Synthese (2021) 199:1577–1597
https://doi.org/10.1007/s11229-020-02819-5

Replies

Catherine Z. Elgin[1]

Published online: 14 August 2020

I'd like to begin by thanking Christoph Jäger and Federica Malfatti for editing this collection. An enormous amount of work was involved. I'm grateful that they were willing to do it, and honored that they thought it was worth doing. I'd also like to thank the editors of *Synthese* for agreeing to publish this collection of papers. Finally, I'd like to thank the contributors to this volume. I expected criticism of my position. I got it. But, to my surprise, writers also came up with a variety of ways to extend my ideas into areas I had never considered. I'm delighted to have prompted so many promising offshoots.

Several of the papers take up my anti-veritistic stance. I will begin by discussing them, then go on to discuss papers on other themes. Following Goldman (1999), I take veritism to hold that truth is necessary for epistemic acceptability. As Frigg and Nguyen point out, *True Enough* focuses on literal truth. I contend that if only literal truths are acceptable, epistemology cannot accommodate scientific understanding. Because the strength of my position and of objections to it depend on the conception of truth at issue, it might be worth making a few preliminary remarks about truth.

I have no monopoly on the concept of truth. All I can do is explain and perhaps justify the conception I had in mind. I do not think there is anything particularly unusual about that conception. But there are alternatives. Some of my critics advocate them.

Whatever the correct theory of truth, it must support disquotation. Tarski's T-sentences, such as

'Snow is white' is true ≡ snow is white

must come out true (see Tarski 1956). Somehow, the liar paradox must be avoided. These are standard requirements. I take truth to be bivalent. Any truth-apt representation that is not true is false. So under my conception, partial truths and approximate truths being not entirely true are strictly false.

✉ Catherine Z. Elgin
Catherine_Elgin@Harvard.edu

1 Harvard University, Cambridge, MA, USA

This, as Frigg and Nguyen point out, does not exclude figurative truths. Indeed, I have argued (see Elgin 1988), that metaphorical statements can be true under their metaphorical interpretations.

‘Love is blind’ is true ≡ love is blind

satisfies Convention (T) as much as

‘Snow is white is true ≡ snow is white

does.

Still, literal truth has epistemically attractive features that distinguish it from figurative truth, as well as from approximate truth, partial truth, and felicitous falsehood. We should acknowledge them, and I suggest, not lose them in a melange resulting from what Warenski characterizes as a more ecumenical conception of truth. Some stem from propositional logic. The conjunction of any two literal truths is, as Dellsén emphasizes, literally true. A valid inference with true premises is truth preserving. But the conjunction of two figurative truths may fail to be figuratively true. And the conjunction of two felicitous falsehoods may fail to be felicitous. According to Frigg and Nguyen's account, ‘The nucleus of an atom is a liquid drop’ is figuratively true; so is ‘The nucleus of an atom has a rigid shell’. But it is at least hard to see how ‘The nucleus of an atom is a liquid drop and the nucleus of an atom has a rigid shell’ could be figuratively true. The properties the first conjunct ascribes to the atom are antithetical to those the second conjunct ascribes. Nor is it evident how the conjunction affords an approximation to the truth about the shape of the nucleus. Pretty clearly, the conjunction is infelicitous.

Eternal sentences—tenseless sentences that contain no indexicals—retain their truth values permanently, and can be exported to new contexts without changing truth values. Tensed sentences can be converted to tenseless ones by substituting tenseless verbs combined with exact dates and times (sometimes augmented with ‘before’, ‘after’, ‘at the same time as’ and so forth). And, apart from sentences with essential indexicals (see Perry 1979), sentences containing indexical expressions can be converted to sentences without them by replacing indexical expressions with co-referential, non-indexical terms. So with a little syntactical maneuvering, such sentences can be converted to eternal sentences as well. There is no expectation, of course, that exported sentences will retain the relevance they had in the original context. But their truth values will be preserved. Figurative truths, approximate truths, partial truths and the like are not freely exportable. They tend to be contextually circumscribed. For certain purposes, and within certain limits, it is reasonable to represent the nucleus as a liquid drop. For other purposes, and outside those limits it often is not.

Because of the valuable epistemic affordances that the foregoing logical and linguistic features provide, I consider it inadvisable to loosen the restrictions on truth, or to blur the distinction between literal and figurative truth. But to say that literal truth has a distinctive profile that epistemology should acknowledge is not to say that literal truth has overarching epistemic importance. Quite the contrary. I argue that a variety of representations that are not literally true are integral to scientific understanding, as well as understanding of other sorts. Hence, I maintain that many approximations, partial truths, figurative truths, and the like are felicitous falsehoods. So are some models that are far more distant from the truth. Because they are not wholly, literally true, I consider them falsehoods. Because the respects in which they diverge from literal truth or the degrees to which they diverge from truth are negligible, I consider them felicitous.

I do not contend that epistemic agents do or should *believe* felicitous falsehoods. Working scientists are well aware that their models are not accurate. Still, they consider the models acceptable. In my view epistemic acceptability is not a matter of belief. To accept a proposition is being willing and able to use it in inference or action when one's ends are cognitive (see Cohen 1992). It is, as Dellsén says, a policy. I turn now to individual papers.

1 *Mirrors Without Warnings* by Roman Frigg and James Nguyen

In *True Enough* I argue that scientific models cannot plausibly be construed as true. Roman Frigg and James Nguyen correctly point out that I should have said that they cannot plausibly be construed as *literally* true. Models might nevertheless be true, if construed figuratively. Moreover, I should have considered whether figurative truth might do the job that I claim felicitous falsehoods do. They think it can. They say that by construing models as figuratively true they can 'retain much of the philosophy of science (and epistemology) based on veritism.' That has to be shown. As I've indicated, the logic of figurative truth is more complicated than the logic of literal truth. Confirmation, induction, and theory choice are likely to be equally complicated. Another concern is this: on their view, *Pride and Prejudice, The Illiad,* even *Winnie the Pooh* should be considered true. They are, after all, figuratively true. I am happy with this result, but I suspect that hard-nosed veritists and scientific realists will balk.

In Frigg and Nguyen's DEKI model, the 'K' stands for 'key'. A model, like a map, they maintain is accompanied by a key to its interpretation. My position is close to theirs, but I contend that the model exemplifies a feature that is to be projected onto the target. It highlights that feature. I do not require that the model provide a key. Frigg and Nguyen acknowledge that models do not and need not always provide explicit keys. Still, they maintain, the keys are implicit. The way models in their discipline relate to their targets is something students internalize during their professional development. To a considerable extent, this is surely right. When we learn to read maps, we learn that unless the key explicitly indicates otherwise the top of the map represents north. This convention is so standardized that it would be otiose for each map to indicate it. Even so, I think that counting on the key to secure a proper interpretation is a bit problematic. For models sometimes reveal something that was utterly unanticipated. Competence with models is not exhausted by knowing how to make the correlations that the key (whether tacit or explicit) specifies. A model can be a vehicle for discovery when it exemplifies something about the target that we had not expected to find.

A very important point that Frigg and Nguyen make, almost in passing, is that more work needs to be done on the nature of scientific representation—indeed, I would say, the nature of representation in general. Criteria for representational correctness need to be identified and justified. Types of representation and the affordances they provide need to be differentiated. It is not obvious that all representational correctness is best cashed out as some sort of truth.

2 *Veritism Refuted? Understanding, Idealization, and the Facts* by Tamer Newar

I do not claim to have refuted veritism. My arguments seek to establish that it is excessively epistemically costly, since it fails to do justice to scientific understanding. That's far from a refutation. Tamer Newar proposes several ways that a veritist might reply to my challenge. As I said above, the conception of veritism I rely on is Goldman's (1986, 1999). He maintains that 'people's dominant epistemic goal is to obtain true belief, plain and simple' (1999, p. 24). This, I contend, is at odds with science, as well as with other disciplines willing to sacrifice truths for systematicity, explanatory power, inferential power, and so forth.

As I said above, in my view, a theory of objectual understanding should take acceptance rather than belief to be the core epistemic attitude (see Cohen 1992). I do not, of course, advocate, believing falsehoods. Moreover, epistemic agents should be aware of the semantic character of the propositions they accept. To accept a proposition as true is different from accepting it as a felicitous falsehood, just as to accept it as literal is different from accepting it as figurative. Epistemic acceptance is concerned with the premises an agent is prepared to *use*, not the propositions she merely assents to. Scientists are well aware that they are using simplified (hence strictly false) models. I contend that they are entitled to use them in circumstances where the models' divergence from truth does not matter. And they are entitled to understand the phenomena via those models when the divergence from truth does not matter.

Nawar suggests that the veritist can easily accommodate the inaccuracy of, e.g., '$PV= nRT$' by replacing it with 'approximately $PV= nRT$' or '$PV = nRT$ is a partial description' or 'as an idealization, $PV= nR$T'. These may be worthy of belief; but they will not do as replacement premises in the thermodynamic inferences for which '$PV= nRT$ ' is used. In some circumstances, '$PV \approx nRT$' is a suitable replacement, but not always. The others are non-starters.

Nawar points out that I have not demonstrated that felicitous falsehoods are ineliminable. He is right. But (a) since we do not know how to eliminate them, they are at least currently ineliminable; and (b) scientists show no inclination to eliminate them. Paul Teller goes further. A simple example illustrates his point. The velocity of sound in dry air depends on the temperature, the pressure, and the proportions of the different sorts of gas molecules that compose the air. Although we can give averages, we are in no position to precisely specify the temperature, pressure, and proportions of the different sorts of molecules in a sample of air. Nor, Teller maintains, is this just an epistemic problem. He says, 'No real world sample of air has such precise values, if only because the values would vary from place to place. So at best one is talking about the speed of sound in some idealized condition, not in the real world'.[1] Because the

[1] Teller discussed this case in an early draft of his (2018). To my regret, he did not include it in the published version. The quotation is from his 2016 manuscript.

temperature, pressure, and proportions of different sorts of gas molecules vary slightly from one region of the sample to another, we cannot say truly that sound travels with exactly velocity *v* throughout the sample. To say something true here we need to idealize, treating the magnitudes as uniform across the sample. This is what we do and what we should do. But we do not thereby mirror the actual (fluctuating) speed of sound across the air sample. Teller's point is not just that contemporary science *contains* idealizations that are strictly false; rather, science *consists* of such idealizations (Teller 2018). If he is right, there is no hope of eliminating idealizations.

In my defense of felicitous falsehoods, I appeal to Grice's second maxim of quantity: 'Do not make your contribution more informative than is required' for the purposes at hand (1989, p. 26). Nawar correctly points out that Grace's maxim concerns conversation or communication, not understanding. I suggest, however, that the reason the maxim is valuable for communication is that it captures something important about understanding. By eliminating or marginalizing irrelevancies, we are better able to grasp relations among relevant factors. We understand better when we respect the maxim; we do not just communicate better. The reason Grice's maxim is a conversational maxim is that it is an epistemic maxim.

Goldman's veritism holds that truth alone has final epistemic value. Nor is Goldman alone in thinking this. As Nawar says, this would still leave room to recognize that felicitous falsehoods have instrumental epistemic value. If so, they are epistemically useful tools because they promote the emergence of true beliefs. Maybe so, but I do not believe that this comes close to exhausting their epistemic value. General Relativity has instrumental value because it serves as a springboard for further research. But it has final epistemic value as well. It embodies a rich and deep understanding of the physical world. If veritists have to deny this, they are missing something important.

3 *Epistemic Norms: Truth Conducive Enough* by Lisa Warenski

Lisa Warenski contends that veritism should be ecumenical about epistemic norms. It should acknowledge all 'truth-oriented values' where 'a truth-oriented value is one that bears an essential relation to truth'. She does not say what it takes to bear an essential relation to truth. This is a problem, since on the face of it, all propositions bear an essential relation to truth, since being a proposition is a necessary condition for being true. But it is unlikely that veritists want to countenance false propositions. We need to know which essential relations matter.

Warenski misconstrues my position, thinking that I advocate giving up truth entirely and believing falsehoods. I do neither. Truths as well as felicitous falsehoods figure in understanding by being elements of networks of epistemic commitments in reflective equilibrium. My theory of understanding focuses on what we should accept for cognitive purposes, not what we should believe. She challenges my, or more properly Goldman's (1999), conception of veritism. She takes it that under a less restricted view about truth, veritism could be preserved. This might be so, but the costs would be higher than she acknowledges.

Much of her argument turns on the Central Park Five Case. She provides a description which makes it clear that the case was a manifest, egregious miscarriage of justice.

Five young Black men were convicted of a horrific crime that they did not commit. Something went badly wrong.

For our purposes we should set aside moral, legal, and procedural wrongdoings unless they bear on epistemological issues. Like Warenski, we should distinguish between the legal case and public opinion.

First, some background about US criminal law: In a criminal case, the criterion for a verdict of *guilty* is *beyond a reasonable doubt,* where that verdict was reached entirely on the basis of evidence admitted in court. If jury members have independent reason to consider a defendant guilty or innocent, they must disregard it. If evidence presented in court is declared to be legally inadmissible, the jury must disregard that too. There are strict laws about what sort of evidence can be admitted. Hearsay, illegally obtained evidence, evidence about the defendant's prior criminal record are ruled out. Care must be taken to insure that the evidence has not been tampered with. And so on. The goal is not simply truth. Recognizing that juries are not infallible, the legal system incorporates a preference for false negatives over false positives. This is captured in Blackstone's ratio: it is better that ten guilty persons go free than that one innocent person be convicted. The ratio is not something the jury need worry about. *Beyond reasonable doubt* is their criterion. Rather the ratio informs the rules of evidence. Those rules are biased so as to avoid a particular sort of error. As a result, it is possible for members of a jury to be epistemically justified in believing that the defendant is guilty and yet in good conscience find him innocent because the admissible evidence (which was not all the evidence the jury members had) was not enough to convict. A criminal trial thus is not a good model for what ought to be believed or accepted when one's ends are purely cognitive.

Warenski takes it that my position cannot explain what went wrong in the trials of the Central Park Five. I disagree. As she describes the trials, my theory yields the result that the case against the five suspects was epistemically untenable. First, coerced confessions are untrustworthy. So the fact that the five teenagers confessed under duress is not reason to think that they committed the crime. Second, the confessions were mutually inconsistent. The crime could not have been committed the way all five said it was. Third, the claim that the defendants both assaulted the female jogger at time t and place p and robbed the male joggers at time t' at place p' was untenable, given that the defendants could not feasibly get from one crime scene to the other in the interval between t and t'. These constitute sufficient reason to reject the contention that the five young men were guilty. The internal incoherence of the story presented in court discredits the case. This conclusion does not require a commitment to the truth or falsity of any particular bit of evidence. It requires a recognition that the various contentions presented as evidence are not cotenable. I am not saying, of course, that it does not matter whether the defendants were guilty. I am saying that the way courts seek to establish guilt is not veritist. Nor is it obvious that ecumenical veritism would do better.

What of public opinion? Warenski points out that from the outset the newspapers represented the five as guilty. She says that that this accorded with the racist views of white New Yorkers. She intimates that I have no grounds for criticizing this bias. She says, 'The story as initially reported was coherent—at least to white New Yorkers'. I am not convinced. She may well be right that many white New Yorkers considered it coherent. But believing that an account is coherent does not make it so. In any case, my criterion for acceptability is not coherence per se. Rather, I maintain that a network of commitments is acceptable only if it is at least as reasonable as any available alternative. By 1989, there was plenty of evidence available to white New Yorkers that racist assumptions are untenable. They had access to the alternative narrative that Black New Yorkers had. They lived through the Civil Rights movement. And, I would add, they had plenty of evidence that jumping to conclusions on the basis of insufficient reason is epistemically untenable. The fact that a lot of people default to a racist stereotype is epistemically unfortunate, morally reprehensible, and socially tragic; but doing so is not warranted by my position. We all have overwhelmingly good reason to reject racist assumptions.

It is not clear to me how ecumenical veritism would help with this case. To insist that there were some truths, some partial truths, some approximations to the truth, and that there were some falsehoods, some claims that were partially false, some claims that were not even approximately true, although correct, does not get us any further. Racist assumptions are untenable. So are incoherent accounts. So are unreliable methods of extracting information. That is what we need in order to explain, epistemically speaking, what went wrong here.

4 *Partial Truth versus Felicitous Falsehoods* by Soazig Le Bihan

Soazig Le Bihan characterizes my view of veritism as maintaining that ‘considerations of truth are necessary, even if not always sufficient, when assessing epistemic value; when other considerations apply, truth is always of central concern.’ That’s not quite right. Although I maintain that, even for a veritist, not all truths are equally epistemically estimable, I do not revise Goldman’s requirement: Truth is a necessary condition on epistemic acceptability. It follows for veritism that, whatever its other merits, a contention is unacceptable if it is not true. It is not clear to me what the relation between epistemic value and epistemic acceptability is supposed to be. But the conjunction of Goldman’s characterization of veritism with my contention that truth is bivalent, there is a problem. Partial truths, being partial falsehoods are not true. So, according to the veritist, they lack epistemic value. I hold that although models, idealizations, and thought experiments are not strictly true, they have epistemic value. Hence, I am not a veritist. Either Le Bihan and I disagree that truth is bivalent or we disagree about what it takes to satisfy a necessary condition.

Le Bihan might, however, insist that the models she considers partly true are epistemically valuable because they are truth conducive. If so, they have instrumental epistemic value. There are a couple of problems with this position. First, there is reason to doubt that all the models we consider epistemically valuable are truth conducive. If the truth about a phenomenon is noisy, then a model that restricts itself to difference makers distances itself from truth. Modifying it to bring it into closer accord with the truth is not desirable, because it would reintroduce considerations that impede understanding (see Dennett 1991). An unrealistic model (e.g., one that represents a population as infinite) may be epistemically preferable to a more realistic one. A second, related point, is that on my view, the considerations we accept are considerations we reason with. We represent the phenomena to ourselves in terms of them and reason accordingly. Thus, inaccurate models that facilitate reasoning, are not, as Le Bihan says, 'second best' when the alternative is an intractable truth. They are preferable if we seek to understand the phenomena and reason about it fruitfully.

5 *Scientific Understanding and Felicitous Legitimate Falsehoods* by Insa Lawler

Insa Lawler assumes that all scientific understanding is a matter of grasping the content of an analysis or explanation. In understanding, she maintains, we extract the true content, and leave the felicitous falsehoods behind. I have a number of worries about this position. The first is that I see no reason to think that all scientific understanding involves explanation. Lipton (2009) provides two compelling examples to the contrary. Someone's understanding of the apparent retrograde motion of the planets may be captured in her ability to properly manipulate an orrery, even if she cannot explain retrograde motion. Her demonstration shows what is going on. If we want to insure that she doesn't get it right by luck, we can ask her to show what would happen if things had been different—if, for example Mercury had been a bit further from the sun. If her displays continue to be correct, we have, I think, no reason to deny that she understands the phenomenon. Galileo demonstrated that gravitational acceleration is independent of weight by showing that the alternative leads to a contradiction. He provided no explanation of the independence. He simply enabled us to understand that it could not be otherwise. Explanation is one way to demonstrate understanding; illustration is another.

A second problem is that even on Lawler's view grasping the content of an explanation is insufficient for understanding. One can grasp the content of a caloric explanation of heat transfer without understanding the phenomenon. One could even grasp the content, being aware that the explanation is false. Lawler might dismiss this case, arguing explanation is factive. If explanations must be true, then caloric explanations are no worry. The problem though, is that one might grasp the content of a thermodynamic explanation in exactly the same way that one grasps the content of a caloric explanation. Even if the thermodynamic explanation is true, unless one had reason to think that it is true, one would not understand the phenomena. And the content itself does not ensure that it is true. It has to link up to the phenomena. This is where the felicitous falsehoods enter. It is via felicitous falsehoods (the gas models) that the content of the explanation connects to the facts.

Lawler allows that the models play an enabling role. But she thinks that understanding consists in extracting the truths that constitute the explanation from the falsehoods that scaffold them. She maintains that we should decouple representation from explanation. The question of how felicitous falsehoods afford epistemic access to the phenomena is, on her view, independent of the question of how they figure in the content of the explanation of the phenomena. Addressing the content question, then, need not answer the question of why models are empirically successful. This is far from obvious.

In any case, it is not clear that we actually can perform the extraction. Our only access to the phenomena that provides any insight into their behavior may be via the models. If we 'kick away the ladder', we seem to have no reason to believe the explanations that result. They might be true, but we have no reason to credit them. Decoupling seems a chimera. Felicitously false models are not just routes to understanding or enablers of understanding; they are partially constitutive of understanding. We understand in terms of our models, not despite them.

Lawler evidently takes epistemic commitment to be a matter of belief. I do not. Nor do I think it is a matter of taking the content of the commitment to be true. Because acceptance is a policy of being willing and able to use a commitment as a basis for inference or action when one's ends are cognitive, mutually inconsistent commitments need not clash. They may be isolated from one another by the constraints on the inferences or actions one is willing to let them figure in. To accept a model is consonant with recognizing that it is felicitously false, hence that it holds only in some contexts and within some limits. Where it holds, it exemplifies something that bears on the phenomena in question. Thus, for example, the model of the nucleus as a liquid drop exemplifies some features of the nucleus. To also accept a different model—say, one that represents the nucleus as a rigid shell—is to recognize that it too exemplifies some features of the nucleus—not the same ones as the previous model. It can be fruitful to accept both. No contradiction is involved, since to accept them is not to be committed to the view that both are literally true of the nucleus.

We regularly accept figurative representations that would lead to contradiction if we considered them literally true. Consider the case of metaphor. We might describe someone as a golden retriever to characterize him as displaying a sort of indiscriminate enthusiasm for which we have no literal label. He is, of course, a human being. No contradiction emerges since the predicate 'human being' applies to him literally and the predicate 'golden retriever' applies metaphorically. Consider the case of allegory: In *The Republic*, Plato included the allegory of the cave. The sun is the allegorical counterpart of the Good. But although the sun is visible and material while the Good is invisible and abstract, there is no contradiction; for the allegory does not represent literally. This reinforces Frigg and Nguyen's point about the importance of attending to the nature of representation. We need to understand how various modes of representation relate to each other, how they reinforce or interfere with each other's functions.

6 *Understanding Realism* by Collin Rice

Collin Rice seeks to vindicate scientific realism. Like Lawler, he favors extraction. On the sort of realism he advocates, scientists extract true modal information from inaccurate models. The worries about extraction that I expressed in my reply to Lawler arise here too. Even if the extracted modal information is true, it is not obvious that, independent of the inaccurate models, scientists have any reason to believe or accept it.

Rice maintains that models are valuable because they provide access to modal information. They enable scientists to figure out what would have occurred or would have been the case if things had been different. This is important. But we need more. In particular we need to secure the relation of the model providing the information to the phenomena it bears on. On the basis of our knowledge of pigs and a bit of aerodynamics, we can correctly infer 'If pigs had wings, they could fly'. But that does not yield any understanding either of porcine locomotion or of aerodynamics. Rice's claim that 'the more modal information one grasps about the possible states of the target phenomenon, the better one understands the phenomenon' is probably false. Much of the modal information about the target phenomenon, like much of the non-modal information is trivial. That the relations of pressure, temperature, and volume reflected in '$pV = nRT$' would hold in a world where bicycles had never been invented does not improve our understanding of gas dynamics.

I maintain that exemplification is the link between models and their targets. Rice disputes this, focusing on my discussion of the Hardy–Weinberg model. He is in good company in thinking that I could have (or anyway should have) been clearer. So let me give another try. The distribution of genes in a population is affected by multiple factors including selection, migration, mutation, and genetic drift as well as the distribution of alleles in the previous generation. The Hardy–Weinberg model concerns how alleles would redistribute if nothing but redistribution played a role. (This is the sort of counterfactual Rice is interested in.) I contend that the model exemplifies a feature which is properly projected onto the target. That projection is not supposed to yield the actual distribution of alleles. Rather it yields the contribution that one factor makes to the overall distribution. Factor analysis figures in the way the model functions.

Rice discusses universality classes, which consist of (perhaps) heterogeneous systems that display similar patterns of behavior. One can construct a simple, idealized model to represent that behavior. His point here is important. Because, for the purposes at hand, the differences between the systems are negligible, the model is effective. I would say that the model exemplifies a pattern it shares with the phenomena. That pattern should hold in suitable, delineated counterfactual circumstances. Rice is suspicious of exemplification. Whether or not his suspicions are warranted, I believe, philosophers of science should take universality classes seriously.

7 *Elgin on Understanding: How Does It Involve Know-How, Endorsement, and Factivity?* by Emma Gordon

Emma Gordon addresses three themes from *True Enough*. She then makes an intriguing proposal. I'll say a little about each.

Her major worry concerning know-how concerns my discussion of rule following. She worries that on my conception we cannot accommodate novelty. I am not so sure. Constitutive rules of a practice define the practice. They specify what it takes to engage in the practice, what the goals of the practice are, and what sorts of actions are permissible in pursuing those goals. The constitutive rules of chess, for example, stipulate that checkmate is the goal of the game. They specify what it takes to checkmate one's opponent, how the various pieces are permitted to move, and so on. Someone who flouts the rules is not playing chess. But as is well known, chess is a game of strategy.

There are multiple, highly creative gambits, all consonant with the rules, that can be used in attempting to win. New strategies can be invented. Knowing how to play chess is a matter of knowing how to exploit the opportunities that the rules afford.

Gordon's example is Ulrich Salchow, the first figure skater to attempt a rotating jump. She notes that the jump deviated from normal standards of figure skating. Two possibilities are open here. I do not know which is historically accurate. The first is that the jump did not violate the rules of competitive figure skating. Salchow did something remarkable and unanticipated that accorded with the rules. If so, he presumably constrained his skating so as to not violate the rules. His rule following was no accident. In this case, he was like a chess player who invents a new gambit. The second is that he violated the rules that were then in effect. His remarkable jump displayed a new kind of figure skating excellence, thereby prompting authorities to revise the rules so that competitive figure skating could henceforth embody that excellence. In that case, even if he was disqualified for failing to follow the rules, he left his mark on the sport. Rules permit some actions and prohibit others. There is thus typically room for novelty within a rule-governed practice (see Rawls 1955).

In challenging my claim that understanding involves reflective endorsement, Gordon presents a variant of Lackey's (2008) creationist teacher. Stella competently teaches the theory of evolution although, being a creationist, she does not believe or accept it. Gordon suggests that Stella nonetheless understands evolution. I agree that she understands something. But the target of her understanding is the theory itself, not the phenomena the theory accounts for. So her understanding is like that of a literary theorist's understanding of *Paradise Lost*, or a historian of science's understanding of phlogiston theory. Because Stella does not reflectively endorse the theory of evolution, she effectively understands the theory as she would a fiction. She does not understand how the diversity of life on earth came about.

Gordon's third issue concerns my non-factivism. Contrary to what she says, I do not maintain that understanding involves false beliefs. I insist that accepting something is different from believing it (see Cohen 1992). Like Kvanvig (2003), she hopes to consign falsehoods to the periphery of understanding. She uses Copernicus as an example, arguing that although Copernicus was committed to the idea that the sun had an absolute position, this was not central to his view. Hence, a factivist can construe him as understanding the motion of the Earth anyway. This might work for the commitment about absolute space. But this is not Copernicus's only deviation from the facts. Copernicus maintained that the Earth's orbit is circular. It is not. Given that there was a strong commitment in fourteenth century astronomy that celestial motions had to be circular, it is harder to consign this error to the periphery.

Perhaps more serious than spats about the relative centrality of commitments that their advocates take to be factual are worries about models and idealizations that are not even supposed to be factual. In many cases, there is no plausible way to marginalize them. If contemporary scientists are to understand gas dynamics at all, it will be in terms of some highly idealized, hence strictly false, model. The strictly false gas laws are central to their understanding.

Gordon's final point extends my ideas in a new direction. Williamson (1996) convinced many philosophers that knowledge is the norm of assertion. If so, knowledge that p is necessary and sufficient for appropriate assertion that p. I have long doubted the necessity claim. If it is correct, we should probably assert a lot less than we do. Gordon takes up the sufficiency claim. Drawing on Lackey (2008), she considers the case of a physician who learns from an intern that test results show that a patient has pancreatic cancer. The physician has not seen the results, much less evaluated them for herself. She simply takes the intern's word for it. If the knowledge norm is correct, the physician should have no qualms about telling the patient his diagnosis. This seems irresponsible. We have the sense that the physician ought not assert that he has pancreatic cancer until she has seen and evaluated the results for herself. Gordon suggests that in a case like this assertion requires expertise, which requires the sort of understanding I describe. As an expert, the physician should know how to wield the commitments that bear on the topic. She should have the information at hand, so that she knows how to answer the patient's questions, which are apt to be fine-grained, to be targeted to his own condition, and to involve a variety of hypothetical scenarios. The patient has a legitimate expectation that she base her diagnostic assertion on the exercise of that know how. This strikes me as an important insight. It not only illuminates some of the discomfort we feel in the Lackey case, it also points in a promising direction for research into expertise.

8 *Catherine Elgin on Peerhood and the Epistemic Benefits of Disagreement* by Kirk Lougheed

Recent discussions about disagreement focus on disagreements among epistemic peers, where epistemic agents count as peers just in case they share the same evidence, background information, reasoning abilities, and incentives. The list of qualifications varies slightly from one paper to the next, but these qualifications are fairly standard. The rationale for the qualifications is to underwrite the assumption that neither party is more likely to be mistaken than the other. Papers on peer disagreement tend to leave it mysterious why epistemic peers disagree. The answer is easy enough in Christensen's restaurant case (2007). In disagreeing about how much each person should pay, at least one of the parties miscalculated. But many cases are not so straightforward. One reason peers sometimes disagree, I argued, is that they favor different reasoning styles (Elgin 2018). Peers, I said, are not clones.

Lougheed takes issue with this. He maintains that different reasoning styles are bound to result in different track records. And if Jen has been right more often than Ken, they are not peers. I am not convinced. The easiest cases come from mathematics. Analytic geometry provably reduces to algebra. If Bill prefers to solves problem geometrically while Phil prefers to use algebra, and each is as good as the other at his chosen form of mathematical reasoning, then so long as they attempt the same number of problems and solve equally difficult problems equally often, they strike me as mathematical peers. We needn't even assume that both answer all and only the same problems correctly. If both are right, say, 90% of the time, they seem to be peers. It might seem that mathematics is a special case. I do not think so. For two people to be epistemic peers, they need not have given all the same answers in the past. Presumably there would be considerable overlap. But there might be divergence as well. This does not affect their status as peers, so long as they are right about equally significant questions, wrong about equally significant questions, and suspend judgment about equally significant questions an equal proportion of the time.

Lougheed introduces a different conception of epistemic peerhood. Rather than attempting to specify what it takes to be peers, he emphasizes the fact that each of the agents has no independent reason to think that he or his counterpart is in an epistemically stronger position than the other. This is a lovely idea. I do not think that it would eliminate the question of what it takes to be epistemic peers, but it would push that question to the periphery. When faced with a disagreement, an epistemic agent needs to decide whether she should conciliate or hold fast. That requires determining whether she has good reason to believe either party is in a better position to know than the other. She rarely if ever has the rich array of background information that familiar discussions use to define peers. Lougheed's suggestion that the reason should be independent of the disagreement itself is reasonable. To incorporate the very fact that we disagree as a tie-breaker seems question begging.

9 *Elgin's Community-Oriented Steadfastness* by Klaas Kraay

According to Klaas Kraay, the steadfastness that I maintain is permissible in cases of peer disagreement can help communities of inquiry achieve their collective epistemic goals. He is right about this. But I worry that he makes the rationale sound too altruistic. Suppose Hal and Mal disagree about whether Neanderthals buried their dead. Hal thinks that they did and that his reasons for thinking they did are as good as Mal's reasons for thinking they did not. If both remain steadfast and, accept the hypothesis that they consider correct, pursue their inquiries accordingly, one may eventually come up with conclusive reasons. Then the other will conciliate. The reliance on the community is a means to discovering something Hal wants to know, in the same way that getting a stronger telescope is a means toward discovering something an astronomer wants to know. The community provides him with evidence that he could not so easily get by himself. But Hal need not be altruistic in his continued pursuit of evidence to support his hypothesis. He certainly need not think that he is (or even might be) sacrificing his own epistemic good for the good of the community. He accepts the hypothesis because he thinks it is correct. The community benefits because diversity protects them against prematurely rejecting hypotheses that may turn out to be correct (see Kitcher 1990).

Kraay construes me as an epistemic rule-consequentialist. I do not think that I am a consequentialist of any kind. Of course inquiry has ends. Inquirers aim to discover things or find things out or figure things out or whatever. But it does not immediately follow that their stances (including their stances on disagreement) are determined by their aims. If they were, demon worlds would pose a problem. In a world where none of our methods succeed, the consequentialist has to say that there is no basis for favoring one method over others. But even in such a world, I believe, agents should consider scrupulous, rigorous, evidence-based inquiry preferable to jumping cavalierly to conclusions.

I have developed a responsibilist, quasi-Kantian position according to which members of epistemic communities function as legislating members of realms of epistemic ends. They formulate, justify and occasionally modify the rules, criteria standards that they think best serve their collective epistemic ends. Very roughly, on my view Hal is epistemically permitted to remain steadfast when, according to the collective judgment of his community of inquiry, the aims of their joint enterprise are fostered by steadfastness in cases like his. This distances epistemic permissibility from questions of likeliness, hence from Hal's having to assess whether his steadfastness is apt to have any effect at all. It allows for the recognition that most inquiries fail. That being said, I'm happy to learn that the sort of steadfastness I discuss is also available to epistemic rule-consequentialists.

10 *Some Fallibilist Knowledge: Questioning Knowledge-Attributions and Open Knowledge* by Stephen Hetherington

Stephen Hetherington has long maintained that familiar internalist and externalist conceptions of knowledge are untenable, but that better conceptions are available. In this paper he addresses the epistemic discomfort we feel with concessive knowledge attributions—statements of the form 'I know that *p*, but I could be mistaken that *p*.' Such attributions sound perilously close to Moore's paradoxical statements. In making them, we seem to be giving our assurance in the first clause and taking it back in the second. It's not a happy situation. Or so it seems.

Hetherington is more optimistic. He introduces the idea of *open*-knowledge. 'Someone would claim to know that *p,* for instance, even while asking whether she might unwittingly be mistaken as to *p*—hence mistaken as to whether she really has the knowledge that *p* that she is taking herself to have.' This is a description of a self-attribution of knowledge. We get to open-knowledge itself by saying that she both does know that *p* and does wonder whether she might be mistaken as to whether *p.* Open knowledge is a self-reflective, questioning stance. It not only recognizes our fallibility, it makes a virtue of it. It strikes me as an excellent epistemic state to be in. It evades Kripke's knowledge puzzle (2011) without defaulting to dogmatism. It also orients the knower to both further evidence vis à vis *p,* and to higher order considerations about the nature and strength of her evidence for *p.* It puts her in a position to learn more both about *p* and about what it takes to know things like *p.* Rather than closing her off to further inquiry by declaring the question whether *p* settled, it invites and directs further inquiry.

There is a cost. Craig (1990) and Hannon (2019) maintain that the point of knowledge is to enable us to identify good informants—people whose word we can safely take in acquiring information that we do not already have. Concessive knowledge and open knowledge seem not to provide that. To see the difficulty, we need not suppose, as Craig and Hannon do, that providing a mark of a good informant is *the* point of knowledge. It is enough that it serves as a desideratum. Suppose I ask someone whether, for example, the brownies at the bake sale contain nuts.[2] She replies, 'I know they do not, but I could be wrong about that'. I would be reluctant to take her word, particularly if

[2] I think this is a variant of an example presented by Brian Weatherson. Unfortunately, I do not know where he gives it, or whether I learned of it from a good informant.

I am allergic to nuts, or fear that I will transmit her information to someone who is. Her concessiveness seems to mark her as a poor informant.

This is not to deny that we should recognize and value open-knowledge. It is, however, to suggest that it does not provide everything we want in a conception of knowledge. Perhaps our everyday concept of knowledge involves a variety of desiderata that are not cotenable. We can achieve some of them only by sacrificing others. Hetherington can readily agree.

11 *Mistakes as Revealing and As Manifestations of Competence* by Felipe Morales Carbonell.

In *True Enough*, I argue that mistakes are epistemic achievements. Felipe Morales Carbonell delves deeply into the question of how this can be so. One dimension of the issue, he maintains, concerns epistemic policy. If acceptance were irrevocable, it would make sense to be epistemically risk averse. If it is revocable, we can afford to take more risks. In that case, however, we need to determine what risks are worth taking. Plainly we should devise policies to identify errors. But, as Morales Carbonell argues, we should do more. We should devise ways to manage our epistemic commitments so that errors (or potential errors), once identified, will be informative and fruitful. Dissection of an error should highlight something significant about the phenomena and/or our current understanding of it.

Ideally, once we have found an error, we should correct it. But, as Morales Carbonell points out, how to do so is not always obvious. If a mistake is integrated into a network of commitments, extracting it will not leave the rest of the network intact. 'Distinct policies [of error correction] will be more or less conservative depending on the size of the changes they will entail.' More than size is at issue. Strategies of revision require epistemic value judgments about which related commitments we should be more or less open to revising. These depend not only on local concerns—those that figure in 'the smallest set of beliefs that yield the inconsistency', but also on more global concerns about the topic and our access to it.

Morales Carbonell goes on to discuss competences. Sosa (2007) maintains that knowledge is true belief due to the competence of the believer. Epistemic competence is construed as an ability to characteristically get things of the appropriate sort right. My view of epistemic competence is different. It can be displayed in getting things fruitfully wrong. A problem for Sosa's view, as Morales Carbonell points out, is that if the task is hard, even a competent agent could regularly fail. Sosa's athletic example is archery: he takes it that the competent archer usually hits the target but, due to environmental interference, occasionally misses. Maybe this is true in archery. But if we switch sports, competence has a different profile. Even the best professional baseball players usually strike out. Those with batting averages of .300—that is, those who only strike out only about 70% of the time—are considered excellent batters. Arguably something similar holds at the cutting edge of inquiry. Whatever it takes to be a competent string theorist, it probably isn't being characteristically right.

One final point concerns environmental interference. Would any epistemic agent be competent in a demon world? Probably not, at least if the criterion is something like believing truly on the basis of ability. This is why I prefer to evaluate epistemic agents on the basis of how well they use available epistemic resources. Even in a demon world, Newton would be epistemically admirable.

12 *Rational Understanding: Toward a Probabilistic Epistemology of Acceptability* by Finnur Dellsén

Regardless of how scrupulously an investigator vets each of her claims, in an extended work it is overwhelmingly probably that at least one of her claims is false. Having been careful, she has reason to believe or accept each of them, but to reject their conjunction. But understanding is holistic. It consists in grasping or appreciating how a multiplicity of claims collectively shed light on their subject matter. And it increases as the agent accepts and integrates additional information into her framework of commitments. The problem is not just the fallibility of the epistemic agent. It is a principle of probability theory that the probability of a conjunction of mutually independent claims, each with a probability of less than one, decreases with each additional conjunct.

Finnur Dellsén argues that this provides reason for the theory of understanding to reject the standard probabilistic picture of justification. In its place, he devises an optimality model. Rather than asking how probable a network of commitments is in absolute terms, we should consider how it compares with its rivals. If *N* is significantly more probable than the other accounts that purport to provide answers to exactly the same questions, *N* is acceptable. This is so even if *N* is, in absolute terms, quite improbable. This is a terrific idea. As Dellsén says, it yields something close to my reflective equilibrium model. That two accounts, starting from such different premises arrive at something close to the same conclusion is evidence that they are onto something.

Still, there are a couple of divergences worth noting. It is not clear to me whether they favor Dellsén's account or mine. The first concerns ties. On Dellsén's view, an account is acceptable only if it is significantly more probable than its rivals. On mine, an account is acceptable just in case it is at least as good as the best of its rivals. Thus two accounts of the same phenomena, each as good as the other, and better than any other alternatives, could be acceptable. They could not, however, be conjoined. In *Considered Judgment* (1996), I consider two theories—one describing light in terms of particles, the other describing it in terms of waves. I maintain that each could be as good as the other. But they cannot be conjoined because in a world where light consists of particles, there is no room for waves; and in a world consisting of waves, there is no room for particles.

Because Dellsén's account prohibits compartmentalization, he would say that neither account affords an understanding. His example of compartmentalization involves two theories, each of which accounts for part of the domain. A history of emigration from Southern Europe and a history of emigration from Northern Europe, taken separately, do not afford as good an understanding as a single, comprehensive history of emigration from Europe. He has a plausible case against the sort of compartmentalization that arises from partitioning the domain. I do not know whether it is equally persuasive against compartmentalization that arises from different ways of conceptualizing the same domain.

It would be straightforward for me to tighten my standards to arrive at Dellsén's position. In that case, where two accounts are tied for first place, neither affords understanding. It would not be easy for Dellsén to relax his standards to align with mine. His consistency requirement mandates that any two acceptable claims can be conjoined.[3]

A second worry, which Dellsén could easily accommodate, is that there should be some minimal threshold on acceptability. We probably do not want to say that the deliverances of the oracle at Delphi afforded understanding of the future, even if its predictions were significantly better than the predictions of competing oracles.

A third concern is that Dellsén's model is highly idealized. It enables us to specify what it takes for an account to be acceptable, but it does not equip an agent to tell whether an account is acceptable. One might say the same about my theory. To determine whether an account is at least as good as any available alternative requires being able to identify, access, and assess the alternatives. That is hard; but it may not be as hard as making the comparisons Dellsén would have us make.

However these details work out, Dellsén's optimizing strategy strikes me as an important contribution to the epistemology of understanding.

13 *Reflective Equilibrium and Understanding* by Christoph Baumberger and Georg Brun

Christoph Baumberger and Georg Brun explicate my account of reflective equilibrium and suggest ways to revise it. They go on to consider how reflective equilibrium figures in my theory of understanding. They are right that even with their emendations, my view does not maintain that reflective equilibrium is sufficient for a person to understand a topic. More is required.

Drawing on Goodman (1973) and Rawls (1971), I contend that in constructing an account (or a theory in Baumberger and Brun's sense), we begin with what I call *initially tenable commitments* (see Elgin 1996). These are considerations that have some (perhaps small) measure of acceptability prior to and independent of our current theory construction. By revising, augmenting, extending and systematizing them, we arrive at a network of commitments in reflective equilibrium. The link to initially tenable commitments, I maintain, assures that the account is reasonable in light of our antecedent views on the topic. Baumberger and Brun correctly point out that in the process of adjudicating competing commitments we may introduce further commitments, which they call 'input commitments'. Such commitments, they maintain must be independently credible. I disagree. The introduction of the positron into the model of the atom had no independent credibility. It was justified by the commitment to the electron (for which there was evidence) and by the commitment to symmetry. In light of those two commitments, it was reasonable to introduce a commitment to the positron (a positively charged counterpart to the electron) even when there was

[3] I thank Samuel Elgin for explaining why the requirement cannot easily be modified.

no independent evidence of such a particle. Still, Baumberger and Brun are right that many input commitments are independent credible and that this contributes to their epistemic utility. It is because, for example, the evidence from an experiment is credible that it is properly integrated into an emerging theory. Baumberger and Brun maintain that input commitments are distinct from initially tenable commitments, because they were not introduced at the outset. But the similarities are striking. Both initially tenable commitments and input commitments have an epistemic foothold that is independent of the current construction. Both can be extremely tenuous. Both play the same role in providing a measure of independent support for the developing account. In my view, input commitments are initially tenable commitments. Perhaps the term ‘initial’ was ill-advised. (I followed Goodman 1972; Scheffler 1986 in using it.) Maybe I should have said something like ‘independent provisional commitments’. The main issue is not whether they are tentatively accepted at the outset or later on; it is that their status is in some measure independent of their contribution to the theory under construction.

Baumberger and Brun worry that the tolerance for revisions built into the process of equilibration needs to be constrained. Otherwise there is a danger of changing the subject. They therefore maintain that within the theory under construction there is an important difference between input commitments and inferential commitments. They say, ‘In contrast to input commitments, purely inferential commitments may not be used to assess whether the resulting position can be seen as providing an understanding of the original subject matter’. It is not obvious to me that there is, or ought to be, a sharp line between a revision that does and one that does not change the subject.

Suppose we had input commitments that included a commitment to the effect that gratuitously inflicting pain on people is pro tanto wrong, and a commitment to the claim that pain hurts. In developing a moral account we connect the two commitments to arrive, via inference, at the commitment that gratuitously inflicting pain on people is pro tanto wrong because pain hurts. We then notice either (a) that many non-human animals feel pain, or (b) that ‘on people’ plays no role in our inference. Either prompts the further inference that gratuitously inflicting pain is pro tanto wrong. This leads to expanding our developing moral theory to recognize obligations to animals.

It is not obvious to me that the reliance on inferential commitments shows that we have changed the subject. I would probably say that they showed that the subject was broader than we originally thought. In any case, it seems to me that in theorizing we regularly change the subject. We begin with a rather vague, inchoate, and motley collection of ideas and regiment them into a sharp-edged account that (more or less, but not exactly) accommodates the views we began with. It is close enough.

A critical question is how reflective equilibrium relates to understanding. Here a bit of disambiguation is in order. Consider a case. The theory of evolution affords *an understanding* of the diversity of life on earth. This is so because the theory of evolution consists of a network of commitments in reflective equilibrium, where that network is grounded in fact, is duly responsive to evidence, and underwrites non-trivial inference and action regarding the phenomena it pertains to. Such an understanding is encapsulated in books, journals, web sites, and so forth. By itself this says nothing about rational agents. For an epistemic agent *to understand* the diversity of life on earth via the theory of evolution, she must accept it. That is, she must be willing and able to use it as a basis for inference and action when her diversity-of-life-relevant interests and actions are cognitive. Baumberger and Brun are right to insist that the fact that the theory is in reflective equilibrium is not enough. For an epistemic agent to understand phenomena via a theory requires something of her as well as something of the theory.

They go on to sketch some of the criteria they think scientific understanding should possess. Following de Regt (2009) they suggest that 'since it is often possible to cheat one's way to a mathematical solution without really understanding it and the underlying theory, understanding a subject matter… requires (a_2) being table to solve qualitative problems by drawing consequences of the theory without performing exact calculations'. I do not share this bias against mathematics. It seems plausible that an understanding of a subject could be purely quantitative. To block 'cheating', the agent would have relevant cognitive dexterity, so that she grasps what would happen if things were different. But, as far as I can see, this could be purely mathematical.

They go on to say 'being able to use a theory also requires (b) that one can explain aspects of the subject matter in terms of the theory'. Here too, I disagree. As I mentioned above, a science student who can use an orrery to show why the inner planets seem to reverse direction in their orbits *displays* an understanding, even if she cannot put that understanding into words (see Lipton 2009). That explanation is the only way to display scientific understanding is a bias that I do not share.

Baumberger and Brun quite properly maintain that the theory of reflective equilibrium is underdeveloped. They could have pointed out that the theory of objectual understanding is too. More work needs to be done.

14 *The Social Fabric of Understanding: Equilibrium, Authority, and Epistemic Empathy* by Christoph Jäger and Federica Malfatti

In this paper, Christoph Jäger and Federica Malfatti discuss how and why to integrate the position I develop in *True Enough* into social epistemology. Let me begin with a point of clarification. The distinction they draw between authority and expertise is not the one I was working with. Authority, as I understand it, derives from institutional role. Someone in authority is authorized by an institution to perform certain functions that others, who are not so authorized, cannot perform. For example, as the professor, I have the authority to determine when assignments in my course are due. My teaching fellow, even though he has a better appreciation of the demands on students' time, lacks that authority. A traffic cop has the authority to issue a ticket for speeding. Even though I was well aware that the driver in the red Porsche was speeding, I have no authority to issue a ticket. Authority in this sense is not an epistemic notion. Expertise, however, is epistemic. Jäger and Malfatti accept Goldman's characterization:

S is an expert about domain *D* if and only if (A) *S* has more true beliefs (or high credences) in propositions concerning *D* than most people do, and fewer false beliefs; and (B) the absolute number of true beliefs *S* about propositions in *D* is very substantial (2018, p. 5).

I reject this for several reasons. First, the focus on true and false beliefs, and the idea that we can tally them seems wrong-headed. This is a major theme in *True Enough*. Second, at the cutting edge of inquiry it is likely that even those we would consider experts do not have a substantial number of true beliefs about their subject. Experts on string theory or dark matter understand their subjects better than most, but probably do not have many significant true beliefs about them. These concerns could be addressed with relatively minor modifications to Goldman's characterization.

My main reason for rejecting Goldman's view is that I think that expertise is more widely distributed than his characterization allows. New parents often seek out the advice of more experienced parents about such matters as how to soothe a teething baby, when an infant is likely to sleep through the night, when it is safe to start a child on solid foods, and so on. I consider the experienced parents to have the expertise needed to answer such questions. (I call such expertise 'playground expertise', since consultations between novice and expert often occur on playgrounds.) I think that inhabitants of a city often have expertise about how to avoid traffic jams. And so on. In my view.

S is an expert for *T* vis à vis *p*, just in case *S* has sufficiently greater understanding of *p* than *T*.

So someone could be an expert about a topic for one person but not for another.

This just clears up terminology. It still leaves the problem that Jäger and Malfatti identify. If *S'*s expertise is to be useful to *T, S* must be able to convey it to *T* in a way that *T* can understand and use. Merely having greater understanding, or greater knowledge, or more information is not enough. As Jäger and Malfatti maintain, epistemic empathy is required. I agree. I characterized an epistemically empathetic informant as one who grasps the inquirer's interests and abilities and tailors his responses accordingly. To do this he must be articulate (in a language that the inquirer can understand). He must frame their responses in a way that the inquirer can understand and use (see Elgin 2019). To capture this, we might introduce another term. Let us say that *S* is an *epistemic source* for *T* vis à vis *p* just in case *S* has the appropriate expertise vis à vis *p,* and *S* is epistemically empathetic. To be an epistemically empathetic informant requires not just expertise about the subject matter, but also an appreciation of the inquirer's situation. Only if he has both is he in a position to provide useful information. As Jäger and Malfatti argue, this is a social-epistemic virtue. It is also a moral one.

References

Christensen, D. (2007). Epistemology of disagreement: The good news. *Philosophical Review, 116,* 187–217.

Cohen, L. J. (1992). *An essay on belief and acceptance*. Oxford: Clarendon Press.

Craig, E. (1990). *Knowledge and the state of nature*. Oxford: Clarendon Press.

De Regt, H. (2009). Understanding and scientific explanation. In H. de Regt, S. Leonelli, & K. Eigner (Eds.), *Scientific understanding* (pp. 21–42). Pittsburgh: University of Pittsburgh Press.

Dennett, D. (1991). Real patterns. *Journal of Philosophy, 88,* 27–51.

Elgin, C. (1988). *With reference to reference*. Indianapolis: Hackett Publishing Co.

Elgin, C. (1996). *Considered judgment*. Princeton: Princeton University Press.

Elgin, C. (2017). *True enough*. Cambridge: MIT Press.

Elgin, C. (2018). Reasonable disagreement. In C. Johnson (Ed.), *Voicing dissent: The ethics and epistemology of making disagreement public*. New York: Routledge.

Elgin, C. (2019). The mark of a good informant. *Acta Analytica*. https://doi.org/10.1007/s12136-019-00418-9.

Goldman, A. (1986). *Epistemology and cognition*. Cambridge: Harvard University Press.

Goldman, A. (1999). *Knowledge in a social world*. Oxford: Oxford University Press.
Goldman, A. (2018). Expertise. *Topoi, 37,* 3–10.
Goodman, N. (1972). *'Sense and certainty' in his Problems and Projects* (pp. 60–68). Indianapolis: Hackett Publishing Co.
Goodman, N. (1973). *Fact, fiction, and forecast*. Indianapolis: Hackett Publishing Co.
Grice, P. (1989). *'Logic and conversation' in his Studies in the way of words* (pp. 1–143). Cambridge: Harvard University Press.
Hannon, M. (2019). *What's the point of knowledge?*. Oxford: Oxford University Press.
Kitcher, P. (1990). The division of cognitive labor. *Journal of Philosophy, 87,* 5–22.
Kripke, S. (2011). *'Two paradoxes of knowledge' in his Philosophical troubles* (pp. 27–51). Oxford: Oxford University Press.
Kvanvig, J. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.
Lackey, J. (2008). *Learning from words*. Oxford: Oxford University Press.
Lipton, P. (2009). Understanding without explanation. In H. de Recht, S. Lionelli, & K. Eiger (Eds.), *Scientific understanding: philosophical perspectives* (pp. 43–63). Pittsburgh: University of Pittsburgh Press.
Perry, J. (1979). The problem of the essential indexical. *Noûs, 13,* 3–21.
Rawls, J. (1955). Two concepts of rules. *Philosophical Review, 64,* 3–32.
Rawls, J. (1971). *A theory of justice*. Cambridge: Harvard University Press.
Scheffler, I. (1986). *'On justification and commitment' in his Inquiries* (pp. 293–302). Indianapolis: Hackett Publishing Co.
Sosa, E. (2007). *Apt belief and reflective knowledge*. Oxford: Oxford University Press.
Tarski, A. (1956). *'The concept of truth in formalized languages', in his Logic, Semantics, and metamathematics* (pp. 152–278). Oxford: Clarendon Press.
Teller, P. (2018). Measurement accuracy realism. In I. F. Prichard & B. C. van Fraassen (Eds.), *The experimental side of modeling* (pp. 273–298). Minneapolis: University of Minnesota Press.
Williamson, T. (1996). Knowing and asserting. *Philosophical Review, 105,* 489–523.