



The Exchange of Words: Replies to critics

Richard Moran

Harvard University, Cambridge, MA, USA

I am grateful to David Hills, Lizzie Fricker, and Adam Leite for the careful attention they have given to my book and their stimulating comments at our APA session in February 2019. Before I get to my responses, I would like to spend some time giving a kind of overview of "*The Exchange of Words*," going over some of its main themes and how I see them hang together. This will put some of my specific replies in context. In doing so, I will run the risk of both being sketchy and occasionally sounding dogmatic, so I beg the reader's indulgence.

When I wrote my first paper on testimony, I felt that the fact that human testimony is a linguistic affair was not being given the right emphasis in the discussions that I was familiar with. Not that this was ignored, of course, but it seemed to me that not enough was being made of the fact that speaking, and hence testifying, is a social practice, dependent on the institution of language, and that the way a person's claim that P can be a reason for her audience to believe that P depends on facts about speech that have no equivalent in the ordinary confrontation with evidence, such as the tell-tale footprint on the sandy path. In considering the person's utterance or the footprint as a reason to believe something, we may, in both cases, be said to be dependent on what some person *does*, but, in the case of being *told* something, we are dependent on the speaker in distinctive ways. Regarding the footprint in the sand, we may draw various conclusions on our own; there are many things we might learn from it, but when a speaker testifies to some fact she commits herself to, the truth of the particular proposition is affirmed.¹ Her speech is not a brute phenomenon, which might mean anything or nothing, but rather the assertion of a particular proposition. And the act of telling or testifying is a particular kind of act of the person, one that the speaker specifically *means* to count as a reason to believe for her audience; she means it to be specifically informative. In this, it is different from her other actions, which are not addressed to anyone but from which an observer may be able to draw interesting conclusions. In writing the book, I wanted to understand the phenomenon of one person *addressing* another, and what that has to do with commitment in speech. And, from the beginning, I was struck by the fact that, as speakers, we each have our own perspective on what it is to make our words count for others, to make the difference between simply saying some words and actually *telling* someone something. In any act of telling, there are two perspectives at play: that of the speaker and that of her audience. But of course each of us adopts these different perspectives in turn. Every speaker is also an interlocutor. This duality of coordinated perspectives already seems different from the ordinary confrontation with evidence like the footprint, which is a phenomenon that has no perspective on itself or its status as a reason to believe something. To me, this seemed related to another puzzle about testimony: If my act of telling can be my interlocutor's reason for believing something, a genuine epistemic reason, could that same act be my own reason for believing the same thing? If it is a genuine epistemic reason, like the footprint in the sand, can it not count in the same way for both of us on the same occasion? If it was I who made that footprint on the sand, I could inspect it and draw conclusions from it in just the same spirit as any other observer. The fact that I was the one who produced it does not put me in a different epistemic position with respect to it than anyone else. Why does it seem different in the case of my claiming of that P? Can I not rely on my own claim the way I expect my interlocutor to do? If

it makes no sense to speak of "believing oneself" the way you can believe another person is that simply because there is no *need* for me to do so since I must be assumed to already be in possession of the knowledge in question for me to make the claim in the first place, or is there some deeper incoherence here? And either way, this suggests that even while I provide my audience with genuine epistemic reasons for believing P when I tell him that P, the two of us bear essentially different relations to the reasons created by my statement to him. And that seems importantly different from either of our relations to the footprint in the sand, even when I am the one who made it. The topic of self-other asymmetries has been important to my thinking since beginning working on my first book, *Authority and Estrangement* (Princeton 2001), and the topic of speech and testimony seemed to open up new ways to think about such asymmetries, in the social context of speech and not only philosophy of mind.

In his classic 1990 book, *Knowledge and the State of Nature* (Oxford, Clarendon, 1990), Edward Craig makes a distinction that is also basic to my outlook.

"There are informants, and there are sources of information. Or, to arrange the terminology differently, among the various sources of information there are on the one hand informants who give information; and on the other there are states of affairs, some of which involve states of human beings and their behavior, which have evidential value: information can be gleaned from them. Roughly the distinction is that between a person's telling me something and my being able to tell something from observation of him." p. 35

The informant, as opposed to the mere source, is assumed to *know* the information that she provides. The person as informant is not an "unwitting" source of information. Human linguistic communication is a distinctive way in which knowledge and information circulate among human beings. We gain true and false beliefs from what people say and do in all sorts of ways, many of which have nothing to do with anything a person means to be communicating. But it is central to human linguistic communication that the speaker *means* or *intends* to communicate something to someone. It is true that in our speaking as in our actions generally, much of what we "give off" or communicate to those around us is, as it were, in spite of ourselves, part of the general informational field that we all live in. But communication in speech, in the central cases of saying, asking, or asserting, is characterized by the speaker *meaning* to communicate something and meaning for this to be accomplished by *being understood*. Again, not all learning from others takes this form; not everything we learn from others depends on understanding what they are doing or on their meaning for us to learn anything from them at all. For human linguistic communication, however, the notions of meaning and understanding are essential and not accidental.

The speaker as informant takes herself to be communicating something to the other person deliberately, directly, and not merely in spite of herself. And she takes the fact that she means what she says, that she means to be providing her audience with a reason to believe something, to itself be an element in how her act of speaking actually *is* a reason for her audience to believe what she says. The speaker does not take the intentional character of his act of speaking, the fact that she understands what she is saying, to be *incidental* to its character as a reason to believe something. Rather, she takes it to be essential to the epistemic import of what she is doing, not only that it is something she is doing intentionally and with understanding of its significance but also that her audience must himself understand her to be acting intentionally and with understanding in order for him to gain the kind of reason for belief that she takes herself to be presenting in her speech. Human linguistic communication aims at people understanding each other.

Sometimes, of course, it is actually true that the speaker's own understanding of her act is irrelevant to some epistemic import it has for another person. When my accent reveals that I am American and not British or French and also when I say to someone "I am American," it is the same thing believed in the two cases. But, in the latter case, the epistemic import of the utterance to the belief in question depends on the *content* of what is said, whereas, in the first case, the content of the statement is irrelevant to its status as a reason to believe that I am American. I could have revealed the same thing about myself by saying "I'm the King of Siam." Ordinary dialogue

and communication of knowledge through testimony depend on the particular content being claimed as true. And, in either belief or disbelief and agreement or denial, there is an identity of content between what is asserted and what is believed or denied. By contrast, in the purely evidential or symptomatic relation to someone's words, as when my regional accent gives me away, the observer learns something that need have no relation at all to the content of my assertion. His conclusions are detached from the claim being made. The observer here is outside the situation of dialogue and not related to the speaker in the way of two conversational partners where the respondent can say "yes" or "no" to the speaker and express agreement or disagreement with the very thing claimed by the speaker. For an observer or detective like this, the speaker's statement or claim is the *occasion* for believing something, which may or may not be related to what he said. Here the speaker's claim is not the very thing believed or disbelieved.

(This is what I mean by the speaker and the audience or observer being out of harmony with each other, when they are not participating in the same practice of assertion and belief, denial, or counter-assertion with respect to the same content. They are not out of harmony with each other when the audience distrusts or doubts the speaker's claim or knowledgeability (I will say more about this in my reply to Lizzie Fricker).

The dependence of the topic of testimony on the understanding of illocutions like claiming and asserting has not, I think, been given the central place it deserves in the epistemology of testimony, and this is a part of what I meant earlier when I said that the specifically *linguistic* nature of testimony had not be fully appreciated. And I think this is in part due to the observer's perspective that comes so naturally to the epistemologist. The classic topics of epistemology emerge naturally out of an imagined situation of observation, the one-way confrontation with some phenomenon or body of evidence, raising issues of what inferences can legitimately be drawn from them, the proper conditions for their assessment. The objects of epistemic assessment are themselves "states of affairs" in Craig's sense, without a perspective on themselves or their own epistemic significance. When it comes to the topic of testimony, philosophers naturally fall into discussing it purely from the point of view of critical consumers of testimony and leave out of consideration how each of us is also a producer of testimony. But these two perspectives are not separable from each other. Epistemically, the speaker has to understand what she is doing in asserting that P as the sort of thing she herself could understand and accept as a reason from someone else to believe something. This is true even when she is lying. A purely epistemological story about testimony tends to obscure this mutual dependence of the two roles, as though we could understand testimony purely from the consumer's point of view (which will be a critical or skeptical point of view, as with any confrontation with evidence), and could leave out of consideration one's own speaker's point of view on one's own act of assertion. As speakers ourselves, we have a perspective on how we stake ourselves in speech when we make a simple statement of fact. When we do so, we produce an utterance in words that could, on another occasion, be merely mentioned rather than used, or could be used to express a variety of different speech acts other than assertion. The speaker herself has to determine how she is presenting her utterance if it is to have any epistemic significance concerning its content. As participants in this practice of speech, we all know how to do this. The philosophical understanding of testimony must account for the *participant's perspective* on the practice as a whole.

If there is no issue of believing testimony until an utterance has the specific form of a claim or statement, then what makes this possible? Here, I give a lot of attention to the nature of illocution in speech, its difference from the perlocutionary dimension of speech (Austin). Briefly, asserting, promising, and renouncing are illocutionary speech acts. For an utterance to count as a claim that P, rather than wondering whether that P, or simply mentioning the words of a philosophical example, is for the speaker to commit herself to her audience in a particular way. She makes no promises or claims when talking in her sleep or pronouncing the words of a language she does not understand. And when asserting that P, she knows that this is what she is doing in her utterance, rather than wondering whether that P, and she knows this in the way she knows her forms of commitment in action. She does not come to know this by observing her own behavior, nor can she look to her interlocutor to tell her whether she is asking her or telling her. Rather, she knows her illocution as a piece of practical self-knowledge. Here, as interlocutors, we defer to the speaker. It is up to her and not her audience to determine whether her utterance should count as *claiming* that P or

wondering whether that P, as a promise or as a mere piece of advice. All this is quite different from what Austin calls the perlocutionary dimension of speech. In addressing my conversational partner, I may in fact be boring him, surprising him, or offending him. I might be aware that I am doing some of these things, and I might be doing some of them intentionally, but my awareness of what I am doing along these lines need not matter at all to whether I am doing them or not. These may be results I am aiming at, but whether I am doing them or not is not up to me. I speak with no special authority here, and another person may know quite better than myself that I am only succeeding in boring or offending my audience. If I *do* know that I am doing these things by speaking, there is a sense in which my knowledge is contingent and external to the facts themselves. By contrast, I argue that the speaker's awareness of the illocutionary act of claiming or promising is internal to it and not external. My knowledge of what I am doing when I am seeking to persuade or surprise my audience will play an enabling role in achieving success, but I might also achieve that end accidentally or unknowingly. But, with illocutions such as claiming or promising, the speaker's understanding and awareness of what she is doing is part of what makes these the particular acts that they are. The utterance of the words, "I promise to meet you tomorrow at three" fails to be a promise at all if produced under hypnosis or by someone asleep or without understanding of what she is saying. The failure of these conditions for the act of asserting or promising do not simply impair it as a good instance of its kind but annul its character as an act of that kind altogether.

I will not be able to cover all the points raised by David Hills, Lizzie Fricker, and Adam Leite, but they have given me plenty to think about. I will begin with a few corrections and clarifications:

Early in his reply, David Hills warns against thinking that "what's done deliberately in committal assertion can't also be done inadvertently, and that what's free is in no way constrained." I quite agree, and I do not think anything in my book denies the possibility of a person *being* responsible for careless or unintended consequences or suggestions of what they say outright. We incur and assume responsibilities in life in all sorts of ways, some of them just by being born in a certain place at a certain time in history. In the book, I am interested in how the practice of language makes it possible for speakers to deliberately and explicitly assume responsibilities that they would not have but for their overt, public assumption of them. Perhaps Hills did not mean to suggest it, but I do not think I am guilty of what Bernard Williams refers to as making a "fetish" of assertion. This comes up in Williams' book, *Truth and Truthfulness* (Princeton 2002), when he is discussing whether someone can be said to have lied when they say something, which is literally true but which they know will be understood by their audience to mean something false (107). Williams argues that it would be making a fetish of assertion to claim that, in such cases, the demands of truthfulness have been observed, and I fully agree.

About promising and explicit commitment, Hills compares Raz's discussion of Harry, who says he advises but will not promise, with Moore's Paradox. He says Harry gives no indication of his strange combination of attitudes and hence cannot be construed as cooperative in the Gricean sense. I agree with Cavell, Scanlon, and others² that the obligations of a promising do not require a special verbal formulation, but Raz's case seems a comprehensible instance of declining to assume a certain obligation. Harry is someone who wishes to be helpful so long as it is at no cost to himself and does not constrain him in any way. This may not be a model of benevolence, but it is hardly incoherent and is perfectly common. Harry has his plans for tomorrow, and if he can help out someone without changing his plans and, more importantly, without being *obligated* to follow through on his announced plan even if his situation changes, then he is happy to do so. He would like to help out if it is convenient for him, and he sincerely believes that it *will* be convenient for him. Hills asks: "Harry invites John to count on a lift from him, and what more could the straightest of straight promises achieve than that?" (10) Well, I think there is this basic difference: When you promise someone a ride, you are now bound to that person. You are no longer free as before to give the ride or not; your plans for tomorrow are now fixed around the time you promised the ride, and only the person you made the promise to can now release you from that obligation. Harry does not want all that, but he still fully expects to be in a position to help and wants his friend to know that. I do not see any incoherence here.

And, in any case, my point in taking up the Raz case was to illustrate the idea of a "normative power" and how particular illocutions such as claiming or promising can be seen as ways that speakers have to deliberately assume or

delineate their responsibilities. We are certainly responsible for more than we deliberately take on, more than we may be aware of. No one denies that. But it is also possible for human beings to publicly and deliberately take on or decline specific additional responsibilities, one's they would not have but for their public declarations. That is what they do in making and refining their claims, in promising and other speech acts.

Both Leite and Hills take up my discussion of Casaubon in George Eliot's *Middlemarch*. I am not sure if I am in any real disagreement with Hills here, but I want to underscore my own point in bringing up this story. I claim that overtly addressing another person about some matter is importantly different from getting them to see something by other means. I believe there are various philosophical pressures that lead people to overlook, deny, or downplay this difference, and that is indeed a theme in the book. Here, toward the end of the account of overt intersubjectivity in the penultimate chapter, I use this example from *Middlemarch* to suggest that we can see what the distinctive importance of outright telling and overt intersubjectivity are by considering the situations where that sort of engagement with another person is precisely what someone is seeking to avoid, where it is a line they do not want to cross. The thought is that we can shed light on its distinctive importance by thinking about the motives a person may have for refusing precisely that overtness, where that is something different from mutual knowledge or mind-reading. Here is what I say in the book about some possible motives:

"There could be many reasons for this avoidance on the part of Casaubon. He does not want to invite a response from his wife Dorothea, or he does not want to make it appear as though he is appealing to her in any way, that he is asking for some sort of response from her, or he does not want to expose himself to further inquiries from her, or to any open display of pity or concern. Or he does not want the facts about his illness to be "common ground" in any conversation between them, such that each of them will now be entitled to assume this truth in their overt dealings with each other. Or he wants to find a way to let her know, perhaps a way that involves his saying something, and perhaps even in a way that involves her recognizing his intention, but he does not want to "talk about it" with her; that is, he does not want to face her about it because that would involve being faced by her about it. He would like to install the correct belief in Dorothea's mind, but he does not want a conversational partner. His informative interests here are purely unilateral. He wants to produce a certain belief in his audience, and he has come to the conclusion that something like the full Gricean mechanism is the best means to accomplish this. If Dorothea understands Casaubon's attempts to bring her to know about his illness and his fears in the same unilateral way as we are imagining that he understands it, then her ability to "get the message" will indeed depend entirely on her skills as a mind reader and an interpreter of the evidence. That is the position she will have been relegated to, since she is not a participant in a conversation with him. It is only in communication that is "essentially avowable" in Strawson's sense that the people in question would have any use for a second-person pronoun or a first-person plural, a "you" or a "we," to address each other and express their presence to each other in the conversation. By contrast, the full Gricean calculation we imagined being rehearsed by Casaubon in his desire to let Dorothea know about his condition can be fully laid out without any use of the second person, since the very conception of "informing" that he is operating with is a unilateral one in terms of acting upon the beliefs of another person." 184-5

Regarding Hills's discussion of the photograph by Lewis Hine, we do seem to be looking at things differently here. As I see it, what matters to the status of a photograph as evidence of the working conditions in the factory has everything to do with whether the photograph was in fact staged or doctored in any way. And we often (though not always) trust the photographer not to have manipulated the image in misleading ways. Hine's photo appears to show a girl at an industrial cotton loom, with an adult person in the background. Kendall Walton's point, which I think is right, is that the photograph would be evidence for the girl and the adult in the background, regardless of what Hine's own beliefs were and what his purposes were in producing or showing the photograph.³ That is, this is what the photograph would show, would be evidence for, even if Hine himself thought the child in the picture was a young boy

and not a girl, and even if Hine did not notice the adult figure in the background. The photographer's beliefs are simply not relevant to whether the photograph is evidence for an adult in the background. That is why the photograph can have documentary value, even for the photographer himself, and can correct his own memory of the scene. If our attribution of the photograph to Hine turned out to be wrong, and it was in fact the factory owner who took the shot, we could still see what it shows about the factory conditions. But when a *speaker* tells us "There's also an adult figure in the background," we can only learn something about the scene in the factory from this utterance if we assume such things as that the speaker means the same thing by these words that we do, that she is not speaking in a trance, is using rather than mentioning these words, and intends her utterance as a statement of fact rather than some other illocution or a philosophical example. This is how the spirit in which an utterance was produced and presented matters to its epistemic import, in ways that it does not matter to what we can learn from a photograph that has not been doctored. And even when a photograph is doctored, it is the doctored itself that matters to its epistemic status and not the beliefs or other attitudes of the photographer.

I have learned much from Lizzie Fricker's work on the epistemology of testimony, and her comments give me a chance to clarify certain basic elements of my story. We do in fact agree about non-linguistic cases of Gricean meaning, and much of Chapter Six concerns cases that fall somewhere along the spectrum of showing and saying (e.g., my discussion of Sperber and Wilson). And we agree that it is not "merely out of a purely interpersonal relation between a speaker and an audience" that the responsibilities connected with telling are assumed but that these apply only within a socio-linguistic practice, which I frame in terms of Thomas Reid's notion of "social acts of mind." At the same time, I appreciate Fricker, noting that I agree with her (Fricker 2006) that "overhearers are not in principle worse off epistemically vis-à-vis a piece of testimony than the addressee." (139) "[T]hat there is such an entitlement available depends on the fact that the utterance is an audience- directed speech act" (2006, 598).

First a clarification and a correction. At one point Fricker characterizes my view in the following terms"

"The gist of Moran's discussion of evidence is this: the epistemic significance of a linguistic act of telling entirely depends on the intentions with which it is produced; but this means that it cannot be evidence for belief in what is told, since evidence, Moran wants us to believe, must have its epistemic force regardless of how it was produced."

On my view, the epistemic status of an act of telling does not entirely depend on the intentions with which it was produced. This epistemic status of any assertion will naturally depend on such things as the speaker's knowledgeability and truthfulness, as well as the wider social-linguistic practice. Secondly, the illocutionary status of her utterance does not so much depend on the speaker's *intentions* as on how she is *presenting* her utterance within that practice. The picture I develop in the book is not a psychological one but the description of a social practice. And finally, when I say that the epistemic status of something like a footprint in the sand is independent of the stance toward the footprint on the part of the person who made it, I do not mean that the footprint's status as evidence that someone was walking on the sand is independent of *how it was produced*, for of course the apparent footprint may have been produced mechanically, from the air or by a device resembling a foot, and that would indeed impair its status as evidence that someone was walking there. As with my remarks to David Hills about photographs, doctored evidence is misleading evidence, and what makes it misleading is the actual manipulation of the evidence, not the photographer's own attitude toward it. What makes a phenomenon evidence in the first place is independent of the stance toward it of the person who produced it. And this is different from how someone's utterance gains the status of reason to believe something, when the speaker, playing her role in the practice, determines its illocutionary status as assertion or act of telling.

As to "evidentialism" about testimony, my own use of the term "evidence" may indeed be idiosyncratic, but, in a moment, I will try to say why I use it the way I do. I take it that, even under "evidentialism," there will be broader and narrow understandings of the term. Often what we mean by "evidence" is a body of empirical phenomena. But if "evidentialism" is presented as a general epistemic principle, then presumably it should also include such things as

logical entailment or inference to the best explanation, which can also provide good reasons for belief. Perhaps these can be seen as “evidence” in some broader sense, but, for this and related reasons, I prefer her later formulation in terms of “truth-related considerations” and leave the term “evidence” for the narrower notion. I do take beliefs gained by testimony to be based on “truth centered” or “epistemic reasons,” and, in Chapter Five, I distance myself from the idea of “epistemic partiality” or believing a friend out of loyalty, etc. (at one point David Hills characterizes my account in terms of the hearer being offered practical reasons for belief in what he is told, but that is not my view).

Here is what I mean by the contrast of evidential relations with the participant's perspective on the act of telling and knowledge communicated through testimony. In general, I take “truth-centered” or “knowledge-grounding” reasons to be a broader category than evidential reasons. For instance, in *Authority and Estrangement* (Princeton 2001), I defend the idea that a person is in a position to know her own mind or know her own future actions in ways that are not observation- or evidence-based but are nonetheless truth-centered and sources of knowledge. As with observation-based knowledge, my practical knowledge of what I am about to do can be communicated to another person and treated as ordinary information in their own planning. There are a variety of forms of knowledge with a basis other than evidential.

In my discussion of testimony, I am thinking of “evidence” in a way that is impersonal, something that is a reason for belief, and is so in a way that is independent of anyone's attitudes with respect to it. By contrast, a speaker's utterance, as a verbal performance, is only a potential testimonial reason for belief when the speaker presents her utterance as an assertion or similar speech act. From the speaker's point of view, her bare utterance has no epistemic significance regarding its content unless she is presenting it specifically as a claim or assertion. As a verbal performance, she could be saying the words by way of quoting someone else's assertion or mentioning the words rather than using them, etc. But the social practice of speaking puts her in a position to present those very words as a statement of fact to her audience, and if her audience recognizes what she is doing, then she will have *told* her audience something.

If her audience relates to her utterance as evidence for something, as a reason for believing something that is independent of the speaker's own stance toward it when she gives it the illocutionary status of a claim, then the audience is not relating to her utterance as she meant it. This is how an observer may discern my nationality by hearing my regional accent, which is evidence for my origins independently of any significance I mean to give it. The audience here would be treating my utterance as an ordinary kind of evidence for something, which, like a footprint in the sand, has that status independently of any assumptions about the stance toward the footprint on the part of the person walking in the sand. For the speaker herself, the question of whether her utterance is to count as a claim or something else is not something she settles by observation, speculation, or evidence. There is no fact of that matter prior to her committing herself or declining to do so.

For her audience to take the speaker's utterance to be evidence for the truth of its content independently of the speaker's own stance toward her utterance is to fail to recognize the kind of reason the speaker takes herself to be offering here. In terms of Craig's distinction, one would be failing to relate to the speaker as informant. Now, once it is recognized by both parties that an actual *claim* has been made by the speaker and that this claim depends on how the speaker has committed herself in her utterance, then we can say that the *fact* of her having asserted that P is a reason to believe that P indeed is *evidence* that P. But that epistemic status of her utterance is parasitic on something that is not evidential: The speaker is giving the illocutionary status of claim or assertion to her utterance. After the fact, when a speaker as informant has made her claim, we can say that her having made her claim is a “state of affairs,” or some behavior of hers, which we can treat as evidence for some fact (perhaps the very one she claimed as true). But that attitude is completely derivative of the primary, participant's relation to her exercise of a normative power within the social practice of assertions and other illocutions.

Finally, I hope these remarks make it more clear what I mean by certain philosophical constructions of testimony construing the speaker and hearer being “out of harmony” with each other and that this is not the same thing as the hearer taking a critical, even skeptical view of the speaker's claim. At one point Fricker says, “Thus I disagree with

Moran's thesis that forming belief in what one is told in a way that is aptly sensitive to evidence of trustworthiness would involve a kind of disharmony; would mean that a speaker's word was never taken in the spirit in which it was offered." But I do not think we really are in disagreement here. Of course we often respond skeptically and critically to what people tell us and are right to do so. When I speak of "disharmony," I am thinking of the situation where the audience responds to the speaker's utterance as having its epistemic import wholly independently of the speaker's own stance toward her utterance or what she takes herself to be speaking informatively about (as when my claim about the weather reveals my American origins). When I claim that P, the way my utterance of P counts as a reason for P depends on what illocution I am performing, whereas when that same utterance provides evidence of my nationality, that is, so quite independently of how I am staking myself in the utterance. Here my audience takes my utterance to be a reason to believe something quite independent of the actual content I am asserting, and in this way we are out of harmony with each other in the speech situation. The interlocutor's responses of "yes" and "no" have no application apart from an assumption of identity of content and reason-giving with respect to it.

Many of Leite's examples concern cases where people are only dimly aware of what they are doing or the import of their own actions. Naturally I agree that life is often like that, but I think the philosophical interest and import of such cases depends on a background picture of intentional action and speech. He may disagree with me here, but, although we do things both intentionally and accidentally, inadvertently, and unconsciously, I believe that the concept of *intentional* action is the central defining idea of human action as such. Perhaps this is not controversial, but I also believe (with Elizabeth Anscombe and others) that intentional action is constitutively related to the agent knowing (without observation) what she is doing under some description or, more simply, that it is not accidental that a person can answer a relevant "what" and "why" question about her own current action. But, human life being what it is, there will be problematic and interesting cases where we want to say the person acted intentionally, but confusedly, or perhaps with only repressed awareness of the import of what she is doing. But I take it those cases are significant and interesting only as departures from the norm and that we should not conclude from this that all we can say is that sometimes people are aware of what they are doing intentionally and sometimes not. The psychological and philosophical questions these cases present only make sense against the background of an understanding of intentional action and self-consciousness being internally related to each other.

I would say similar things about illocutions and failures of understanding such as when one updates software. The law and social life generally create pressures to make such "agreements" mechanical and superficial (just as the expression of an apology or thanks can be empty rituals). But, even here, it still matters (legally and otherwise) that the user is being asked to *affirm that she has read and understood* the lengthy and incomprehensible document she is signing. We could not conclude from these examples that understanding is only contingently related to agreement.

Along with David Hills, Leite describes cases in which a person takes on a certain responsibility without clearly being aware that this is what they are doing. But here again, my claim in the book is not that deliberate, conscious acts were the only way to incur various responsibilities. Rather, I am interested in how we can understand what is special about what makes an utterance into a particular illocution by attending to what is self-conscious in the act and that this helps to understand the distinction between the illocutionary and the perlocutionary. So, in a sense, I can agree with Leite here: "Nothing about the idea of reconfiguring your normative relation to a hearer simply in speaking – thereby modifying your commitments and responsibilities to her and the forms of assessment and address from her to which you are thereby open – requires that you always have to know that you are doing it or what shifts in your relations are being thus enacted." Indeed: I can shift my normative relation to someone by inadvertently insulting them or by accidentally misleading them in some important way.

What is the importance of "self-conscious agency" as I see it? I assume we agree that people can pronounce words whose meaning they do not understand and that when they do, we do not count them as having asserted, claimed, or promised anything, similarly for words produced during sleep or hypnosis. In these situations, we can agree that the person is not participating in the practice of asserting or claiming and that none of the ordinary responses to a claim having been made apply here (such as "counter-assertion," agreeing or disagreeing with what was said, etc.). Similarly, for a young child to make promises or to forgive or apologize, it is not enough for her to

know the words, but she has to be initiated into the whole social practice and to understand something about the forms of human relation where these acts take place, the particular importance and consequences of promising or apologizing. This takes time.

If we do agree, then what accounts for the fact that we do not take the relevant illocutions to be performed in these circumstances? That is what I try to give an integrated account of in the book. In my remarks about “illocutionary authority,” I have sometimes invoked the difference between simply mentioning some words and using them to perform some speech act. But, in this context, it is important to remember that even the use/mention distinction can be abused. Today we are only too familiar with comedians and others playing with slurs or other inflammatory language, with the alibi that no one has really been maligned because, in the context, the topic is really just the word itself and not the people it supposedly applies to. We often have no trouble seeing right through such pretense. Often it is merely the brazen exploitation of how easily shifting contexts can be, and, in some situations where such language is being deployed, it is simply the enforced presence of the word in the discourse that is the harm being aimed at, inserting it into the discourse, and the difference between use and mention pales in significance [I take this up in my paper “Seeing and Believing: Metaphor, Image and Force,” in my collection *The Philosophical Imagination* (Oxford, 2017)]. And yet, none of this undermines the use/mention distinction itself, and there can be no critical or philosophical reflection on language and discourse without it.

Concerning Leite's example of the person “asking for help,” I am inclined to say what I say about the Raz case in response to Hills. I think it can be perfectly legitimate for the speaker to want to draw the line there and *not* actually ask for help, even while venting and complaining. And it can be a form of respect for the person's autonomy to leave that matter to her explicit presentation when she is good and ready to ask for help. (Naturally there is plenty of room for neurotic conflict and bad faith in these situations. The angry retort: “I never *asked* for your help!”) About such cases Leite writes: “These examples are not merely cases in which someone's mental states manifest themselves in their behavior, as when we can read someone's emotions in their expression and gestures. The speaker is doing something – criticizing, asking for help – and they are not aware that that is what they are doing.” I am not as confident as Leite seems to be about the simple “fact of the matter” here. Instead, characterizing what someone is doing as “asking for help” will often sound to me like the insistence a therapist or a partner will make for someone to *change* what they are doing, not simply to admit the supposed fact of what they are already doing but rather urging them to make the move from revealing their need and desire for help to actually *asking* for help. The very fact that the person is resisting this move is a reason to see it as a distinct form of expression and possibly as a therapeutically important step for them to take. If it were already straightforwardly the case that the person well and truly succeeded in asking for help, there would be no need for any such intervention. Often the right thing to say in such a situation is: “Why don't you just ask for my help? What's holding you back?” To simply insist that the plain fact of the matter is that the person really *was* asking for help all along is to miss the importance of the step they are currently unwilling or unable to make and thus what therapeutic intervention might help them to make it.

Leite's discussion of “Mr. K.” raises similar issues. First, I make it clear throughout the book, and in disagreement with other philosophers, that I think the notion of “verbal communication” is much wider than the particular form of intersubjectivity that I am focused on in the book. So I have no problem seeing this as a case of verbal communication. And I spend a good part of Chapter Six saying what I think is right about the idea (in Sperber and Wilson, among others) that there is a spectrum of cases from “showing” to “telling” and arguing why nonetheless this does not dissolve the distinctive importance of open, explicit telling. So we can talk about such cases as “forms of telling,” and I have no wish to legislate our choice of words here. But here is the problem: If we just go ahead and say that Mr. K. *told* his analyst that he “felt cast out by seeing the other patient yesterday,” then we miss the point of the analyst's intervention. For the meaning of the intervention was to help Mr. K. do something that he was not quite able or willing to do the previous day. That is what the analyst picked up on, and today he makes this intervention (“You're trying to let me know ...”) the point of which is to help the patient make a new move in their therapeutic relationship. There is a particular *need* of the patient being addressed here, and we would miss it if we simply said that he already *told* the analyst that he felt cast out.

Toward the end of his comments, Leite returns to the Casaubon case and raises some questions about my picture of sincere expression and the meaning of mutuality. Briefly, I do make a distinction between two senses of “expression” and claim that, in the “personal” sense, one may be said to express the belief that P or gratitude for a favor even in the absence of the belief or any feeling of gratitude. And I believe a person can speak sincerely while being self-deceived (that is one of the difficulties of dealing with self-deception, in oneself or in others). I take Leite's point that mutual recognition can fail in cases of explicit self-presentation, for the sorts of reasons that he subtly explores toward the end of his comments. His considerations open up new questions for thinking about the very meaning of “mutual recognition,” which I do not have the space to explore here. But I agree that “there are recognizable forms of fully intersubjective openness that involve interpersonally-directed non-linguistic emotional exchange.” I briefly discuss some of these in Chapter Six (responding to a joke or a metaphor or impressions of the seaside), but these will often be cases where there is no determinate thought or “mental state” in the offing and hence where I do not think we will be helped by the requirement of a match between expression and “the very mental state in question.” The Casaubon case is one of a failure or refusal of openness and explicitness, and I wished to understand the nature of the ordinary forms of open expression that are being refused there. At the same time I think it is clear that there is no point in a *generalized* demand for explicitness in human affairs. Another side of intimacy includes the desire to be understood without having to *say* anything, without having to spell everything out. In the book, I am interested in the equally everyday phenomenon whereby we take it upon ourselves to inform others by outright telling.

ENDNOTES

- ¹ As in the book, except where the context would make it confusing, I will tend to use feminine pronouns to refer to the first speaker and masculine ones for second speaker and interlocutor.
- ² Cavell, *The Claim of Reason* (Oxford, 1979), chapter 11, and Scanlon *What We Owe To Each Other* (Harvard, 1998), chapter 7.
- ³ ‘Transparent Pictures: On the Nature of Photographic Realism’, *Critical Inquiry* 11, no. 2: 246-77.

How to cite this article: Moran R. The Exchange of Words: Replies to critics. *Eur J Philos.* 2019;27:786–795.
<https://doi.org/10.1111/ejop.12491>