Awareness, Apperception and Understanding

Catherine Z. Elgin

Abstract

Although neither knowledge nor understanding requires apperception, without it a subject is epistemically vulnerable. Apperception enables a subject to assess her epistemic commitments, opportunities, and obstacles. If an apperceptive subject takes an agential stance, she can correct and extend her commitments, fostering the enhancement of understanding. I argue that communities as well as individuals are capable of apperception and thus of epistemic agency.

Conditions on Reflective Endorsement

Without conscious experience we would not know or understand anything. So we might think. This claim is ambiguous. It might mean that if we were not conscious organisms, we would not know or understand anything. This is likely true, even trivially true. 'What is it like to be a turnip?' would probably be neither as long nor as interesting as 'What is it like to be a bat?' (see Nagel 1974). Or it might mean that each particular bit of knowledge or understanding must be grounded in or secured by conscious awareness of some specific datum that bears on the object of her awareness. This may seem plausible, but it is false. Because it is false, the question of the epistemic importance of conscious awareness needs to be investigated. I will begin by arguing for the falsity of the plausible claim. Then I will go on to discuss the relations between conscious awareness and understanding.

Let us begin with some terminology: In what follows I use the term 'cognizance' to refer to conscious awareness in general. The term 'conscious' is used to characterize mental states or

attitudes the agent is aware of having. Although I will be concerned only with contentful mental states – beliefs, desires, emotions, and the like – it is worth recognizing that the term 'conscious' as I use it also applies to moods one is aware of being in. 'Subliminal' refers to mental states or attitudes the agent has but is unaware of having. Repressed attitudes are subliminal. So are consequences – even obvious consequences – of a person's occurrent attitudes if she has never worked them out. The term also applies dispositional beliefs – like the belief that giraffes can't speak French – that she is unaware of harboring, having never had occasion to bring them to mind. Thus some currently subliminal attitudes can be easily brought to conscious awareness. The term 'apperception' refers reflective, reflexive awareness. Apperception is multi-faceted. It involves harboring a propositional attitude – belief, desire, suspicion, or whatever; awareness of the content of the attitude – that the cat is on the mat, for example; awareness that the attitude pertains to a particular propositional content – that her suspicion is of the cat's being on the mat; and awareness of oneself as harboring that attitude toward that content. This is roughly Leibniz's conception of apperception – 'reflective knowledge of [an] internal state'. (Leibniz, 1981, § 4, p. 208).

Knowledge

Four positions dominate current theory of knowledge. None holds that knowing that *p* requires conscious awareness of *p* or of one's attitude toward *p.*

Reliabilism insists that it does not. According to reliabilism, *S* knows that *p* just in case *S*'s belief that *p* is produced or sustained by a reliable process or mechanism (Goldman 1986). *S* need not be aware that her belief satisfies this requirement, that the particular process that secures *p*'s epistemic status is reliable, or even that she harbors the belief that *p.* That knowledge that *p* is consonant with such obliviousness is touted as an advantage of reliabilism. Because

cognizance is not required, the position readily accommodates perceptual knowledge, dispositional knowledge, and knowledge consisting of reliably produced or sustained subliminal beliefs. It also acknowledges that animals who are incapable of reflecting on their own mental states still know.

Evidentialism is cagier about the issue. According to evidentialism, *S* knows that *p* only if she has a non-fortuitously justified true belief that *p*, where satisfaction of the non-fortuitousness requirement insures that the belief's justification and its truth-maker align (Adler 2002). For our purposes, the precise connection between the justification and the truth-maker is irrelevant. What matters is that evidentialism does not require that *S* be aware she has the belief that *p,* aware that her belief is sufficiently justified, or aware that what justifies it is what makes it true. All that is required is that the conditions in fact be satisfied.

Virtue theoretic epistemology holds that for *S* to know that *p,* her belief that *p* must display her exercise of epistemic virtues (Sosa 1991). She need not be aware that she believes that *p.* Nor need she be aware that her belief is epistemically virtuous. Indeed, if epistemic humility is one of the relevant virtues, she probably should be unaware of its manifesting epistemic virtue.

According to *knowledge first epistemology,* an epistemic state is luminous if the agent is aware of having it. But luminosity is not required for knowledge. If *p* is not luminous, then even if *S* knows that *p,* her knowledge is opaque to her (Williamson 2000). She knows because she stands in an appropriate relation to the facts. Since knowledge is held to be more fundamental than belief, if she knows that *p,* questions about belief do not even arise.

That knowledge is consonant with obliviousness is unsurprising. Two familiar factors combine to explain the possibility of subliminal knowledge. The first is the widespread

recognition that the KK thesis is false. It is possible to know that *p* without knowing that one knows that *p*. Were this not so, to know that *p, S* would have to know that she knows that *p;* to know that she knows that *p,* she would have to know that she knows that she knows that *p;* and so forth ad infinitum. We are not that smart.

The second factor is that is common to have beliefs without being aware that one has them. This is our ordinary condition when mundane beliefs do not reach the threshold of conscious awareness. On 'automatic pilot' when her mind on other things, *S* stops at the traffic light, not even consciously registering that she does so. All four accounts of knowledge would claim that *S* knows that the light is red. Moreover, it is possible to have beliefs one sincerely denies having. According to psychotherapists, this occurs in cases of repression; according to Marxists, it occurs in cases of false consciousness. There is no reason to deny that of these beliefs sometimes satisfy the conditions on knowledge – that some patients who harbor repressed memories of a childhood trauma subliminally believe and thus know that they were abused; that some workers employed by a seemingly generous employer subliminally believe and thus know that they are exploited. If anything close to contemporary theories of knowledge is correct, knowledge does not require cognizance of one's doxastic or epistemic state.

Understanding

Turn now to understanding. Following Kvanvig (2003), let us distinguish between propositional and objectual understanding. Propositional understanding is expressed in a sentence like '*S* understands that *p*' or '*S* understands why *p',* where the direct object of the sentence expresses a proposition. It is understanding of an individual matter of fact. *S* understands that rattlesnakes are dangerous. *T* understands why her computer crashed. Objectual understanding is understanding of a topic or subject matter. *S* understands

biochemistry; *T* understands the theory of forms; *U* understands the candidate's campaign strategy.

Propositional understanding is a form of knowledge. *'S understands that p'* seems just another way of saying *'S knows that p'*. That being so, understanding that *p* no more requires cognizance than knowing that *p* does. What of understanding why *p*? It might seem that understanding why a proposition is true is sufficiently complex that cognizance would be required. Surely, one might think, an epistemic agent can't understand why the house is on fire without being consciously aware of the cause of the fire. But the issue is not so clear. Suppose a firefighter rushes into a house to extinguish a blaze. He does not turn on his fire hose; he sprints past an available water source; at considerable danger to himself he gets close enough to the fire to throw a flame-retardant blanket over the fire, depriving it of oxygen. Given his expertise, he immediately understands what sort of fire it is; he recognizes that it is a grease fire, which would be exacerbated rather than extinguished if doused with water.[1] He understands why that particular fire is burning. We should probably concede that he is consciously aware of the fire. Nevertheless, in extinguishing it he is on auto-pilot in much the way the driver who stopped at the light was. His understanding, it seems, is embedded in beliefs that do not, and in the circumstances need not, rise to the level of cognizance. People who regularly confront emergencies – first responders, battlefield medics, triage nurses, and the like – regularly display this sort of immediate, subliminal understanding of why an alarming incident is happening. Even if, as Grimm (2014) and others maintain, propositional understanding is knowledge of causes[2], and even if knowledge of causal connections is more complex than knowledge of (some)

1 This is a variant on Pritchard's example where the house burns down due to a short circuit. I do not know enough about short circuit fires to make my case using his example.

2 I disagree with Grimm about this, but I will not dispute the matter here.

brute facts, propositional understanding does not always involve cognizance of the facts one understands, or of one's epistemic state as a state of understanding them.

Objectual understanding is holistic. I have argued that an understanding of a subject is grasp of a systematically linked body of information in reflective equilibrium, where that body of information is grounded in fact, is duly responsive to evidence, and enables non-trivial inference and action regarding the phenomena the information pertains to (see Elgin 2017). This certainly seems complicated enough that cognizance would be required to achieve it. I think not. If we focus on cases like understanding biochemistry or understanding the theory of forms, it seems plausible that cognizance is mandatory. But if we consider a basketball player who adjusts his play to accommodate the moves of the opponent he is guarding, the idea that someone with objectual understanding must have real-time conscious awareness of the factors that constitute his understanding looks suspect. The player, it seems, viscerally understands his opponent's moves. Someone like the fire fighter or the basketball player who has merely subliminal understanding may be at a loss if asked, 'why are you doing that?' even though he is doing something that is manifestly responsive to the subtleties of the circumstances and effective in them. Objectual understanding can be embedded in know-how that resides below the threshold of cognizance.

This raises the question: What does conscious awareness add? Is an agent epistemically better off if she is cognizant of the phenomena her attitudes bear on? Is she epistemically better off if she is reflexively aware of her epistemic state vis à vis those phenomena? Here I will focus on objectual understanding. I suggest that to the extent that an agent's attitudes are subliminal, she is at their mercy. Even if she on objectively solid ground, she is subjectively vulnerable. She is in no position to recognize the strength or structure of her network of cognitive

commitments. Nor is she in a position to subject that network to assessment. After the fact, the firefighter could probably explain his action, grounding it in his understanding of the different profiles different sorts of fires display. He could, that is, answer the question 'why *did* you do that?' by bringing to conscious awareness something he tacitly drew on when acting. But if the action was grounded in a subliminal understanding of the fire, then at the time he acted, he would have been hard pressed to answer 'why *are you doing* that?'. Moreover, his explanation would likely be relatively sparse, failing to reflect the intricacies of the understanding that his action displayed. His accommodation to the vicissitudes of the fire may have been as fine-grained as those of the basketball player. Even if a belief or constellation of beliefs is sufficiently supported by evidence without conscious awareness, cognizance seems mandatory if one is to provide an explanation. Similarly, if one is to provide an assessment. Critical scrutiny requires cognizance of whatever one is scrutinizing.

Objectual Understanding

To make my case, I need to sketch my conception of objectual understanding. An understanding of a topic consists of a network of epistemic commitments in reflective equilibrium: its elements are reasonable in light of one another, and the system as a whole as reasonable as any available alternative in light of the epistemic agent's antecedent commitments. Such an equilibrium is labeled 'reflective' because together the two features make it a network of commitments that the agent can, on reflection accept.

When it comes to subliminal networks (or subliminal regions of networks) the issue is a bit tricky. The problem comes with the word 'can'. Clearly the agent cannot on reflection accept a consideration if it is inaccessible to her reflection. She cannot survey it, cannot assess it, cannot examine it at all. But if the criteria for reflective equilibrium are satisfied, she

understands the topic. We can accommodate this worry by introducing modal gloss to 'on reflection': if she could reflect on it, she would accept it. There is nothing wrong with her network of commitments per se; what is missing is her capacity to appropriately access that network. This explains why subliminal understanding is often effective. The elements of the network are mutually supportive, and the system as a whole is reasonable. Moreover, it explains why in epistemically resistant cases of false consciousness the network of commitments does not qualify as an understanding of the situation. In resistant cases, if the subjects could reflect on what they subliminally believe, they would not accept it. They would reject it in favor of the view they consciously hold that their employer is benevolent. This would be a type of confirmation bias.

Understanding is open-ended. A subject's fabric of commitments concerning a topic expands as it incorporates new propositions, new standards, or new methods. It deepens when she integrates new connections among commitments she already holds. A geometry student deepens her understanding of triangles when, having already accepted a proof of the Pythagorean theorem, she incorporates a new proof into her network of mathematical commitments. She understands the theorem's role in mathematics better when she sees how the new proof displays a different network of mathematical relations. Understanding is refined when, on the basis of the commitments the agent already accepts, she comes to draw finer distinctions, or redraw the boundaries between disparate elements. It is strengthened via contraction when problematic commitments are rescinded. Because it is fallible and open-ended, understanding is flexible and dynamic. It alters its constitution and contours over time as it accommodates or rejects new inputs. A network of commitments ought never be construed as complete.

Understanding can be tacit. To the extent that it is, the agent relies on it as a basis for

inference or action without appreciating that or why she does so. This, I suggested, leaves her vulnerable. Even though she changes her network of commitments in response to new inputs, adjusting it in the face of whatever impediments she encounters, she does so blindly. For all she can tell, the restored equilibrium – if it actually is an equilibrium – is unstable. Revisions are not always improvements. An adjustment that enables her to escape an immediate pitfall may give rise to more serious difficulties downstream. Moreover, if the agent cannot survey her network of commitments, she cannot entertain alternatives, spot vulnerabilities or opportunities, or make corrections.

The epistemic situation of someone whose understanding is subliminal is less than ideal. Unsurveyability results in epistemic insecurity, even if the agent is unaware of it. When her understanding is tacit, an agent cannot identify the strengths and weaknesses, opportunities and obstacles that figure in her thinking about the topic. She may be on safe ground, but she would be ill-advised to think she is, particularly if the stakes are high. Moreover, she cannot exploit her understanding. Inasmuch as the network is flexible and open-ended, it is responsive to new inputs. But she is in no position to seek out the sorts of inputs that would be fruitful, or to safeguard her understanding by preventing the incorporation of inputs that might lead her astray. She can either strike out blindly or wait for whatever happens to impinge on her cognitive system. She cannot leverage her current understanding to gain a better understanding.

Nevertheless, it is possible that certain strands in one's understanding should be, and should remain, subliminal. Perhaps networks or portions of networks consisting of what Gendler (2010) calls 'aliefs' are valuable precisely because they operate below the threshold of cognizance. Perhaps they can function as effective heuristics because they skirt the added burdens that come with cognizance. Jumping to conclusions is not always ill-advised. The

integration of aliefs into a network that contains beliefs may be valuable. It enables agents to deploy fast and frugal heuristics, which work well enough often enough. The mere fact that aleifs embed falsehoods is not an objection. Propositional elements in a network are acceptable if they are true enough (see Elgin 2017). The issue that concerns us here is not the fact that the contents of aliefs are apt to be false; it is that aliefs are apt to be subliminal. Subliminal commitments exact an epistemic cost. Perhaps they convey benefits that compensate for their cost. But they are risky. To the extent that the understanding of a topic is subliminal, it is not susceptible to the sort of intentional correction or improvement that conscious understanding is.

The Spectatorial Stance

Cognizant understanding of a topic is an explicit awareness of how the elements of a constellation of commitments relate to one another and how that constellation answers to the evidence that supports it. We can distinguish two stances a cognizant agent might take toward her understanding of a topic – the spectatorial and the agential stance.

The spectatorial stance is third personal. A spectatorial cognizer can apprehend the warp and weft of her tapestry of commitments. She can identify her various commitments, recognize their interconnections and lacunae, locate redundancies and idle wheels, as well as putatively relevant bits of data that do not seem to mesh. She can survey and explore the network. So far, this is purely descriptive. Her epistemic stance is like that of a cartographer, mapping the terrain, identifying and classifying what he finds.

She can do more. She is in a position to explain how and why various commitments hang together. She not only appreciates that the network constitutes a mutually supportive system, she can recognize what supports what, and why some seemingly relevant items are not included. Perhaps, for example, her survey reveals that, given the standards of evidence the network

incorporates, anecdotal evidence must be excluded. Then she can appreciate why what might have seemed like a relevant bit of information ought not be integrated into her understanding of a topic.

The spectatorial cognizer can be critical. She can assess her network, identifying gaps and flaws in her understanding. A constellation of mutually supportive commitments may be incomplete. Perhaps her current understanding of the topic enables her to see that there are questions she ought to be able to answer but cannot, or that there gaps in the information that she lacks the resources to fill. She may be able to identify errors – commitments that she previously incorporated into her evolving understanding that do not stand up to scrutiny. She can identify weaknesses. Perhaps some supporting links in her network are tenuous. So long as they hold, the network is in reflective equilibrium. But she may appreciate that her reasons for accepting them are relatively fragile. Perhaps, for example, her epistemic support for a commitment stems entirely from the testimony of a single source or evidence of a single experiment. That now strikes her as risky. She can discern vulnerabilities. If a proposition she has integrated into her network has a significant margin of error, she may see that despite the antecedent support, and despite its contribution to the network, it still a weak element. It does not contribute much. She can recognize obstacles. Perhaps there are limitations in her inferential commitments that prevent her from drawing certain inferences that she would like to be able to draw.

The fruits of her survey are not all bitter. She can also recognize strengths in her network. She can appreciate how and why it constitutes a solid take on the topic. It explains things she wants to explain, provides the resources for inferences she wants to make, draws lines between cases where it seems important to draw them. She can see how her current understanding supplies resources to extend her network of commitments, and to engage in

effective actions.

If the spectatorial cognizer had access to it, she could survey someone else's understanding of the subject in exactly the same way. She could identify the strengths and weaknesses, opportunities and obstacles to his understanding of the topic. He ought to be able to explain, infer, extend, exploit, and so forth. He is in a position to avoid these pitfalls, draw these inferences, extrapolate from those findings. He ought to shore up this commitment and revise that one. Inasmuch as the spectatorial stance is third-personal, this is not surprising.

The spectatorial stance is the stance of a film critic. A good critic describes the work in sufficient detail that the important features stand out. In effect, she maps the terrain of the work, pointing out the significant features. She may contextualize the work, perhaps locating it in a tradition, situating it in the director's oeuvre, identifying relevant artistic conversations it is a part of, and/or relevant socio-cultural influences that it incorporates and trends it reacts against. She identifies its the strengths and weaknesses. She says, for example, that the film had strong characters but a weak plot. She maintains that it was a mistake to cast an actor whose metier is comedy in a wholly serious role. She points out that the use of flashbacks was a brilliant way to overcome the obstacles posed by the passage of time. She is not impressed by the lighting, and thinks that some of the camera angles were ill-advised. She regrets the missed opportunities which could have been exploited. In doing this, she does more than say 'thumbs up' or 'thumbs down'. She enables her viewers to better understand the film.

Although she can provide a judicious description and assessment of the film, she is powerless to correct the flaws or exploit the opportunities she finds. Even if she is right that the film would have been better if the fade outs had been more gradual, and even if she is right that the director could easily have instructed the camera man to make them more gradual, she is still

an on-looker. She is in no position to intervene. Nor is this only because the film is already complete before she sees it. If she were a drama critic reviewing a play that was still running, the most she could do is make recommendations to the director, the actors, the lighting crew, or stage manager. She could not make the corrections she advocates. This is the fundamental limitation of the spectatorial stance. It can recognize where and how things might be better; it can provide cogent reasons for the verdicts it reaches. But as a third-personal, outsider stance, it cannot implement the needed changes.

An anecdote of John Perry's brings this out. He says,

> I once followed a trail of sugar on a supermarket floor, pushing my cart down the aisle on one side and back the aisle on the other, seeking the shopper with the torn sack, to tell him he was making a mess. With each trip around the counter the trail became thicker. But I seemed unable to catch up. Finally it dawned on me. I was the shopper I was trying to catch. I believed at the outset that the shopper with the torn sack was making a mess. And I was right. But I didn't believe that I was making a mess. . . . When I came to believe that, I stopped following the trail around the counter, and rearranged the torn sack in my cart. (1979:3).

Perry's worry is not mine. He is concerned with essentially indexical beliefs. But the situation he describes fits my purposes as well. Initially, Perry had a largely spectatorial view of the mess-making. He thought that the trail of sugar was likely caused by a shopper with a leaky sack in the cart. He inferred that the shopper would probably not want to be spilling sugar and that the shopper would therefore want to know that he was making a mess. He did not follow the trail merely because he wanted to know where it led, but because he sought to intervene to prevent further mess-making. To that limited extent his stance was agential. But because Perry

did not know that he himself was making the mess, the best he could hope to do was to catch up with the mess-maker and recommend that he adjust the leaky sack. In this he was like the drama critic who makes recommendations about a play that is still running. It was only when Perry realized that he himself was making the mess that he could reposition the torn sack so that it no longer spilled sugar onto the floor. That is, once he recognized that he was the mess-maker, having taken an agential stance, he could resolve the problem.

The Agential Stance

Agential understanding is fundamentally first-personal. An epistemic subject whose stance is agential is in a position to take responsibility for her networks of cognitive commitments. Like the spectatorial cognizer, she can survey and assess her understanding of a topic. She can identify its strengths and weaknesses, recognize the opportunities and obstacles to further understanding it presents, identify errors and vulnerabilities embedded in it. But unlike the purely spectatorial cognizer, she can do something about what she finds. She can amend, rescind, or strengthen commitments that she now considers under-supported or ill-advised. There is, of course, no guarantee that she will thereby improve her understanding. That depends on the circumstances she is in and the changes she chooses to make. But she is capable of making intentional changes in response to perceived defects, and taking advantage of what she sees as opportunities. If, for example, she recognizes that a consideration or technique that strengthens support in one area could be effectively used to shore up another, she can import it.

She can fill gaps in her understanding. Recognizing that her mastery of ancient Greek does not enable her to decide whether 'episteme' means 'knowledge' or 'understanding', she can run through her memory of actual uses of the term and see which translation fits best. Failing that, she can run through her list of acquaintances with greater expertise in Greek and see what

they think.

Rather than just waiting for bits of data to conveniently coalesce into an epistemically useful package, she can marshal evidence. She can bring together seemingly diverse considerations to support a hypothesis. Her evidence may be scattered across her network of commitments. Professor White wonders why Bill, whom she considers an excellent student, performed so badly on the final exam. She knows that he had been ill, and had missed several lectures. She recognizes that he is proud, so it is unsurprising that he would be loath to ask for help. She heard that because his car broke down, he had to work extra hours to pay for repairs. As long as these bits of information are unconnected, she finds his performance mysterious. But once she puts them together, they coalesce into a pattern that enables her to see why he did badly (see Lewis 1983).

The agential cognizer can not only recognize tensions in her commitments, she can take steps to alleviate them. On surveying her political commitments, Jess comes to recognize that her commitment to freedom of expression are in tension with her views about hate speech. She thinks that people should be free to express their opinions, and also thinks that people should not be subject to hateful, derogatory rants criticizing them for their race, gender, sexual orientation, or religion. If she were a purely spectatorial cognizer, she would simply conclude that the tension between her commitments is a regrettable weakness in her understanding of political rights. As an agential cognizer, however, she can and should ask herself: 'What are you going to do about this tension?' She can assess her resources for resolving the dilemma. If she thinks that she currently has no acceptable way to relieve the tension, she can consider where she might go to gain additional resources. Should she read more about the issue? If so, what? Should she consult with others? If so, who? An agential cognizer thus can, in principle, recognize the limits

of her current understanding, and make moves directed to extending her epistemic range.

Know Thyself (or anyway understanding thyself)

Thus far, I have discussed the ways cognizance enables an agent to fruitfully engage with her network of epistemic commitments. This is not the whole story, for cognizance is Janus-faced. To be able to exploit the resources her understanding provides, an agent needs to be self-aware. She needs to recognize that she is capable of taking steps to improve her epistemic lot. This involves recognizing that she herself is in circumstances where certain actions are called for, and recognizing that she herself is in a position to undertake these actions. The 'she herself' locution is designed to underscore that her stance toward her situation must be first-personal. It cannot be the disinterested perspective of thinking merely that someone should do something about the problem. The epistemic agent qua agent requires apperception.

Think back to Perry's predicament. He was cognizant of the fact that someone was making a mess and consciously supposed that the perpetrator probably did not want to make a mess. He realized that he suspected that the perpetrator, if alerted, would take pains to cease his mess making. That is, Perry was aware of his own views about the spilled sugar and the sugar spiller. He also recognized that he himself was aware of mess making in his vicinity – aware that, for example, he was not watching the mess making on zoom. He recognized that he himself was in a position to alert the perpetrator. He was not paralyzed or mute. And he thought that a good way to find the perpetrator was for him, himself to follow the sugar trail. In order to undertake the task, he had to be cognizant of his own situation, with its affordances, powers, and limitations, not just the perpetrator's.

To be in a position to exploit her resources, the cognizant agent should also be able to take a normative stance toward her network and her resources for improving it. I said earlier that

unlike the cognizant spectator, the cognizant agent can do more than assess her commitments, as a film critic would; she can also take steps that she has reason to think will improve her epistemic lot. Both a spectator and an agent have the resources to recognize that their network includes an epistemic infelicity – perhaps invalid inference license that permits affirming the consequence. Both can recognize that this is a defect. But the cognizant agent is in a position to devise and test remedies. To do that, she needs not only the capacity to be critical, but to be self-critical. She needs to be able to take a normative stance not just to her understanding, but to her self as someone who understands.

What a person understands and how she understands it cannot be prized apart. What an agent considers established depends on what she considers evidence, what she counts as defeaters, how much evidence she takes to be required, and so forth. So her understanding of her epistemic norms infiltrates her understanding of a topic. The growth of understanding of a topic goes hand in hand with a growth in understanding how to understand a topic, and a refinement in one's views about what the understanding of that sort of topic consists in.

Apperception is far more fine-tuned than this description suggests. It involves an appreciation not only of one's take on things, but also of the scope and limits of one's agency. Perry, following the sugar spiller on foot, had different options from those available to Jerry, who was tracking the miscreant via a surveillance camera. To take advantage of his options, an epistemic agent needs to be aware of what he can do in the circumstances. If he is to act responsibly to improve his understanding, he needs to be aware of what alternatives it would be epistemically and morally responsible to pursue.

Collective Apperception

So far, this paper has focused on an individual epistemic agent. But elsewhere I have

argued that understanding is in large measure a collective epistemic achievement (see Elgin 2017). This raises a question: Are communities of inquiry capable of apperception? One might think that the answer must be 'no'. Cognizance, as I described it, is a matter of conscious awareness. Since apperception requires cognizance, it might seem that to say that communities are capable of apperception requires saying that groups are conscious – that is, that they have phenomenal states. This is something we should avoid. 'The mood of the crowd' should remain a metaphor.

Although it is plausible that an individual's conscious awareness has a phenomenal element, what is doing the epistemic work is not phenomenal. Whether or not there is a felt quality to conscious awareness makes no epistemic difference. Such awareness is epistemically valuable for the individual because it highlights aspects or regions of her network of commitments, making them available for scrutiny and assessment. But there is no reason to think that only a state with phenomenal properties can do that. My conception of cognizance thus extends beyond the phenomenal. It consists in a sensitivity that renders what it registers available for scrutiny and assessment by the entity that registers it. Thus the fact that a community does not have phenomenal states over and above those of its members does not preclude the community's being cognizant. Nor does it preclude the community's being cognizant of things that no individual member is.[3]

I have argued that a community of inquiry is a realm of epistemic ends (see Elgin 2017). Its members constitute their joint practice by making and reflectively endorsing the rules that bind them and devising the procedures that define and circumscribe their collective epistemic ends. Because they consent as legislating members of a realm of epistemic ends, each member

3 I suspect that this happens in scientific communities where there is a considerable division of cognitive labor. But here my point is simply that the possibility is not ruled out by my conception of cognizance.

recognizes that the consent is mutual. Each one's consent is grounded in the recognition that the others also consent, and each recognizes her epistemic responsibilities to her compatriots. Moreover, because the practice is devised precisely to pursue certain ends, the members design it and its constituent commitments to be subject to public scrutiny, correction, and development. It thus makes sense to say that if the community is a realm of epistemic ends, it is consciously aware of the commitments that constitute it, and that constrain their practice. Its awareness is not always complete. Some of its commitments may be subliminal. This is how Marxists describe false consciousness. The workers constitute a community that collectively fails to recognize that they are being exploited. Arguably, scientists who endorsed Newtonian physics constituted a community that subliminally assumed that mass is independent of acceleration. A practice is grounded in subliminal commitments is vulnerable. The commitments may be mistaken, or limited in scope. Once they are identified, they can be examined for tenability.

The constitutive commitments of an epistemic practice are susceptible to scrutiny and assessment. The community that constitutes the practice evaluates its ends and means, correcting, amending, and extending it as it sees fit. Because the community is capable of joint reflective endorsement of its commitments, I suggest, it exhibits collective apperception. Its apperception, like that of an individual agent, is Janus-faced. The community reflects on its epistemic domain, and on itself as capable of generating an understanding of the domain, adjusting its understanding of each in light of the other. The more it learns about the domain, the more it learns about its capacity to fruitfully investigate that domain and the more refined its understanding of the domain. Thus, for example, a scientific community that once depended on unaided observation, learns to create devices like telescopes and microscopes that extend its epistemic range. It learns how such devices can distort and how to protect against distortions. It

learns to understand how those devices extend its range, and that in turn equips it to further extend its range – to devise, for example, scanners that augment perception rather than merely sharpen it. Depending on what those devices reveal, it learns more about what sorts of detectors might enable it to further extend its range. The process is familiar and uncontroversial. A community of inquiry advances its understanding and self-understanding in tandem. In so doing it exhibits and relies on collective apperception.

That collective stance is agential.[4] The practice involves self-monitoring and self-correction. The responsibility to enforce the epistemic norms, and extend or correct the commitments may be distributed. Distributed cognition is commonplace. There is no suggestion that all members of the community play the same epistemic role. Thus, for example, first-order norms – such as those bearing on experimental design – may be the province of a particular individual or a single lab. Within a single lab investigators with different areas of expertise, may make different contributions, and hold themselves answerable to different standards. One group assesses the experimental design; another runs the analysis; a third takes charge of the calculations. The deliverances of the lab are assessed by peer review. The responsibility to assess and correct the standards for peer review may be distributed more broadly. Revisions of standards may require deliberation across a wider swath of the community. Should hypotheses be preregistered? Should journal articles be blind reviewed? Should the review be double blind? Etc. The exact distribution of responsibilities is apt to vary from one community to the next. But because a community of inquiry polices its own epistemic activities and products, its stance is agential. And because it recognizes, assesses, and publicly acknowledges the basis on which it does so, the community manifests collective apperception.

4 I am grateful to Kareem Khalifa for pushing me to acknowledge this.

A different community might take a spectatorial stance toward the same material. This is what we see in the history, sociology, and philosophy of science. A historical or sociological stance is descriptive. It maps commitments, showing connections and dependences inherent in the scientific practice it investigates. A philosophical stance adds a normative dimension – disclosing errors, faults, strengths, and weaknesses. But like the film critic, the philosopher functioning as such could not implement corrections. To do that, she would have to become a member of the community whose work she studies.

Conclusion

Although knowledge and understanding may be subliminal, I have argued that the intentional, epistemically responsible advancement of understanding requires conscious awareness of the constitution and contours of one's network of epistemic commitments. It also requires conscious awareness of one's own powers, limitations, affordances, and resources – that is, recognizing oneself as an epistemic agent engaged with a particular topic. Epistemic achievement is not an endowment; it must be earned. Apperception involves recognizing oneself as an epistemic agent and being cognizant of the constitution of one's network of commitments provide a basis for the expansion and revision of that understanding. But whether the additions and revisions constitute advances, declines, or plateaus in understanding does not depend on apperception and cognizance alone. It depends on whether the agent uses the insights they afford well.

References

Adler, Jonathan (2002). *Belief's Own Ethics*. Cambridge MA: MIT Press.

Elgin, Catherine Z. (2017). *True Enough.* Cambridge MA: MIT Press.

Elgin, Catherine Z. (1996). *Considered Judgment.* Princeton: Princeton University Press.

Gendler, Tamar (2010). 'Alief and Belief' *Intuition, Imagination, and Philosophical Methodology.* Oxford: Oxford University Press, pp. 255-281.

Goldman, Alvin (1986). *Epistemology and Cognition*. Cambridge MA: Harvard University Press.

Grimm, Stephen (2014). 'Understanding as Knowledge of Causes' *Virtue Epistemology Naturalized: Bridges Between Virtue Epistemology and Philosophy of science (Synthese Library),* ed. A. Fairweather. New York: Springer, pp. 329-346.

Kvanvig, Jonathan (2003). *The Value of Knowledge and the Pursuit of Understanding.* Cambridge UK: Cambridge University Press.

Leibniz, Gottfried (1989). 'Principles of Nature and Grace' *Philosophical Essays,* ed. R. Ariew and D. Garber. Indianapolis: Hackett Publishing Co.

Lewis, David (1983). 'Truth in Fiction: Postscript' *Philosophical Papers* volume 1, Oxford: Oxford University Press, pp. 276-208.

Nagel, Thomas (1974). 'What is it Like to be a Bat?' *Philosophical Review* 83:435-450.

Perry, John (1979). 'The Essential Indexical' *Nous* 13:3-21.Sosa, Ernest (1991). *Knowledge in Perspective.* Cambridge UK: Cambridge University Press.

Sosa, Ernest (1991). *Knowledge in Perspective.* Cambridge UK: Cambridge University Press.

Williamson, Timothy (2000). *Knowledge and Its Limits.* Oxford: Oxford University Press.