

SOSA'S VIRTUE EPISTEMOLOGY

JONATHAN L. KVANVIG
Department of Philosophy
Baylor University
Jonathan_Kvanvig@baylor.edu

SUMMARY: Ernest Sosa's latest epistemology remains a version of virtue epistemology, and I argue here that it faces two central problems, pressing a point I have made elsewhere, that virtue epistemology does not present a complete answer to the problem of the value of knowledge. I will press this point regarding the nature of knowledge through variations on two standard Gettier examples here. The first is the Fake Barn case and the second is the Tom Grabit case. I will argue that Sosa's latest virtue epistemology fails to handle either case acceptably, and that as a result, cannot explain the value that knowledge has over that of the sum of any of its proper subparts.

KEY WORDS: defeasibility, Gettier problem, value of knowledge

RESUMEN: La última epistemología de Ernest Sosa continúa siendo una versión de epistemología de las virtudes, y siguiendo con una idea que ya planteé en otra parte, aquí argumento que afronta dos problemas centrales: que la epistemología de las virtudes no presenta una respuesta completa al problema del valor del conocimiento. Insistiré en esta idea sobre la naturaleza del conocimiento mediante variaciones de dos ejemplos estándares tipo Gettier. El primero es el caso del granero falso y el segundo es el caso de Tom Grabit. Argumentaré que la última epistemología de las virtudes de Sosa no maneja ninguno de estos casos de forma aceptable, y en consecuencia no puede explicar el valor que tiene el conocimiento por encima del de la suma de cualesquiera de sus propias partes.

PALABRAS CLAVE: derrotabilidad, problema de Gettier, valor del conocimiento

Ernest Sosa is the father of modern virtue epistemology, having first introduced such an approach near the end of “The Raft and the Pyramid” (Sosa 1980). His most recent, and most developed, version of the view is contained in *A Virtue Epistemology* (2007). The view relies on a distinction between animal and reflective knowledge, attempting to understand both kinds of knowledge in terms of the notion of apt belief. An apt belief is, according to Sosa, one that is true because competent, and it is a virtue epistemology because of this use of the notion of competence. To have a competence is to have an intellectual virtue, and thus knowledge is, at bottom, understood in terms of the intellectual virtues.

As with other versions of virtue epistemology, this version faces several kinds of problems, and here I will focus on two of them. I will press a point I have made elsewhere, that virtue epistemology does not present a complete answer to the problem of the value of

knowledge.¹ In order to have provided a complete answer to the problem of the value of knowledge, any given virtue epistemology will have to clarify the nature of knowledge in terms of the virtues, and then explain how virtuous true belief is more valuable than true belief itself. I have argued that the latter point is correct (that a true belief that is a display of an intellectual virtue is more valuable than a true belief that is not), but I have argued as well that virtue epistemologies fail to provide an adequate account of the nature of knowledge.

I will press this point regarding the nature of knowledge through two standard Gettier examples here. The first is the Fake Barn case and the second is the Tom Grabbit case.² I will argue that Sosa's latest virtue epistemology fails to handle either case acceptably, and that, as a result, cannot explain the value that knowledge has over that of the sum of any of its proper subparts. I take up each case in the next two sections.

1. *The Fake Barn Case*

The fake barn case runs as follows. You are driving through some rural area, perhaps some part of Wisconsin. The locals, bored with ordinary farm life, have decided to play a trick on visitors, and so have tried to replace all the barns in the area with fake barns. They inadvertently leave one real barn in place. As you are driving through the countryside, you take notice of various objects: houses, cars, horses, cows, pigs, fields of corn and other crops, etc. You only notice one barn-like object, and it happens to be the only real barn in the locale. You believe it is a barn, and your belief is an ordinary perceptual one, and one that is true. But because of the activity of the locals, you do not know that it is a barn.

At first glance, this example is a problem for Sosa-style virtue epistemology. The person in question possesses an intellectual virtue: he or she is competent perceptually, and is using this competence in conditions that are appropriate for its exercise. The result of such a display of competence is a true belief. It thus appears to be a case of virtuous true belief, but not a case of knowledge.

A careful reading of Sosa's claims about knowledge may make one wonder whether there is room for maneuvering to avoid this first impression. Sosa says, "We can distinguish between a belief's

¹ I argue this point at length in Kvanvig 2003.

² Carl Ginet is the inventor of the fake barn case, but it was made famous by Goldman 1976. The Tom Grabbit case is found in Lehrer and Paxson 1969.

accuracy, i.e., its truth; its adroitness, i.e., its manifesting epistemic virtue or competence; and its aptness, i.e., its being true because competent" (2007, p. 23). One might wonder whether this appeal to the *because of* relation can help avoid the fake barn case, and it is worth noting that one of the primary defenders of Sosa-style virtue epistemology, John Greco, makes just such an appeal. Concerning the barn belief, Greco says,

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. (Greco 2009, p. 318)

The idea, then, is to distinguish between a case in which a true belief arises out of one's cognitive abilities or virtues or competences and a case in which a belief is true because of the display of competence in question. Both could legitimately be called examples of virtuous true belief, but, according to Greco, only the second is to be identified with knowledge.

I have argued against this proposal elsewhere (Kvanvig 2009), and it is interesting to note that Sosa agrees that this proposal will not work. One of the difficulties with Greco's proposal is the problem of testimonial knowledge in young children. In the fake barn case, Greco focuses on the accidental character of being correct, and on the counterfactual point that focusing on a different part of the landscape would have resulted in a false belief. Being right, in the barn case, owes too much to other factors, Greco holds, to legitimately be creditable as an achievement to the believer because of a display of an ability or competence. It is precisely at this point that the problem of testimonial knowledge is most pressing, and it is worth noting that Sosa spends considerable time in *A Virtue Epistemology* arguing for the view that partial credit for true belief is credit enough and achievement enough to count: "Testimonial knowledge can therefore take the form of a belief whose correctness is attributable to a complex social competence only partially seated in that individual believer" (Sosa 2007, p. 97).

If so, however, the fake barn cognizer is going to pass scrutiny as well, as Sosa recognizes. Sosa holds that the fake barn belief is, in fact, an apt belief: it is a belief that is true because competent. If testimony early in life leads to beliefs that are true because of the

competences of young children, in spite of the fact that the total explanation must appeal to competences and practices that reach far beyond these individual competences, then a belief about a barn in the landscape must be judged similarly. It too is a belief that is true because of the competences of the cognizer in question.

Sosa thus parts company with other defenders of Sosa-style virtue epistemology. Though he does not explicitly connect his discussion of testimony with the fake barn case, the connection is fairly immediate and direct, and fits well with what he does say explicitly about the fake barn case. Surprising as it may seem, Sosa grants that the belief in question is an instance of apt belief — “surprising” because such an admission appears to doom his virtue epistemology. Sosa believes, however, that the language of demise is premature, for he thinks that the standard response to the fake barn case, the response that denies that the cognizer in question knows that the object in question is a barn, is mistaken. His response to the fake barn case, then, is not that of crafting an account of knowledge immune to this counterexample, but rather to explain away the sense that the true belief in question isn’t knowledge. He says, in a footnote,

Our account does help to bring out, however, how not all Gettier cases are created equal. In some cases, such as Gettier’s two actual examples, and such as Lehrer’s Nogot/Havit case, the subject does not attain so much as animal knowledge: apt belief, belief that gets it right in a way sufficiently attributable to the exercise of a competence in its proper conditions. However in other similar cases, what the subject lacks is rather reflective knowledge. Our kaleidoscope perceiver, in Lecture 2, is a case in point. The Ginet/Goldman barns example arguably belongs with the kaleidoscope case. (Sosa 2007, p. 96, fn. 1)

The strategy is thus to explain away, rather than accommodate, the fake barn case. Goldman and others are mistaken in thinking that the cognizer in question lacks knowledge. According to Sosa, the reason for the mistake is that they are confusing animal knowledge with reflective knowledge.

We must ask, then, about this distinction and how it helps with the barn example. Sosa’s discussion of the distinction and its application to given cases proceeds using the kaleidoscope example referred to in the passage just quoted: a case in which you see a red surface, but where the lighting is controlled by a jokester who might just as easily have presented you with a white surface/red lighting situation as with the actual red surface/normal lighting situation. After noting

that his account allows such a belief to be a case of apt belief and hence of knowledge, Sosa remarks:

Recall, however, our distinction between two sorts of knowledge, the animal and the reflective. Any full account would need to register how these are matters of degree. For present purposes, however, the key component of the distinction is between apt belief *simpliciter*, and apt belief aptly noted. If K represents animal knowledge and K^+ reflective knowledge, then the basic idea may be represented thus: $K^+p \leftrightarrow KKp$. (Sosa 2007, p. 32)

By identifying reflective knowledge with iterated animal knowledge, Sosa preserves the unity of his theory, with the notion of apt belief being the unifying factor. The idea, then, is to explain away the kaleidoscope example and the fake barn example in the same way. Both cases are cases of apt belief and hence of knowledge, but they are not cases of apt belief aptly noted, and hence not cases of reflective knowledge. They are not cases in which one knows that one knows.

This point alone, however, will not explain away the cases. Most simple perceptual beliefs are cases in which one has knowledge but also in which one fails to know that one knows. We simply don't reflect that often on simple perceptual beliefs, and when we don't we won't have reflective knowledge. The fake barn case is thus no different in this respect from ordinary, unreflective perceptual beliefs, and yet, as Sosa recognizes, it is awkward to say the same thing about both ordinary perceptual and perception in the fake barn case. So how does the distinction between animal and reflective knowledge help?

Sosa's further discussion reveals that it is not the simple distinction between animal and reflective knowledge that is employed to explain away the fake barn case but rather a distinction between what kind of knowledge can obtain and can't obtain in a specific case. One might say, for example, that in ordinary perception, reflective knowledge is possible, but in the troubling examples, it isn't. Sosa notes that "in ordinary perceptual belief, one can aptly presuppose, or take it for granted, that the relevant competence and conditions are in place" (Sosa 2007, p. 35), but that in the kaleidoscope case, "we fall short of reflective knowledge [...] because the jokester precludes the aptness of our implicit confidence that our perceptual belief is apt" (Sosa 2007, p. 36). Such meta-confidence couldn't be apt because "[h]is being in control makes it too easy for us to be confident in that

default way, in normal conditions for the exercise of our perceptual competence, while still mistaken” (Sosa 2007, p. 36). It is thus not the distinction between animal and reflective knowledge that does the explanatory work, but a distinction between what kinds of knowledge are possible.

The remarks just quoted, however, are nonetheless perplexing. The idea is that in ordinary perception, we legitimately assume that conditions are normal, and that in the kaleidoscope case, this ordinary assumption is too easily mistaken. While true, the point doesn’t help unless somehow it is a requirement for reflective knowledge that all the assumptions made for animal knowledge must themselves be apt assumptions.

There are two problems with the claim that such assumptions must be apt in order for reflective knowledge to be possible. The first is that nothing can be apt except a mental state of some sort, and it is not clear that assumptions and presuppositions are mental states at all. We sometimes talk of implicit commitments and the like, and perhaps assumptions are such things. But I doubt it. Remember that animals and small children have a great deal of knowledge, and for many such individuals, they lack the conceptual resources needed even to entertain or consider anything about normalcy or aptness or perception. How then can they be in any mental state, whether explicit or implicit, that involves such concepts?

The second problem is that the inference made is faulty. The conclusion is that confidence that our perceptual belief is apt can’t itself be apt in the kaleidoscope case. The reason given for this conclusion is that, in that case, it is too easy for us to be confident “in that default way” and yet be wrong. The difficulty arises from considering which confidence Sosa has in mind. It can’t be confidence in the perceptual judgement itself, since that belief has already been judged to be apt. Is it the possible belief that the first-order belief counts as knowledge? That is not very plausible, since this belief is no more susceptible to being held while false than is the first-order belief. The jokester isn’t trying to mess with the reflective capacities that would be used to form the meta-belief, so the meta-belief is not a good candidate here. Yet, Sosa maintains that the fragility of the first-order belief doesn’t imply that it is inapt, so there would be no reason for thinking that any fragility inherited by the meta-belief would imply that it is inapt either. More likely is that Sosa is thinking of what is taken for granted in ordinary perceptual cases. What is taken for granted in ordinary circumstances, and aptly so

according to Sosa, is that the circumstances are normal for the kind of competence in question. The idea, then, is that what is aptly taken for granted in ordinary circumstances isn't aptly taken for granted when the possible belief is the belief that one knows that the surface is red when the jokester is present.

This result becomes clearer when we attend to Sosa's discussion of the connection between aptness and safety. The official account of aptness involves a belief being true because of a display of competence,³ and Sosa maintains that this account leaves open the possibility of a belief being unsafe but apt. The connection between safety and aptness is more limited. According to Sosa, it is only basis-relative safety that reveals aptness. Safety requires that a belief would have been held only if likely to be true (Sosa 2007, p. 25), where basis-relative safety is the property a belief has of having a basis which is such that it would likely have had only if true (Sosa 2007, p. 26).

In the kaleidoscope case, the belief that a given surface is red is, thus, unsafe (since the belief could easily have been false) and yet basis-relative safe (because the particular basis in question involves normal conditions for color vision, including "conditions of lighting, distance, size of surface, etc., in conditions generally appropriate for the exercise of color vision" (Sosa 2007, p. 31)). Yet, similar remarks seem initially in order about one's reflective belief that one knows that the color of the surface is red. That belief is unsafe (because the jokester easily might have rendered the belief false), but it is nonetheless basis-relative safe. It is basis-relative safe because the belief, imagined now to exist, is the product of a competence to reflect and distinguish what one knows from what one merely believes. That competence is displayed in conditions that are normal for its operation: the cognizer is not on drugs, is not emotionally conflicted in such a way as to be unable to reflect well or carefully, etc. In short, any of the ordinary ways in which conditions need to be in order for reflection to operate competently are in place. The jokester doesn't affect those conditions at all, though he might have so acted that the belief in question could easily have been false. That only shows, however, that the belief is unsafe, not that it is basis-relative unsafe.

³ "Aptness requires the manifestation of a competence, and a competence is a disposition, one with a basis resident in the competent agent, one that would in appropriately normal conditions ensure (or make highly likely) the success of any relevant performance issued by it" (Sosa 2007, p. 29).

So what Sosa must have in mind is something concerning the presupposition or assumption that the conditions for color vision are appropriate, and that it is this presupposition or assumption that is inapt and thus prevents reflective knowledge that the surface is red (prevents, that is, the knowledge that one knows that the surface is red). In fact, he says as much:

The perceiver would there [in the kaleidoscope case] be said to have apt belief, and animal knowledge, that the seen surface is red. What he lacks, we may now add, is *reflective* knowledge, since this requires apt belief that he aptly believes the surface to be red (or at least it requires that he aptly take this for granted, or assume it or presuppose it, a qualification implicit in what follows). (Sosa 2007, p. 32)

The thought pattern underlying these claims would seem to be as follows. Reflective knowledge is just second-order knowledge: it is knowledge of first-order knowledge. But first-order knowledge is understood in terms of apt belief, so to have second order knowledge is to aptly believe that one's first-order belief is apt.

Sosa then traces the consequences of requiring apt belief that one's first-order belief is apt:

Consider now the kaleidoscope perceiver's belief that he aptly believes the seen surface to be red. We are assuming that the competence exercised in that meta-belief is a default competence, one which, absent any specific indication to the contrary, takes it for granted that, for example, the lights are normal. Because of the jokester in control, however, the exercise of that competence might then too easily have issued a false belief that the lights are normal. (Sosa 2007, p. 34)

The route to a denial of reflective knowledge thus proceeds as follows. Sosa first identifies reflective knowledge with meta-knowledge: knowing that one knows. Since knowledge has already been identified with apt belief, Sosa then requires that one knows that one knows only if one aptly believes that one aptly believes (or at least that one can take for granted or presuppose or assume that one aptly believes). Now, the aptness of the first-order belief requires a default competence displayed in taking it for granted that the lighting conditions are normal, so an apt meta-belief that one's first order belief is apt requires aptly taking it for granted that the lighting conditions are normal. But because of the jokester, no such taking for granted can be apt.

This pattern of reasoning is repeated in Lecture 5 as well, where Sosa returns to the kaleidoscope case and expands on the account first given in Lecture 2. He says,

Nevertheless, the kaleidoscope perceiver and the ordinary perceiver are still dramatically different epistemically. They differ in whether they can know their respective conditions to be appropriately normal for the exercise of their perceptual competence. The jokester precludes the kaleidoscope perceiver from knowing this; but [...] the ordinary perceiver is not similarly affected. (Sosa 2007, pp. 105–106)

This explanation, I will now argue, runs contrary to the account of reflective knowledge with which Sosa began. That initial account clarifies reflective knowledge in terms of iterated animal knowledge. The explanation in the passage just quoted of failure of reflective knowledge in the kaleidoscope case requires a different account of reflective knowledge, one in terms of which one can know reflectively only if one knows, or is in a position to know, how the requirements for first-order knowledge have been met. On this alternative account, reflective knowledge requires more than iterated animal knowledge—it requires an embedding of first-order knowledge in a system of understanding or knowledge that provides an explanatory basis for knowing how one has achieved first-order knowledge, a system reminiscent of certain types of coherence theories. As I will argue, once we are careful to apply the apt-belief account carefully, there will be no good reason to deny reflective knowledge in the kaleidoscope case once we have granted animal knowledge in that case.

We can begin to develop the argument for this last claim by considering Sosa's more developed account from Lecture 5 of what it is to base knowledge on some indication or sign of truth. Sosa formulates the following necessary condition for basing on indications of truth:

(I') S has animal knowledge that p based on indication $I(p)$ only if either (a) $I(p)$ indicates the truth outright and S accepts that indication as such outright, or (b) for some condition C , $I(p)$ indicates the truth dependently on C and either (i) S accepts that indication as such not outright but guided by C (so that S accepts the indication as such on the basis of C), or else (ii) C is constitutive of the appropriate normalcy of the conditions for the competence exercised by S in accepting $I(p)$. (Sosa 2007, p. 105)

Applied to the question of animal knowledge in the jokester case, you know because condition (b)(ii) is satisfied. The condition C concerning the normalcy of lighting conditions may not guide belief formation, but that condition is “constitutive of the appropriate normalcy of conditions for the competence exercised by S ” in forming a belief on the basis of how things look.

Note, however, that (I') gives one no resources for claiming that reflective knowledge requires knowing or aptly taking for granted that condition (b)(ii) is satisfied by the first-order belief. Reflective knowledge is nothing more than meta-knowledge, on Sosa's official account of it, so the question is simply the question of whether, in the kaleidoscope case, one can know that one knows that a given object is red. On some accounts of reflective knowledge, one could know that one knows only if the particular story about how one has first-order knowledge is known to obtain. But Sosa's account is not such an account. It is a unified account of reflective and animal knowledge in terms of apt belief, with reflective knowledge being nothing beyond iterated animal knowledge.

So suppose we test such iterated animal knowledge to see if it passes scrutiny by condition (I'). The claim in question is thus not the claim that the surface is red (p), but that the individual in question knows that the surface is red ($K(p)$). To know that $K(p)$, (I') will require basing one's meta-belief on some indication $I(K(p))$ which indicates the truth in question outright or dependently on some further condition or by that condition being constitutive of normalcy for the competence exercised. Presumably, the source of one's meta-knowledge will be reflection itself, reflection perhaps about which of one's beliefs count as knowledge and which do not. So suppose, in the jokester case, you reflect. Conditions for such reflection are perfectly normal, since the jokester is operating at first-order level, not meta-level: his shenanigans concern the color of the surface, not your ability to reflect on which of your beliefs count as knowledge and which don't. You form the belief that you know that p , based on the indication of reflection. Conditions for reflection are as normal as they ever are. The content of your reflection might be as simple as this: “what in my present experience is experience of things regarding which I have knowledge? Well, my present experience is made up of things I believe from my senses, and they are good sources of knowledge, as well as things I believe based on less laudable sources, such as mere prejudice, hunches, and guesses of various sorts. So among the things I know are things got from my senses, which includes eyesight, and that means I know that the object in question is red.”

That is a perfectly ordinary kind of reflection, and it is occurring as much in its home territory as perception itself was occurring in home territory regarding the color of the surface.

I will belabor this point a bit, to make the point as clear as possible. Abnormality at the meta-level must be abnormality for the competence exercised, which is reflection, and so the operation of the jokester has no role to play in ascertaining whether the conditions for reflection are normal. Abnormality can arise in a number of ways: it could be drug-induced, temporary insanity, psychological impairments of the sort that prevent some people from being able to reflect honestly on certain issues, etc. But the jokester's activities *are not relevant to reflection as such*, but are relevant only to the operations of eyesight. Hence, the activity of the jokester has no effect whatsoever on whether the reflection in question passes scrutiny by (I').

It is important to note that (I') gives only a necessary condition for basing, and thus it would be a mistake to conclude that if the reflection in question passes scrutiny by (I'), it is home free. The point, however, is that the theoretical resources of the unified theory of animal and reflective knowledge do not sustain the high demands Sosa places on reflective knowledge in the kaleidoscope case. To sustain those high demands, we need a different account of reflective knowledge. For example, suppose knowledge is justified true belief. Then, on one account of reflective knowledge, one could have reflective knowledge only if one knows that these conditions are satisfied in a given case. One would thus need to know that one believes the claim in question, that the claim in question is true, and that one's belief in that claim is a justified belief. If reflective knowledge is nothing more than animal meta-knowledge, no such requirements on reflective knowledge are appropriate. Instead, one should apply the usual tests for animal knowledge, altering only the propositional content from p to Kp . In general, testing for animal knowledge involves finding a competence to explain the presence of the belief in question, not looking at the truth conditions for the claim in question or at implications of the claim in question to see if they are themselves known to obtain.

One might wonder whether this last point can be exploited to show that Sosa's pattern of reasoning about reflective knowledge can be sustained. The link would be a closure principle of some sort, perhaps a principle that claims that you know p only if you are in a position to know that any implication of p is true. Then, if the apt belief account of knowledge is correct, one would need to be in a

position to know that the requirements for aptness are satisfied in the kaleidoscope case. Wouldn't that be enough to rescue Sosa's denial of reflective knowledge?

No, it wouldn't. Consistent with Sosa's attempt to provide a unified account of reflective knowledge, to be in a position to know that the relevant clause of (I') is satisfied would require powers of detection and reflection to notice that the conditions involved in believing that the surface is red are conditions that are constitutive of normalcy for the operation of perception. Ordinary perceivers are in such a position, and so are in a position to know that the first-order belief that the surface is red is an apt belief and satisfies the conditions required for the belief to be apt.

So the conclusion stands. There are requirements for reflective knowledge that support Sosa's conclusions about the kaleidoscope case, but those requirements are not elicitable from Sosa's official account of reflective knowledge in terms of iterated animal knowledge. Moreover, adopting an alternative account of reflective knowledge that emphasizes aspects more at home in coherence theories of knowledge would both threaten the unity of the account of knowledge, and force one to withdraw the claim that all knowledge can be understood in terms of apt belief.

2. *The Tom Grabit Case*

This difficulty with the fake barn case, seen as treatable in the same way as the kaleidoscope case, is reinforced by consideration of other examples in the Gettier literature. In particular, virtue epistemologies have difficulty with the range of possibilities surrounding the Tom Grabit case. The skeleton for this range begins this way: Tom is your best friend, and you see him running from the library with a book, in conditions that make it obvious to you that he is stealing the book (e.g., the security guard is running after him and shouting for him to stop, etc.). Added to this beginning is a police report, with the police searching for Tom at his home. When they arrive Tom's mother says, "Tom couldn't have stolen the book; he's in Hawaii on vacation; it must have been his twin brother Buck, who doesn't live here and whom I haven't seen for years." The mother's testimony is thus a defeater of your basis for believing that Tom stole the book, but it is a defeater of which you are unaware.

Whether this defeater prevents you from knowing that Tom stole the book depends on how the skeleton above is fleshed out. The default setting, of course, is that the defeater in question prevents

you from knowing, but this default setting can be changed depending on what additional details are added to the case. Lehrer and Paxson, for example, add the details that Tom's mother is a pathological liar recently released from a mental institution, that there is no twin, and that all of this is well-known to the police. In such a case, the defeater in question is a misleading one, and doesn't undermine one's knowledge, according to Lehrer and Paxson.

Some may be unpersuaded by this assessment. If so, other details can be added that make clear how it is possible for the defeater in question to be misleading. Suppose, for example, that Tom's mother is rehearsing for a play that evening, a play in which one of her lines is the sentence above: "Tom couldn't have stolen the book; he's in Hawaii on vacation; it must have been his twin brother Buck." Suppose further that the police hear her practicing, knowing full-well what the situation is and that she is just rehearsing for the play. In such a case, the potential for defeat is blocked by the additional details, so that your knowledge is not undermined.

The point to note is that there is a continuum of detailed specifications that can be added to the basic outline of the story, a continuum running from cases in which the mother's testimony obviously prevents you from knowing that Tom stole the book to cases in which it obviously does not (even though in all of the cases, the mother's account is totally fabricated). Where the precise cut-off is between fully elaborated cases of knowledge and fully elaborated cases of a lack of knowledge, or whether there even is such a precise cut-off, is not of concern. What matters is the range of fully elaborated cases, with clear cases of knowledge on one end and clear cases of a lack of knowledge on the other end. In all such cases, what Tom's mother says is not true, but in some of these cases, the ruse is elaborate and well-conceived, by a fully competent person gifted at such ruses. It is a sad fact of intellectual life that people better positioned on certain matters than we are have the power to prevent us from knowing things that we would have known apart from their subterfuge. I am not making any claim as to how easy it is to do so, or how frequently it occurs, but one of the lessons from the Tom Grabit case is that it is possible.

This implication of the Tom Grabit case fits into a larger framework concerning the epistemology of disagreement that yields the same general picture, a picture about the way in which our knowledge depends in important ways on what others believe and say. In the epistemology of disagreement, the issue is one of internal defeat, and the degree to which evidence of disagreement undermines

the rationality or justification of one's beliefs. Examples such as the Tom Grabit case simply externalize the defeaters in question, placing them outside of the cognizer's awareness, relevant then to whether the cognizer has knowledge. I believe there are important lessons to be learned in the epistemology of disagreement literature from such Gettier examples, since some are tempted in the disagreement literature to exaggerated positions in which the known presence of an equally or more competent inquirer with roughly the same evidence is a defeater of rationality that cannot be overridden. One might argue against such a position from the inside, as I have done elsewhere (Kvanvig forthcoming), but one might also argue against the position from the outside, noting that the analogous position with respect to Gettier situations would yield an improper division between versions of the Grabit case that imply a lack of knowledge from those that leave knowledge intact.

There is no need to pursue that issue here, but I raise it to point out that Grabit cases are not some isolable backwater in Gettier lore, but are instead examples of a phenomenon that is a large and important aspect of intellectual life. Any temptation to dismiss Grabit cases as unimportant or uninteresting anomalies for a preferred theory should be resisted precisely because of the centrality of social context to a full understanding of when knowledge is present and when it is not.

In this regard, it is worth noting how the two topics interact. Suppose one were inclined toward the heroic path of denying that there are any Grabit cases in which one's knowledge is threatened. (It is worth noting in passing here that it isn't even remotely plausible here to explain away the Grabit cases by saying that the defect is not one concerning knowledge, but one concerning meta-knowledge: the possibility of knowing that one knows is not threatened in any way by the Mother's testimony that doesn't threaten knowledge itself.) The analogous heroic path concerning the epistemology of disagreement would be that disagreement itself never provides a defeater for one's present opinion. That position is among the least plausible positions to adopt in that arena, and its implausibility is one consideration counting against trying the heroic path. Instead, we should acknowledge in our attempts to understand cognitive achievements that, once socially positioned, a complete explanation of our cognitive achievements can't ignore or dismiss the role that others in the social arrangement play. The presence of others can both hinder and help us, epistemically speaking, and the point above about the continuum

of Grabit cases is just one particular example of this more general phenomenon.

It is this possibility that places limits on the role virtue epistemology can play in explaining the nature of knowledge, and on Sosa-style virtue epistemology in particular. For your belief about Tom is a perfectly ordinary perceptual belief from perceptual abilities as competent as one might care to specify, generated in perfectly normal perceptual circumstances. The presence or absence of knowledge is explained, not by anything having to do with the local facts concerning belief formation, but rather on the basis of social facts concerning events at another place and time. Nothing about your competence is threatened by these distant facts, no disparagement of aptness is implicated, and no loss of safety, whether basis-relative or unrestricted, is incurred.

The threat to knowledge is not from the inside, but rather from the outside. Those familiar with the inability of purely internalistic accounts of knowledge to survive scrutiny by Gettier and his cohorts will recognize the pattern here immediately. Put in the language of defeat, purely internalistic accounts of knowledge fail because there are two kinds of defeaters. Internal defeaters undermine the internal rationality of belief, but knowledge answers not only to internal defeaters but to external ones as well. Just so in the case of apt belief, involving displays of competence, ability, or power. Knowledge can be undermined by an internal failure of the sort that happens when the belief in question is true and yet still inapt, but knowledge answers to external factors as well. That is the lesson of the Grabit cases.

3. *Conclusion*

These results have predictable results for the issue of the value of knowledge. A full defense of the value of knowledge will explain how its value exceeds the value of its proper subparts.⁴ Part of the attraction of virtue epistemology is that it gives a plausible explanation about why virtue-based true belief is more valuable than true belief itself. Apt beliefs are accomplishments of a certain sort, displays of skill or ability at detecting truth, and the products of such displays are valuable both because they are true and also because they are ability-based accomplishments. Virtue epistemology thus has much to contribute to the question of the value of knowledge. It cannot, however, be a complete answer to the problem of the value of

⁴ For a defense of this claim, see Kvanvig 2003 and Kvanvig 2009.

knowledge, since it cannot do that with getting the nature of knowledge right in the first place. As the discussion above shows, there is something more to knowledge than can be explained solely in terms of the language of aptness and competence that Sosa's epistemology employs. A full explanation of the value of knowledge still eludes us.⁵

REFERENCES

- Goldman, A., 1976, "Discrimination and Perceptual Knowledge", *The Journal of Philosophy*, vol. 73, no. 20, pp. 771–791.
- Greco, J., 2009, "The Value Problem", in Haddock, Millar and Pritchard 2009, pp. 313–322.
- Haddock, A., A. Millar and D. Pritchard (eds.), 2009, *Epistemic Value*, Oxford University Press, Oxford.
- Kvanvig, J., 2009, "Response to Critics", in Haddock, Millar and Pritchard 2009, pp. 339–352.
- , 2003, *The Value of Knowledge and the Pursuit of Understanding*, Oxford University Press, Oxford.
- , forthcoming, "The Rational Significance of Reflective Ascent", in T. Dougherty (ed.), *Evidentialism and Its Critics*, Oxford University Press, Oxford.
- Lehrer, K. and T.D. Paxson, Jr., 1969, "Knowledge: Undefeated Justified True Belief", *The Journal of Philosophy*, vol. 66, no. 8, pp. 225–237.
- Sosa, E., 2007, *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*, Oxford University Press, Oxford.
- , 1980, "The Raft and the Pyramid", *Midwest Studies in Philosophy V*, University of Minnesota Press, Minneapolis, pp. 3–25.

Received: March 25, 2009; accepted: May 13, 2009.

⁵ For help with previous drafts, I express my thanks to John Turri and members of various audiences to whom I have presented these ideas, including quite helpful discussion with students in my graduate epistemology seminar.