Misunderstandings

False Beliefs in Communication

Georg Weizsäcker

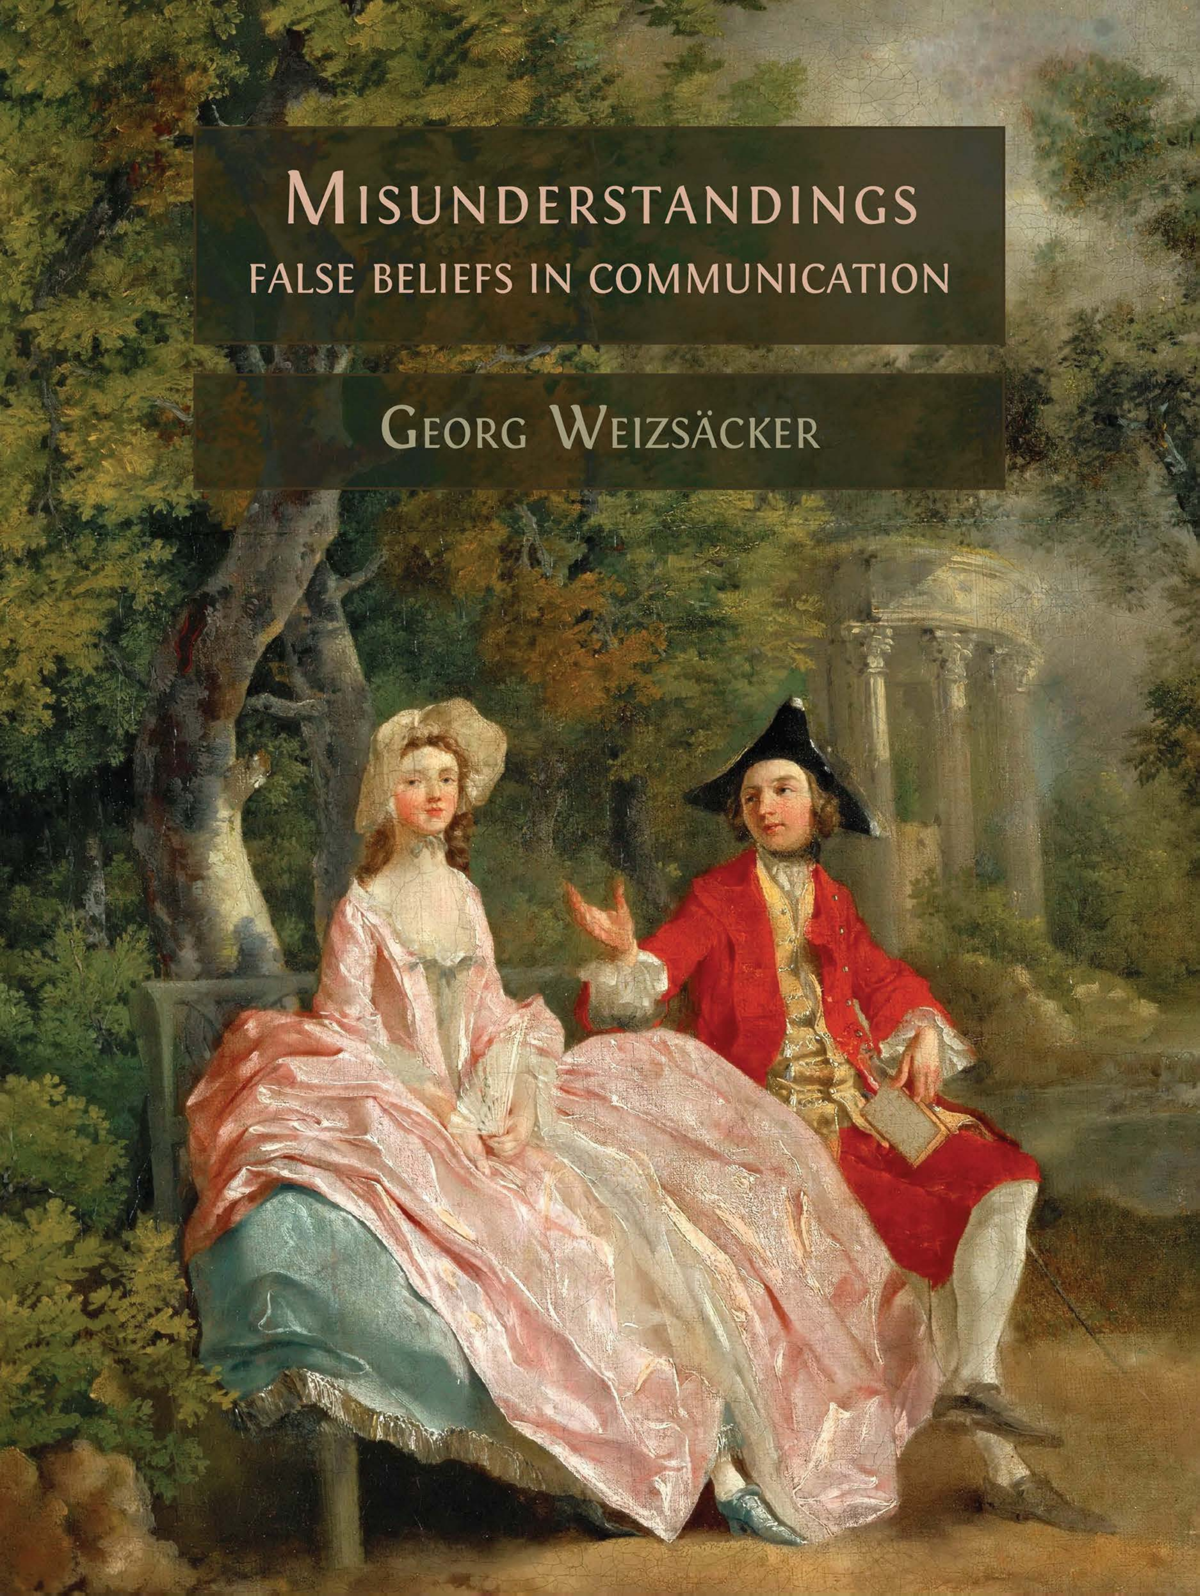

https://www.openbookpublishers.com/

©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All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at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
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 3

Seeing what we don't see

This chapter asks the first three empirical questions about beliefs. They address what we expect ex ante, before the conversation begins.

Since we do not care about their type *per se* (within the limits of the conversation), the relevant uncertainty is about two unknown aspects: their actions and the state of the world. Question 1 and Question 3 address the former and Question 2 addresses the latter.

In what way are the questions in this chapter about ex-ante beliefs, as opposed to ex-post beliefs? There are two answers to this. First, at the beginning of the conversation, it may not be clear whether they or we will listen or talk, respectively. Chapter 4 will discuss the beliefs for the case that they listen and we talk, and Chapter 5 for the case that they talk and we listen. Here, in Chapter 3, this is still open.

Second, Chapters 4 and 5 will proceed by conditioning the analysis on the possible statements: what happens if we say X, and what if we say Y? Here, in Chapter 3, we do not yet consider how beliefs are conditional on statements. In this sense, the beliefs discussed in Chapter 3 are prior beliefs, and those discussed in Chapters 4 and 5 are posterior beliefs. Yet, as we now discuss, even prior beliefs are conditional.

Question 1: Do we underappreciate that every person is different?

We have to gauge our partner in conversation. We have to judge their situation and predict their actions. Every person is different and, hence, every

© 2023 Georg Weizsäcker, CC BY-SA 4.0 https://doi.org/10.11647/OBP.0367.03

person's actions are different. If we think that we already know how they will act, then we might regard the conversation as unnecessary.

One possible goal of the conversation is that we want to find out about their information. Do we know what they might tell us?

Another reason for wanting to know their actions is that our utility depends on the combination of their actions and our actions. Accordingly, what we believe our best actions to be depends on what we believe their actions to be.

Their actions, in turn, depend on their type. We have some information about their type. Do we use it well? This question thus asks how our belief about their actions *conditions on our information about their type*. Does this conditioning move our belief far enough?

Researchers who work on belief elicitation often distinguish two properties of beliefs: discrimination and calibration. Applied to our context, discrimination is the extent to which we differentiate at all between different people, in the sense that we expect different actions from them. Calibration measures the extent to which these conditional beliefs are correct.

'Do we underappreciate that every person is different?' is, by and large, a question about discrimination. The subsequent 17 questions in this book are also of this type. They all ask whether our conditional beliefs show a particular deviation from their target: we may differentiate too little.

Stating the question more formally, from person i's perspective:

$$\text{Is } P^i_{a^j}(\cdot|I^i_{\theta^j}) \text{ too close to } P^i_{a^j}(\cdot)?$$

Rachel's belief about the governor's way of leading the conversation is ignorant of his type. She has met the governor before and she is aware of his political record, but she fails to draw any inference from the information that she has about his type.

The governor, in turn, was briefed about the people around the table and anticipates that Rachel may use the opportunity to push a new research agenda, instead of continuing the previously-started inititative. Also, he is always prepared to receive requests for funds (and is unlikely to engage with them).

The book repeatedly uses the wording "too close to" to describe the empirical possibility that two subjective probability distributions – here, $P^i_{a^j}(\cdot|I^i_{\theta^j})$ and $P^i_{a^j}(\cdot)$ – are closer to each other than the actual distributions of their targets, i.e., the random variables that they refer to. Given what person i knows about θ^j, she may conceivably predict person j's action better than she actually does. Her conditional belief would be more accurate if she anticipated better that certain types choose certain actions and if she combined this insight with her information $I^i_{\theta^j}$, which indicates something about the type that she faces.

The conditional belief can err in many other ways, too. Beliefs are probability distributions – assigning a probability to each possible value of an unknown variable – and their distortions may tend into many different directions. The expression "too close to" alludes to a particular one: person i does not differentiate person j's types as much as the true data generating process.

We also notice the vagueness of the expression: a precise formal expression for "too close to" would require a proper distance measure between distributions. This can be defined with some effort, but is skipped here as it would not substantially aid the discussion. The reader will have a good idea what "too close to" means, even without a formal definition.

Taking this good idea for granted, one can notice a further property of the formulation: it asks about a directed hypothesis. The question can be answered with a simple "yes" or "no" – and the answer indicates a particular bias. It is also noteworthy that the question does not attempt to clarify why the answer is what it is. The question is purely empirical.

Later chapters of the book will provide candidate explanations of why the answer is what it is. For now, the focus lies on the methods for answering Question 1. How does one measure this?

The easiest case is that the to-be-predicted action, a^j, has only two possible values that the placeholder $\cdot$ in person i's belief can take on; call them A and B. The belief assigns a probability to A and a probability to B. The two probabilites sum to one and the belief is therefore fully described by a single number, the subjective probability of A. (As the reader noticed, our simplified discussions of Rachel, Dimitri and Steve view the actions as binary, wherever possible.)

In this case, a suitable measurement of person i's conditional belief requires *(i)* observing the information $I^i_{\theta^j}$ that person i has about person j and *(ii)* eliciting the belief about the probability of the event that $a^j = A$ from person i.

Measuring these things is straightforward. A key advantage of a laboratory study is that the experimenter can control the information that person i has about person j, $I^i_{\theta^j}$, or at least important parts of it. A second key is that the experimenter can ask whatever she wants. She can, for example, ask person i:

"Is person j more likely to choose A than B?"

The answer indicates whether or not $P^i_{a^j}(A|I^i_{\theta^j})$ lies above 50 percent. Voilà, a bound on the conditional belief.

Or, the experimenter can ask for numeric values:

"Expressed in percentage probabilities, how likely is it that person j chooses A?"

Or, if the experimenter expects that the participants are not sufficiently experienced with probabilities, he or she can also refer to frequencies:

"Out of 100 repetitions of this experiment, in how many cases would person j choose A?"

Another possible method is that the experimenter may use a graphical tool to let the participants allocate probability mass. Which of these methods is best depends on the situation and on the eloquence of the instructions. Each of the methods has been shown to work.

But what is the *unconditional* probability $P^i_{a^j}(\cdot)$?

This is harder to measure because one usually observes person i with only one value of information, $I^i_{\theta^j}$. One does not observe what person i would predict with any other value of her information, or without her information.

One possible solution is to consider another group of participants in the role of person i who receive no information about person j. The belief of these participants may be taken as a $P^i_{a^j}(\cdot)$.

As a second possible solution, the experiment can match each participant in the role of person i with a different partner in the role of person j. For each pair, person i receives some information about person j and the experimenter asks person i what they expect person j to choose. This elicits the conditional belief $P^i_{a^j}(\cdot|I^i_{\theta^j})$ for different information sets $I^i_{\theta^j}$ and the experimenter can calculate the population-wide unconditional belief $P^i_{a^j}(\cdot)$ as the aggregation of these conditional beliefs, averaging over the values of $I^i_{\theta^j}$.

Both solutions work in principle. There is, however, a fundamental way in which the second approach is preferable: it takes the fact that *some* information is given as given. Conditionality – here and throughout the book – should be understood as the belief being conditional on the *value* of the information and not on the fact that information is available.

So far, so good – but still not there. The experimenter also needs to measure the "actual distributions" that the beliefs are all about. What is person j likely to do?

Towards this measurement, the experimenter uses the second group of participants, acting as person j. The experimenter measures two proportions in this group: the proportion of A-choosing participants among all participants in the group (this is the actual frequency that the unconditional belief $P^i_{a^j}(A)$ tries to predict) and the proportion of A-choosing participants among those participants for whom a given information $I^i_{\theta^j}$ applies (...that the conditional belief $P^i_{a^j}(A|I^i_{\theta^j})$ tries to predict).

The difference between these proportions can be compared to the analogous difference in beliefs, yielding the answer to Question 1.

`A famous economic decision-making experiment is the ''Trust Game''. In its basic version, one player decides whether to END the game immediately or to SEND some funds to the other player. To simulate the productivity of an economic interaction, the experimenter multiplies the sent amount by a factor greater than 1. The second player therefore either receives nothing (if the first player chooses to END) or receives more than the amount that was sent (...SEND). The game's main point of interest is that the second player can now PAY some funds back to the first player, or RUN and keep it all. Depending on the first player's belief about the` second player's PAY/RUN choice, it may therefore be optimal for the first player to END or to SEND. Binzel and Fehr (2013) play this game in a poor neighborhood of Cairo, Egypt, and each participant faces two pairings: they play either with a friend or with a stranger. In each case, the participants acting as the first player know whether their partner is a friend, with whom they arrived together at the experiment, or a person who is randomly chosen from the other participants. The frequency of PAY is statistically larger for friends than for strangers (72% versus 55%) but the beliefs of the participants in the role of the first player are not: when playing with a friend or a stranger, respectively, 40% and 49% of these participants say that they expect their partner to choose PAY. That is, they tend to trust friends less, not more, than strangers. (This belief difference is statistically insigificant.) We also see a general tendency to be pessimistic: calibration is imperfect in the sense that the unconditional, average belief in PAY is too low. This book's focus, however, lies on the degree of discrimination. In the present game, participants do not realize the sizable increase in PAY frequencies when playing with friends rather than strangers.

Question 2: Do we underutilize our knowledge of the world?

The second uncertain aspect is the state of the world. Knowing it is important because we can choose actions that adapt to this knowledge.

But how much can we know? The state of the world is high dimensional. What are the dimensions on which we need more information? What are dimensions that influence our utility? What are dimensions on which the conversation can inform us? And conversely, what are dimensions that help us predict the conversation?

This asks about what dimensions we *should* focus on. What we *do* focus on is another matter. In many cases, it is a matter of accessibility. Some parts of the state of the world are easy to think about. We have words for them, or lively memories. Other parts are more elusive. Are the dimensions that are accessible to us also the ones that are relevant for us?

This leads to the distinction between things that are known to us and things that we *could* know if we were to focus on them. Let us say that these things are *manifest* to us.

In a rich world, a lot of things are manifest to us when we enter the conversation. Most of these things are unknown to us, in the sense that our actions and statements do not reflect them, but nevertheless they are manifest: we could be made aware of them and our actions and statements could reflect them. This is important in a conversation, not least because some of these manifest things may be pointed out by our partner in the conversation.

The distinction between known things and manifest things resonates with the book's differentiation between information and beliefs. A person's information describes the totality of things that are manifest to her. For person i, it is given by the sets I^i_ω and $I^i_{\theta j}$. In the book's remaining parts, the full information that includes all of these manifest things will often be called the *circumstance* of person i. It is a rather sophisticated description – everything that person i could know.

In contrast, the beliefs of person i, denoted by P^i_ω and $P^i_{\theta j}$, are likely to be less sophisticated: they describe what she actually considers in her choices and statements. They differentiate only along the dimensions that person i pays attention to.

(In an important way, beliefs go further than information: they specify probabilities for the uncertain states of the world. But beliefs may nevertheless be rather unsophisticated.)

We thus re-phrase our question: instead of asking whether the manifest things are the right ones, we ask whether our beliefs differentiate between the right dimensions. Would our utility increase if our beliefs differentiated more along other dimensions?

Having accurate beliefs is not just about differentiating but it is also about getting the quantities right (recall: discrimination and calibration). It requires that we find the right measure of weighing different dimensions against each other. How important is a given piece of additional information, relative to what we believe without considering this piece of information?

It is convenient to state this question, too, through conditional beliefs. Incorporating a new piece of information means forming a conditional probability, where the conditionality refers to the new piece of information.

$$\text{Is } P^i_\omega(\cdot|I^i_\omega) \text{ too close to } P^i_\omega(\cdot)?$$

Dimitri does not yet have a clear view about how the disappointing events of today's competition affect Agniezka's and his collaboration. He jumps to a tentative conclusion, vaguely describing a course of action that he believes to be available to them: "We will hit back." Given how little he knows, it may have been better to not conclude anything, and/or to ask Agniezka for information.

Agniezka enters the conversation with a more pessimistic view about the collaboration with Dimitri, and with more knowledge about the issue. In secret, she already hedged her bets during the last few months, by establishing contact with the competitor team, about the possibility of joining them. Before entering the elevator, she realized that today's competition has greatly reduced the options for her team with Dimitri.

To give an empirical answer to this question, the experimenter needs to measure, or control, the information I^i_ω. Experiments are well suited to do this because the controlled design can govern the information flow.

In its extreme form, an experiment can introduce a stylized set-up where the entire information structure is fully under control. The experimenter describes prior probabilities of ω's possible values, and explains how additional information is generated. For example, that it comes from random draws from one of several urns that have known proportions of balls of different colors. The state ω is the unknown identity of the urn from which a ball is drawn, and the information I^i_ω is the color of the drawn ball. Upon observing the ball's color, the participant can update about ω.

In such a stylized set-up, the unconditional belief $P^i_\omega(\cdot)$ comes for free: the experimenter may take it for granted that the participants agree to the prior probabilites of ω's possible values, as long as these prior probabilities are clearly stated in the experimental instructions. The experimenter only needs to measure the conditional beliefs $P^i_\omega(\cdot|I^i_\omega)$, for the different values of I^i_ω.

In other experiments, prior beliefs may be "homegrown" and do not relate one-to-one to a stylized set-up. For example, the beliefs may refer to real-world events. In these cases, the unconditional belief $P^i_\omega(\cdot)$ is harder to know but it can still be measured.

A similar methodological discussion to that in Question 1 applies here. The unconditional belief $P^i_\omega(\cdot)$ is best measured as the population-average belief, averaging over the values of I^i_ω. This method works irrespective of beliefs being homegrown or not.

For example, let ω be a binary event with values A and B. Information about this event may be homegrown in the sense that different participants have different knowledge about it even before the experiment begins. In addition, the experimenter can give half of the particpants a piece of information that is labelled "high" (which, let us say, indicates that $\omega = A$ is relatively likely) and give the other half a piece of information that is "low" (...unlikely). The population-average belief about A's likelihood is the desired "unconditional" belief $P^i_\omega(A)$: the midpoint between $P^i_\omega(A|\text{"high"})$ and $P^i_\omega(A|\text{"low"})$; both of them can be observed in the experiment.

A fully controlled environment (like balls from an urn) has the advantage that Bayes's rule gives analytic predictions for the beliefs' targets, i.e. for the true conditional probabilities. Question 2 can be applied directly to these benchmarks: is the updating from the signal too weak, relative to the Bayesian prediction? In contrast, an experiment that measures beliefs about naturally occuring events – and no Bayesian prediction – requires that the experimenter can sample the events sufficiently well to estimate their true distribution.

Experimental economists and cognitive psychologists have a long tradition of recording failures in the process of updating from new information. The evidence points to a particular stylized pattern: we update too much from small samples (or weak information), and too little from large samples (or strong information). In an early demonstration of this effect, Griffin and Tversky (1992) instruct their experimental participants that a particular coin is either biased by yielding 60% Heads and 40% Tails, or biased by yielding 60% Tails and 40% Heads. They also specify that the two possible biases occur with equal probability. In the experiment's main part, Griffin and Tversky show their participants sample throws of this coin and ask them to guess, conditional on a given sample, how likely it is that the coin is biased towards Heads. In cases where samples contain only few throws of the coin, the participants make guesses that are too extreme, i.e. too close to 0 or 1. For example, if observing three throws that all land on Heads, the typical response is to say that it is 85% certain that the coin is biased towards Heads, whereas the true Bayesian posterior is only 77%. In contrast, cases where the samples contain more throws induce the participants to hold beliefs that are too close to the ex-ante likelihood of one half. For example, if respondents observe that 11 out of 17 throws land on Heads, the typical assessment is that it is 65% certain that the coin is biased towards Heads, whereas the actual Bayesian posterior is 88%.

Question 3: Do we underappreciate the context when considering their action?

The third question combines the two previous ones. Do we use our information about the state of the world when predicting their actions?

We should certainly do so: their actions (including their statements) are far more predictable if we use our knowledge about the world.

This question covers the third and final conditional belief that arises before the conversation starts. Recall that only two aspects of uncertainty enter our utility directly, their action and the state of the world, and that we have two sets of information, about their type and about the state of the world. In principle, each of the two sets of information helps to predict each of the two aspects of uncertainty. This generates four conditional beliefs that are of interest. But by assumption, types are uncorrelated with the state of the world (see Chapter 2) and our information about their type is therefore not indicative about the state of the world. This leaves us with three relevant conditional beliefs.

Let us use a well-known term to re-phrase the question: do we appreciate the *common ground*? They and we live in the same world, we make similar experiences and obtain similar information. Do we account for this information sufficiently well when predicting their future behavior?

Appreciating the common ground sounds so easy. And yes, the common ground is very important for a good understanding of communication. But clearly, their information is different from our information. (How common is the ground?) Moreover, their inferences from their information may be hard for us to predict – and so may their actions.

This leads back to the distinction of things that we believe in, versus things that are only manifest to us. When predicting their actions, we need to ask what things *they* believe in, and what things are only manifest to them. We will enter this discussion in later chapters of this book. For now, the analysis's first step is to formulate a question without any mentalizing about their beliefs. Instead we consider our information – the things that are manifest to *us* – and ask how this information correlates with the actions of our partner in conversation. (We may happily observe that this is quite a simple question.)

As with the previous questions, we formulate a directed hypothesis: we may discriminate too little.

$$\text{Is } P^i_{a^j}(\cdot|I^i_\omega) \text{ too close to } P^i_{a^j}(\cdot)?$$

In the few seconds before the conversation with Ralph, Steve recognizes the context's exceptionality. Something is making Ralph very unhappy, to the extent that he effectively broke down in an empty street. Steve understands immediately that Ralph would hate to talk about what depresses him. He also understands that the context may induce a whole array of unusual actions of Ralph. On the one hand, Ralph may punish Steve for the intrusion. On the other hand, the revelation of his vulnerability may make Ralph relatively friendlier, softer, towards Steve. Overall, the context creates a wide set of utility consequences for Steve. He concludes that he must tread with care. We notice that this conclusion is fully consistent with the description of Ralph's possible types that appeared in Chapter 2. That is, Steve uses his information well, in light of this type description.

For Ralph, the dominant issue is his own acute problem. Before arriving at the scene, he had witnessed an act of domestic violence at his home, which he cannot forget. The fact that a smaller schoolkid, Steve, enters the scene is an additional disturbance. But it is difficult for Ralph to focus on the interaction with Steve and he does not know what to expect from it.

The flip side of using the state of the world to predict their action is to find a correct interpretation of a given action. Do we see their actions in context? This asks about attribution: do we view the person's type as the driver of their action, or do we view the context as the driver?

In many conversations, the answer is tricky because the context is endogenous. People choose their context, at least to some degree. Certain types are more likely to be found in certain contexts.

A key feature of experimental research comes in handy here: random role assignment. It guarantees that personal characteristics are uncorrelated with the context in which the participants find themselves in.

This feature creates a straightforward null hypothesis. In any experiment with random role assignment, how participants assess the personal characteristics of other participants should be independent of context and of the role that they take on.

`The ‘‘correspondence bias’’ describes the tendency to attribute the reasons and motivations of actions to the person committing them, and not to the context. A classic experiment on the correspondence bias is by Ross et al. (1977) who ask pairs of participants to simulate a quizmaster situation. The two roles, quizmaster and contestant, are assigned to the participants at random, at the beginning of the interaction. The quizmaster then asks the contestant several quiz questions. In one condition of the experiment, CONTROL, the questions are taken from an existing set of questions that were formulated by other quizmasters in previous experimental sessions. In the TREATMENT condition, in contrast, the quizmaster herself writes the questions during the present session. In all sessions, quizmasters are under the instruction that questions should be ‘‘challenging but not impossible’’. All of this is known to the contestants. After the quizzes, both participants rate each other’s level of general knowledge, on a scale from 0 to 100. The main result is that in TREATMENT, where quizmasters write the questions themselves, the contestants judge the quizmasters to be more knowledgable than themselves in 23 out of 24 pairs. Conversely, the quizmasters’ ratings of the contestants are balanced and 12 out of 24 quizmasters rate themselves higher than the contestants, on average. In CONTROL, where the quizmasters do not formulate the questions themselves, such a balanced result appears in both roles. Overall, only the constestants in TREATMENT violate the null hypothesis of balanced assessments: they believe that quizmasters are knowledgable because they come` up with difficult questions. On average, this is a false belief because the role of quizmaster is, as described above, chosen at random. The contestants' mistake is that they fail to incorporate the contextual information that the quizmasters are under the explicit instruction to formulate challenging questions.