Misunderstandings

False Beliefs in Communication

Georg Weizsäcker

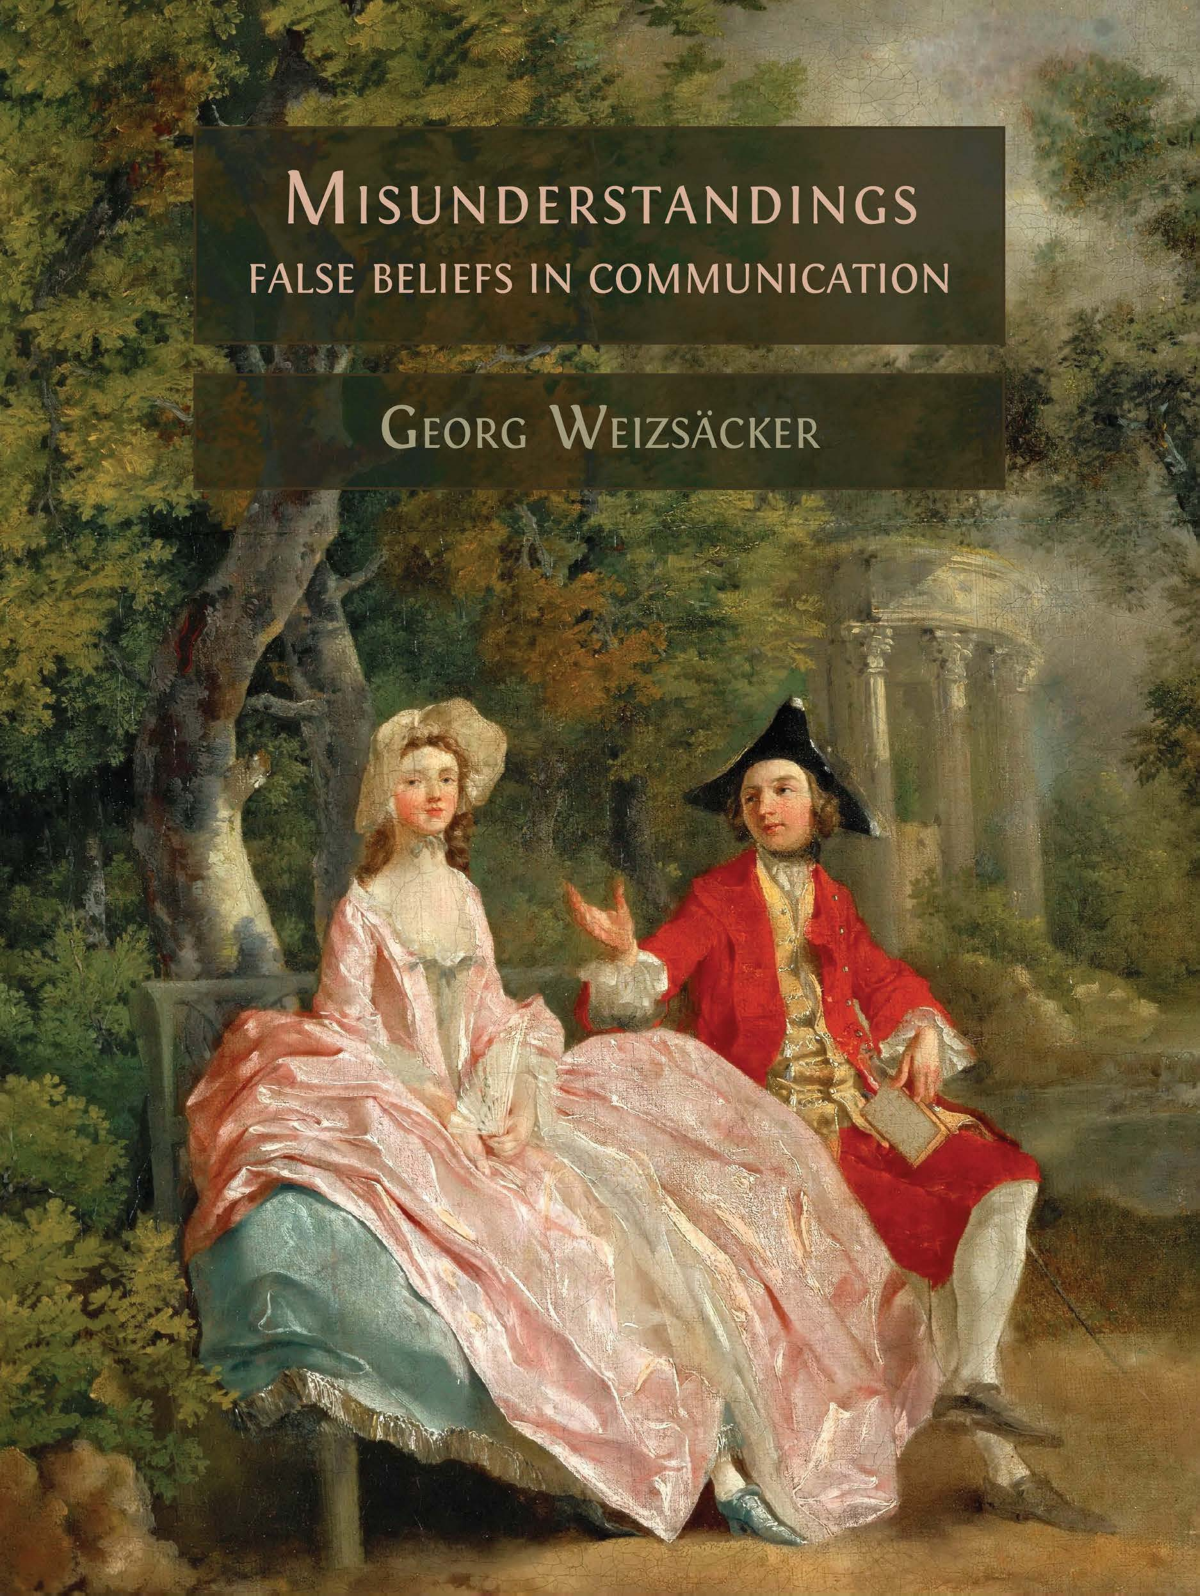

https://www.openbookpublishers.com/

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at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 5

Listening

Through listening, we update our beliefs about two utility-relevant aspects of uncertainty: the other person's future actions and the state of the world. Both co-vary with the other person's statement. If we listen well, we can understand what the statement indicates about the two unknowns. This chapter deals with different ways in which the updating may be inaccurate.

The chapter formulates each of its questions only once although it can be asked in two variants, one of them about the other person's future actions and the other about the state of the world. The two variants are fully analogous and it therefore suffices to ask each question only once.

Question 7: Do we listen too little?

The previous chapter illustrated how our talking is a contingent plan. When in the role of the talker, we say something under a specific circumstance and we say something else under another circumstance. We know that they, when in the role of the talker, do the same.

That is, our interpretation of their statement follows the logic of a language. What would they say if their circumstance was X and what if their circumstance was Y? (And so forth.) With a language in mind, and hearing what they actually say, we learn about the circumstance that they are in.

At least, we do so in theory. What we actually learn is an empirical question about our beliefs, just like all the other questions in this book. Other chapters of the book ask in more depth why they say what they say. This chapter contributes something more superficial but nonetheless important:

 https://doi.org/10.11647/OBP.0367.05

it asks directly – in the sense of: at face value, without a deep interpretation – how our belief changes when we hear their statement.

This is *our listening belief*: for each of the two relevant uncertain aspects (their future actions and the state of the world), for each possible statement that they may make, and for each circumstance that we may be in, it specifies our belief about the aspect, conditional on hearing the statement.

As before, the focus lies on conditional probabilities; here, the conditionality refers to their statements and to our circumstance. This is a pretty large set of conditions – there are many, many combinations of possible statements and circumstances and, no doubt, our listening belief is therefore inaccurate. We cannot ever precisely guess a large set of numbers. But not every inaccuracy is equally important. What we care about is whether our listening belief is on target in the relevant dimensions of uncertainty: the dimensions that allow us to increase our utility through optimal reactions. For example, we may only be interested in predicting a particular future action of theirs, but not everything else that they might do or say.

Actually, we have many good excuses for updating badly when hearing their statement. One such excuse is that their statement is not our only input. In a sea of information that is manifest to us, a single statement's significance is easy to miss.

We must, however, not ignore the sea of information. Even if we find most of it uninteresting per se, it may be relevant for the interpretation of the statement. Questions 8 and 9 describe ways in which we may mis-judge the circumstance of the statement. The present question asks only about how much we condition on the statement itself.

An especially easy-to-miss element is that the talker may have had the *choice* to talk. As mentioned in earlier chapters, their statement carries a different meaning depending on whether they could have chosen to say nothing.

This caveat cuts both ways. If we observe that they do not say anything, we may update from this observation. Conversely, if we observe that they do say something, then we may update from this fact alone. We may, in particular, believe that the fact that they talk is indicative of the relevance of what they have to say.

The same logic applies to the choice of topic that they do, or do not, talk about. If we notice that they avoid a certain topic, this may inform us about this very topic. Conversely, if they do address a certain topic, then we may regard it as especially relevant and update our belief about it more than about other things.

An important formality is that listening beliefs are different in nature from talking beliefs: they relate to different uncertain aspects. Talking beliefs are about the other person's reaction to our statements. Listening beliefs are about what they do after their statement and about the state of the world. We therefore cannot compare the two sets of beliefs with each other. When measuring them, we determine their accuracy by comparing each of them to its own target, i.e., the actual probability (or frequency) that it refers to.

Finally, another note on notation: here and in the formulation of Questions 8 and 9, the symbol x stands in for either of the two unknown aspects of uncertainty $\{\omega, \tilde{a}^j\}$: ω is the state of the world and $\tilde{a}^j$ is the future action of person j, which is still uncertain after j said a^j.

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

Agniezka listens carefully to Dimitri and tries to assess his future actions. He had drawn her into this project, to form a team. Now, he is talking about "blood" and the question arises whose blood it is that he refers to. Would she, Agniezka, be in harm's way if she continues to be on the team, more than if she leaves the team? A noteworthy aspect of Dimitri's statement is that he does not imply in any way that he will blame her if she withdraws from the project now.

Measuring listening beliefs is tricky business. A talker makes only one statement. The listening belief, however, includes the posterior beliefs after every possible statement.

Once they have heard the statement, it would be difficult to ask the experimental participants about their belief for the counterfactual case that the talker would have said something else. It may therefore be better to ask for the listeners' conditional belief reports before the talker talks: what they believe in case that the talker makes one possible statement, likewise for the next possible statement, etc.

Another difficulty, a variant of which came up in Chapter 4, is that it is unclear how to compare the talking belief to its corresponding unconditional belief. What is $P^i_x(\cdot)$? Surely it is a weighted average of the conditional beliefs $P^i_x(\cdot|a^j)$, for different values of a^j – but with what weights?

In Chapter 4, we answered an analogous question only for the binary case of a statement that is either A or B. Now, we consider a more complete answer for the case of a multi-valued statement, in two steps. First, an abstract answer that works fine in theory – but requires a lot of data. Second, we consider how a realistic experiment would deal with this case.

The first, abstract answer is actually quite simple: simply weight each value of a^j by the frequency with which it occurs. In a large experiment with many participants in the role of person j, these frequencies are readily available: the experimenter observes the distribution of chosen statements. He or she can thus use these frequencies as weights to calculate both the belief $P^i_x(\cdot)$ and its target (the unconditional distribution of the unknown x that the belief is about). We note that this set of weights is appropriate in the sense of payoff relevance: the frequencies reflect the actual importance of each statement for person i's utility.

The second observation is about the practical implementation. Here, things are more difficult. If a^j has many possible values, then the data set must be so large that for every single one of these values, there are sufficiently many observations to estimate the average conditional belief with high precision. Otherwise, some of the conditional frequencies cannot be trusted and one may need advanced statistical methods to account for this sampling error.

Moreover, the expresssion "too close to", which describes a comparison of probability/frequency distributions, is perhaps too vague for a large set of values for a^j. If the experimenter is really interested in a large set of such values, then he or she may also be interested in some specific comparison of the relevant distributions and the expresssion "too close to" may not capture this well.

In actual practice, an experimenter will therefore likely proceed as in Chapter 4: if he or she can rely on the simplification that person j makes a binary statement that is either A or B, then the experimenter will do so and answer Question 7 (like Questions 4, 5, and 6) in a simple way. For a binary statement, the question boils down to asking whether $P^i_x(\cdot|A)$ is closer to $P^i_x(\cdot|B)$ than the corresponding actual distributions of x.

Importantly, an *ex-post* binarization is possible for experiments with a large set of possible statements: the experimenter may classify each of them into A or B, thereby enabling an analysis that is equivalent to the case where only these two statements exist.

Many statements amount to being a promise, or other announcement of one's own future behavior. Charness and Dufwenberg (2006) conduct an experimental game similar to the Trust Game that was described in Chapter 3 but with the additional feature that one player can make a promise, thereby potentially enticing the other player to trust him. A promise works as follows in this particular study. The second player announces that, if trusted, he will roll a die whose (random) outcome determines a monetary payment to the first player -- and payoff rules are such that if the second player actually rolls the die, then it is highly likely that the first player benefits from it. But the announcement is non-binding, i.e., the second player does not have to actually roll the die. The first player hears the announcement and her task is to decide whether to end the game early and thereby earn a low but safe payment, or to trust the second player and let him make the decision of rolling versus not rolling the die. The game also involves payments to the second player that give him an incentive to mislead the first player: if she trusts him, his payment increases and he can increase it even further by deviating from the promise and choosing not to roll the die. The experimental instructions are clear about the fact that the second player is obliged to say something, and the experimenters can classify the statements that the participants make into two categories, those that effectively promise the friendly ''I will roll'' versus those that do not do so, which we may understand as announcing ''I won't roll''. The data show that participants who promise ''I will roll' are indeed much more likely to actually roll the die. The respective frequencies of rolling are 79% for the participants who promise ''I will roll'' versus 33% for those who do not promise it. But when participants acting as first players are asked what they believe about the probability of rolling after hearing the statement, they reveal listening beliefs that do not discriminate enough: the reported average expectations of actual rolling after hearing statements ''I will roll'' and ''I won't roll'' are 64% and 51%, respectively.

Question 8: Do we underdifferentiate talkers?

We form our expectations based not only on what they say, but also on who they are. The same statement, when uttered by different people, means different things. We therefore condition our listening belief on our information about their type.

When predicting their future actions, two effects arise that require such differentiation of talkers. First, different types do different things in the future, and second, different types make different announcements (and then do different things in the future).

When predicting the state of the world, the first of these effects is irrelevant but the second effect remains valid: different types say different things. We therefore learn in a type-dependent way from their statements.

A possible complication is that we may also learn something important *about* their type when we hear their statement. They say something and we may realize that our prior belief about their type was wrong. Why should we, then, condition on our prior information?

The answer is that we should use all information. Neither should we condition only on the prior information nor should we use only the new information that the statement contains. The relevant belief is a posterior belief – forming it well requires that we use our prior belief, but with good measure.

This requires a good meta-accuracy about our prior. If we are too certain about their type, then our interpretation of the statement may suffer. Likewise, if we give too little weight to our prior view of their type, then our posterior is inaccurate, too. We may, e.g., enage in too much temporal extension, meaning that we overgeneralize the momentary impression that we receive from the present statement.

Recall also that we do not care about their type per se. Chapter 2 ruled it out by assumption. Assessing their type is important to us only because it helps us learn.

This is a good moment for pause. *Why* does our listening belief depend on our information about their type?

A very rational reason is that their type describes their incentives. We anticipate that they choose their statement optimally, from their perspective. For instance, they may want us to learn something. They may also simply enjoy saying what they say.

Recall our slow, piecemeal approach: this chapter does not yet consider the other person's perspective in any detail. Presently, we only recognize that different types will choose different statements. This may indeed have to do with their incentives, in a large number of ways. Yet, it may also stem from less sophisticated reasons. We may be differently perceptive with different people. We may attribute a different degree of informativeness to them, for reasons other than incentives (including prejudices). We may be in a different mood depending on who is the talker. Knowing who they are may make us more, or less, curious about what they have to say. The information may mislead us by installing an *idée fixe* about their future actions or the state of the world. In any case, our listening belief varies, for rational or irrational reasons, with what we know about the talker.

The discussion of Question 6 included a definition of the perceived relevance of an information that person i has. This was in the context of person i being the talker, and the information was about the state of the world. Here, for Question 8, we notice that the concept of perceived relevance also applies to person i being the listener (not the talker) and to cases where the relevant information is about the other person's type (not the state of the world). We skip over the technical specifications of all these variations of relevance – they would be fully analogous to the one in Question 6 – and merely remark on the generality of the concept: perceived relevance may apply for the talker and for the listener, and for all kinds of information. All variants of perceived relevance will also be important later, in the second half of the book where we will ask about the interlocutors' beliefs about what is relevant from the perspective of the other interlocutor.

To summarize the discussion of type-specific listening: in the case of person i being the listener who has information $I^i_{\theta^j}$, the information is potentially relevant because it helps her to anticipate better the subsequent action of person j and/or to better update about the state of the world. The empirical question is whether she does it to the right extent, which we express – as usual – in a directed question about her beliefs.

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

The governor listens well enough to put an end to the discussion; he understands that he can proceed in the meeting by giving praise to Rachel, strong but vague, not promising her anything. He does not notice, however, the possible significance of scientific research about the provenance of art collectibles. Rachel's remark about it establishes certain facts (she is a scientist, after all) and implies the possibility that these facts may have relevant political consequences. Yet, the governor only refers to other issues.

In other words, he has a sophisticated updating about future actions and the state of the world on those dimensions that he was prepared for: Rachel as a political stakeholder who may ask for funds. He fails to engage with her expertise, and therefore with the information that is new to him.

How can an experiment create a meaningful variation in person i's information about person j's type? Most populations of experimental participants are quite homogeneous. For example, university students are all of similar age and educational status. Highlighting the differences between them requires some effort. Moreover, the highlighting itself may distort the measurement. The experimental participants may pay attention to things that they would otherwise not pay attention to – an instance of "experimenter demand", meaning that the participants react to what they believe is the purpose of the experiment.

One possibility is to use an existing set of statements that were uttered by a heterogeneous pool of people in a different environment and that were recorded with a different purpose.

`Belot et al. (2012) report on an experiment using video material taken from a TV game show. Participants in a decision laboratory watch actual episodes of a show where the contestants play a simple game of cooperation: two contestants simultaneously decide to either cooperate or not, with monetary rewards. Each player has a unilateral incentive to avoid cooperation but earns far more if the other player cooperates. Before the cooperation game commences, the game show contestants can talk to each other. They use this pre-play communication phase to make announcements about their future cooperation -- which may induce the other person to cooperate, too. In the decision laboratory, the video screening stops after the communication phase and the experimental participants are asked to report their beliefs about whether or not the contestant whose statements they have heard will actually cooper-`

`ate. The pool of contestants contains people with many different characteristics, at least one of which is known to the experimental participants: gender. The lab participants should therefore condition on the contestant's gender in their interpretations of what they hear. They do so, but too little. While promises made by women are more predictive of cooperation than promises made by men by a difference of 19 percentage points, the lab participants only predict a difference in predictiveness of eight percentage points.`

Question 9: Are we too impressionable?

A discussion similar to that in Question 8 applies to our use of manifest information about the world – the context. We know something before they say something. We have to put the two together. Do we condition on context to the right extent?

We condition on context for three reasons. First, their talking strategy differs for different contexts. Second, for a given strategy on their side, we update about the state of the world differently for different contexts – because we have different prior beliefs about it. Third, we also update differently about their future actions – we expect them to be context-dependent in addition to the possibility that statements are context-dependent.

The first of these considerations highlights that we need to investigate the mind of the talker. Just as we use our knowledge, they use theirs, too. Our information is not the same as their information but it is indicative of it and we may thus appreciate the context in which they talk. The next chapters will continue this thread.

The second and third consideration takes the talker's talking strategy as given and focusses on how we interpret it. Our information about the world may help us to see that their statement makes certain states of the world, or certain future actions, more or less plausible.

All of these considerations can be expressed by a single question about our own listening belief: does it react well to our knowledge about the state of the world? Just as all previous questions, it is an empirical question whose answer does not require a full understanding of underlying mechanisms. The statement that we hear has a certain statistical informativeness and it is in our interest to recognize this informativeness as accurately as possible.

Challenges arise, once again, from the multitude of dimensions: what parts of the context are we aware of, and to which of them does our listening belief react? Or, to use the term relevance again: what context information is perceived relevant for us as listener?

A feature may help here: that we hear their statement and can reconsider the context's relevance in light of the statement. That is, we may listen to their statement not only by updating about the state of the world, but also by shifting our attention to a new background.

We may, however, still neglect the objectively relevant parts of the context. For instance, we may be overly impressed by what they say. Their statement may come with a lot of eloquence, evidence, or other persuasive force.

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

Ralph's listening is naive. He accepts all-too-easily that the possibility of a ball game is the relevant topic. (Perhaps understandably so – recall that Steve carries a ball, which is a visible cue.) Ralph misses what else Steve could have addressed and what would have been far more relevant, namely Ralph's state of despair. Interpreting this context as the topic of conversation and showing a violent reaction to Steve's statement – or, more gently, telling a cover-up story – may have enabled Ralph to better defend his position in the schoolyard pecking order.

Making monetary payments as part of the experiment is commonly done in experimental economics but not so in experiments of other fields. Why is that?

An ill-meaning explanation is that economists are money fetishists. They regard money as an important motivator, whereas other scholars do not. A less ill-meaning explanation is that economists want maximum control over the experiment. They prefer to restrict attention to an analysis where one dimension is well understood and controlled, even if other dimensions are left out.

This is no coincidence. The thinking of economists often follows their leading theory of decision making, expected utility theory, which describes utility as a one-dimensional real number. Money is one-dimensional, too, which is part of why economists like to analyze it.

In this book we use the same simplification. Recall the payoff mapping introduced in Chapter 2, $u^i : A^i \times A^j \times \Omega \times \Theta^i \to \mathbb{R}$; it is important that u^i is a one-dimensional summary of person i's welfare, about which we can assume that person i maximizes its expected value.

(The maximization may be misguided because person i's beliefs are subjective, i.e., potentially false. This is the whole point of the book. Notice that this point, too, could not be made equally well if we did not assume that person i maximizes u^i given her belief.)

u^i is not money, though. It is the utility that a person receives from a combination of the realized values of the three uncertain aspects. The mapping u^i can take on many forms and the analysis works under a whole range of different assumptions about it. For instance, economists sometimes suppose for simplicity that u^i is a linear scaling of person i's own money payments. In this case, the earnings of others, and the risk arising from the variability in the uncertain aspects, do not bother person i. She is selfish and risk neutral. Alternatively, the utility u^i could be assumed to be a non-linear transformation of money. This allows capturing risk preferences. Or, u^i could depend on person j's payments as well, which would capture social preferences.

The experimenter may make monetary payments not only for the actions (or statements) in the experiment, but also for the participants' belief reports. The experimenter can thereby give appropriate incentives to give honest reports.

Paying for belief reports – how does this work? Imagine that person i is asked to predict a binary action of person j, $\tilde{a}^j \in \{A, B\}$. Then person i's belief about this action is a number between zero and one – the probability that $\tilde{a}^j = A$. Incentivizing the belief report amounts to paying money such that person i maximizes her expected utility if she reports her true belief.

A simple mechanism achieves this: person i receives a fixed payment F with a probability that decreases quadratically in the distance between her reported belief and the ex-post realized truth state of A (which is 1 if person

j chooses A, and 0 otherwise). That is, if person j actually chooses A and person i predicted that this would happen with probability r, then she earns F with probability $1-(1-r)^2$. If, on the other hand, person j actually chooses B then person i receives F with probability $1-r^2$.

This mechanism has the property that reporting one's (subjective) expectation of j's behavior is the optimal response for person i, no matter what her risk preferences are. Other, simpler mechanisms have other theoretical properties. E.g. if the reward that a participant gets for her belief report decreases linearly in its distance from the truth, then the mechanism makes it optimal for person i to report the median of her subjective distribution.

Most participants will likely find the exact version of the monetary incentives not too important. Indeed, there is very little evidence in the literature that the exact rule matters, given that an experiment rewards accuracy in one way or another. There is, however, some evidence that it matters (mildly) whether one uses monetary rewards at all. The belief reports have been shown to discriminate better and be more consistent with one's own actions if money is paid.

Eyster et al. (2018) run experiments where people sometimes forget their prior beliefs and react to other people's statements far too much -- they are too impressionable. A large group of participants listen to each others' statements, where each statement is an estimate of a sum of numbers. The participants act in sequence, after each of them was endowed with a privately known number that is randomly drawn from a mean-zero distribution. When called upon to make her statement, a participant's goal is to guess the sum of all privately-known numbers of the participants who appeared previously in the sequence, including her own number. She reports her guess of this ''target'' and is rewarded for accuracy: her monetary reward decreases (here, linearly) in the distance between the guess and the true sum of numbers, calculated up to her own position in the game. The game-theoretic, rational prediction is that everyone hits the target and earns the full payoff: since all guesses are public, later participants can use the previous guesses to infer what underlying numbers each of the previous participants must have seen. If they figure the previous numbers out, they can add their own number to it. But figuring out the previous numbers may be difficult. In one treatment, participants announce their guesses one after the other, in a slow and transparent manner. Each person's best guess is, here, simply to add one's own number to the previous guess. Participants are mostly successful in tracking the correct sum of numbers in this way. They, however, mildly over-react to previous participants' guesses, and would earn more on average by shading their guess towards the ex-ante expected value of their target (which is zero) by roughly one third. In a second treatment, participants act in lumps. At each point in time, four participants have to announce their guesses. This makes the learning process more complicated as one cannot simply add one's number to the previous guess. (The game-theoretic prediction still uses simple arithmetics, but in a lengthy way.) Participants make systematic mistakes in this treatment, by following the direction of previous guesses to an extreme extent. Early guesses are incorporated again and again, with far too much weight: they are, effectively, multiple-counted, leading everyone's guesses to go far astray. On average, participants should move closer to the target's ex-ante expected value by no less than 98% in this treatment. The effect is so strong that participants who act in the second half of this treatment would earn twice as much money in the experiment if they were to ignore all previous guesses and simply reported only their own numbers.