evidence for the influence of informativeness on the decision to mislead. If too much information would be revealed with a true but misleading statement, people are more inclined to be uninformative. Surprisingly, we did not find an effect of suspicion. What might be going on here?

One possible explanation is simply that our manipulation was ineffective. Perhaps changing the number of deceptive crew members was not salient enough or required too much effort for participants to interpret or keep in mind.

Experiment 2 was identical to Experiment 1, but rather than have people infer how suspicious the explorer might be from the composition of the crew, we instead gave participants explicit information about the explorer’s beliefs. In the LOW SUSPICION condition, participants were told that the explorer suspected that the hint came from a teammate, whereas in the HIGH SUSPICION the explorer suspected the hint came from an opponent. The experiment was similar in all other respects except for a control condition in which participants were asked to help the explorer (used to identify participants who were not trying or did not understand the task). There were also a number of filler items in which there was no obviously misleading option.

Participants were 98 adults recruited via Amazon Mechanical Turk and paid \$1.25USD for 15 minutes participation. Data from 22 participants who failed to demonstrate a sufficient understanding of the experiment were excluded from subsequent analysis.² The remaining 76 participants were 46% female and aged 20-63 (median age 28.5). Participants faced 30 trials in all: 10 map sets (six experimental, four fillers) were presented in each of the three condition blocks.

²These exclusions were of participants who failed to select the HELPFUL message in a CONTROL condition (where the goal was to help) on at least 40% of trials. We also excluded those who selected the HELPFUL message in the LOW SUSPICION condition on 40% or more trials (where the goal was to hinder, and double bluffing was unreasonable). There was no difference in the significance of our findings if these people were included.




Figure 5: **Experiment 2: Informativeness.** The y-axis shows the number of times each type of hint (given by column label) was selected as a function of the number of treasure maps excluded by the misleading hint (x-axis). As in Experiment 1, people were sensitive to informativeness: when the misleading hint was more informative, people were less likely to mislead. Data from the control condition are excluded.




Figure 6: **Experiment 2: Suspiciousness.** The y-axis shows the number of times each type of hint (given by the column label) was selected as a function of the implied suspicion level of the explorer. The x-axis reflects whether people were told that the explorer was expecting a hint from a member of another team (the HIGH SUSPICION condition) or from a teammate (the LOW SUSPICION condition). When participants knew that the explorer was apt to be suspicious of them, they were less inclined to be misleading, opting instead to be uninformative.

Results and discussion

As before, our first question was whether people were sensitive to informativeness when choosing which hint to provide. Once again, there was a significant effect of informativeness ($\chi^2(4, 912) = 18.04, p = 0.001$). As Figure 5 shows, the more maps the misleading option excluded, the less inclined people were to select it, favouring instead the uninformative option.

In light of the null effect in Experiment 1, a perhaps more interesting question is whether people were sensitive to the suspicion level of the explorer when deciding what to tell them. As Figure 6 reflects, when the suspicion level of the explorer is made more obvious, people are indeed sensitive to it. Although the UNINFORMATIVE hint was still the most popular overall, the MISLEADING option was chosen far more when the explorer was expecting a hint from a trusted teammate ($\chi^2(2, 912) = 85.95, p < 0.001$). This suggests that people acting as senders are indeed attentive to the level of trust presumed by the receiver; although, taken together with the results from Experiment 1, tracking suspicion may be too cognitively challenging where it is not especially salient.

Towards a computational model

Experiments 1 and 2 manipulated two important factors: the information content of the messages that deceivers could

choose and their beliefs about the degree of suspicion with which their messages would be received. Taken together, our results show that both factors were important considerations. In this section we present a computational model whose goal is to aid our understanding of these results and generate new testable predictions. While an in-depth analysis of the model is beyond the scope of this paper, here we briefly describe the relevant features.

A convenient starting point for a model of the deception game employed here – and for communication in general – is rational inference (e.g., Goodman & Frank, 2016). In this framework, a receiver faces the challenge of updating her beliefs on the basis of information disclosed by a sender. The sender, for his part, selects information designed (according to his goal) to help or hinder the receiver in her efforts.

We first evaluate things from the perspective of the **receiver**, who is confronted with a hint x (or, more generally, an utterance). The receiver is assumed to update her beliefs h according to:

$$P_{\text{RECEIVER}}(h|x) \propto \sum_{s \in \mathcal{S}} P_{\text{SENDER}}(x|h,s)P(h)P(s) \quad (1)$$

where s represents a sampling strategy employed by the sender and \mathcal{S} represents the range of such strategies considered. As a simplifying assumption, we assume that the receiver considers the sender’s sampling strategy to be independent of the true hypothesis.

This inference thus depends on the **sender**, who selects information according to a sampling strategy:

$$P_{\text{SENDER}}(x|h) \propto (P_{\text{RECEIVER}}(h|x))^\alpha \quad (2)$$

where α reflects the goals of the sender, and $P_{\text{RECEIVER}}(h|x)$ the sender’s assumptions about how the receiver updates her beliefs. A sender who wishes to reveal the truth to the receiver (i.e., to increase the receiver’s posterior probability for the correct hypothesis h) will have an α with a positive value; one who wishes to conceal the truth has a negative α ; one who behaves somewhat randomly has an $\alpha = 0$. There are other ways to capture conflicting goals, like assigning separate utility functions for the sender and receiver with regard to truth-predicated action, but we chose this for its relative simplicity.

To capture the patterns observed in our deception game, both equations have to be considered simultaneously. That is, both sender and receiver must recursively consider the assumptions and strategies used by the other party. Importantly, from the receiver’s perspective the inferential potential of a message depends not only on the information as such, but also on the “sampling strategy” of the sender, which reflects the sender’s goals and assumptions about the receiver. For example, sampling procedures that follow the principle of cooperation and the Gricean maxims have a stronger inferential potential (e.g., Bergstrom, Moehlmann, & Boyer, 2006; Shafto, Goodman, & Griffiths, 2014; Voorspoels, Navarro, Perfors, Ransom, & Storms, 2015). Crucially, the receiver

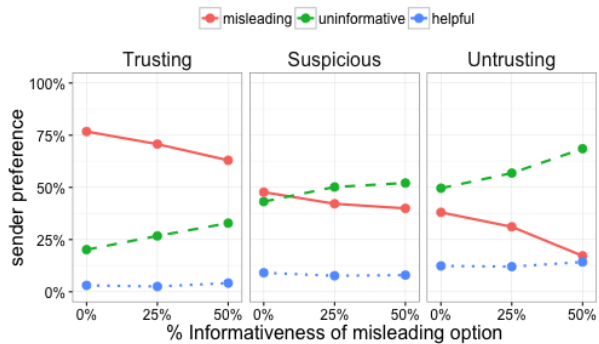


Figure 7: **Model predictions for sender actions.** Model predictions for the preference of a sender in the deception game for the misleading, uninformative and helpful message options. From left to right, the panels present scenarios with increasingly suspicious receivers (modelled through different kinds of inference about the sender’s goals and assumptions). The x -axes indicate how informative the misleading option is (in terms of the proportion of hypotheses excluded by it). The model predicts a decrease in preference for the misleading option as it becomes increasingly informative, as well as an increase in preference for misleading when the receiver is less suspicious.

not only updates her beliefs about what is true, but simultaneously makes inferences about the sender’s sampling strategy: learning whether the sender is helpful and knowledgeable play a critical role in epistemic vigilance, and has a substantial impact on how rational agents reason (e.g., Shafto, Eaves, Navarro, & Perfors, 2012).

In Equations (1) and (2), the universe of sampling strategies \mathcal{S} evaluated by the receiver is defined in terms of two things that she presumes about the sender: what does the sender assume about her (reflected by $P_{\text{RECEIVER}}(h|x)$), and what are his goals (reflected in α). Many scenarios may be modelled in this way, but here we consider three. If the receiver is TRUSTING, this means that she is performing inference over two possibilities: either the sender is trying to be helpful ($\alpha = 1$), or he is inattentive and thus not selecting information with care ($\alpha = 0$). If the receiver is UNTRUSTING, this means that she believes that the sender is trying to conceal the truth from her ($\alpha = -1$), under the mistaken assumption that he is trusted. Lastly, if the receiver is SUSPICIOUS, this means that she is performing inference over whether to be trusting or untrusting.

How well does this approach capture the main qualitative patterns in the deception game? To answer this, we simulate outcomes for the three scenarios we have outlined. In the leftmost panel of Figure 7, the receiver trusts the sender, but is not sure how attentive he is: he may be acting helpfully ($\alpha = 1$) or he may be providing poor but not actively misleading data, perhaps due to lack of motivation, attention, or information ($\alpha = 0$).

If the receiver updates her beliefs (concerning the true treasure map) at the same time as her assumptions about the helpfulness of the information received (α) then there is reason for the sender to choose a hint that seems informative. That

the message appears informative supports the receiver's assumption that it has been carefully selected, which further fuels inference. This recursive process may lead the receiver to draw a misleading conclusion if the sender is not actually helpful (as in our experiments). However, as the information content of the hint increases, so too does the risk that the receiver will inadvertently arrive at the truth. Consequently, the model captures the fact that the sender's preference for the misleading hint declines with its information content.

In the rightmost panel of Figure 7, in contrast, the receiver is certain that the sender is not to be trusted. If the sender is aware of this, there is little to be gained by attempting to mislead, and so the uninformative hint is preferred. The extent of this preference is, once again, moderated by the information content of the misleading option.

In many situations, a receiver will not be predisposed to regard the sender with complete trust, nor complete distrust, but rather will remain open to either possibility. We model this case by assuming that the receiver is performing inference about whether the sender should be trusted ($\alpha = 1$ or 0) or not ($\alpha = -1$). The preferences of an antagonistic sender facing a suspicious receiver are shown in the center panel of Figure 7. The two conflicting forces are most pronounced here, dividing the sender's preference between the two strategies. On the one hand, the sender may convince the receiver that he is actually trying to help by appearing informative, yet the (real) information can be used by the receiver to rule out previously plausible (but false) hypotheses.

Overall, there are two clear patterns that were found in our experiments and were also predicted by our model. Firstly, as the information content of the misleading option increases, there is an increasing preference for choosing the uninformative hint. Secondly, the more trusting the receiver is assumed to be, the more popular the misleading option becomes. This pattern of results is consistent with the idea that people may be performing some kind of recursive inference over how suspicious their interlocutor is when deciding how to deceive. Furthermore, our results are consistent with the notion posited here and elsewhere (e.g., Goodman & Frank, 2016; Shafto et al., 2012), that receivers (from the sender's perspective at least) perform joint inference over the goals of the sender and the truth of the matter at hand given the information received.

Conclusion

"There is nothing more deceptive than an obvious fact."
— Arthur Conan Doyle

In two studies we have demonstrated that people's preference for a deceptive strategy hinges on their assumption of whether cooperative norms are expected to apply. In situations where high levels of trust and cooperation are warranted, deceivers are more inclined to actively mislead than to simply withhold information. In this scenario, the deceiver seeks to leverage the inferential boost of cooperative communication. In contrast, when the deceiver believes the false implication will not

be inferred — when the receiver is already suspicious — then preference shifts towards limiting the information disclosed.

Acknowledgments

We thank Paulien Van Rijmenant for help with data collection. KR was supported by an Australian Government Research Training Program Scholarship. WV is a postdoc at KU Leuven. AP was supported by ARC grants DP110104949 and DP150103280. DJN received salary support from ARC grant FT110100431.

References

- Aquino, K., & Reed, I. (2002). The self-importance of moral identity. *Journal of Personality and Social Psychology*, 83(6), 1423-1440.
- Bergstrom, B., Moehlmann, B., & Boyer, P. (2006). Extending the testimony problem: evaluating the truth, scope and source of cultural information. *Child Development*, 77(3), 531-538.
- Blandon-Gitlin, I., Fenn, E., Masip, J., & Yoo, A. (2014). Cognitive-load approaches to detect deception: searching for cognitive mechanisms. *Trends in Cognitive Science*, 18(9), 441-444.
- Bond, C., & DePaulo, B. (2006). Accuracy of deception judgments. *Personality and Social Psychology Review*, 10(3), 214-234.
- Bond, G. D. (2012). Focus on basic cognitive mechanisms and strategies in deception research (and remand custody of 'wizards' to Harry Potter movies). *Journal of Applied Research in Memory and Cognition*, 1(2), 128 - 130.
- DePaulo, B., Lindsay, J., Malone, B., Muhlenbruck, L., Charlton, K., & Cooper, H. (2003). Cues to deception. *Psychological Bulletin*, 129(1), 74-118.
- Dynel, M. (2011). A web of deceit: A neo-gricean view on types of verbal deception. *International Review of Pragmatics*, 3, 139-167.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11), 818 - 829.
- Grice, H. P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Horn, L. (2004). Implicature. In L. Horn & G. Ward (Eds.), *Handbook of pragmatics* (p. 2-28). Blackwell Publishing.
- Mills, C. M. (2013). Knowing when to doubt: Developing a critical stance when learning from others. *Developmental Psychology*, 49(3), 404-418.
- Montague, R., Navarro, D., Perfors, A., Warner, R., & Shafto, P. (2011). To catch a liar: The effects of truthful and deceptive testimony on inferential learning. In L. Carlson, C. Holscher, & T. Shipley (Eds.), *Proceedings of the 33th annual meeting of the cognitive science society*. Austin, TX: Cognitive Science Society.
- Rogers, T., Zeckhauser, R., Gino, F., Schweitzer, M., & Norton, M. (2014, september). *Artful paltering: The risks and rewards of using truthful statements to mislead others* (HKS Working Paper No. RWP14-045). Harvard Kennedy School. Retrieved from <https://ssrn.com/abstract=2528625> doi: <http://dx.doi.org/10.2139/ssrn.2528625>
- Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science*, 15, 436-447.
- Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by and learning from examples. *Cognitive Psychology*, 71, 55-89.
- Sperber, D., Clement, F., Heintz, C., Mascaro, O., Mercier, H., Origg, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*.
- Voorspoels, W., Navarro, D. J., Perfors, A., Ransom, K., & Storms, G. (2015). How do people learn from negative evidence? non-monotonic generalizations and sampling assumptions in inductive reasoning. *Cognitive Psychology*, 81, 1-25.
- Vrij, A., & Granhag, (2012). Eliciting cues to deception and truth: What matters are the questions asked. *Journal of Applied Research in Memory and Cognition*, 1(2), 110-117.