

Epistemic Territory

Presidential Address, APA Central Division

Denver, CO, February 22, 2019

Jennifer Nagel, University of Toronto

A Presidential Address, if I understand the genre, is an opportunity to make bold claims without the ordinarily cooling prospect of a question period. While it is a pleasure to be granted this freedom, I must confess some anxiety about exercising it, an anxiety heightened by recollection of previous addresses, not least last year's *tour de force* from Charles Mills. Going further back, I am haunted by the memory of a Presidential Address I heard 23 years ago in Pittsburgh, on a topic directly relevant to the one I have chosen today. In his address to the 1996 meeting of the Central Division in Pittsburgh, Michael Friedman launched a scathing attack on philosophical naturalism, where naturalism is understood as the thesis that empirical knowledge is our only kind of knowledge, and that 'philosophy, as a discipline, is (...) best understood as simply one more part—perhaps a peculiarly abstract and general part—of empirical natural science' (Friedman, 1997, 7). Friedman drew on the history of physics to develop a strong case for the idea that pure mathematics and philosophy have sharply distinct structural roles in the development of human knowledge, particularly at moments of 'fundamental conceptual transformation' such as the early modern period, and the dawn of the twentieth century (1997, 18). At these moments, Friedman argued, philosophy does not 'proceed in splendid isolation' from scientific developments, but somehow finds motivation from the course of these developments, while operating at quite another level (1997, 18).

Notwithstanding the power of Friedman's arguments, the question of how to conceive of the relationship between philosophy and natural science—or how to map their proper epistemic territories—remains somewhat unsettled, especially in branches of philosophy and science far from the relatively clean structure of physics. As an epistemologist, I am interested both in squarely philosophical questions about the nature of knowledge itself, and in some messy empirical questions about how knowledge is attributed, live and in real time, especially in contexts of social interaction. Following Friedman, I am convinced that these questions are not on the same level, but also that they cannot be investigated in splendid isolation from each other: there is something of great epistemological interest to be gained in studying the empirical details of what is happening in ordinary knowledge attribution. These empirical details are studied in a great range of different subdisciplines in psychology, linguistics, and sociology, and the epistemologist who becomes curious about how knowledge is actually attributed

will need to work hard even just to make sense of the disparate terminology and methods of these disciplines. While clarifying terms and methods might help one to summarize what is known empirically, such summaries will not on their own answer core philosophical questions about knowledge, not least because each of the relevant empirical disciplines brings heavy (and sometimes dubious) philosophical presuppositions to its study of knowledge attribution. However, if Friedman is right that philosophical progress can be motivated by the empirical without being subsumed by it, then it should be possible for philosophers to engage with empirical work on knowledge attribution in a suitably critical spirit. If we do not see philosophy as subservient to science, our aim can be not simply to uncover the philosophical commitments of the relevant empirical work, but also to criticize these commitments and see this empirical work, and knowledge itself, in a new light.

I take the expression ‘epistemic territory’ from the sociologist John Heritage, whose work in conversational analysis will form my empirical point of departure today. Conversational analysis studies verbal interactions of all types, including phone calls to distress hotlines, televised news interviews, doctor-patient exchanges, courtroom cross-examinations, and casual gossip. Particular attention is paid to aspects of interaction not commonly studied within traditional linguistics, such as the timing of switches between speakers, the frequency and nature of stumbles and repairs in speech, and the significance of interjections such as ‘oh’, ‘um’ and ‘uh’ (for an overview, see Sidnell & Stivers, 2012). Within conversational analysis, Heritage has been a pioneer in the subfield of ‘epistemics,’ which investigates the significance of what is known, and what is taken to be known, by the various participants in a conversation. Philosophers may be more accustomed to working with ideas from Grice (1975) or Stalnaker (2002) in this general area; however, there is value in picking up a somewhat different, if often overlapping, set of theoretical tools, and in exploring data that may be novel to at least some of us.

As an example, we could start with the conversational particle ‘oh’, which might at first seem to be an insignificant element in language. In his landmark 1984 article on this topic, Heritage observes that traditional linguists and philosophers have had little to say about ‘oh’. He quotes Charles Carpenter Fries, the author of *The Structure of English*, classifying it as a mere ‘signal of continued attention’, and elsewhere as a ‘noncommunicative’ utterance (Fries, 1952, as quoted in Heritage, 1984, p. 337); on the side of philosophy, Wittgenstein is cited as remarking that ‘oh’ is well described simply as a sigh (Wittgenstein, 1974, 67, as quoted in Heritage 1984, 337).

Heritage argues that these authors have underestimated ‘oh’. Finding an interesting set of patterns in its use, he contends that it functions as an epistemic change-of-state marker.¹ Heritage’s original (1984) idea was that ordinary conversational use of ‘oh’ is ‘used to propose that its producer has undergone

¹There are a few other functions served by ‘o(h)’, such as the poetic vocative use inherited from Greek and Latin; these are relatively rare in contemporary spoken English, and will be set aside in what follows (for an overview of the full variety of functions of ‘o(h)’, see Heritage 2018).

some kind of change in his or her locally current state of knowledge, information, orientation or awareness' (1984, 299). Subsequent elaboration makes it clear that this change has a positive character, and in what follows I will argue that 'oh' more specifically signals a moment of gaining knowledge, or regaining awareness of stored knowledge. This is a strengthening of Heritage's original line: in describing what 'oh' signals, he gives weaker disjunctive formulations, speaking of 'knowledge or information', rather than knowledge. I fear that in doing so he is lumping together the situations which truly underpin the function of this signal with situations in which the signal misfires. The tail flash of the white-tailed deer signifies danger, not some disjunction of danger-or-apparent-danger, even if in fact these animals suffer many false alarms. Suggesting that the core function of 'oh' is to mark knowledge gain is a friendly amendment; I hope to persuade you that this will actually be a better fit with Heritage's overall framework for understanding conversation, which is explicitly structured in terms of knowledge. I trust it will also be clear how much I am indebted to Timothy Williamson's (2000) knowledge-first epistemology in trying to make sense of this material.

In any event, the most striking use of 'oh' indicates surprise, or sudden and unexpected knowledge acquisition, and this may be what first comes to mind when we introspect on what it means, but most uses of 'oh' are more subdued. The changes in epistemic state marked by 'oh' can come either from an 'extra-conversational contingency' which is then taken up into talk ('oh, that T-shirt reminded me...'), to mark something 'sparked off' or brought to awareness by some feature of the conversation, or in response to being informed ('oh, ok'). In fact, sometimes what brought to current awareness is your own prior and independent knowledge of the point that your interlocutor has just made, which is why 'oh, I know' is such a strong form of agreement, above a generic agreement marker like 'yeah' (Heritage & Raymond, 2005). When produced after the response to a polar question, the 'oh' constitutes a kind of conversational receipt for what one has learned, a crucial success signal in the joint action that conversational partners are cooperatively undertaking (Heritage, 1984, 304). We will keep in mind that even if it functions as a signal of knowledge acquisition or revival, 'oh' can be triggered by something which only seems to the producer like knowledge, just as our use of other words (like 'diamond') can be triggered by lookalikes. It is also possible to be duplicitous with this signal, strategically presenting oneself as just now learning something one in fact knew all along. When a conversational partner seems excited to inform you of something that happens to be old news to you, rather than responding with 'yeah', and an elaboration which will make it clear that you are the more knowledgeable party on this topic, you have the option of saving face for them with some 'oh'-prefaced assessment ('oh dear', 'oh, wow'). The possibility of faking a naïve position in this manner is clearly parasitic on a general understanding of 'oh' as a marker of a positive change of epistemic state. Although it plays an important role in conversational dynamics, this tacit understanding is not necessarily available to introspection, which does something to explain why Fries and Wittgenstein could have advanced such dismissive explicit theories of 'oh'.

In my view, the norm to produce ‘oh’ as a knowledge acquisition marker is a member of a special class of norms, the class that Sarah Murray and Will Starr are describing when they write, ‘Many of the norms that govern our interactions are not principles we would endorse if asked about. They are heuristics of social cognition that we absorb from our social environment without being explicitly formulated or taught’ (Murray & Starr, 2018, 218). They argue that such norms ‘play a key and inadequately appreciated role in communication’, enabling coordination among agents with conflicting aims (2018, 219). They further argue that the coordinating function of communication between such agents should lie at the core of any theory of what is happening in conversation. Narrower theories such as Grice’s are problematic, they argue: if you are satisfied that communication is fundamentally about the mutual recognition of communicative intentions, you haven’t thought hard enough about what could ever motivate perfectly rational Griceans with possibly conflicting aims to express themselves, and get others to recognize their states of mind. According to Murray and Starr, norms like the Gricean maxims of Quantity and Manner—and here I would add the norm to produce ‘oh’ on acquiring or re-viving knowledge—are not followed because we are rational, but ‘because they are part of our cultural inheritance’; such norms are ‘self-fulfilling expectations about what agents like us [do] in particular circumstances’, ‘operating in the shadows of our unconscious minds’ (Murray & Starr, 2018, 219).

I think Murray and Starr are absolutely right to focus on the ends served by conversation, and to set themselves the task of explaining utterance force—what is happening when we assert, question or command—in a way that will ‘show that this is a stable and reproduced way of coordinating’ (Murray & Starr, 2018, 230). Their own theory argues that conversation should not be understood in the self-contained terms of updating what we are mutually assuming for the purposes of conversation, but as something that serves the function, and has been selected to serve the function, of coordinating the real-world beliefs and intentions of the interactants. As animals in the world, we can achieve much more by working together in a coordinated fashion, and synchronization of our intentions and beliefs in conversation could doubtless aid our joint action. However, if we are going to think about structure of conversation as being explained by its benefits, I wonder whether alignment of belief is the best we can do. Some beliefs are better than others, and we might worry about the desirability of a neutral drive towards alignment, in which we could all end up running in perfect synchrony with each other, but out of line with the world, like the proverbial lemmings off the cliff. So, I would like to pick up the spirit of Murray and Starr’s approach, but try to experiment instead with the idea that the epistemic side of human communication is fundamentally about achieving coordination through the sharing of knowledge. (There is a motivational side of communication as well, focused on coordination of desire or the pursuit of the good, but I will be setting it aside in what follows.) Shared knowledge is important because we ideally align not only with what other agents are thinking, but also with states of the world.

Returning to our particle ‘oh’, given how much we learn about the world

from talking to others, it is not surprising that we frequently acquire knowledge during conversations. It remains an interesting question why we so frequently broadcast this type of change, in the course of conversation, with an audible expression.

Sociologists situate ‘oh’ in what they call the ‘epistemic backchannel’ of conversation, along with a variety of other signals and devices. This channel is heavily used: studies of spoken English conversation suggest that roughly half of conversational turns are launched with ‘something besides a constituent of a grammatical unit’—most commonly ‘yeah’ and ‘oh’ (Norrick, 2009, 871). While these interjections initiate many utterances, they additionally occur in medial positions and as free-standing conversational turns. Today I will be using English-language examples, but I should mention that there is a large body of cross-linguistic work in this area. There are some points of deep similarity across languages; for example, in the generic repair initiator ‘huh?’ which is pronounced in very similar fashion globally (Dingemanse, Torreira, & Enfield, 2013). There are also interesting some points of difference: Finnish, for example, has several distinct particles that divide the work done by ‘oh’ in English (Koivisto, 2016). Notwithstanding these differences, it is widely understood that some common functions are served by the epistemic backchannel, for speakers of all languages. The backchannel, it will turn out, plays a key role in helping us track our relative epistemic territories, where these territorial calculations are a crucial determinant of what we are actually doing at any point in a conversation.

Looking at the history of conversational analysis, Heritage traces the theoretical development of the notion of ‘epistemic territory’ back to the 1970s (Heritage, 2012b, 4), starting with an observation made by William Labov and David Fanshel. They noted that asymmetries of knowledge sometimes have an interesting impact on the structure of conversation. In particular, curious effects arise in discussion of events to which only one party to the conversation has privileged access. In their terminology, where A and B are a conversational dyad, A-events are those known to A but not B (or as we will later say, in A’s territory and not B’s), and B-events are known to B but not A. (There are several other types of events, but they will not concern us at present.) With this distinction in hand, Labov and Fanshel formulate the following rule: ‘If A makes a statement about B-events, then it is heard as a request for confirmation’ (Labov & Fanshel, 1977, 100). One way they tested this rule was in a series of structured interviews with residents of New York, about their experiences of life in the city. When any subject mentioned a burglary, the interviewer’s script prompted them to say flatly, ‘And you never called the police.’ Even with both declarative syntax and falling intonation, subjects all responded to this ‘B-event’ statement as if it were a question (1977, 101).

The move of asking by declaring can be made in various ways. In her classic treatment of ‘fishing devices’, Anita Pomerantz distinguishes between two grades of knowledge, the direct (or ‘type 1’) knowledge that one typically has of one’s feelings, of events in which one was the ‘subject-actor’, and personal data such as one’s name or marital status, and the indirect (or ‘type 2’) knowledge that one has through inference and testimony, and through witnessing the

actions of others from some distance. To fish, speakers can make declarative statements expressing their type 2 knowledge to recipients with type 1 knowledge of the same topics: Pomerantz's examples include, 'You were in Room 252 for a long time this afternoon,' and 'She said that uh you guys were having a party Friday' (Pomerantz, 1980). The speaker's presentation of their lower grade of knowledge is treated as an occasion for the recipient to elaborate on this topic from their type 1 authoritative perspective. By fishing rather than questioning more directly through interrogative syntax or intonation, Pomerantz argues, the speaker demonstrates some respect for the privacy or authority of the more knowledgeable recipient who is spending unusual amounts of time in room 252 or apparently planning a party.

In his account of the development of the theory of epistemic territory, the last major figure that Heritage credits is the Japanese linguist Akio Kamio, who introduced the concept of 'territories of information'. Kamio replaces Pomerantz's two grades of knowledge with a continuous gradient from 0 to 1, where these numbers encode how deeply the topic is situated within one's territory; any given event can take any value on this line for each speaker and each hearer. Speakers deploy different linguistic forms to signal their relative territorial possession of a topic, relative both to the hearer(s) and to the upper and lower bounds of the scale. This theory aims to explain, for example, the conditions under which a speaker will say, 'Summer in Alaska is beautiful', as opposed to 'I hear summer in Alaska is beautiful.' Japanese has particularly rich resources for epistemic territorial marking, but Kamio argues that the same broad functions are performed in different ways English-language conversations as well, with subtle features of intonation contour playing a role in both languages alongside more conspicuous evidential markings in enabling speakers to signal their relative position at one point or another on the gradient (Kamio, 1995).

Heritage's own theory of epistemic territory adopts the basic idea of the gradient from Kamio, but with a qualitative rather than numerical scale, situating speakers relative to each other on a spectrum between more knowledgeable (K+) and less knowledgeable (K-) positions with respect to the topic at hand. Heritage distinguishes epistemic stance (what a speaker signals about their status, at a given moment in conversation) from the underlying epistemic status itself. As an example of how to stake out different stance positions on the gradient, Heritage (2012b, 6) contrasts three different way a doctor might ask her patient about his marital status, in the course of taking his medical history:

1. Are you married?
2. You're married, aren't you?
3. You're married.

As utterances directed to someone presumed to have authoritative knowledge—a person's own marital status is squarely within his instinctively recognized epistemic domain—these are all naturally heard as requests for information, but the first, with its straight interrogative syntax, strikes a more deeply

‘unknowing’ stance (a strong K-), in contrast to the second, where the tag question requests confirmation, indicating that the speaker takes the positive answer to be likely but not certain. The third strikes a milder still (K-) position, but will still serve to motivate some reinforcement on the part of the addressee. Of course, conversational participants can adopt an epistemic stance which differs from their actual status, including hostile practices of deliberate deception, but also innocent mistakes, and more constructive practices such as exam questioning; I will set these problem cases aside for the present with a promise to return to them.

Heritage argues that in order to understand what social actions our conversational partners are engaged in, we invariably need ‘a fine-grained grasp of epistemic domains and relative epistemic status within them.’ The question of whether another person is reminding you of something or asking you something, for example, cannot be resolved simply by whether this person is using interrogative syntax: with that syntax, you still need some grasp of the speaker’s epistemic domain in order to tell the difference between a rhetorical question and a genuine question. Without that syntax, as in the ‘You’re married’ example, a speaker might still be requesting information: indeed, a recent corpus study of information requests in English found that only a quarter of them were marked with *wh*-question words, only half of declarative questions had strongly rising intonation, and many had falling intonation (Stivers, 2010). Meanwhile, declarative syntax might or might not be heard as making an assertion, with rising intonation inherently ambiguous between questioning and continuing, and assigned to one of these categories in part by calculation of epistemic status.

The chronic need for calculation of epistemic status is very closely linked to what Heritage sees as one of the fundamental purposes of conversation, which is to pool the knowledge of the participants. Here is a core principle that supports the pooling of knowledge: whatever a speaker’s actual epistemic status might be on a given point, she can motivate conversation by taking an epistemic stance that contrasts her (presumed) epistemic position with that of her audience. There are two possibilities:

First, speakers can position themselves in a relatively unknowing (or K-) position relative to others concerning the matter at hand, thereby initiating sequences by inviting or eliciting information from a projectedly more knowing (or K+) recipient. Alternatively, knowing (K+) speakers can simply initiate talk concerning the matter at hand, thus launching a sequence, finding a warrant for this conduct by projecting their recipients to be in a relatively unknowing (K-) position. (Heritage, 2012a, 33)

K- initiations prompt replies. There a cooperative norm of answering questions, and this norm meets strong compliance in practice, with one recent study finding fewer than 5 percent of conversational questions going unanswered (Stivers, 2010), or fewer still if we include not only direct but also ‘transformative’ answers, which for example correct a mistaken presupposition (Stivers & Hayashi, 2010). K+ initiations also prompt interaction: the story-teller seeks receipts

and assessments ('oh my goodness') from his audience, and will struggle, repeat and refine his contribution if he does not find it (Bavelas, Coates, & Johnson, 2000). The dynamic here is subtle, because either type of initiation can prompt the addressee to reject the status projected upon them, for example, by responding to a K+ initiation with 'yeah', and an elaboration, reversing the presumed epistemic gradient by replying with more details of what has just been presented by the other party as news. But whether the epistemic stances taken are accepted or challenged, disparities in perceived status function as the 'hydraulic' engine of conversation: 'any turn that formulates a K+/K- imbalance between participants will warrant the production of talk that redresses the imbalance'; Heritage maintains that either type of gradient will start 'an epistemic seesaw motion that will tend to drive interactional sequences until a claim of equilibrium-for-all-practical-purposes is registered by the person who had previously assumed (or was assumed to be) the K- position' (Heritage, 2012a, 49). (Note the asymmetry here: it is the apparently K- person who must be satisfied that equilibrium has been reached, a point that will matter when I turn to discuss skepticism.) This satisfaction can be registered in various ways, down to a simple nod or, 'oh, ok'. Standing cooperative norms of conversation give us the sense that something is wrong if a question is left hanging, unanswered, or a reply (or story) is left unacknowledged. In summation, Heritage writes, 'giving and receiving information are normative warrants for talking, are monitored accordingly, and are kept track of minutely and publicly. It could, in principle, be different, but it is not' (Heritage, 2012a, 49).

Even if it is essential to meaningful communication that there is some flow of information between the parties, this flow could in principle exist without the running commentary on how it is going; in fact, for all nonhuman animal species, as far as I know, this is how it works. Other animals certainly share knowledge, alerting each other to everything from nectar sources to predators, but they do so without monitoring or broadcasting a public record of the shifting epistemic gradient between them. Other animals have superficially interactive duets and contact calling bouts, but not genuinely interactive conversations, where contributions are elicited and questioned, repair signals are sent if something is not quite heard, corrections are given, and confirmations or receipts are issued for successful transmission (Enfield, 2017, ch. 3, 9). The representation of epistemic gradient enables human beings to take an active part in managing the flow of information between them, for better and for worse. Representations of epistemic gradients enable correction and cooperation, but also tactical deception, an arena where humans vastly outperform other animals (Martin & Santos, 2014).

It is time to distinguish the better and worse cases. Note that if any represented epistemic gradient between speaker and audience motivates talk, then there are two satisfactory stopping points: both parties can move up to a shared K+ position or down to a shared K- position. As soon as we are in a position where both of us take ourselves to know the relevant fact, or a position where both of us agree that we are ignorant, then it is clear that there is nothing more to be gained by sharing what we take ourselves to know on this point in

conversation (although the position of recognizing our mutual ignorance could motivate other actions, such as efforts at inquiry and inference). Conversation can improve our epistemic position through either type of levelling: most obviously, when we move up to shared knowledge, but also when a speaker who had mistakenly taken herself to be in a K+ situation is corrected by a hearer, and, assuming that the hearer lacks knowledge on this point, both parties level down to a K- stance that accurately reflects their lack of knowledge.

Conversational levelling is not always a good thing, epistemically: where stance departs from status, we can worsen our status by aligning our stances. When people who are actually ignorant present themselves as knowledgeable, we can worsen our epistemic position by moving to accept what they say; or, if we start out knowing and are inappropriately challenged, we can lose knowledge by levelling down, which is presumably what's happening when we meet the gaslighter and the skeptic (more on them later).

If conversation is motivated by representations of epistemic gradients, then conversation's capacity to improve our epistemic position depends on the quality of those representations. We need some ability to gauge what others will know, especially in cases where their presumed knowledge differs from our own. One might expect the calculation of relative epistemic status to be computationally intractable, or at least extremely difficult in the ordinary course of conversation, but Heritage contends that we generally seem to solve this problem in real time: 'While it may be thought that the notion of epistemic territory introduces a contingency of daunting difficulty and complexity into the study of interaction, in fact relative access to particular epistemic domains is treated as a more or less settled matter in the large bulk of ordinary interaction' (Heritage, 2012b, 6). He suggests that this apparent general agreement on relative epistemic status is enabled in large measure by common fixed habits of territorial recognition, for example, expecting individuals to know more than outsiders about their own thoughts, plans, feelings, friends, jobs, pets, and families (*ibid.*), and to follow some common rules about the adjudication of evidence types and territorial disputes (Heritage & Raymond, 2005).

Indeed, it is possible to trace the core of our initial recognition of epistemic territory much further back, to the mechanisms for tracking gaze and perspective that we share with nonhuman primates, mechanisms which give us an immediate sense of what lies within another creature's field of view, and what is hidden from them (Bräuer, Call, & Tomasello, 2005). The possession of epistemic territory is not exclusive—for example, we can share visual access to what lies within our common field of view—but it is nevertheless helpful to track the different boundaries of privately held epistemic territory. In strategic games, apes and monkeys remember whether competitors do or do not know where food has been hidden: they selectively take advantage of ignorant competitors (Hare, Call, & Tomasello, 2001; Marticorena, Ruiz, Mukerji, Goddu, & Santos, 2011) and make appropriate inferences from the choices of competitors who they recognize as being in an epistemic position superior to their own (Kaminski, Call, & Tomasello, 2008). Mindreading has clear payoffs in anticipating the actions of competing animals over time, and in using those other animals' signs

of epistemic access as a guide to parts of the world that they cannot directly see for themselves.

I stress competition here, because nonhuman primates are surprisingly weak at tracking mental states in cooperative contexts. In a friendly situation where trainers are pointing out the location of hidden food, chimpanzees fail to distinguish knowledgeable from ignorant helpers (Povinelli, Rulf, & Bierschwale, 1994), despite the fact that they can draw these distinctions for similar rewards when competing (Kaminski et al., 2008). Mindreading capacities in competitive contexts are now well-established in experiments that aim to control for shallower explanations such as reading another animal’s outward behavior as a guide to action. Something deeper is already happening in primate recognition of epistemic territory.

Specifically human tracking of epistemic territory runs deeper still, because we track not only what lies in or out of the epistemic territory of others as compared to ourselves, but we also track how well they are tracking our territorial recognition. One of the defining features of human social intelligence is our capacity for joint attention, in which subjects attend to the same object in the awareness that this attention is socially shared. Other primates will often align on the same object, for example by following sightlines; only humans deliberately engage the eyes of another to direct attention onto an object, and look back at each other to check our mutual alignment. Human beings deploy their social intelligence not just in competing with each other, but also in cooperating. Michael Tomasello observes that this fact decisively changes the value of outward evidence of our mental states: human beings have much richer social intelligence because, in our capacity as cooperators, we need to broadcast those states. He puts the point as follows: ‘Whereas during competition individuals read the minds of their competitors against the competitor’s will (when we are competing, I want to conceal my mental states from you), in cooperation and coordination individuals want their partner to read their minds (when we are cooperating and coordinating, I do everything I can to display or advertise my mental states to you to facilitate the process)’ (Tomasello, 2018, 7). These advertisements are enabled in a number of ways, starting with anatomical differences between us and our closest animal relatives. We have evolved the most ‘readable’ eyes of any primate, with our contrasting white sclera and elongated eye shape making it exceptionally easy to see which way a human is looking, even at a distance (Kobayashi & Kohshima, 2001). Broadcast of epistemic states then continues up into conversational epistemic signals like ‘oh’.

These devices help us to calculate shifting epistemic territory in live conversation, enabling the social division of epistemic labour. Communal pooling of knowledge would be hampered both by widespread distrust and by complete credulity, so it makes sense that generally cooperative creatures like us instinctively attribute knowledge to each other concerning our local expert domains of our experiences, plans, pets, families and so on. However, these instincts of ours generate some initial difficulties we need to tackle, even before we tackle problems arising from rogue speakers who defect from the generally cooperative system to engage in deception.

Differences can arise between a sincere person’s actual epistemic territory (more simply, their knowledge) and their perceived epistemic territory (definable in terms of the set of questions whose answers the relevant audience will take on trust from this speaker, instinctively seeing the speaker as knowing). Sincere speakers can make innocent mistakes about their own plans and pets, and they can know—and want to communicate—things that go well beyond the safe zone of what other people will instinctively see them as knowing. I take the first kind of problem to be structurally similar to the problem we already have as individuals with fallible sensory systems: trusting an innocently mistaken speaker is not so different from trusting one’s innocently mistaken eyes. As for the second kind of problem, if a speaker wants to communicate some knowledge which her audience will not simply take on trust, she can turn to explicit argument, and she can do this even if what she knows was not originally formed on the basis of an argument. Much of what we know is unreflective, not formed on the basis of prior sequential consideration of evidence; but even unreflective perceptual judgments—for example, reports of having seen something surprising—can be defended argumentatively upon challenge.

My understanding of argument is partly shaped by the theory that Hugo Mercier and Dan Sperber have been advocating in their recent work, in which explicit reasoning has evolved as a solution to our cheap signalling problem (Mercier & Sperber, 2017). Unlike bees, who can only communicate an innately restricted set of facts about nectar source quantity, distance and direction, or peacocks, who can send the message that they are fit only by maintaining extremely costly plumage, we can send a signal saying almost anything at virtually no cost. But unless our signals are credible, there is no value in transmitting (or receiving) them. Mercier and Sperber argue that speakers and hearers who have the same instincts about what argument forms are valid can exploit those shared instincts to their mutual advantage (and they also produce extensive empirical evidence that even without any training, humans have strong instinctive capacities for logical reasoning, but these are optimized only in argumentative contexts). If speakers can show that their claims follow from premises already accepted by the hearers, or premises that will be accepted on trust, they can radically extend the scope of what they can successfully communicate (Mercier & Sperber, 2017). Like Murray and Starr, however, Mercier and Sperber focus on the value of belief coordination, or more precisely, on the value for a speaker of persuading the audience to accept what she believes, where this belief might or might not be true. Without denying that reason can be used to persuade someone of something falling short of knowledge, the existence of reasoning in our species would be better explained by its contribution to the expansion of shared knowledge, as opposed to shared belief. If the presence of reasoning in our species is to be explained by its advantages, the advantages of shared knowledge for both parties to communication are much more striking than the advantages of synchronized belief. Furthermore, mere belief synchrony might be achieved by devices much simpler than argument, such as deferral to the dominant animal.

Initial calculations of epistemic territory can guide me in deciding whether to

simply tell you something or work at persuading you, although I can also switch to argumentation when my effort at telling fails to secure uptake. My subsequent choice of premises is all about territorial calculation: I need to think about what you will know, and what you will take my word for. We think of argumentation in terms of territorial capture. Reasoning soundly from known premises can expand our jointly held actual epistemic territory, if you recognize my initial territory and follow my argument. Alternatively, I could engage in sophistry and reason on the basis of what I have calculated you will falsely perceive me as knowing, even to conclusions I know to be false, but there are risks in taking that path. If I need you to fool you with something that is going to look like knowledge, I am going to have to keep track of the lies I have told you to ensure their consistency, and to the extent that maximizing knowledge is the prosocial thing to do, I should fear the discovery of my campaign of deception, and your subsequent classification of me as a malevolent actor, which will be a problem if I want you to take my word for anything in the future, even on a matter well within the territory you would ordinarily assign to me instinctively (Clément, 2010; Sperber et al., 2010). If I try to expand my perceived epistemic territory illegitimately, through sophistry, and I am caught, I can end up shrinking my socially acknowledged territory below its original starting point.

The reduction of recognized territory is not always a bad thing. Unlike the white-tailed deer who are automatically startled into flight by every white flash, we have some flexible choice about the K+ initiations we hear from others, even when they are speaking of matters clearly within the territory we would ordinarily instinctively grant them. I might know that there's something strangely tricky about that thing that looks like a cup of water, right in your line of sight, and might thus be able to correct your false presupposition as you are about to drink from it. In these cases, cooperative groups can outperform isolated individuals in gaining knowledge, through a socially distributed safety network; we can to some extent cover each other's generally recognized territory, saving each other from cases in which we had only the appearance of knowledge, and not the real thing.

However, there are also cases of hostile reduction in perceived territory. The gaslighter gets into the evil business of (for example) denying our knowledge of states of affairs we have experienced (Abramson, 2014). To the extent that we can recognize their malevolent agenda—the people who assaulted you clearly have something to gain in denying your claims to know what happened—we can sometimes find a way of resisting gaslighters' efforts to 'tell us what really happened', their misrepresentation of the epistemic gradient (although I would not want to underestimate the power of another person's expressed doubts in making one begin to doubt oneself).

We turn now to skepticism. It may seem strange that I am dwelling on the cooperative mechanisms of oral conversation in laying the groundwork for my explanation of what is going wrong with us in philosophical skepticism. Surely Descartes confronted an admirably strong form of the skeptical problem while thinking in silence and isolation, in his stove-heated room. But it is striking that even Descartes creates the character of the evil genius to personify the

skeptical problem; he imagines fetching lines of dialogue for this character. It seems to me vastly easier to motivate skepticism by imagining being confronted by questions, following an effort to make some assertion about the world, rather than attempting to (for example) construct a pure monologue aimed at deriving the result that no one knows anything from some first principles about the nature of knowledge. Efforts in that latter direction seem to me more likely to raise doubts about those principles than to raise confidence in the conclusion of the attempted derivation.

In real (and imagined) conversations, the skeptic does something unusual, and apparently disinterested. If we have a cooperative system of gradient levelling, which generally functions to maximize knowledge among prosocial and largely epistemically competent creatures who can be assumed to want knowledge, the skeptic perversely chooses to run it in the opposite direction, as if his desire were to minimize knowledge. He exploits our usual cooperative conversational mechanisms throughout: if I start by claiming for example that there is water in this glass, he does not simply disregard this claim, or declare that knowledge is always impossible. The skeptic gets under my skin by granting me some patch of uncontested epistemic territory: for example, he grants that I know how things visually appear to me. That is a premise the skeptic can choose to take on trust, and it seems relevant to my original claim about this being a glass of water, so it could in principle work as the foundation of some argument we could pursue together. He is not just shutting down conversation in a way that would let us step out of the seesaw: we feel the usual hydraulic pull of the epistemic engine. The skeptic then poses the question of whether what seems to me to be water might not be another, visually indistinguishable liquid. By adopting this K- stance on the question of whether there is water in the glass, the skeptic has control of the conversation: his position will motivate further conversational exchanges until he grants that equilibrium has been reached. We recall the asymmetry of sequence completion: it is the person who had taken the K- position who must be satisfied.

Meanwhile, through his active line of questioning, the skeptic has also put me in a new frame of mind about the contents of the glass. In normal circumstances, I could have made the unreflective judgment that it contained water, and a normal interlocutor, sharing visual access to this scene and to my line of sight, would readily have made the same unreflective judgment and credited both of us with shared knowledge. When the skeptic refuses to grant my customary epistemic territory, I am stuck trying in vain to get us together into a K+ position, while restricted to using arguments that start from premises my conversational partner will accept as known. My ensuing failure of argument can start to feel like a failure of knowledge, not least because one generally good sign that something is known, that it actually falls in my territory, is that others are recognizing my territorial claim to it. It is in general a good feature of our cooperative knowledge-gaining system that we respect the rights of others to challenge our epistemic standing, not least because it is rarely in anyone's interest to refuse to take our knowledge on board, setting aside malevolent characters like the gaslighter and the coal-funded climate change denier.

Far from appearing as malevolent, the pure skeptic can come across as straightforwardly impractical, in a way that seems to radiate philosophical purity. He can even present himself as in an odd way pro-socially committed to a deep and very important project, the investigation of the nature of knowledge itself. If I am right, then the skeptic does actually bring to light an interesting feature of how knowledge is instinctively defended under pressure, but his methods of relentless questioning are not well designed to uncover the larger nature of what knowledge actually is. The skeptic is missing something important about how epistemic territory is initially gained and instinctively granted, even though he himself has to dole out these grants at least about appearances in order to keep the conversation going (and with his insatiable questions, he's a genius at keeping the conversation going). He is playing us, at a game we are instinctively driven to play, but as philosophers, we can step back and take a larger perspective on this whole game.

Today I have made some bold and under-argued claims that the natural function of human conversation is to share knowledge, rather than just to synchronize beliefs. It is tempting to end on a jolly note, praising the conversations we have as humans with each other, especially those at the APA, at the reception we are about to enjoy. But my parting words are going to take us in quite another direction. If we have a cognitive and cultural system whose natural function is to share knowledge, this does not guarantee that we will end up sharing knowledge. The unconscious drives supporting our epistemic trust and vigilance evolved in small-scale societies, where judgments were typically based on first-hand experience, and testimony was delivered face to face. We are now immersed in a complex media environment whose algorithm-driven dynamics are barely understood even by those who created it. We are like white-tailed deer who have wandered into a forest full of loudspeakers broadcasting startling noises and hunters who can spot us better every time we flash our tails, clicking on one reaction symbol or another. The same instincts which support knowledge gain in some environments could in other environments turn us into hostile and polarized teams of lemmings, increasingly detached from each other, and from reality. If we had some deeper epistemic self-understanding, we could perhaps do more to protect ourselves, as a species, from the new forms of epistemic pollution that are emerging globally. Just as learning to represent epistemic gradients opened up new possibilities for active knowledge sharing (and for deliberate deception), so also, moving up a level and learning to represent how we represent epistemic gradients also opens up new possibilities, both to protect our pursuit of knowledge and to threaten it like never before. I think it is an open question whether philosophers motivated to defend the pursuit of knowledge will develop the kind of epistemic self-understanding we need, in time to save our species. As President, I declare a state of emergency.

References:

- Abramson, K. (2014). Turning up the lights on gaslighting. *Philosophical Perspectives*, 28(1), 1-30.
- Bavelas, J. B., Coates, L., & Johnson, T. (2000). Listeners as co-narrators.

Journal of Personality and Social Psychology, 79(6), 941.

Bräuer, J., Call, J., & Tomasello, M. (2005). All great ape species follow gaze to distant locations and around barriers. *Journal of Comparative Psychology*, 119(2), 145-154.

Clément, F. (2010). To trust or not to trust? Children's social epistemology. *Review of Philosophy and Psychology*, 1(4), 531-549.

Dingemanse, M., Torreira, F., & Enfield, N. J. (2013). Is "Huh?" a universal word? Conversational infrastructure and the convergent evolution of linguistic items. *PloS one*, 8(11), e78273.

Enfield, N. J. (2017). *How we talk: The inner workings of conversation*. New York: Basic Books.

Friedman, M. (1997). Philosophical Naturalism. Paper presented at the Ninety-Fifth Annual Central Division Meeting of the American Philosophical Association in Pittsburgh, PA, on April 25, 1996 *Proceedings and Addresses of the American Philosophical Association* 71:2, 7-21.

Fries, C. C. (1952). *The Structure of English*. San Diego: Harcourt Brace.

Grice, H. P. (1975). Logic and conversation. *Syntax and Semantics*, 3, 41-58.

Hare, B., Call, J., & Tomasello, M. (2001). Do chimpanzees know what conspecifics know? *Animal Behaviour*, 61(1), 139-151.

Heritage, J. (1984). A change-of-state token and aspects of its sequential placement. In J. M. Atkinson & J. Heritage (Eds.), *Structures of Social Action: Studies in Conversation Analysis* (pp. 299-245): Cambridge University Press.

Heritage, J. (2012a). The epistemic engine: Sequence organization and territories of knowledge. *Research on Language & Social Interaction*, 45(1), 30-52.

Heritage, J. (2012b). Epistemics in action: Action formation and territories of knowledge. *Research on Language & Social Interaction*, 45(1), 1-29.

Heritage, J. (2018). Turn-initial particles in English: The Cases of 'Oh' and 'Well'. In J. Heritage & M.-L. Sorjonen (Eds.), *Between Turn and Sequence: Turn-Initial Particles Across Languages* (Vol. 31, pp. 155-190). Amsterdam: Benjamins.

Heritage, J., & Raymond, G. (2005). The terms of agreement: Indexing epistemic authority and subordination in talk-in-interaction. *Social Psychology Quarterly*, 68(1), 15-38.

Kaminski, J., Call, J., & Tomasello, M. (2008). Chimpanzees know what others know, but not what they believe. *Cognition*, 109(2), 224-234.

Kamio, A. (1995). Territory of information in English and Japanese and psychological utterances. *Journal of Pragmatics*, 24(3), 235-264.

Kobayashi, H., & Kohshima, S. (2001). Unique morphology of the human eye and its adaptive meaning: comparative studies on external morphology of the primate eye. *Journal of Human Evolution*, 40, 419-435.

Koivisto, A. (2016). Receipting information as newsworthy vs. responding to redirection: Finnish news particles *aijaa* and *aha* (a). *Journal of Pragmatics*, 104, 163-179.

Labov, W., & Fanshel, D. (1977). *Therapeutic discourse: Psychotherapy as conversation*. New York: Academic Press.

- Martcorena, D. C. W., Ruiz, A. M., Mukerji, C., Goddu, A., & Santos, L. R. (2011). Monkeys represent others' knowledge but not their beliefs. *Developmental Science*, 14(6), 1406-1416.
- Martin, A., & Santos, L. R. (2014). The origins of belief representation: Monkeys fail to automatically represent others' beliefs. *Cognition*, 130(3), 300-308.
- Mercier, H., & Sperber, D. (2017). *The Enigma of Reason*. Cambridge, MA: Harvard University Press.
- Murray, S. E., & Starr, W. B. (2018). Force and Conversational States. In D. Fogal, D. Harris, & M. Moss (Eds.), *New Work on Speech Acts* (pp. 202-236). Oxford: Oxford University Press.
- Norrick, N. R. (2009). Interjections as pragmatic markers. *Journal of Pragmatics*, 41(5), 866-891.
- Pomerantz, A. (1980). Telling my side: "Limited access" as a "fishing" device. *Sociological Inquiry*, 50(3-4), 186-198.
- Povinelli, D. J., Rulf, A. B., & Bierschwale, D. T. (1994). Absence of knowledge attribution and self-recognition in young chimpanzees (Pan troglodytes). *Journal of Comparative Psychology*, 108(1), 74-80.
- Sidnell, J., & Stivers, T. (2012). *The handbook of conversation analysis* (Vol. 121): John Wiley & Sons.
- Sperber, D., Clement, F., Heintz, C., Mascaró, O., Mercier, H., Origg, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, 25(4), 359-393.
- Stalnaker, R. (2002). Common Ground. *Linguistics and Philosophy*, 25(5-6), 701-721.
- Stivers, T. (2010). An overview of the question-response system in American English conversation. *Journal of Pragmatics*, 42(10), 2772-2781.
- Stivers, T., & Hayashi, M. (2010). Transformative answers: One way to resist a question's constraints. *Language in Society*, 39(1), 1-25.
- Tomasello, M. (2018). How children come to understand false beliefs: A shared intentionality account. *Proceedings of the National Academy of Sciences*, 201804761, 1-8.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wittgenstein, L. (1974). *Philosophical Grammar* (A. Kenny, Trans. R. Rhees Ed.). Oxford: Basil Blackwell.