

Responses to Critics

I begin by expressing my sincere thanks to my critics for taking time from their own impressive projects in epistemology to consider mine. Often, in reading their criticisms, I had the feeling of having received more help than I really wanted! But the truth of the matter is that we learn best by making mistakes, and I appreciate the conscientious attention to my work that my critics have shown.

My responses here will focus on two main areas of contention by my critics. In the first part, I'll address some of the questions and criticisms concerning the nature of understanding and the quasi-factive account I gave of it; and in the second part, I'll talk more about the Meno problem and the prospects for a successful response to it.

Prior to doing so, however, I want to address briefly Riggs' remarks about the value of truth. There is a growing sense among some epistemologists (e.g., Sosa,¹ Roberts and Wood,² DePaul and Grimm,³ and perhaps now Riggs as well⁴) that some value for truth applied uniformly and universally to all truth may not be defensible. Instead of counseling the love of truth for inquirers, such critics recommend a more circumspect love for deserving truths only. Perhaps they are correct, but I do not think anyone has provided the necessary argumentative support for the view, and I'll explain why here briefly.

Let me start with an analogy. I think pain is bad, even though pain often contributes to well-being. When I put my hand on a hot stove, the pain causes me to move my hand quickly. If you leave everything the same and remove the pain, things would be much worse.

There's no good reason to explain such examples by saying that the pain in question is good.

Instead, the language of defeat helps us understand the case. Pain is a bad-making feature of a case, always and everywhere, but sometimes it occurs in circumstances where its intrinsic badness is overridden or defeated by other factors. We often use the term 'prima facie' to capture this defeasible character of the badness in question.

The view that truth is valuable should be understood in a similar, defeasible way. When philosophers point out that the truth about the number of grains of sand regarding a particular stretch of beach is hard to see as worthy of discovery, that point by itself doesn't undermine the prima facie value of truth. Instead, we can explain this perception as involving defeat of the prima facie value in question with other features: the claim in question doesn't do much explanatory work for us, it doesn't help our overall understanding of the universe much, taking the time to learn it will interfere with other learning, etc.

So it is not hard to see how to resist the claims of lovers of limited truth, but it would be nice to have a strategy for getting around the impasse thus created. I believe a strategy can be outlined that makes it very difficult to defend the limited view. The strategy is to address the value of a posteriori truth first by imagining that you are an angel in Plato's heaven.⁵ You have available to you all the information that doesn't depend on concretia, and you are trying to determine which of the truths about concretia one ought to attend to once plucked from Platonic heaven and placed in the world of concrete reality. I think the answer to this question is that you'd have no reason to classify certain truths about concretia as unimportant, prior to placement in the world. You wouldn't be able to classify some as more basic from an explanatory perspective or more important for a complete understanding of any given phenomenon, for these kinds of claims depend on the nature of the concrete universe, not

on some features detectable from what is known in Plato's heaven. Once so placed, however, you'd learn things that would immediately generate a list of priorities in terms of important and unimportant truths. That's how defeasibility functions, however, and leaves intact the view that all truths of the sort in question begin from the same level of importance.

If I'm right about this thought experiment, the only route left is to argue that some of the truths available to one within Plato's heaven are themselves unimportant. It is interesting to note, however, that none of the critics of the universal value of truth use examples of truths that would be available in such a setting. Instead, they always appeal to empirical truths about the concrete world that, from our perspective, are fairly worthless (memorizing the phone book, counting grains of sand, etc.). I suspect, then, that any argument for a limitation on the value of truth won't be extendable to the truths available in Plato's heaven. On the simplest picture of the truths here, there are logical truths and metaphysical truths (i.e., necessary truths that aren't logical truths). Which of these are unimportant? The logical truths are all derivable from no premises whatsoever, and are interderivable; the metaphysical truths are truth in all metaphysically possible worlds and mutually entail each other. For any system in which some such truths are proposed as basic, there is another system in which other truths are proposed as basic. What, then, could account for the importance of some such truths and the unimportance of others? I can think of one such property some might wish to appeal to: some such truths are more complex than others. I have no idea, however, why anyone would think that such a property is a mark of importance or unimportance. If no such account can be found, however, then an account of the value of truth in universal yet defeasible terms is able to withstand any of the purported counterexamples to the view, just as the similar view about the badness of pain can withstand counterexamples to it.

None of my critics here, including Riggs, actually endorses the arguments for the limited perspective, however, so I won't pursue this issue any further here. Instead, I want to focus on the questions concerning understanding and the precise nature of the Meno problem. I turn first to the issues surrounding the nature of understanding.

ON UNDERSTANDING

Let's call the view I defend the "quasi-factive view." The form of understanding that I am interested in, and which I claim escapes the Meno problem, is objectual understanding, the kind of understanding in which the content of the attitude is an object of some sort (person, theory, part of reality, etc.). Such understanding is related, presumably constitutively, to various pieces of information, and on the quasi-factive view, the pieces of information that are central to the understanding in question must be true.

My critics raise important concerns about the defensibility of the quasi-factive view, but in my opinion, many of these objections can be explained away in one of two ways. Some concerns can be explained away in terms of honorific attributions of understanding and others in terms of the vagueness and ambiguity of ascriptions of understanding. Let me say a few words about the latter issue first. Concerning vagueness, the lesson is clear. On any view, there will have to be some connection to the facts in order for objectual understanding to obtain, but it will always be vague precisely what the connection needs to be. On the quasi-factive view, vagueness is found twice over. First, there is the vagueness of exactly what amount of information can be tolerated as false and yet understanding be

retained. Second, there is vagueness in the distinction between central and peripheral pieces of information. When falsehoods are found within the noetic system of a particular person, it will often be vague whether that falsehood is sufficient to undermine understanding. This result is just as it should be, since, as Aristotle informed us, we should not look for more precision than the subject matter allows. More relevant to the cases presented by my critics, however, is that ambiguity is found as well in attributions of understanding. I think something like this is occurring in Riggs' example concerning his wife.⁶ He claims to understand his wife, but when you look at the example, it looks to me like the object of his understanding is his wife's psychological constitution. He understands her moods, her fears, her motives, how she thinks, what her interests are, and the like. For such understanding, it is not a central piece of information precisely how she came to have these features. Had the claim to understanding been one about his wife's life story, her history, I'm inclined to say that the story doesn't sustain that attribution, since what appears to be a major event in her life is misunderstood. Perhaps not, however, since what caused a central feature of her psychological makeup need not be central from the point of view of her life's story. I suspect that is not the case, since any decent biography would seem to need an account of how she came to have the particular aversion in question. In any case, once we disambiguate the ascription of understanding, we can say correctly that Riggs understands his wife's psychological constitution even though he doesn't have a good grasp of some central biographical facts about her.

The other point of note in responding to apparent counterexamples to the quasi-factive view of understanding is to recognize that epistemic terms are often used in an honorific sense. The clearest example of such a use is when we talk about the present state of scientific knowledge about a certain

phenomenon. If pressed, most such ascriptions would be retracted: we speak this way to express some laudable epistemic standing for a certain conception of the phenomenon in question, but we would deny in many cases that such a conception is literally known to be true. Such honorific uses can be found using the notion of understanding as well. On such case is Elgin's example of the second grader's understanding of evolution. Elgin notes rightly that I will treat such cases on the model of language describing the current state of scientific knowledge, and she objects that such extended usages fall outside the scope of epistemology, on my view. She says,

When we construe such a take on a subject as understanding, Kvanvig believes, we use the term 'understanding' in an honorific sense, just as we use the term 'knowledge' in an honorific sense when we speak of 'the current state of scientific knowledge', while conceding that some of what belongs to the current state of scientific knowledge is false. Such honorific usages of epistemic terms are, he believes, extended usages that fall outside the scope of epistemology. ⁷

I think this objection is not quite accurate, however. If I were to offer a theory of knowledge, I would not expect it to answer to locutions involving the current state of scientific knowledge, and when we are talking of understanding, we should not expect the theory to answer to honorific ascriptions of understanding of evolution to 2nd graders. This point doesn't mean, however, that honorific uses of epistemic terms are not suitable items of epistemological theorizing. They would fall under what I would metaphorically refer to as the pragmatic dimension of epistemic terminology, and it is a worthwhile epistemological project to determine exactly what epistemic reality underlies such uses. My view isn't that such uses fall outside the scope of epistemology, but rather only that we should not confuse real

knowledge and understanding with whatever epistemic reality underlies such honorific uses.

This treatment of understanding, as falling in some safe harbor protected by appeal to honorific ascriptions on the one hand and the phenomena of vagueness and ambiguity on the other also helps with the case of scientific understanding that Elgin raises. The heart of the objection, Elgin points out, involves the idealization characteristic of scientific models. Since these models are essential for scientific understanding, and since idealization results in models which are inaccurate in central ways, it appears that scientific understanding is not quasi-factive in the way I claimed. Quoting from her paper,

But there is another aspect of science that is even more troublesome for Kvanvig's view. That is science's penchant for idealization. . . . Elimination of idealizations is not a desideratum. Nor is consigning them to the periphery of a theory. . . . I suggest that effective idealizations are felicitous falsehoods. That they are false is evident. They are felicitous in that they afford epistemic access to matters of fact that are otherwise difficult or impossible to discern. Idealizations are fictions expressly designed to highlight subtle matters of fact. They do so by exemplifying features they share with the facts. . . . These felicitous falsehoods are not fictions. Fictive sentences neither are nor purport to be true. They function in other ways. ⁸

I think Elgin is onto something quite important here, but it is not, I think, the idea that understanding is not quasi-factive in the sense I maintained. The issue here concerns the object of understanding. One might understand the model or theory itself, as when one understands phlogiston theory. One does not thereby understand combustion, however. Understanding the world scientifically is not simply a matter of understanding the given model but involves, rather, some relationship between

the model and reality. Scientific understanding of the sort in question consists in the possession of a model and a realization of the extent to which it is an idealization and what aspects of reality the model is intended to shed light on. Here Elgin agrees, I believe: she says that the falsity of the idealizations are “evident”. Of course, she might mean that the falsities in question are evident only to enlightened epistemologists, but I suspect she realizes that scientists themselves are aware of the idealizing that is occurring and that the understanding made possible by the models they construct involves grasping the nature of the connection between these models and reality. What follows from these points is that the understanding made possible by the sophisticated models of modern scientific inquiry cannot be identified with the information in the model, but must advert to the modeling relationship itself and details about it, including information about the extent to which the model idealizes reality.

Notice as well that this last point is one that Elgin quite correctly stresses. Idealizations are felicitous in virtue of “exemplifying features they share with facts.” So the understanding generated by the model depends not only on an understanding of the model itself, but the ways in which the model mirrors, and does not mirror, the facts. Once we move past the model itself as being the body of information, however, the point about idealization doesn’t yield non-factivity for understanding, but rather affirms it. There is much more to be said here about the precise nature of scientific understanding, but we don’t need the details to see the point that the quasi-factive view of understanding tells the right story here once we recognize the fact that the object of understanding is not simply the model itself but some more complex thing involving a relationship between the model and reality.

ON THE PROPER UNDERSTANDING OF THE MENO PROBLEM

The other major issue I want to respond to concerns the attempts by my critics to undermine the claim that there is a Meno problem of the sort I describe. Greco, for example, takes the heroic line that there simply is no presumption that knowledge is more valuable than its parts, and Riggs holds that a suitably subtle rendering of the problem leaves room for a decent answer to the problem in terms of the notion of credit for true belief. He says, “Another reason is that I hope some version of the credit theory of knowledge (e.g.—the one defended today by John Greco) will be able to answer Kvanvig’s charges and account for the value of knowledge.”⁹

Before addressing these issues, however, let me clear up a couple of issues that Greco and Riggs talk about that do not accurately reflect the position I argued for. Greco worries that there is some ambiguity in what I wrote about what the Meno problem really is, and Riggs is concerned that what I say about understanding conflicts with how I describe the Meno problem. There is some truth in both charges, but a proper understanding of the position I was arguing for will show that there is nothing theoretically worrisome here.

A brief summary of the view I defended will help avoid some side issues. First, the Socratic worry about the value of knowledge over true belief, transposed into the key of contemporary epistemology, is the worry about the value of knowledge over any proper subset of its parts. To avoid repetitiveness, I often put the issue in the book using less precise language, but that was for stylistic effect only. At no point did the work focus on any other question. In characterizing the worry in this way, I did not endorse the idea that knowledge really is, in fact, more valuable than any proper subset

of its parts. Instead, my approach was as follows. The history of epistemology focuses on the nature and extent of knowledge (and elements such as justification taken to be partially constitutive of knowledge). Such a focus involves a narrowing of focus from the more general issues concerning cognitive achievements and excellences of a purely theoretical sort, and I wondered why an interest in such achievements and excellences should focus so exclusively on one particular kind of excellence or success. The beginning of an answer to this question points out that we want to get to the truth but we want more besides. Think of a purely hypothetical meeting on marketing strategies for, say, some pulp journalistic publication. The marketers are trying to decide which human motivation to tap into. One marketer suggests belief itself he's a fan of fundamentalism and remarks that holding beliefs and being certain of them is so satisfying to human beings that they should use the slogan, "Inquiring minds want *opinions*." A second marketer scoffs, suggesting instead, "Inquiring minds want *the truth*." The third and fourth marketers are former students of Feldman and Sosa, respectively, and they resist both suggestions, offering instead, "Inquiring minds want to justifiably believe the truth," and "Inquiring minds want to virtuously believe the truth." Talk of firing the marketing firm begins to circulate, since these proposals are so hopeless, until the bright young former Gettier student (turned marketer) exclaims, "I know what we should use! We should use "Inquiring minds want to *know!*""

The publication makes millions relying on that ad campaign.

What is instructive about the example is not only the intuitive superiority of the slogan to its competitors and the central role that concept of knowledge plays in it, but the natural way in which the proposal is first offered. Notice the force of saying, "I know what we should use!" Why appeal to the concept of knowledge here? This is something we do all the time, by the way, and doing so serves as

an inquiry-stopper. We either reject the claim to knowledge or we acquiesce with the suggestion. Notice that we don't do that for any of the other suggestions in the neighborhood. If someone says, "I have a true belief about what we should use as a slogan," such a remark doesn't function in this way. Nor does adding that the belief is either justified or virtuous in the appropriate way. An appeal to knowledge does.

When you conceive of yourself as knowing a given claim, you will be puzzled by any counsel to investigate the matter further. "Do you know who that is on the stage talking?" "Yes, I do; it's Kvanvig." "Shouldn't you check to make sure?" "Excuse me?" Knowledge is the kind of thing that licenses closure of inquiry from a purely theoretical point of view. Mere belief does not license such closure, and neither does mere true belief. Moreover, neither does justified true belief, as is shown by the lottery paradox. One can be justified in believing that one's ticket will lose without that justification legitimating the closure of inquiry. It is for reasons such as this that I claim that knowledge is ordinarily thought of, or assumed to be, more valuable than its proper subparts.

This position on the unique value of knowledge is buttressed by some fairly reasonable proposals about the connections between knowledge and assertion. According to Williamson, for example, knowledge is the norm of assertion, formulable pithily as the advice not to say what you don't know to be true. Slightly weaker is the view defended by Max Black, G. E. Moore, Robert Shope, and Peter Unger that, in asserting a claim, one represents oneself as knowing that the claim is true. By forming an acronym of the last names of this group of philosophers, we can affectionately refer to this position as the BUMS view. Whether one agrees with the knowledge norm view or the BUMS view, it is clear that these philosophers are onto something important about our ordinary conception of

knowledge. One of the platitudes about the functional role of knowledge ascriptions is that it is a legitimator of inquiry closure. Nothing similar can be said about belief, true belief, or justified or virtuous true belief. In each case, such theorists could quite naturally display the phenomenon of metalinguistic negation at such proposals: if you suggest that in assertion you represent yourself as believing correctly, a natural response by one of the BUMS would be, “no you don’t; you represent yourself as *knowing!*”

The argument that there is a Meno problem concerning the value of knowledge derives from these considerations. It is because attributions of knowledge function to legitimate the closure of inquiry that it ought to have a value that exceeds that of its proper subparts, for these parts do not legitimate such closure. It is banal to remark that Bo believes something that he ought to investigate further, or even that Joe ought to investigate further even though his believing is correct, a display of cognitive excellence, and justified or rational. Things are different, however, when we utter the perplexing “Bo knows . . . but ought to check further.”

One thing that seems to bother Greco here is the idea that ordinary folk haven’t thought enough about knowledge to think of it in this way.

Suppose also that the “plus” part amounts to some minor, technical adjustment to the traditional idea that knowledge is justified true belief, and that this further condition adds no further value to knowledge over justified true belief. Would these suppositions conflict with ordinary thought? Would this show “that knowledge does not have the kind of value it is ordinarily thought to have”? It is hard to see how that could be the case. Most people are not at all aware of Gettier problems, and we can suppose that virtually no one was before 1963.

But then how could ordinary thought include the idea that knowledge is more valuable than justified true belief?¹⁰

To sum up, there is no pre-theoretical conviction that knowledge is more valuable than any subset of its constituents, and therefore it is not appropriate to require that a theory of knowledge explain why knowledge has that sort of value. Kvanvig's criterion for an adequate solution to the value problem is too strong.¹¹

Greco voices two criticisms here. The first is about how the ordinary conception of knowledge could include features that seem to require awareness of the Gettier problem. This issue arises in part because I wrote carelessly regarding the value that knowledge is “ordinarily thought to have.” Two points are worth recognizing here, however. The first point is a subtle one about the relationship between thoughts, ideas, assumptions, presuppositions, and the like. Our behavior in ascribing knowledge can reveal assumptions and presuppositions about knowledge that are not in our thoughts at all, just as one's behavior can reveal that one is a racist even though one would sincerely disavow it. Our linguistic behavior can show that we are assuming or presupposing that knowledge has some special value even if no one had ever reflected on the issue of the value of knowledge.

One might resist here on grounds that assumptions at least have to be non-occurrent mental states akin to dispositional beliefs, and that ordinary folk not only have never thought about the Gettier problem, they are in no mental state whatsoever that involves a content concerning that problem. I doubt, however, that this view is sustainable. I have two reasons for doubt here. The first questions the claim that assumptions and presuppositions need to be mental states. Why wouldn't it be enough to assume something that one be strongly disposed to believe it, as opposed to believe it already (albeit

only dispositionally)? It is one thing to believe a claim dispositionally, and a different thing to be disposed to believe it. In order to sustain the claim that assumptions must be mental states, we'd need an argument against the suggestion that dispositions to be in a relevant mental state could be enough to count as having made an assumption.

The second point makes the stronger point that even if assumptions are mental states, they are not plausibly thought of as dispositional beliefs. In many cases where our assumptions are pointed out to us, we experience chagrin at the realization. Upon thinking about the particular propositional content in question, we do not embrace it. Instead, we reject it. So what explains the chagrin? To experience chagrin in such an immediate fashion, it would seem that we need some propositional attitude in place that runs contrary to the assumption in question. Suppose, for example, that you are strongly averse to racism, but respond on a given occasion in terms that you agree are accurately described in terms of assuming that people of a particular race are more dangerous. You experience chagrin upon having this fact pointed out to you.

Your chagrin depends upon your aversion to racism, and such aversion involves, I would expect, cognitive commitments. You believe, or are committed to, lots of claims, including the exact opposite of the assumption underlying your response. But in general we don't want to try to make sense of ascribing both the belief that p and the belief that not-p to an individual (except in cases where modes of presentation explain away the absurdity in question). Nor do we want to say, I think, that some of your beliefs about the races somehow went out of existence during the period of your assumption. Neither will appeal to degree or strength of belief help, so long as we are still willing to countenance the reality of the distinction between beliefs and non-beliefs (even if the threshold varies by

context).

The above is a bit of philosophical overkill, since most of it isn't needed to make the basic point that the fact that most people have not considered the Gettier problem and have not reflected on the value problem regarding knowledge does not show that there is no assumption or presupposition that knowledge is more valuable than any proper subset of its parts. So when Greco claims that there is no "pre-theoretical conviction" that knowledge has some special value, we should perhaps grant the point but deny its relevance. If convictions are conscious, occurrent mental states, then I agree that there is no such generally shared conviction. That leaves untouched, however, the point about assumptions and presuppositions, as well as the data about the functional role of knowledge ascriptions in terms of legitimating the closing of inquiry.

Even worse, if Greco were right that there is no assumption that knowledge is more valuable than any proper subset of its parts, the focus in the history of epistemology on the nature and extent of knowledge would be downright indefensible. Recall that the more general issue for reflection is one concerning cognitive successes and excellences from a purely theoretical point of view (i.e., from a point of view abstracting away from other purposes such as purposes which are practical, moral, aesthetical, religious, or political in nature). Why should the history of reflection on such successes and excellences focus so centrally on knowledge? If knowledge had no special value, there would be no hope of vindicating the history of this philosophical subdiscipline.

This point ties directly into Riggs' worry that what I say about the Meno problem may not fit well with what I say about understanding.¹² He is right about this, but the result is one that should be expected, given the nature of the project I undertook. It is important here that the characterization of

the Meno problem with which we begin is only a working hypothesis. It involves a plausible characterization of the special value knowledge is assumed to have, and it characterizes a viewpoint that has the power to substantiate the obsession the history of epistemology shows concerning the concept of knowledge. In the end, I do not think that this assumption will survive scrutiny. The point is, rather, that one should approach the theory of knowledge with the working hypothesis that knowledge is more valuable than its proper subparts. If I'm right, this hypothesis cannot withstand careful scrutiny, so it should be no surprise if, in defending the special value of understanding, I say things that do not fit well with the working hypothesis with which I began. The working hypothesis is false precisely because the myopic focus in the history of epistemology on the nature and extent of knowledge cannot be defended. The result may be that some special value can be found for knowledge, but it won't be a special value of the sort that would be needed to justify the singular attention to the concept of knowledge that the history of epistemology displays.

So there is a Meno problem and it needs a good solution. Perhaps Riggs and I are right that a better understanding of the problem will show that its proper construal is not enough to justify the history of epistemology, but we should look first to see whether a construal of the problem that has hope of vindicating that history has any chance of success. Here Greco's work is important, for he thinks his virtue account can give us everything we want, and Riggs indicates some agreement on this point. As I argue in the book, we should expect such heroism to fail either on the issue of the nature of knowledge or on the issue of the value of knowledge. In my book, I argued in favor of the idea that beliefs formed through virtue display or indicate a special kind of value, and hence that virtuous true belief is more valuable than true belief. What I argued is that such a view can't explain the value of

knowledge, since knowledge is more than virtuous true belief.

This point plays out in expected fashion in some of Greco's own remarks. Greco's own proposal to the Meno problem is as follows:

Knowledge is a kind of success through virtue. And in general, success through virtue is more valuable than either success without virtue or virtue without success. In particular, virtuously produced true belief is more valuable than true belief that is not virtuous and virtuous belief that is not true. Neither subset is intrinsically valuable, or constitutive of what is intrinsically valuable, in just the way that knowledge is.¹³

Note that Greco and I agree that success through virtue is more valuable than success without virtue. That is, we agree that virtuous true belief is more valuable than mere true belief. What I deny is that this point solves the Meno problem. To do that, the account would have to be an adequate account of the nature of knowledge, and it's not. Familiarity with the Gettier literature reveals a variety of cases to which it succumbs: Ginet's fake barn case, Harman's assassination case, the Lehrer/Paxson Tom/Buck case, and even Gettier's original inferential cases. In the fake barn case, the traveler's belief that the thing in the field is a barn is a product of excellent perceptual abilities. In Harman's case, Jill believes what the reporter tells her in virtuous fashion, and the reporter reports sincerely what he saw, where the perceptual equipment in question can be as superior as one cares to stipulate. In the Tom/Buck case, you know that Tom stole the book because you saw him do it, and the facts about belief production remain constant whether or not the testimony of Tom's mother amounts to a misleading defeater. In the bad case, the police have no record of contact with her, and so the testimony implies a lack of knowledge on your part. In the good case, the police have a long file on her

and don't need to take the story seriously, and so the testimony counts only as a misleading defeater. What is telling is that the facts about the nature of your belief and whether it counts as an instance of success through virtue remain the same across the cases, so success through virtue cannot account for the possible variants of the case. Finally, in Gettier's original inferential cases, there is no reason to suspect that there is some defect of intellectual character or failure to display it. On any ordinary reading of the notions in question, success is achieved through a display of positive character traits.

Greco holds, however, that such an understanding of the key notion of success through virtue is not the one he intends, and uses the fake barn case as a test case for clarifying that notion in a way that provides some hope of escape from the apparent counterexamples above. He says,

[I]t is unclear how perception can count as an ability relative to S's environment and yet be unable to distinguish real barns from fake ones. But suppose we agree for the sake of argument that S's belief is not only true but also formed from a virtue. Still, the belief is not true *because* it is formed from a virtue. Put more carefully, the belief's being so formed does not explain why S has a true belief rather than a false belief. On the contrary, S believes the truth because she happens (accidentally) to be looking at the one real barn in the area. If she had been looking anywhere else nearby, excellent perception or no, she would have a false belief.¹⁴

There are two key claims here. The first sentence suggests the possibility of taking a truly heroic path of trying to explain away every Gettier case as involving a person who simply does not have the requisite abilities or virtues in question. Jill's believing that the President was assassinated couldn't count, on this path, as a display of cognitive abilities; your belief that Tom stole the book couldn't

either; etc. The trick to this approach is the relativization move: it is not success through virtue that is to be identified with knowledge, but rather success through virtue-relative-to-environment. It is clear that the relativization move can be used to carve off all the cases of knowledge from the cases on non-knowledge: just gerrymander the individuation of environments enough to get that result. The problem, however, is that such gerrymandered individuation will not be useful in the project of trying to explain the value of knowledge but will instead undermine any such account. That, as I have argued, is the price of gerrymandering and ad-hoc-ery.

To see how this charge plays out in a given case, consider variants on the Tom/Buck case. In one variant, you know that Tom stole the book because the police have this long track record with Tom's mother, so that her testimony counts only as a misleading defeater. In another variant, you don't know that Tom stole the book because the police have no such track record and have to take her story seriously. The question for the heroic path here is how the relativization is supposed to work. The answer is that the two "environments" have to differ, in spite of the fact that they can be as far away from you geographically as you please. There is no interesting intuitive notion of 'environment' on which the facts about the prior relationship between Tom's mother and the police constitutes any part of the environment in which you observed Tom steal the book and flee from the library. One can, of course, concoct a technical notion to get the desired result, but one does so only at the expense of drawing distinctions that clearly have no relevance whatsoever to the value problem, leaving unexplained why one should value knowledge over its lack in the case in question.

It is not clear, however, that John intends to take the heroic path, for after suggesting it he quickly moves to a different approach. On this approach, the crucial issue is a causal/explanatory one.

On it, we acknowledge the obvious point that in a broad range of classic Gettier examples, a true belief results from a display of cognitive virtues or excellences. What is missing, according to Greco, is the idea that the true belief is achieved *because of* the virtues, the virtues do no *explain why* the person in question has a true belief rather than a false one.

I think there are three fundamental difficulties with this proposal, and I'll take them in increasing order of importance. The first point concerns the evidence that Greco cites for concluding that the virtues do not explain, in the fake barn case, why one has a true belief rather than a false one. Notice that John infers that the belief's being the product of a virtue does not explain why the belief is true rather than false from the counterfactual claim that if she had looked anywhere else nearby with the same virtues, she would have had a false belief. This counterfactual claim, however, is, in general, false. It is only true when we restrict the claim to what she takes to be a barn. Regarding trees and shrubs, cows and horses, and people and houses, everything would be fine. So the counterfactual has to be very specific to the content of the belief in question to support the conclusion Greco wishes to draw, and even then the claim need not be true because of the possibility of Frankfurt cases. Let the demon be interested in your having true beliefs about barns. The landscape is still littered with fake barns, but the demon will do nothing unless you direct your attention to a fake, in which case he'll zap it into a real barn. But you don't look at the fakes, you look at the one real barn. The demon does nothing. You don't know, however, that the object in question is a real barn, and this even though had you looked at any other barn-like thing in your environment you would have had a true belief. In short, one cannot use the truth or falsity of the counterfactual in question as decisive evidence for the explanatory claim Greco uses to try to avoid the fake barn case.

Moreover, the strategy instanced here for the fake barn case is hardly capable of being exported to all the other cases. In particular, it won't be any help in the variants of the Tom/Buck case noted earlier. The difference in the two variants concerns differences in the past interaction between the police and Tom's mother, but it is simply implausible to think this difference implies any difference in the explanation of why your belief is true rather than false.

The second point is that it is unclear that there is any notion of explanation that will generate the desired result. In the fake barn case, the perceptual abilities of the observer do explain, in part, why a true belief is achieved. Those abilities are not the only factors relevant to a complete explanation, but the obvious fact is that our abilities never constitute a total explanation of why we get to the truth rather than fail. Only infallible abilities would be good enough to yield a complete explanation on the basis of the abilities alone. So the critical question such an account must face is why it will focus on one part of the complete explanation in one case but not in another. Here the philosophical investigation of the term 'the cause' is probably instructive. The result of this investigation is pretty clearly the following: what counts as the cause of a given event is not any set of semantically invariant features. Instead, what counts as the cause of an event seems to vary with the interests and purposes involved in a given context or conversation. For example, what was the cause of the forest fire, the campfire not completely put out, the trail of trash strewn by the bear, or the decade-long drought that allowed the smoldering campfire to move to nearby leaves, then to the trash, and finally to the entire forest? Environmentalists will cite the campfire and Republicans will favor the drought, while the scientist will, in all likelihood, abstain on the project of selecting a distinguished factor from the set of causal factors underlying the forest fire.

The pessimistic conclusion grounded in the literature on the notion of the cause is that there is no interest-independent notion here to be found.¹⁵ If that conclusion is correct, then the best Greco could hope for is that our language of explanation tracks our interest in knowledge over non-knowledge, and even if we are interested in knowledge, there is no reason to suppose that the language of explanation will always favor this interest over other interests. Since that assumption is surely implausible, Greco's proposal will have the result that attributions of knowledge have to track whatever interests are dominant, *even when these interests are clearly trumping our interest in knowledge*. That's interest-relativity in the theory of knowledge that is indefensible.

The third problem with Greco's proposal is that it is hard to reconcile with the variety of kinds of knowledge that are possible. One kind is testimonial knowledge of a mundane sort, where you simply take the word of someone about a given topic. Children learn in this way much before they acquire the critical perspective to assess their sources, and any suitably nuanced attempt to sort explanatory factors properly should put more stress on the virtues of the source instead of on the virtues of the child in explaining why the child has acquired a true rather than false belief. There is, of course, the possibility of stipulating some notion of explanation that places greater stress on the child's virtues, but by now we should be familiar with the objection that results: the degree of stipulation is inversely proportional to the plausibility such an account possesses for addressing the value problem.

An even more difficult possibility than knowledge by testimony, however, is the possibility of innate knowledge. It may be that we have no innate knowledge, but in some epistemologically interesting sense, the possibility of such knowledge exists. Any account of the value of knowledge should be able to explain the value of innate knowledge as well, but the idea that innate knowledge

involves some kind of success through virtue, at least on the part of the individual in question, is difficult to imagine. As before, I'm sure we can stipulate some understanding of success through virtue to yield this result, but stipulations don't explain value.

Conclusion

I believe, therefore, that there is still an unsolved Meno problem and that future attempts to solve the problem as I have stated it will fail. I have no compelling argument for the prediction, however, and the discussion here is a good illustration of the direction discussion must take about this prediction. Those less pessimistic than I will continue to propose solutions, and those sharing my pessimism will continue to try to demonstrate that the solutions do not work.

About the quasi-factive view of understanding I defended, I believe there is much more to be said both by way of elucidation and by way of criticism. In one way, this conclusion is wholly unsurprising, since there has been so little discussion in the history of epistemology concerning the nature of understanding. I will end by expressing the hope that future work in the field will remedy this deficiency.

Endnotes

1. Ernest Sosa, "The Place of Truth in Epistemology," in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, edited by Michael DePaul and Linda Zagzebski (Oxford University Press, 2002).
2. Robert Roberts and Jay Wood, *Virtues that Deliver the Epistemic Goods*, (Oxford University Press, forthcoming 2007).
3. Stephen Grimm and Michael R. DePaul, "Review Essay: Kvanvig's The Value of Knowledge and the Pursuit of Understanding," *Philosophy and Phenomenological Research*, forthcoming 2007.
4. Wayne Riggs, "Getting the Meno Requirement Right," pp. 3-6, manuscript. I should note that Riggs says, on p. 6, "I am not endorsing any of these views..."
5. I would hope that the following goes without saying, but to prevent any misunderstanding, I will point it out anyway: the thought experiment in question doesn't assume that there is such a thing as Plato's heaven or even that it is possible that there is such a thing. It is a rhetorical device only.
6. "Getting the Meno Requirement Right," p. 9, manuscript.
7. Katherine Elgin, "Is Understanding Factive?" p. 7, manuscript.
8. Elgin, pp. 9-10, manuscript.
9. "Getting the Meno Requirement Right," p. 3, manuscript.
10. John Greco, "The Value Problem," p. 6, manuscript.
11. "The Value Problem," p. 8, manuscript.
12. "Getting the Meno Requirement Right," pp. 12-13, manuscript.
13. "The Value Problem," pp. 12-13, manuscript.
14. "The Value Problem," p. 16, manuscript.
15. See, for example, J.L. Mackie's *The Cement of the Universe*, as well as the papers in the collection edited by Ernest Sosa, *Causation and Conditionals*, (Oxford University Press, 1975) as well as the more recent collected edited by Collins, Hall, and Paul, *Counterfactuals and Causation*, (MIT Press, 2004).