Misunderstandings

False Beliefs in Communication

Georg Weizsäcker

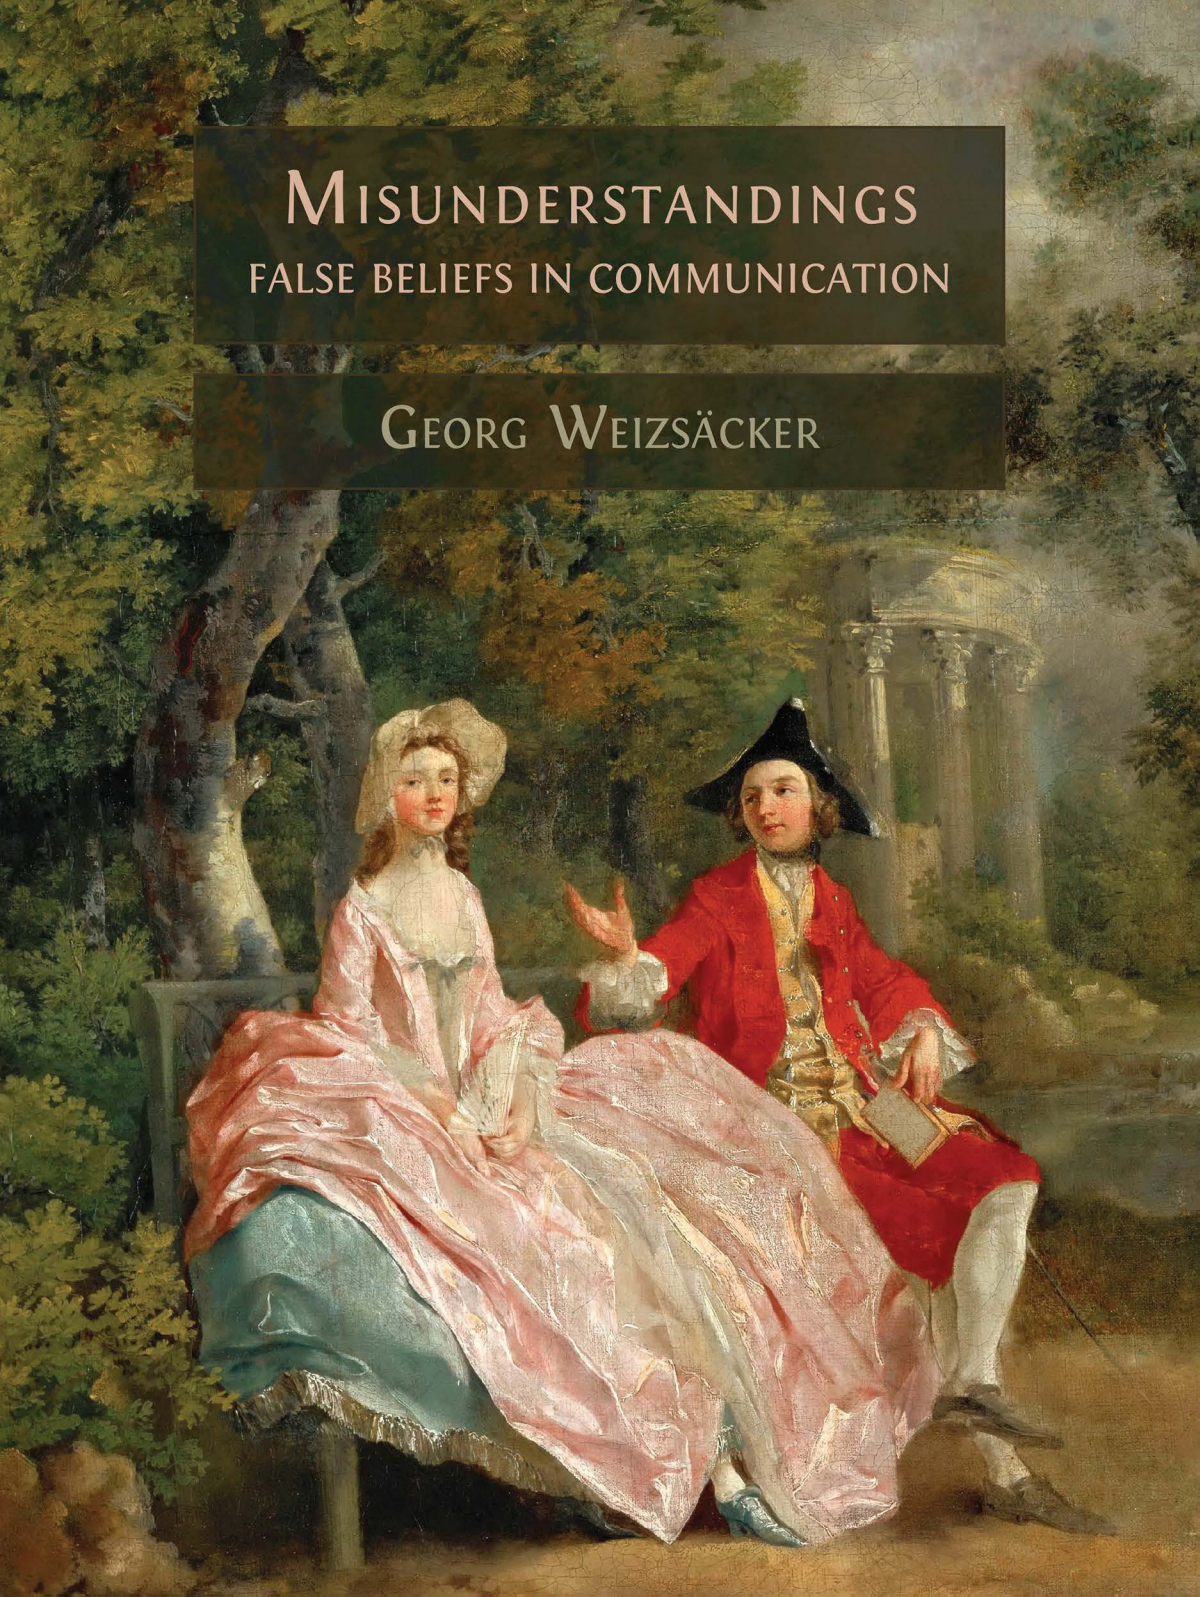

https://www.openbookpublishers.com/

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 8

Perceiving how they listen

Recall from Chapter 4 that we anticipate how they may react to each of our statements – as described by our first-order talking belief. We now aim to support this belief, by turning the table on Chapter 5 and asking how we perceive their listening belief: how we understand their belief change upon hearing the statement. This belief change translates into a reaction, i.e., what we care about.

(A nasty note of caution over two paragraphs, which the reader may want to ignore: let us always keep track of who has what information. When considering how they listen, we condition on their possible information sets. There are many of them, and we have to consider them all. All of *our* beliefs depend on our own information. In particular, our beliefs about their beliefs also depend on our own information.

Nevertheless, we do not need to re-consider each of their circumstances separately for each of our possible circumstances. It suffices to discuss the questions that arise in this chapter simply for one of our possible circumstances, namely the one that we are in. This shortcut stems from the fact that this chapter only asks about the inner layer. It considers what we believe about how they update their listening beliefs; these updates depend only on their, the other person's, information sets, not on ours.)

 https://doi.org/10.11647/OBP.0367.08

Question 16: Do we think they just don't listen?

They hear our statement and react to it. Our (outer-layer) expectation of this reaction is our first-order talking belief and underpins our talking. We now justify the first-order talking belief, and hence our talking, through a perceived first-order listening belief.

To see how exactly this belief underpins our talking, we proceed with the usual slowness. First, a name. Our subjective expectation about their first-order listening belief is called, no surprise, *our second-order talking belief*. It is accurate if it correctly predicts their first-order listening belief for each possible statement and each information set that they could be in.

With an accurate second-order talking belief, we correctly understand how they would update from each of our statements, in each possible circumstance of theirs. Realistically, we cannot achieve exact accuracy – but we should try. The more accurate our second-order talking belief is, the better we predict what their reaction to each statement is, and the better is our choice of statement. A more accurate second-order talking belief increases our utility.

The correct use of their information is a challenge and makes this question interesting. Just like we use our information when listening – see Chapter 5 – we consider how they use theirs. We imagine how they consider their information about our type and how they consider their information about the context. This results in a quite rich interplay between second-order beliefs and first-order beliefs, in a similar way to that in Chapter 7. Just like we did there, we can ask the questions about meaning (Question 13), about lies (14) and about the richness of language (15).

Indeed, many things are the same from both sides. For another example, recall the importance of aligned interest in Chapter 7, which is equally important here. Whether or not the incentives are aligned influences whether we can both gain from the statement, versus only one of us. Second-order beliefs reflect this, and they thus reflect how much information we, as talker, believe to transmit. Still another analogy to Chapter 7 is that an accurate second-order talking belief does not imply that the first-order talking belief is accurate; and vice versa.

But there are differences, too. In this chapter, *we* are the talker and our statement carries meaning to *them*. We know the connection between our circumstances and our statement because we choose the statement ourselves. They, as listener, need to figure it out. The fact that we already know the connection may help our second-order belief's accuracy, or reduce it. Another difference – one that helps *their* belief – is that they hear the statement that we make and can condition their beliefs on the fact that we make it. We, as talker, have to consider different statements without the additional information that one of them is chosen. For all these reasons, second-order talking beliefs require a different kind of mentalizing than second-order listening beliefs. It may be that when we listen, we have accurate second-order beliefs but when we talk, we fail to have them – or the other way around.

Notice, finally, that in many conversations we talk because we may *want* to change their expectations, for reasons of vanity. We may love to surprise. (It may be the reason why we are in the conversation in the first place.) In other conversations, we may want to keep a secret from them.

At the core of all these cases lies our second-order belief accuracy. How well do we assess the extent to which our statements changes their belief?

$$\text{Is } P^i_{P^j_x(\cdot|a^i)} \text{ too close to } P^i_{P^j_x(\cdot)}?$$

Here, and also in Questions 17 and 18, we let $x \in \{\omega, \tilde{a}^i\}$, in analogy to the description in Question 7.

Dimitri, ill-prepared for the elevator conversation, does not anticipate that Agniezka understands his statement as indicating that he would not blame her if she left the project now. Furthermore, he does not anticipate that describing his vision of the team's possible joint future makes Agniezka's belief about this joint future only more gloomy. In his current state of mind, everything that he talks about would sound negative to her. It may have been more effective for Dimitri to talk about the possibility of the team splitting up.

The illusion of transparency is a direct obstacle to the accuracy of second-order talking beliefs: we tend to overestimate how well they know our circumstance. We may therefore, erroneously, think that we need not describe it.

Conversely, in a situation where we want to hide our circumstance, we may underestimate how well we can do it. We are better liars than we think.

For an experimental demonstration of the illusion of transparency, one need not specify the set of possible statements with much detail. It suffices to measure first-order listening beliefs about the truth status of a statement (whatever the statement is) and compare it with the talker's second-order talking belief.

Gilovich et al. (1998) provide a nice demonstration of the illusion of transparency: one group of participants, the ‘‘drinkers’’, participate one by one and are asked to drink five drinks of equal visual appearance while being watched by a second group of particpants, the ‘‘observers’’. One of the five drinks has a foul taste but the drinkers are asked not to show which one it is. An observer's task is to watch the drinker drink and to identify, from their demeanor, the foul-tasting drink. That is, he or she attempts to detect a lie: the instance of the drinking excercise where the drinker only pretends not to be disgusted. After drinking, the drinkers report their second-order beliefs: they predict how many of the observers have correctly identified their feeling of disgust. The actual detection rate is 20%, so the observers are no more accurate than guessing randomly. But the drinkers predict that a significantly larger share of the observers, 36%, identify the true circumstance.

Question 17: Do we think they interpret our statements and other people's statements in the same way?

We usually do not say as much as we could say. Our language is vague and does not specify all the detail. As discussed for Question 15, the rich variety of contexts makes it even less conceivable that we specify all the detail. Yet, vagueness sometimes increases eloquence and can help to make ourselves clear.

Importantly, a statement's effect on the listener is not only a matter of the words that we use: how we believe to influence the listener depends a lot on what we know about them, and what they know about us. For certain

constellations of types, it may well be that using vague language is our best option.

A well-known game-theoretic analogue of this is the impossibility of informative equilibria in the presence of conflicting interests. If our and their incentives are not well aligned, an equilibrium where we speak in a fully informative way is ruled out by a simple contradiction argument. Suppose that we could talk to them in such a fully informative equilibrium. Then we would attempt to influence them and move their belief away from the truth, at least a little bit, in the direction that helps us the most. This deviation would not help them, by the assumption of conflicting interests. They know this and would therefore not listen to everything that we say. The equilibrium therefore cannot be fully informative.

An equilibrium with vague language, in contrast, can well exist if the incentives are not too disaligned. For instance, we are able to indicate to them a rough summary of our information – whether it falls into one rough category, or another. They cannot infer anything precise from such a vague indication, but at least they learn the rough summary. Note the equilibrium property: Given how they listen, we actually want to make such a vague indication because it is the best we can do. And they, in turn, have no reason to fear that we trick them, given that our language is vague and does not allow micro-managed influences of their reaction.

The reasoning is an equilibrium argument and it therefore involves higher-order beliefs. But of course it also involves beliefs of lower order. So let us consider to what extent we can tell the same story while restricting ourselves to first-order beliefs and second-order beliefs.

To tell the story, we compare the value of a vague statement with that of a precise one. The hypothesis that the vague statement is optimal, expressed as a property of our second-order talking belief, is that we expect that they update weakly more (in a direction that helps us) in response to the vague statement than in response to the precise one.

Why would we expect this? The precise information has all the information of the vague statement, and more.

It is, however, certainly a feasible belief. We may expect them to learn the wrong thing if we give them too much information, i.e., that it misleads them or keeps them from listening to the key part of what we say.

A deeper reason has to do with our type. A long and elaborate statement may change their view of our preferences. It may emphasize our type and our hope for a particular, fine-tuned reaction. That is, a more detailed statement may make them suspicious.

This reasoning goes back to asking about what is (perceived) relevant. Talking has a cost, so if we talk with much detail, we must believe that there is a benefit in doing so. Our expectation must be that they respond positively to the detailed talking. But this expectation requires that we have a detailed view of their type, to the extent that they share our interest even about our detailed utterances. We must regard the details as relevant, and believe that they are relevant for the listener, too. This directs attention to the question of preference alignment: if their preferences are not so congruent with ours, then this fact, too, becomes more salient. They may ask themselves how much they share, or trust, our preference.

A more general observation is that we use *ostension* when we talk. We point to something. We use the fact that the listener knows about our opportunity cost: we could do something else with our time and energy – e.g., say something different. Given this opportunity cost, the fact that we say something specific makes it obvious to the listener that we expect a reaction to this specific statement. This, in turn, sends the listeners on a search for a motivation for the specific statement, i.e., for a context that we perceive as relevant if it is combined with the statement.

Question 18 and Chapter 9 will return to the issue of context for the listener. Anticipating this discussion, we notice that our doubts about our ability for precise language, as expressed in the previous paragraphs, may become even stronger if we try to justify our second-order talking belief by introducing more levels: beliefs of third and higher order. The more we ask about what we believe about each other's type and the beliefs about these beliefs, the more questionable it may appear that our preferences are well-enough aligned to convey much information.

A more language-oriented view of vagueness generates further possible effects. Linguists, psychologists and philosophers have examined that vagueness is far more than just imprecision. A statement that is merely imprecise can conceivably be made more precise, for instance by using a longer and more detailed sentence. But many vague statements cannot even conceivably be subject to such a "precification". (Think of poetry again.) Instead,

they are simply a different category of statement. They may therefore evoke a belief change that is different from anything that a precise statement can evoke.

Even these linguistic aspects of second-order talking beliefs are connected with what they think about our type. We may believe that they, when hearing a vague statement, may think of us as being friendly or wise, or perhaps as useless or offensive. Their interpretion of our statement depends on how they view us and our second-order belief about it, correspondingly, contains large degrees of freedom.

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

Rachel is a scientist who uses fairly precise language. Her statement about having "verified the reports with scientific studies" is not exaggerated. She knows about the reliability of her statement, and she does not like to simplify or boast. Her second-order talking belief specifies that the governor understands her type in this respect. She therefore expects that he understands her words to be literally true (give and take). The governor does not, however, view her type in this way – he does not regard Rachel as an authority for truth, any more than most other people with whom he interacts – and he therefore does not entertain the possibility that her statements describe a significant political risk for him.

Being personally close with the person to whom we talk may make our second-order talking beliefs more accurate. But they may also become less accurate if we know the listener well.

Experimenters can make the appropriate measurement for this question by simply combining measures that were previously discussed in the book. The experimenter can vary the identity of the listener and measure how the second-order listening beliefs varies with it, relative to how the different listeners' first-order listening beliefs change. To demonstrate that the listener's identity creates a bias, it suffices to carry out the exercise for a given statement, or a small set of statements.

A surprising pattern of results in psychology shows the extent of misunderstandings among spouses. Savitzky et al. (2011) conduct a communication experiment where they record the expectation of participants on how well they expect to be understood.

The talkers say a pre-determined phrase that has multiple possible meanings, like ‘‘What a nice pen’’, which could indicate that one is about to grab the pen, or that one would like to know who owns it. Four different interpretations are listed on the experimental instructions, for each statement. The talker is under instruction to convey one of the interpretations via their way of uttering the phrase, using variations in their speaking tone in any way they like. The listeners then indicate one of the four candidate interpretations. Before the listeners' interpretations are revealed, the talkers are asked to predict in how many of ten phrases they expect the listeners to grasp the intended meaning. The experiment, crucially, also varies the identity of the listener. When talking to a stranger, talkers predict that a little more than five phrases are understood in the intended way, whereas in truth a little less than four were understood. When talking to one's spouse, the actual understanding was almost at the same level - roughly four phrases were interpreted in the intended way - but talkers believed that more than six phrases were identified correctly. The results indicate an over-conditioning of second-order talking beliefs, for the case that one's spouse is the listener.

Question 18: Do we fail to see how their interpretations are in context?

Choosing the topic of our statement, we play with the difference between what is merely manifest to our listener, on the one hand, and what they believe to be relevant, on the other hand. We make references to certain dimensions of the state of the world, influencing what is salient to them. Their, the listener's, information about the world – all things that are manifest to them – is more than what they can process. We re-direct their focus and we do this better if we understand how their first-order listening belief reacts to context.

Given the high dimensionality of context, it is quite a miracle that they understand us at all. Yes, we can influence their focus, but how far does this take us? There are so many dimensions of context and, with them, there are so many possibilities of interpretating our statements. But it works, at least somewhat. The human mind's strive towards efficiency has lead to great mastery of language. Talking and listening are highly efficient ways of pro-

cessing information. In a haystack of possible meanings, we and they both have a good idea of where to search for the needle.

This efficient selection of the contexts has, once again, a lot to do with relevance and utility maximization. As discussed for Question 17, we use ostension when we talk, i.e., we point to things. They know that we choose the statement in a way that is best for us. The fact that we point to something means that we expect them to react to this fact – otherwise our ostension would be a waste of time and effort.

A short way of putting it: utility is benefit minus cost and for a context to be relevant for us when we talk, it must therefore be (i) accessible to them (low cost of connecting our statement with it), (ii) accessible to us (ditto), (iii) beneficial to us (their reaction helps us) and (iv) perceived beneficial to them (they believe it helps them). This narrows down the interpretation.

Our second-order talking belief reflects this reasoning. If we believe that their first-order listening belief makes a particular context perceived relevant for them, then we believe that their reaction depends on this context. We, too, can condition our statement on the context and, between the two of us, we can transport meaning more successfully by conditioning on the same context.

That is, the relevance of contexts are likely to be connected between the two interlocutors, at least in our mind. Clearly, our beliefs may be wrong. We will surely not be able to anticipate perfectly how each context is relevant for them, but the closer a fit we obtain, the better it is for us.

Notice again the asymmetry between their role and ours. They, as listener, cannot influence our belief about the relevant context – only we can influence theirs, through our statement. The listener's perceived relevance of a particular context may change with the statement they hear. The ability to influence the perceived relevance of context is, in all this, the talker's alone.

One way of justifying our second-order belief is to link it to a third-order belief: we believe that they condition on a certain context because we think about their second-order listening belief. If we have a reasonable belief about how they respond to their belief about our belief, then we may be more confident in our own second-order belief. In principle, one can carry this further, to justify the third-order belief by fourth-order beliefs, and so on.

But for many considerations, this type of deeper thinking is not necessary: the analysis described before the previous paragraph applies even if we just take our second-order talking belief as the primitive concept. The second-order talking belief describes to what extent we perceive a context as relevant to them. This is all that we need to select our statement. And, in all this, the more accurate is the second-order belief, the smaller is our misprediction about their reaction.

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

Steve's second-order talking belief is off target, by conditioning too much on the context of Ralph's despair. In particular, Steve believes that Ralph believes that with a relatively high probability, Steve will tell other children about what he saw. Steve's statement aims to minimize this probability in Ralph's belief, i.e., it aims to express the smallest possible challenge to Ralph's authority.

Unbeknownst to Steve, the inaccuracy of his second-order talking belief is beneficial – a lucky coincidence. Steve's statement appears to be so unrelated to the context that Ralph does not connect them in his mind. Had Ralph made such a connection, he would have been more pessimistic about Steve's future actions and would have reacted more aggressively.

Experimental design is often guided by the solution of a reference theory, such as the set of Nash equilibria in a game. It is good science to determine precisely what the reference theory predicts for an experiment before actually running it.

The more general the theory is – the weaker the assumptions that it rests on – the better. The strongest experiments often have a clear prediction that is valid for a wide set of assumptions. An unambiguous theoretical prediction gives the experimenter a firm control over the interpretation of the data.

By the same token, the experimenters should stay in control of their theoretical prediction when varying the state of the world in an experiment. In cases where the variation of ω is a pure framing variation, this is usually not a problem since the theoretical prediction mostly does not react to framing.

In cases where material incentives are affected by a change in ω, retaining a clear theoretical prediction may be a challenge. But it may be surmount-

able – e.g., it may be possible to keep the theoretical prediction constant across different values of ω.

In an economic experiment by Agranov et al. (2023), competition is switched on and off between two treatments. In each treatment, there is a single buyer who wants to buy a good only if it is of high quality. In one treatment, NO COMPETITION, there is only one seller who observes the quality and sends a message to the buyer -- ‘‘high quality’’ or ‘‘low quality’’. As usual in buyer-seller games, the seller may have an incentive to lie to the buyer if the true quality is low. Whether or not he has this incentive depends on the credulity of the buyer. In this experiment, the seller may also incur a payoff reduction that reflects guilt aversion and/or lying aversion, and the buyer may incur a disappointment cost if she believes a lie. These possible psychological effects are private information, making the behavior of other players hard to predict. Agranov et al. (2023) formulate a precise parametric theory of these effects, in order to factor them into the analysis and predict how they differ between different contexts. In particular, competition may matter: a second treatment, COMPETITION, introduces another seller who also has a product for sale and can also send a message to the buyer. The buyer then selects a seller and, potentially, decides to buy from the selected seller. In this treatment, lies and disappointment may be evaluated very differently than in the absence of competition. Despite the relatively high complexity of the games played in this experiment, the two treatments have analogous sets of equilibria and the design allows a ceteris-paribus investigation: does competitive pressure induce more lying, and is lying more or less successful under competitive pressure? The results on the buyer’s first-order listening belief show that buyers tend to believe a message that promises high quality in both treatments: the average posterior beliefs in high quality after hearing the message ‘‘high quality’’ are higher than after hearing the message ‘‘low quality’’, by a difference of 50 and 55 percentage points, in NO COMPETITION and COMPETITION, respectively. The sellers’ second-order talking belief show corresponding differences of 53 and 45 percentage points in the two treatments, respectively. That is, if anything, the buyers are mildly more credulous in the treatment where the sellers face competition. The sellers, in contrast, believe that the buyers become slightly less creduluous in this treatment. But the sellers still

believe that the lie is effective, even under competition -- where they are more likely to need it. This is consistent with the observed messages by sellers. They lie more than twice as often in COMPETITION than in NO COMPETITION.