


MISUNDERSTANDINGS FALSE BELIEFS IN COMMUNICATION

GEORG WEIZSÄCKER



<https://www.openbookpublishers.com/>

©2023 Georg Weizsäcker

This work is licensed under an Attribution-ShareAlike 4.0 International (CC BY-SA 4.0). This license allows users to share, copy and redistribute the material in any medium or format; to adapt, remix, transform, and build upon the material for any purpose, even commercially, and only so long as attribution is given to the creator. Attribution should include the following information:

Georg Weizsäcker, Misunderstandings: False Beliefs in Communication.
Cambridge, UK: Open Book Publishers, 2023,
<https://doi.org/10.11647/OBP.0367>

In order to access detailed and updated information on the license, please visit
<https://doi.org/10.11647/OBP.0367#copyright>

Further details about the CC BY-SA 4.0 license are available at
<https://creativecommons.org/licenses/by-sa/4.0/>

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at
<https://archive.org/web>

Digital material and resources associated with this volume are available at
<https://doi.org/10.11647/OBP.0367#resources>

ISBN Paperback: 978-1-80064-881-4
ISBN Hardback: 978-1-78374-051-2
ISBN Digital (PDF): 978-1-80511-138-2
DOI: 10.11647/OBP.0367

Cover: Thomas Gainsborough, Conversation in a Park (1746), Louvre Museum,
[https://commons.wikimedia.org/wiki/File:Thomas_Gainsborough_-_Conversation_in_a_Park_\(1746\).jpg](https://commons.wikimedia.org/wiki/File:Thomas_Gainsborough_-_Conversation_in_a_Park_(1746).jpg)
Cover design: Jeevanjot Kaur Nagpal

The publication of this work was supported by the Open Access Publication Fund of Humboldt-Universität zu Berlin, funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 491192747. The publication of this work was further supported by the CRC/Transregio 'Rationality and Competition: The Economic Performance of Individuals and Firms' (TRR 190), funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 280092119.



Chapter 4

Talking

Now, at last, the actual conversation. This chapter describes how our beliefs may be off target when we talk. The difference to the previous chapter is that we now focus on their *reaction* to our statement.

Our beliefs about their reaction are, conveniently, the only bit of our beliefs that depends on what we say: there are three uncertain aspects (actions, issues, types) but we already know our type, we control our actions, and we cannot influence the state of the world or their type. Our success as talker depends, entirely, on how well we predict their reaction.

A widely held view of communication is that language is a convention. They and we have arrived, through our previous experiences or perhaps through some type of agreement, at a use of signals that allows information to be conveyed.

The key of such conventions is common interest: they and we *want* to communicate because it makes our utilities increase in unison. The talker wants the listener to learn something. Having learnt it, the listener does something that is good for both.

The common interest also fixes the beliefs about language use. The talker expects the listener to stick to the agreed language because doing so lies in the listener's interest. The listener expects the talker to do the same. The convention reinforces itself.

Or does it? The reasoning certainly opens a number of questions. What *is* the interest of our partner in conversation, and is it the same as ours? Are our beliefs really conventional, in the above sense of the word, and do

we agree which convention, out of many possible ones, we use? Do we really know what signals to use, and are these signals suitable for the present purpose? Is our language not far richer than what could possibly be agreed upon?

The book proceeds very slowly and seeks answers in small steps. It does not take any convention as given. Rather, it asks about the individual beliefs that, taken together, may or may not form a convention. As we will see, the interplay of beliefs is indeed far richer than what a convention can achieve.

Question 4: Do we predict a too-small reaction?

Even small steps are complex: notice that our talking is not only what we actually say in the circumstance that we find ourselves in, but also what we would have said in other circumstances. Our talking is a contingent plan. It depends on our type, it depends on what we know about the state of the world, and it depends on what we know about their, the other person's, type.

Our talking is informative because of this property: what we say in a given circumstance would not have been said in other circumstances.

Moreover, not only do we say something different in every circumstance but we could also have said many different things in the same circumstance. Checking whether or not our talking is optimal means checking for each circumstance whether or not the statement that we make gives us a higher utility than everything else that could be said.

The assumptions of Chapter 2 allow translating this into an empirical question in terms of our beliefs. Are they accurate?

We already observed that the other person's reaction is the only uncertain aspect that is moved by our talking. We therefore characterize our talking by our belief about how they react to our statements, in any given circumstance that we may be in.

Using the formal notation, the relevant belief of person i is $P_{a^j}^i(\cdot|\cdot)$. The first placeholder \cdot holds the place for the possible values of person j 's reaction a^j and the second \cdot holds the place for the things that person i conditions

on: her statement a^i and her circumstance, which is summarized by the information sets $I_{\theta^j}^i$ and I_{ω}^i . (Knowledge of her own type, θ^i , is of course also part of person i 's circumstance. The book leaves this out of the discussion, for brevity.)

We consider a whole collection of such conditional beliefs; one for each statement and circumstance. That is, the collection has one belief about a^j for each possible value of $(a^i, I_{\theta^j}^i, I_{\omega}^i)$. This collection is called *person i 's talking belief*.

(Later chapters will re-name it person i 's *first-order* talking belief, but for now the brief version is fine.)

The question of accuracy, then, asks about the predictive power of our talking belief: is it close to the true probability distribution of their reaction, for each of our possible statements and circumstances? Accordingly, this chapter asks three directed questions about our talking belief, accounting for the fact that the belief changes with our statement, our information about their type, and our information about the state of the world.

The first question, Question 4, is the simplest and asks whether our beliefs respond too little to the possible statements that we could make: do we believe that they react less strongly than they actually do?

The opposite effect may be empirically true, too. Instead of under-estimating, we may over-estimate how strongly they react to our statement. They may decide not to answer us at all. We may mumble. They may not identify the words, or not have the capacity to process them. More generally, they and/or we may suffer from constraints in the transmission and if we tend to be unaware of this, then we may over-estimate their reaction. However, in parallel with the other questions in the book, we ask about a possible under-differentiation:

$$\text{Is } P_{a^j}^i(\cdot|a^i) \text{ too close to } P_{a^j}^i(\cdot)?$$

The action that Steve cares about is whether or not Ralph acts aggressively. Steve is quite pessimistic about this but expects that an offer to play ball makes it less likely that Ralph would express anger. Indeed, the push by Ralph, while unfriendly, is a good outcome for Steve. Ralph behaves in

a way that one may call normal, for his standards. In contrast, a statement that enquires about Ralph's state of despair would plausibly have backfired. Steve's belief is accurate in predicting the direction in which his statement influences Ralph, perhaps even underestimating the size of the effect.

For the purpose of measurement, let us examine talking beliefs more closely. $P_{aj}^i(\cdot|a^i)$ is a probability distribution (as subjectively expected by person i) over the actions that person j may pick in response to a^i . It is a different distribution for each possible a^i , and can be measured by asking simple questions. ("How likely is it that they react by choosing A if you say a^i ?" – and similar questions for the other possible reactions of person j .)

These beliefs also depend on person i 's circumstance ($I_{\theta j}^i$ and I_{ω}^i) but for now we can take the circumstance as given. If the experimenter wants to measure only the believed reaction to a^i but not its interaction with $I_{\theta j}^i$ and I_{ω}^i , then she can simply ignore them.

Parallel to the discussion in the previous chapter, a question arises: how can one measure the unconditional belief $P_{aj}^i(\cdot)$?

The answer, also parallel to the discussion in the previous chapter, is that it is best to measure this unconditional belief as a weighted average of conditional beliefs. Specifically, the average of $P_{aj}^i(\cdot|a^i)$, averaging over different statements a^i .

(This average belief is potentially very different from what we would expect person j to do in the absence of a statement. The very act of talking may make person j change their behavior and person i may anticipate this.)

So the experimenter aims to measure a weighted average of the collection $\{P_{aj}^i(\cdot|a^i)\}_{a^i}$ – but with what weights? The problem is that most experimental studies, even if they fix the set of *possible* statements, do not impose the choice of a^i ; it is chosen by person i . This renders it unclear how to weight the possible values of a^i in the aggregation.

Here in Chapter 4, we address this complication only for the simplified, binary case there are only two possible statements, $a^i \in \{A, B\}$, like in our discussions of Rachel, Dimitri and Steve. (Chapter 5 will discuss the more general case.) In the binary case, the weights of the possible values of a^i do not matter for answering Question 4, at least not for answering it in a qualitative way: any weighted average lies between the two conditional beliefs

$P_{aj}^i(\cdot|A)$ and $P_{aj}^i(\cdot|B)$, and the same is true for the empirical choice frequencies of person j that these beliefs refer to. The question can be thus answered by simply asking whether $P_{aj}^i(\cdot|A)$ and $P_{aj}^i(\cdot|B)$ are closer to each other than the corresponding choice frequencies of person j .

In any case, $P_{aj}^i(\cdot|a^i)$ needs to be measured for each participant in the role of player i and each statement a^i , independent of whether or not the participant makes the statement. One can then piece the measurements together and compare them to the empirical reactions of person j , yielding the answer to Question 4.

A key feature of such an experiment is to have well-defined action spaces: the set of possible statements and the set of possible reactions. If both sets are small, a participant in the role of person i can report the entire family of talking beliefs.

The experimenter may also want to control the incentives of the participants, e.g., by paying money for given combinations of statements, reactions, and states of the world. This allows deriving predictions under standard assumptions about the participants' preferences. In particular, the experimenter can determine the game-theoretic equilibria as benchmarks for the analysis. Doing so may well be useful for the understanding of the data patterns – not least because the previously mentioned “conventions” are often described as game-theoretic equilibria, and the experimenter can use the belief data to test the equilibrium hypothesis.

This re-emphasizes, in passing, one of the book's main points: while the equilibrium requires all beliefs to be accurate, the book describes a more general approach where beliefs can be off target. The reader will see that many of the book's concepts are suitable for many possible beliefs.

A final note on conventions: different conventions are differently informative. For example, one can always agree to be entirely uninformative. After such an agreement, no-one will deviate: the listener does not react differently to different things that the talker may say. The talker knows this and may just as well say anything, with no relation to the circumstance. A self-fulfilling prophecy.

Game theorists call this uninformative convention a “babbling equilibrium”. It is an extreme case. Many similar conventions, perhaps with small degrees of informativeness, may also apply.

But talking beliefs are relevant in every possible set-up: even if the situation is not so well-described, or if the theory is incomplete, or if its assumptions are invalid. Question 4 does not require an equilibrium analysis. It asks directly about the talking belief, wherever it may come from.

Sheremeta and Shields (2013) conduct an experiment that measures the talking belief in a minimal version of a salesman pitch. The talker knows the quality of a possible investment (good/bad) and indicates the quality to a listener who then decides about investing. The investment is binary and thus the situation is a simple sequential 2x2 game: the talker announces ‘‘good’’ or ‘‘bad’’ quality, the listener invests or not. If the quality is indeed good and the listener invests, she earns money in the experiment. If the quality is bad and she invests, she loses money. The talker, in contrast, earns money if and only if the listener invests, independent of quality. This creates an incentive to announce ‘‘good’’ quality even if quality is actually bad -- the non-alignment of incentives reduces the effective informativeness of the statements. (In this game, standard game theory predicts that only babbling equilibria exist.) The experiment also includes a belief elicitation task where the talker reports his talking belief: the probability with which he expects the listener to invest for each of the two possible statements (‘‘good’’/‘‘bad’’). The data show that the talking beliefs are quite accurate in this game, and that messages are viewed as informative. The averages of the talker’s stated beliefs about whether or not the listener invests after hearing the ‘‘good’’/‘‘bad’’ message are 69% and 18%, whereas the actual investment rates in these cases are 67% and 15%, respectively. While the talkers’ stated beliefs are spot-on, the question arises whether their quality announcements are consistent with the belief statements: for these beliefs, it is money-maximizing to always announce ‘‘good’’ quality. Do talkers’ announcements show such a pattern? The answer is, only a mild majority of talkers show it: 60% of talkers are willing to lie, i.e. announce a ‘‘good’’ quality also in the case that the actual quality is bad. The remaining 40% of talkers are honest about a bad quality if it arises. A simple possible explanation is lying aversion, in the sense that the talker prefers to not announce ‘‘good’’ quality if this announcement does not match the truth. However, the effect of lying aversion would have to be

quite strong, perhaps implausibly strong, in order to explain the highly frequent truth-telling: the elicited beliefs show that participants expect honest behavior to be quite costly, shaving off a substantial part of one's earnings. A complementary explanation, for which Sheremeta and Shields (2013) find evidence with additional analyses, is that the talkers are "ahead averse" (on top of being lying averse) in the sense of having a preference against earning money at the expense of the listeners. Both explanations are within the scope of assumptions in Chapter 2, corresponding to specific utility functions u_i . Overall, this discussion illustrates both caveats and advantages of measuring beliefs: the elicitation of talking beliefs may not easily result in a definitive explanation for observed talking behavior, but it can help quantifying competing hypotheses.

Question 5: Do we underappreciate different people's differences in reacting?

Refining the previous question, we ask how our talking belief varies for different listeners. We should use our information about their type, as it is indicative of their reaction. If we were in their shoes, our reaction to a statement would clearly be person specific, wouldn't it? Only we, or someone very much like us, would react like we do; or so it seems to us. Well, we are not them and we therefore have to anticipate how they, in their own shoes, may react.

Doing so, we face an informational asymmetry. They know much more about themselves than we do. But we know at least something, on the basis of which we can discriminate (in the belief sense of the word).

A sophisticated belief considers what their different types *could be* and how each of them would react. The details of this conditional thinking depends on the conversation. In some conversations, we care only about how the average person reacts to what we say. In other conversations, we care about the entire set of reactions by different types.

This leads to the "art of questioning". We may want them to reveal their type. We may want them to say something that is specific to their prefer-

ences and restrictions. We may want to get them to talk.

It can be a good idea to ask an *open* question, i.e., one with a wide set of possible answers. Communication scholars have found that listeners often respond to such questions with relatively high frankness and accuracy. Another trick to find out about their type is to ask an *affective* question, i.e. one that addresses their emotions or personality traits directly.

Connecting the idea of open questions to the assumptions of Chapter 2, we note that the restriction of their possible answers is not a technical one. In most conversations, our statement does not literally make their set of possible reactions wider or narrower. (They are free to say whatever they want.) Rather, our statement influences the utility that they get from their different possible reactions. It thereby rules out, or rules in, some of their possible reactions.

A question that is open is, then, one that makes many of their possible reactions potentially beneficial for them. In this sense, openness is close to politeness. Using an expression from the linguistic literature, our statement is polite if it does not threaten their *negative face*: it does not diminish their freedom to make a decision at will.

All of this applies not only when we ask a question. It applies to any statement that we may make. We want to influence their possible reactions. How they react depends on our statement but also on their type. All the more reason for us to anticipate their type-specific reaction well.

Is $P_{aj}^i(\cdot|a^i, I_{\theta j}^i)$ too close to $P_{aj}^i(\cdot|a^i)$?

Dimitri should have given Agnieszka's type more thought. What if she dislikes engaging in conflict? Dimitri knows her too little to predict her next steps. In particular, it is natural to suspect that she may not react positively to his talk about future "blood". Moreover, Dimitri is impolite in that he pushes the question of the team's continuation. He implicitly pressures her to make a statement about it. Her inclination to continue the teamwork may have been more positive if he had simply said nothing.

Open questions – or other prompts for the other person to speak unrestrictedly – are difficult to deal with in an empirical analysis, exactly because

of the lack of restrictions. The talker cannot report beliefs about the listener's possible reactions if the listener has too many of them available.

Instead of asking for talking beliefs directly, one may infer them indirectly, from the statements that a talker makes. The empirical strategy of the experimenter may be to ask, "What beliefs would justify these statements?"

This is, of course, precisely the empirical strategy that the book pursues all along. The question arises, then, whether we can pursue this question equally well, or perhaps even better, if we do not have belief reports but instead rely on indirect inference from the statements.

Such an inference certainly requires assumptions about the connection between statements and beliefs. But this alone makes no big difference: recall Chapter 2, arguing that we need to make such assumptions even if we measure beliefs directly by asking for reports (we may need pretty much the same assumptions for both methods, actually).

What else, then, speaks to the methodological question whether it is better to infer beliefs or to ask for belief reports?

On the one hand, asking for belief reports can be costly and intrusive. Costly because their measurement requires an elaborate survey. Intrusive because the measurement itself may change the observation – it imposes a machinery on the conversation. The interlocutors, when facing the task of belief elicitation, are not in their natural habitats any more. This may affect their way of communicating, or it may affect their beliefs.

On the other hand, as described, a disadvantage of inferring beliefs is that the action spaces are often too large. Analyses of free-flowing conversations are likely to be underdetermined. Too many things could be said by the talker, and for each of them, too many possible expectations about the listener's reactions can provide possible reasons for what is said. The analyst has too many degrees of freedom.

As a feasible way out of the problem of having too-large action spaces, one can artificially limit the scope of the conversation *ex post*, by classifying the available statements into a small set of categories. An *ex-post* text analysis can thereby transform the statements into a suitable data set that describes the talker's choice set in a simplified way. It is unlikely, however, that the interlocutors converse with the same classification in mind. The

method therefore carries the baggage of additional, and unrealistic, assumptions.

Overall, from the experimenter's perspective, allowing a free-flowing conversation and inferring underlying talking beliefs is a possible option, but an imperfect one.

Let us ask this methodological question differently: if we ask for a belief report, is this really the belief that we are interested in? Are the 18 questions of this book not rather aimed at a set of latent beliefs that underlie the choices and statements – what economists call *revealed expectations*?

The answer is that we are indeed interested in these latent beliefs, first and foremost. The elicitation of belief reports only gives an incomplete proxy for latent beliefs. It may or may not be the best available proxy. It, too, is a possible-but-imperfect option. The analyst, when choosing the method, should weigh the pros and cons. This book, aiming for precision in the experimental set-up, makes a choice by largely focusing on the technique of direct elicitation. But the choice is not meant to be definitive about what is the best method.

A final remark on method, especially for the purpose of answering Question 5: one also needs measures of the dependence on $I_{\theta j}^i$. This requires observations for different sets of listeners. As a variation of the methods in Chapter 3, one can describe several hypothetical vignettes to the talker, each of which ask her to imagine a different identity of the listener. The experimenter can then measure whether the (elicited or inferred) talking beliefs react to the information about the listener.

Partial and indirect evidence on Question 5 appears in the empirical literature on politeness. Holtgraves and Yang (1990) take a middle ground in terms of restricting statements: they pre-fix a set of statements and ask the participants to assess the statements in terms of how likely they would use each of them in a given situation. These assessment have a fairly clear interpretation in terms of talking beliefs: judging a statement to be likely to be used indicates that one also judges the statement as fulfilling the goals of the talker's aim better than a less-likely-to-be-used statement. (The reader may dispute this assumption, but let us take it as given.) Holtgraves and Yang's study also lets the identity of the listener vary, by describing the situation as a

pair of vignettes with two different co-workers: one in which the talker and the listener are close to each other, here named Treatment CLOSE, versus one where talker and listener have a more distant relationship (Treatment DISTANT). Moreover, an important part of the research design is that the statements reflect different levels of politeness, ranging from bold (e.g., ‘‘Go get the mail’’) to a more polite, indirect request (e.g. ‘‘Has the mail arrived yet?’’). The results of the experiment show that impolite requests are believed to be significantly more likely to be chosen in CLOSE than in DISTANT, and in the opposite order for polite requests (but insignificantly so). On a scale from 1 to 7, the participants assess an impolite request to be chosen with an average likelihood score of 2.62 in CLOSE, versus 2.38 in DISTANT; a polite request is judged to have likelihood scores of 4.43 in CLOSE and 4.54 in DISTANT. This indicates that talking beliefs differ depending on the information about the listener. Whether the differentiation is sufficiently large, given how listeners actually behave, cannot be judged from this experiment.

Question 6: Do we forget about the world when predicting their reaction?

As indicated in Chapter 2, talking is not only about information. When we talk, we perform. Our statement is an action and it changes the world (a little bit) and certainly our utility.

Linguists and philosophers of language differentiate types of *speech acts*. For instance, a statement about the world is an assertive speech act. Such an assertive speech act is usually not meant to change the world directly. It only conveys information.

A direct influence on the world can be achieved through a declarative speech act like the governor’s utterance that “here and now, we have to move on”, which ends the conversation. A declarative speech act, if made appropriately, changes the world only by virtue of being made.

So what is the state of the world, and what is it not? In Chapter 2, ω was introduced as the collection of things that the people in the conversation cannot influence. But person i ’s statement a^i is chosen by her. How is all

this consistent with the previous paragraph, which asserts that a statement can influence the state of the world?

The resolution of this little puzzle is that ω is unaffected by i 's *present* statement. For any subsequent statement, in contrast, the present statement will already belong to the description of the state of the world, i.e., it will be part of the new ω .

Our analysis, thus, does not differentiate between assertive speech acts and declarative speech acts – in either case a^i cannot affect the state ω that is relevant for the present statement but it does affect the state that is relevant for the next step in the conversation. (This is just a technical observation about this book's analysis. The reader may nevertheless find it helpful to differentiate the different kinds of acts.)

All statements share the property that when we make them, we need to anticipate the other person's reaction. To do this well, we use our information about the context, I_ω^i . Has something relevant been said at an earlier stage? Are they, the other person, in a position to hear and process our statement? Did we make any other experience or observation that helps to predict their reaction?

This shows that the previously-mentioned “conventions” are misleading, or incomplete at best: no convention can cover all contexts. The talker knows this and anticipates a (non-conventional) context-dependent reaction.

Later chapters of the book will focus on how the listener's interpretation, and the talker's anticipation of this interpretation, may incorporate context information. But even now, we can already observe one important element of the reasoning: the listener's incentives to react change with ω , and the talker knows it.

More precisely, when person i talks, she knows that person j 's utility reacts to ω and that his information about it will therefore influence his reaction. Given that i 's information about ω is correlated with j 's, i can use I_ω^i in her prediction about how j will react. The challenge is to do this in the right direction and to the right extent.

The discussion leads to the definition of person i 's *perceived relevance* of information I_ω^i . This is i 's expectation about the utility increase that she herself obtains from taking I_ω^i into account in her own talking.

(For precision, the sentences in these parentheses express the definition more formally: Fix i 's belief about the reaction of j , $P_{aj}^i(\cdot|\cdot, I_\omega^i)$, and first consider the statement that she finds optimal if using this belief – call it \bar{a}^i . Now consider what statement she would find optimal if she did not have the information, i.e. if her belief was $P_{aj}^i(\cdot|\cdot)$, and call this statement $\bar{\bar{a}}^i$. This alternative action will lead to a different utility level. If person i compares the utilities from using the two statements, she will evaluate them with her actual belief in mind, $P_{aj}^i(\cdot|\cdot, I_\omega^i)$. The perceived relevance of the information is, therefore, the difference in indirect utility that she expects to get from making statement \bar{a}^i instead of $\bar{\bar{a}}^i$, using belief $P_{aj}^i(\cdot|\cdot, I_\omega^i)$ in these expectations.)

An information is, thus, perceived relevant for person i if she thinks that she increases her utility by paying attention to it and making her statement accordingly. For later use, notice how this definition uses person i 's belief about person j 's reaction – she expects that he will react in a particular way – but that the definition does not use person j 's information or his beliefs.

Notice also that this definition uses person i 's *subjective* beliefs, when evaluating the utility difference that relates to an information about ω . Her belief P^i may, of course, be wrong in many ways – the relevance of I_ω^i is only perceived. But it is straightforward to also define an objective counterpart: one may consider how much her utility actually increases from paying optimal attention to the information, anticipating the true distribution of j 's behavior.

One may call this utility difference the objective relevance of the information and, under this nomenclature, Question 6 asks whether the perceived relevance is close to the objective relevance of a given information. However, while these are all useful concepts, they are not easily measurable: utility differences cannot be elicited without a large machinery. It is easier, and consistent with the general notation in the book, if we formulate the question in terms of beliefs:

$$\text{Is } P_{aj}^i(\cdot|a^i, I_\omega^i) \text{ too close to } P_{aj}^i(\cdot|a^i)?$$

Rachel is surprised by the fact that the governor does not make even a half-decent commitment to follow up on her request for support. In her view,

she puts forward a convincing statement about the state of the world, and he sits still. Rachel's misunderstanding is that while she finds the scientific state of knowledge about provenance decisive for what is the right course of action, the governor does not generally follow the science. For him, it is a political issue, not a scientific one. The relevance of the scientific truth of the findings is not objective, for this conversation.

To answer this question in an experiment, it is usually best to assign the relevant context information exogenously. (It avoids selection.) The experimenter may, for instance, describe the entire information structure to the participants: that there exist certain states of the world, each of which presents different incentives to the interlocutors, and how the uncertainty about the state of the world is resolved.

In particular, the participants acting in the role of the talker should not only observe their own signals about the state of the world, but they should also learn about how the information structure creates signals for the listener, i.e., what the listener may learn.

The ‘Dictator Game’ has only one active player, the dictator, who chooses the size of a donation that she makes to the other player, the recipient. That's it -- the game is rather simple and this is good for the experimenter: the donation size serves as a straightforward measurement of generosity. The drivers of generosity can then be investigated in variants of the game, e.g., by including pre-play communication in the form of a binding suggestion: a third player, the talker, makes a suggestion to the dictator (the listener) about how much money to donate. The dictator first listens to the talker's suggestion and then either accepts or rejects it. Peltzer (2019) played this game under two conditions: a donation is either not very effective, meaning that the recipient receives exactly the donation amount that is given by the dictator, or it is very effective, in which case the recipient receives the donation amount multiplied by 3. The effectiveness is the state of the world. A signal about the state of the world is given to both the talker and the dictator: if the donation is not very effective, both of them receive a red signal with 0.75 probability and a green signal with 0.25 probability (the same signal for both, which is commonly known). Conversely, if the donation is very effective, both of them receive the green signal with 0.75 probability and the red signal with the remaining 0.25 probability.

ity. The experiment measures talking beliefs by asking what the talker expects about the probabilities of the dictator accepting each of three different suggestions (EUR 2, EUR 4 and EUR 6, in each case out of a maximum donation budget of EUR 8), conditional on the signal being red or green. The experimenter can thus observe how talking beliefs change with the information about the state of the world. The results show how the dictator's reaction to the suggestion does indeed depend on the signal: with a green signal, the acceptance probability is above one half for each possible suggestion, but with a red signal, the acceptance probability ranges from about one quarter for the suggestion of EUR 6 to almost three quarters for a suggestion of EUR 2. The talkers, however, underestimate the signal's influence on the dictator's acceptance: they believe that with a green signal, the dictators would show a similar reaction to the case of a red signal -- e.g., that they accept a suggestion of donating EUR 2 only one third of the time.

