Midwest Studies in Philosophy, XXXII (2008)

Why We Should Prefer Knowledge

STEVEN L. REYNOLDS

I will discuss Plato's question from the Meno: Why should we prefer knowledge that p over mere true belief that p? My answer will be quite un-Platonic. We do and should prefer knowledge mostly to obtain the approval of others.

Plato discusses the question using the example of knowing that a road leads to Larissa as opposed to merely having a true opinion that it does (Meno 97a–98d) and suggests that we should prefer knowledge because it is "tied down" by reasons, and so is not likely to run away.[1] It is not clear, however, as a matter of psychological fact, that nonknowledgeable true opinion is generally less lasting than knowledge. For we tend to forget the grounds of our beliefs even where they have grounds, and we are not likely then to hold the mere opinions more tentatively.

Still, the claim is apparently that the guide for action that consists in knowledge is more reliably available than the guide that consists in the corresponding true belief that is not knowledge. So knowledge is primarily an instrumental good. It is valuable because it enables us to obtain other goods, such as satisfying our desire to go to or avoid Larissa, and it is better than mere true belief because it is likely to persist longer and so guide more successful actions.

1. As historical claim, this (and further attributions to Plato below) is of course lacking in nuance and not based on any genuine understanding of the historical and philosophical background relevant to Plato's views on the topic. I include it against the advice of Tom Blackson, because I think it is useful as a schematic position (or caricature) against which to contrast the view of epistemic value that I mainly want to propose here.

It seems to me that there are two important innovations from Plato's Meno[2] answer in the recent literature.[3] One is that it is allowed that we may desire truth for its own sake, rather than merely as a guide to successful action. The other is that we may look back for a source of value. It is not merely that we shall benefit more in other ways, now or in the future, from having knowledge than we would from having mere true belief, as Plato suggests. There is something that is good-making in the way that we have come to hold the truth, if we thereby come to know it instead of merely truly believe it. For reasons to be explained below, I think the recent discussion still retains too many of Plato's preconceptions about how to approach such questions of value, but let us first take a closer look at how the recent answers improve on Plato's.

The first main difference is that recent discussions allow that we may want the truth for its own sake, apart from any other goals having the truth may help us achieve. It is suggested that we seek knowledge because seeking knowledge is the most reliable or effective way to obtain true beliefs, and obtaining true beliefs is our main cognitive goal (BonJour 1985, 1998; David 2001; Horwich 2006; Williams 1978).

That suggestion, at least understood in the most natural, reliabilist way of taking it, runs into a problem made vivid by Linda Zagzebski, which motivates much of the recent discussion and particularly its strategy of looking backward for good-making features (Zagzebski 2003a, 13–15). If the choice is between a true belief that p and knowledge that p, it seems we still prefer knowledge to true belief, even if assured (by some means that does not amount to giving us knowledge of it) that it will be a true belief. Perhaps God tells us that in the near future she will give us either a true belief or knowledge of the same content, where the proposition to be believed or known is not presently indicated, and we can now choose which to have. Suppose God also assures us that there will be no extraneous net cost or benefit to either choice. It seems that we would still choose knowing over merely having true belief, if thus assured that other things really were equal. But why should we choose knowledge, on the view that it is valuable mainly as a distinguishable, reliable route to truth, since choosing knowledge in this specific case will not make it any more likely that we will have a true belief? If it will be true in either case, we shall have all of the benefits of a true belief whether or not it also counts as knowledge. In Zagzebski's analogy, once we have a good cup of coffee, why should we value it more for having come from a reliable coffee maker? What we want in the particular case is, by hypothesis, good coffee, or a true belief, so once we have it there is no reason to care about how we got it. This problem is sometimes

2. I specify the Meno, because in the Republic (358e) Plato lists knowing as one of the things we desire for its own sake as well as for what comes from it.

3. I will only be considering the discussions that accept Plato's assumption that knowledge is and should be preferred to mere true belief, but of course there are important recent contributions that deny that claim. Jonathan Kvanvig holds that it is really understanding that we (should) value and that we are under the misapprehension that knowledge should be valued only because of its normally close association with understanding (Kvanvig 2003). A different way to support the denial of the presuppositions of Plato's question is Crispin Sartwell's spirited defense of the claim that knowledge just is true belief (Sartwell 1991, 1992).

called "the swamping problem" apparently because the goodness of the result, good coffee, or true belief (the only relevant thing we desire for its own sake), "swamps" any difference in value due to the source.

The most common response has been that knowledge is better than mere true belief because it has good-making historical qualities. Having true belief from a source that yields knowledge may be a virtuous cognitive action (Sosa 1997, 2003, 2007; Zagzebski 1996, 2003a,b), or to the believer's credit (Greco 2003; Riggs 2002). In Sosa's analogy, a shot released by a skilled archer is rightly valued more highly than one released by a lucky archer, even if both strike the center of the target (Sosa 2007, 77ff.). The true belief that counts as knowledge is valued not merely because it happens to be true, but also because it is true as a result of the exercise of cognitive virtues or epistemic skills.

It has been suggested that these answers cannot be quite right because knowledge is often not the result of skills or virtues and also not an achievement or to our credit. Jennifer Lackey alleges a counterexample: One might accept the testimony of the first person one meets in Chicago regarding the way to the Sears Tower, and thereby come to know it (Lackey 2007). This is not an achievement or to the believer's credit (she holds), yet it still counts as knowledge and as such is to be valued over mere true belief. So the historical explanation of value that cites cognitive virtues or achievement cannot account for this case.

But is it really so clear that there are no cognitive skills presumed in the background of this story? If the receiver of the testimony *would* accept it even if it came from a manifestly deranged or deceitful person, then we might be reluctant to grant that she knows, even if the person from whom she received the testimony is not deceitful or deranged. So it is apparently assumed as background for the story that there is no reason in the testifier's observable demeanor or circumstances to doubt her testimony, and that the inquirer has the requisite skills and inclinations to recognize circumstances undermining her testimonial justification were they to occur. Such background assumptions may make the knowledge so obtained seem more of an achievement, just because it will yield knowledge only to a person who has these further epistemic skills, over and above the mere ability to understand what was said in testimony.

A different example may make Lackey's point more forcefully: It is easy to know that the sky is clear, if one has normal vision and is standing outdoors at the time. But where is the achievement in opening one's eyes? I do not find entirely convincing even the claim that this last example is a counterexample. The world class archer who makes an easy shot is still shooting skillfully, even if it is the sort of shot a much less skilled archer could also reliably make. Small achievements are still achievements, and I rather think that there is more skill even in the most ordinary perceptual judgments than Lackey allows (Reynolds 1991).

So I think there is something right about the suggestion that knowledge is due to cognitive virtues or skills. It does constitute an achievement.[4] The suggestion I will make can be construed as a further specification of these proposals. But they

4. This is not intended to minimize differences in the proposals. I am taking a long focus on them, since I think they all suffer from a similar problem.

are not complete explanations of the value of knowledge, since they invite the question why we should so value what leads up to knowledge. Why should we regard this particular way of coming to believe as due to a skill or a virtue or as constituting an achievement or to our credit? If the answer is only that the skill, virtue, etc., reliably leads us to true beliefs, then it seems we have not escaped Zagzebski's swamping problem.

One might hope that skills or virtues are recognizable as such on their own, apart from their leading to such valued outcomes as knowledge (so that they can explain the additional value of knowledge rather than presuppose it), and that they are valued independently of their immediate outcomes, so that epistemic skills (etc.) could be a source of epistemic value independent of the value of true belief and thus avoid the swamping problem. But it is not so clear that we can recognize skills as such apart from their results. For example, not every cultivated way of acting or set of habits constitutes a skill or a virtue, even if it is rather difficult to learn. There are tricks one can do with a piece of cotton lint that are every bit as difficult as the achievements of archery, but we do not value them and do not think that we should.

We have a number of reasons to value archery over tricks with lint. One is that there is in our society an established set of games or competitions for archers, with widely acknowledged standards, where archers can demonstrate their skills and be rewarded for outstanding achievement. A second (and not really independent) reason to value archery is that those who succeed in archery competitions, or in more informal ways demonstrate skill in archery, are praised by others. There are no competitions and no praises for autistic children who spend hours doing tricks with lint. But again that invites the question why we have competitions in archery and why we praise it.

There is an obvious answer to that question. It is not found in the amount by which the intrinsic value of arrows suddenly occupying the centers of targets exceeds that of bits of lint drifting through the air in just the way the autistic child finds so fascinating, but rather in the history of archery. Archery was once of great benefit in hunting and warfare, while lint floating never has had such uses. Keeping one's family fed and protecting them from aggression once made the skills of archery of great practical value. Archery no longer has those uses to any significant extent in our society, but its history helps explain why we now have the archery competitions and the cultural practice of praising skillful archers.

Plato's approach to questions of value discourages us from considering such cultural explanations however. If we are asked what value there is in that arrow now quivering in the center of a target over and above the value of a bit of lint floating in just that way, where both are the result of hundreds of hours of concentrated practice, we are not likely to consider the more remote history of the activities. It is tacitly assumed in the way the Meno question about knowledge is framed that the difference in value will be present and discernible in any case in which we (should) desire knowledge in preference to true (or true justified) belief, such as the case in which we desire to know which road leads to Larissa. So one considers the differences that are salient in those particular cases, seeking some good possessed by the knower over and above the good possessed by the mere true

believer that that is the way to Larissa, a good that would explain why it is reasonable to prefer knowledge. If an apparent difference in the specific cases does not seem valuable enough to outweigh the additional (small) costs of inquiry, it is concluded that it is not the relevant difference.

Another philosophically interesting case where I think a larger view is likely to be illuminating is the mysteriously high value of original visual art over exact copies.[5] I have in mind here the value of the work from the master's hand, as opposed to an attempted exact copy of a given piece, not the distinction between the highly original type of artistic work and the merely derivative (but not copied) work. If we value oil paintings, say, mainly for their aesthetic effects on the viewer, and a good copy is good precisely because it looks just like the original, and hence will have the same visual effects on the viewer, just as a viewer, then it is hard to see why we should value a particular original painting so much more than an excellent copy of it. Attempts to explain the value of the original in terms of actual aesthetically relevant defects of the existing copies only emphasize the general problem: that a really good copy should apparently be just as valuable as the original, but, mysteriously, isn't.

If we consider the society-wide practice of such valuing however, a plausible rationale for praising the original work over even perfect copies is not hard to see. Suppose that our society desires to encourage the production of new art, perhaps for the sake of producing new kinds of aesthetic experience. (There may well be other aims too. The point is not to give a complete account of our reasons in this case—no doubt they are complex—but just to point out that one plausible possible reason is likely to be overlooked on Plato's way of approaching the issue.) One way to do that is to praise original art for being original and not a copy. An artist will then reliably earn that sort of praise for her own productions only by making new art, not by copying others' works. The praise given would thus encourage artists to try to do something new, not merely to reproduce the old, no matter how esteemed the old art is for any of its features that would be reproduced in a good copy.

Although the praise of art as "the original" may thus have as its main purpose or function changing the behavior of those engaged in producing art, it will also naturally affect those who purchase an original work from its previous owners, so that original works will have a higher value in the market than mere copies, however exact. I suggest that it is (in part) the practice of praising particular pieces of art in this way that produces our feelings that only the original piece is truly valuable, rather than the antecedent value of being the original piece that motivates our praise of it. That is not to say that just any praise will lead us to value that which is praised, of course, but praise that has some significant motivation and is widely and consistently offered will tend to produce an attitude of valuing.

If we, as society, praise original art in order to encourage new art, that explains the *practice* of such valuing, not an individual case. Considering a single "clear case" obscures our reasons for preferring originals in general, because it focuses our attention on differences pertinent to that case, instead of its character

5. Thanks to Peter French and Peter de Marneffe for helpful and knowledgeable discussion of this issue.

as an instance of the practice. There is a plausible partial explanation of the practice—our collective desire to encourage with our praise the production of new works of art—but that explanation would not be suggested by any salient difference between, say, Rembrandt's *Night Watch* and its best copies.

I think that the nature of the value of knowledge is similarly overlooked when we follow Plato's method and consider only ourselves in a given case and ask why we should prefer to know, rather than merely truly believe, in that specific case.

But before turning to a cultural explanation of the value of knowledge, I would like to say a few words about praise and criticism, or expressions of approval and disapproval. It seems to me that many philosophers, especially those who do not work in value theory generally, tend to grossly underestimate the influence of the approval or disapproval of others in our daily conduct. We like to imagine ourselves to be free spirits and independent thinkers, and some of us are, in some respects, some of the time. But we are also, like the rest of humanity, constantly evaluating ourselves and others according to standards that are enforced mainly by the approval and disapproval of other people, which we may see in a disapproving look or in an indirect implication of someone's remarks or which, perhaps more often, we infer that they will have toward us from their remarks about absent third parties. Where we do not see particular evidence of their evaluations of our own conduct, we nevertheless constantly expect them to be silently evaluating us, as indeed we are constantly evaluating their conduct, in various standard and expected ways.

To cite some obvious but usually overlooked examples, our standards of grammar, appropriateness in topics of conversation, physical distance from conversational partners, gestures, tone of voice, where to look and when, and the rhythms of conversation that indicate when to make a remark without interrupting, are all taught and motivated by the sometimes visible approval and disapproval of others. This becomes painfully apparent in the social difficulties of high functioning autistic persons, who often strongly desire to obtain the approval of others, but are able to recognize only its plainest manifestations, and so are apt to get all of these matters wrong in what seem to the rest of us to be obvious ways, oblivious to the reactions of the people with whom they are trying to converse.[6] These examples also illustrate how approval and disapproval may play an important role in teaching us basic skills which we may then exercise with perhaps only the very occasional reminder (e.g., the uncomfortable look and backing away when we stand a little too close to someone).

6. I speak from personal experience. My son Andrew, who is eighteen at the time of this writing, has high functioning autism. After much patient coaching he has reached the stage at which people with whom he tries to make small talk merely wonder why his parents never tried to teach him manners. Lack of awareness of the unspoken approving or disapproving reactions of others is likely to be only one contributing factor (but it is a very obvious one) in the difficulties persons with autism have in learning these skills, since they are also relatively uninterested in imitating others and they often have cognitive or perceptual difficulties that interfere with matters of timing. For more on the difficulties of persons with autism and methods effective in remediating those problems, see Koegel and Koegel (2006). A good review essay of the recent general psychological literature on praise and motivation is Henderlong and Lepper (2002).

It is held by a number of recent philosophers that knowledge is the norm of assertion, that is, if we assert that p, when we do not know that p, then we are properly subject to criticism.[7] That suggests part of the view that I want to defend—that we value knowledge because others require that we know what we assert, at least to the extent of disapproving of us, if they think that we do not know that p when we assert that p. Since so much of our life is social, and requires frequently making assertions in conversation, that would give us a reason to value knowledge for anything that we might expect to assert. But it does not yet explain satisfactorily why we value knowledge. For it invites the question why knowledge should be the norm of assertion. Why should we value knowledge in assertions, over the expression, say, of true beliefs or justified true beliefs?

Williamson suggests that no reason can be given for why knowledge is the norm of assertion, because what makes or constitutes a speech act an act of assertion is that it is subject to the norm of knowledge (Williamson 2000, 266–69). As he notes, that raises the question why we have decided to transmit information in the form of assertions, constituted by the norm of knowledge, rather than, say, qu-assertions, which have a norm of true belief. Qu-assertions that p are subject to criticism if made by those who lack true belief that p, but they are not subject to criticism if made by those who do not have knowledge that p, so long as they have true belief that p. In answer to this question he suggests that we can only transmit knowledge through our assertions if we know what we assert. A requirement of knowledge discharges our responsibility to epistemically ensure the truth of the content asserted.

It is not clear, however, that we have such a responsibility, apart from the norm of knowledge itself. To say that I am responsible for ensuring the truth of what I say, where the truth can be ensured only by knowing, seems to be only a variant way of claiming that knowledge is required for assertion.

In previously published work I have argued that the function of saying that people know is in part to express approval of the corresponding actual or potential testimony, rather than approval of assertions more generally (Reynolds 2002). In seeking to testify only in ways that will merit this approval we comply with public standards for testimony and thus, in the long run, improve the average quality of our testimony, making it more likely to be helpful to the recipients. Those who accept testimony offered by those who are trying to comply with a requirement of knowledge are more likely to be successful in their actions or projects than they would be if they relied instead on testimony offered by those who do not attempt to comply with such a requirement.

Testimony may be roughly characterized as that subclass of assertions that are offered as informative, with the expectation or purported expectation that the recipients of the testimony will take the speaker's word for it. The main reason for shifting to talk of testimony instead of assertion more generally is that we often

7. Among those who have defended the view that knowledge is the norm of assertion recently are Peter Unger (1975, 250–71), Timothy Williamson (1996, 2000), Keith DeRose (2002), John Hawthorne (2004), and Jason Stanley (2005).

make assertions in philosophy or while discussing sports or politics (etc.) that do not express our knowledge, and yet it seems that we are not to be criticized for so doing.[8] However, when we testify—expecting others to accept what we say—it seems we are expected to know and will be justly criticized if we do not.

There are a number of points to bear in mind about the thesis that attributions of knowledge function as expressions of approval of testimony. I think these points are obvious on a little reflection on the notion of the function of an artifact, such as a word like "knows." But it seems they are easy to overlook in our anxiety to spot the flaw quickly, an estimable ambition perhaps, but too much admired in our philosophical culture. First, to say that "know" functions to express approval of testimony is not to say that that is its only use, nor that every ordinary competent utterance containing "know" serves to express such approval. A carpenter's hammer may still have the function of driving nails even if most owners of such a tool more often use it for other purposes. Nor does my thesis imply that testimony given with knowledge may not still be legitimately subject to all sorts of other criticisms (e.g., as tactless, pointless, or harmful in a thousand other ways). It does not imply that nonknowledgeable testimony is in fact always so criticized, nor that it should be. Nor does it imply that we may not sometimes excuse people, to whom it appears that they have knowledge, for testifying, even though they do not really know (e.g., the protagonists of the Gettier stories). The occasional criticism or praise of third parties that we hear ("She said it was so, but she couldn't have known") may sufficiently remind us of the requirement of knowledge for testimony, so that we rarely need to hear overt approval or disapproval of our own testimony. It is the approval we mainly desire, not the expression of that approval. That is why it is consistent with my thesis that, once past early childhood, we are only occasionally told that we know or do not know something.

These various points should sufficiently indicate why the simplest sorts of surveys of the usage of "know" (e.g., compiling randomly selected utterances of "know" to determine what percentages are used as expressions of approval) are far too simplistic to be useful tests of my thesis. It is an empirical thesis however, and so should in principle be open to empirical testing. My point here is not to discourage attempts to think of ways to bring it to an adequate empirical test, but just to warn against the simplest sorts of empirical testing, and especially against that "test," irresistible to philosophers because it is apparently available from one's armchair, of trying to recall whether the ways one has recently heard "know" used seem to count intuitively as praise.

I think similar observations indicate why the claim that knowledgeable testimony tends to be more helpful than nonknowledgeable testimony is not put in doubt by the citation of fairly significant classes of exceptions, such as the white lies

8. Williamson allows for such cases as inconsequential breaches of the normal rule for assertion, and cites as a parallel breaches of grammar in animated conversation (Williamson 2000, 258–60). The cases do not seem to be fully parallel however. We would regard the grammatical oddities as breaches of grammar, even though we might think it obstructive of the purposes of conversation to point them out. But I think we are not similarly inclined to admit a fault if it is pointed out to us that we do not really know that p where we have just asserted that p in a discussion.

that spare our vanity, or the alleged benefits of acquiring a high opinion of our own abilities through the flattery of, say, our students and friends. It is no accident that proverbs discourage us from taking compliments as testimony: "a compliment is to be inhaled, not swallowed." It is a large and difficult empirical question whether having a norm of knowledge in fact makes testimony more beneficial to us, not in every case, but on the whole. However, surely it is very plausible that it does.

Turning then from some initial objections to reasons in favor of the view: One reason to hold that "know" functions to express approval of testimony is that it neatly explains why knowledge is the norm of testimony. The word "know" (or its ancient counterparts) was developed for the purpose of indicating acceptable testimony.

Another reason to hold that "know" functions as praise for testimony is that this view explains all of the main features of our concept of knowledge. It thus improves on Bernard Williams' project. He tried to explain the justification and "something for Gettier" features of knowledge on the assumption that we have a goal of true belief, arguing that seeking the other features of knowledge would be the most reliable way to obtain true belief (Williams 1978, 37–45). He assumed that we desire some features of knowledge in order to explain why we would desire the other features, and so explained some aspects of our concept of knowledge in terms of its other aspects. But as I have argued elsewhere, and shall now briefly summarize, the hypothesis that we use "knows" as an expression of approval for acceptable testimony explains why we have a concept in common use that has all of the features of our concept of knowledge.

When a term for praising testimony, say "gnows," was introduced, it would presumably have been used to praise testimony that seemed to people to have been helpful in their projects, such as finding people or things.[9] If Sam told Sally that there were wild onions on the other side of a certain hill, and Sally consequently walked over the hill looking for onions and failed to find any, she might then have criticized what Sam told her as not "gnown." If she did find the onions, however, she might have expressed her approval of Sam's testimony by saying he "gnew" there were onions over the hill. She might have done that whether or not

9. This resembles Edward Craig's argument that our ancestors developed the concept of knowledge and words to express it in order to indicate approved sources of information (Craig 1990). Craig assumes that our ancestors would have been seeking true belief whether p (so far resembling Bernard Williams' idealized inquirer), and considers how they might recognize informants who are likely, through testifying, to give them true belief whether p. Those informants would be designated as knowing whether p. I begin instead with the assumption that our ancestors would create a term with which to praise those who gave helpful testimony, not assuming that they would have any antecedent desire for true beliefs. The stories of how the concept of knowledge would develop from these different starting points are quite different. For example, Craig suggests that our ancestors would seek confidence in their informants as more likely to induce belief in their testimony in those who hear that confidently expressed testimony (Craig 1990, 12–13). (See the text for my alternative.) But the main difference between my view and Craig's is that I focus on the role of attributions of knowledge as encouraging better testimony in the future, while Craig focuses on their conveying information to aid in locating good testimony now. I do not suppose that either of us would have to deny that the area the other focuses on is also an important aspect of our use of "know."

she possessed concepts of truth or belief. Those concepts would not have been necessary for her to judge whether Sam's testimony helped or hindered her project of gathering onions. In time people would notice that someone who had just come from over the hill tended to give helpful testimony about what could be found there, and that those who had not recently been there were less likely to give such helpful testimony. Users of "gnow" would thus gradually come to apply it to those who had a certain kind of history—a history of certain kinds of perception, received testimony, and so on—even in advance of soliciting or acting on their testimony. If our ancestors disagreed with the testimony someone had given, they would expect that testimony to be unhelpful (to others at least, since of course they would not act on it themselves), and so would criticize it as not gnown, and also eventually (because of the disagreement) as not true. Perhaps on observing cases of apparently confident testimony that had not been helpful, although the testifiers were willing to act on it themselves, they would come to regard those testifiers as having been in states that governed their actions in the way that received testimony sometimes does, but that did not amount to having gnowledge, that is, as having mere beliefs.[10] Our ancestors would have recognized that some persons seem to themselves to gnow, but because of errors or omissions in their awareness of their own relevant history or circumstances (as in the Gettier stories), they would not have the sort of history that was required for really gnowing. Those persons might be said to be "justified" in what they testified or would testify (because it reasonably seemed to them that they knew, and so that they were obeying the requirement), but to lack knowledge. If someone seemed to themselves to lack knowledge, but testified anyway, they would be blamed for flouting the norm of testimony, even if it happened that they really gnew. So our ancestors would gradually develop a concept of justification, or at least a practice of refraining from criticizing those who seemed to themselves to know, and they would come to regard justification, or seeming to the subject to know, as a requirement for gnowledge. They would also recognize an absence of belief by the subject as strong but fallible evidence that the appropriate history for gnowledge was lacking. So they would probably come to require belief for gnowledge too. My view thus apparently explains all of the standard requirements of our concept of knowledge—truth, belief, justification, and something for Gettier—without merely assuming an antecedently existing desire for any of them.[11] I do not of course deny that two of these features are more

10. I am of course echoing Wilfrid Sellars' myth of Jones, which explains how our ancestors might have acquired a concept of belief from observations of the testimony and related actions of others. I have extended it to a myth on which it would be natural to expect "knows" to express a more fundamental concept than "believes," as Williamson holds (Sellars 1997 [1956], 102–7; Williamson 2000, 41–48).

11. For further details and attempts to answer some objections, see Reynolds (2002, 149–58). Some recent accounts of knowledge omit the justification condition, replacing it with reliabilist conditions. But they usually allow that there is at least an appearance that justification is a requirement for knowledge. So explaining that appearance would be a point in favor of my view even on these accounts of knowledge. Does my view amount to adopting a kind of "attitudinal theory" of knowledge, perhaps even a kind of noncognitivism or nonfactualism about knowledge, as in Field (1998)? I do not think I am committed to any such view. My idea is that the function of "know" to indicate approval of testimony, and the fact that only certain kinds of testimony have

important than the others: If the testimony is not true, or is not believed, it will not be helpful, though it may still be helpful in the absence of the history required for being knowledge, provided it is in fact true.[12]

Another reason to prefer the view that in using "know" our society encourages better testimony, rather than true belief as such, is that testimony is more under our voluntary control than is belief. We can change our testifying behavior in order to avoid the disapproval of being thought not to know what we have testified, or to obtain the corresponding approval of being thought to know, but it is not so clear that we could thus change our believing in order to avoid the disapproval or obtain the approval (Reynolds 2007).

The view that we say people know in order to encourage better testimony, and discourage substandard testimony, thus suggests a cultural explanation why we value knowledge. Our ancestors found that the sorts of features that now constitute our concept of knowledge that p, when possessed by the testifier, tended to produce helpful testimony that p. So they came to approve of testimony by those who had those features and to express that approval using the ancient equivalent of "know." Knowledgeable testimony had been found to be helpful and nonknowledgeable testimony had been found to be less helpful, on the average and in the long run. Since people approve of our testimony when we speak with knowledge, and disapprove of it when we do not, and since we want to obtain or at least merit the approval, and avoid the disapproval, that gives us a reason to prefer to know whatever we might have occasion to testify.

The function of our practice of classifying people as knowing continues to be achieved, although its history has been forgotten and its function is not clearly understood. We value knowledge over true belief because of the persisting practice of approving only of knowledgeable testimony. My personal preferences do in fact conform to the social practice, because I have been trained to have such a preference by hearing people talk approvingly of knowledge, and disapprovingly of those who testify without it. Even if I do not consciously want such praise (but I do), it has produced in me a preference for knowledge. We are aware that others approve or disapprove of our testimony depending on whether it is knowledgeable, and this motivates us to try to know.

That is not to say that we are conscious of wanting to merit the approval when the question is raised whether we prefer knowledge to true belief in a given case. We certainly do not, as individuals, share the larger social goal of improving testimony. It may seem to us that we just prefer knowledge to true belief, even though no further reason for the preference comes to mind. I have such a preference because, social animal that I am, I have always desired the approval of others, and given our existing institution of approving knowledge, I naturally responded to

tended to be helpful in the past, has led to a fairly definite informational content for the term "know" as we now use it. Given that it has that content now, it is an ordinary factual question whether one knows in any particular case, subject to the usual allowances for vagueness, reasonable contextual variation in standards, and the like.

12. Thanks to Bernie Kobes for emphasizing to me the importance of acknowledging the greater importance of truth among the requirements of helpful testimony.

that approval by developing a preference for it. I do not, however, need to be aware of this in order to be thus influenced in preferring knowledge.

But that may seem not to answer Plato's question about the value of knowledge. So far we have only heard a causal story about how our society, and I qua member of that society, have in fact come to prefer knowledge that p to true belief that p. But now that I am reflecting on that preference, I may still ask Plato's normative question, whether I *should* continue to so prefer it.

There does seem to be on this view a sense in which I *should* want knowledge, although it is perhaps not the "should" of rationality, but rather a "should" that indicates healthy or proper psychological functioning. If I did not prefer knowledge that could only be due to a serious defect in my education or in my psychology, so that either I was not properly socialized into this important cultural practice of ours, or I was pathologically indifferent to the praise and criticism that would have produced such a preference. If I were indifferent to the approval and disapproval of my fellows generally, I would very likely be even more incapable of being trained into my part in a human community than are persons with autism. A failure to learn what counted as knowledge, even if not accompanied by a general indifference to the approval or disapproval of others, would still be a very serious social handicap. Someone who had no desire to determine whether he knew in offering testimony would often give offense and perhaps in time effectively become a social outcast. I am not thinking merely of the incautious retailer of interesting gossip, but of someone who would routinely tell others the time of day, or the presence of Sally in the next room, or the way to the restroom, regardless of whether he thought he knew. Since I could fail to prefer knowledge only by having some fairly serious psychological or social defects, it seems I should prefer knowledge to true belief.

But it may still be said that this account of why we value knowledge over mere true belief is not really an answer to the very specific Platonic question. The real question about our preference for knowledge is: Why should I, as a matter of self-interested rationality, in this case, prefer knowledge over true belief? Plato did not ask why I should want to be responsive to social influences so that I will prefer it. He asked for a reason why it is better for me to have knowledge in the particular case. But instead of answering that question I have only given a naturalistic story about how we have come to regard it as better to know, and perhaps a functional account of why it is better for us that we should be the sorts of beings who would thus respond to our training by preferring knowledge.

It seems to me, however, that the sort of answer I have just given is not so very far from some familiar Platonic themes. The psychological defects that would make me capable of considering the question without having any effective prior preference for knowledge seem to me to be something like the damaged soul that, according to Plato, results from doing the wrong thing, especially if one frequently "gets away with it."[13] That is, they are serious mental deficiencies, although they are

13. *Republic*, 444c–d. I think the epistemic question is easier to answer than the corresponding moral question because it asks only why we should normally prefer knowledge over true belief, when other things are roughly equal. What makes the question why it is best to do what is morally

not internal conflicts of the sort that Plato claimed would occur in those who act unjustly. Someone who could choose not to care for knowledge until shown a personal advantage in each particular case would be very close to someone who just did not care.

But suppose the choice whether to care was only to be offered to me once, and there was no reason to think that making it either way would significantly affect my psychology in social situations generally, or even my future sensitivity to the requirement of knowledge for testimony in particular. Would there be any reason we could give to prefer knowledge then, to someone whose antecedently inculcated preference for knowledge was thus temporarily and otherwise harmlessly suspended? Suppose furthermore that I temporarily did not desire to merit the praise of other people, or even desire my own self-approval as knowing—my lack of a preference for knowledge also includes a complete but temporary lack of concern whether I am to be praised or criticized in this regard.

In such a case I concede that I do not think there would be any self-interested reason for me to prefer knowledge to true belief. But is there any reason to think that only in such a case are we focusing on the genuine value of knowledge over true belief? We might worry in some other sorts of cases that our society may have inculcated in us preferences that we would be better off without, that we have come to hold "false values." That might suggest that only an individualist thought experiment of Plato's sort, one that prescinds from society's influences, could yield correct assessments of values. But the idea that the only way to find the true value of anything whatever is to make oneself consider it from a radically individualistic point of view, such as we are now trying to do, seems wildly off the mark. There is no reason at all to think that there is anything false about our preference for knowledge, just because it has been partly created by our enculturation. Suppose, as we have been speculating, that knowledge would not seem preferable to an otherwise rational person some of whose important human qualities were on holiday. Why should that indicate anything disturbing about our own healthy preference for knowledge?

There is one last related problem for this view about the value of knowledge that I would like to briefly address. It is that my view may seem to posit an odd sort of creation of value. It is natural to think that when we praise something we are reporting that it is good. We cannot make it good just by praising it, like the (fictitious) teacher of small children who is always chirping "Good job!" no matter how the children perform. My account apparently says that we have created a feeling that knowledgeable testimony is good by praising such testimony, and have thereby produced a preference for knowledge where it did not exist before. The good of knowledge is as it were created by our practice of approving it.

Perhaps some of us will not mind that very much. Some of the things that we value, such as having a powerful serve in tennis, seem to be created by fairly arbitrary decisions about how the game is to be played and so about what the

right harder, and a Platonic answer to it harder to defend, is that we seem to be required to prefer doing what is right even when prudential considerations appear to strongly favor a morally wrong course.

players' aims within the game will be. Merely saying something is good does not make it so, but coordinating our actions and attitudes with others to value some things and disvalue others, as we do in creating a game, may have the consequence that something previously neutral becomes really good.

I think, however, that that is not quite what is going on in the case of knowledge. Even on my view it is not arbitrary that we came to praise knowledge, and so to value it over another quality such as mere true belief. There is a good in this system of social control of testimony, a good for our society in the improvement of the average quality of testimony, which motivated the cultural evolution of our way of praising testimony until it became the current practice we have of requiring people to know that p if they testify that p. It was as much a process of discovery of the good of knowing for testimony, and a process of educating ourselves and others to seek that good, as it was the creation of our preference for it.

So the reason we prefer knowledge and should prefer it is that we have been taught to prefer it by our awareness, over many years, that others approve of it. We constantly judge our own actual and potential testimony by this standard. Although we also want knowledge for ourselves for many diverse personal reasons, the main social purpose of classifying ourselves and others as knowing, or not knowing, is to encourage compliance with standards for testimony that benefit all of us in myriad ways. Having been open to that encouragement over the years, and so preferring knowledge now, is a good thing for each of us.

To sum up then: Why *should* I prefer to know that this is the way to Larissa over merely truly believing that it is? I find I just do prefer knowledge, as it were for its own sake. But from a larger perspective it seems that I should prefer it, because given the social practice of approving of testimony only if given with knowledge, I could fail to prefer knowledge, when other things seem to me to be equal, only by having the sorts of serious social or psychological defects that would make me unresponsive to the approval of others. Finally, the social practice that produces this particular preference is good for all of us because it improves the average quality of our testimony, which results in greater success in our projects generally. We do and should prefer knowledge in order to obtain the approval of others.[14]

REFERENCES

BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.

BonJour, Laurence. 1998. In *Defense of Pure Reason: A Rationalist Account of A Priori Justification*. Cambridge, England: Cambridge University Press.

Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford: Oxford University Press.

David, Marian. 2001. "Truth as the Epistemic Goal." In *Knowledge Truth and Duty: Essays on Epistemic Justification, Responsibility and Virtue*, ed. Matthias Steup, 151–69. Oxford: Oxford University Press.

DeRose, Keith. 2002. "Assertion, Knowledge, and Context." *The Philosophical Review* 111(2): 167–203.

14. Thanks to my colleagues Brad Armendt, Tom Blackson, Peter French, Bernie Kobes, and Peter de Marneffe for many helpful comments on an earlier draft of this article.

Field, Hartry. 1998. "Epistemological Nonfactualism and the Prioricity of Logic." *Philosophical Studies* 92(1–2): 1–24.
Greco, John. 2003. "Knowledge as Credit for True Belief." In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, ed. Michael DePaul and Linda Zagzebski, 111–34. Oxford: Clarendon Press.
Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
Henderlong, Jennifer, and Lepper, Mark R. 2002. "The Effects of Praise on Children's Intrinsic Motivation: A Review and Synthesis." *Psychological Bulletin* 128(5): 774–95.
Horwich, Paul. 2006. "The Value of Truth." *Nous* 40(2): 347–60.
Koegel, Robert L., and Koegel, Lynn K. 2006. *Pivotal Response Treatments for Autism: Communication, Social, and Academic Development*. Baltimore, MD: Brookes Publishing Company.
Kvanvig, Jonathan L. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge, England: Cambridge University Press.
Lackey, Jennifer. 2007. "Norms of Assertion." *Nous* 41(4): 594–626.
Reynolds, Steven L. 1991. "Knowing How to Believe with Justification." *Philosophical Studies* 64(3): 273–92.
———. 2002. "Testimony, Knowledge, and Epistemic Goals." *Philosophical Studies* 110: 139–61.
———. 2007. "Against the Goal of True Belief." Presented at the APA Eastern Division Colloquium: Truth and Belief, December 28, 2007.
Riggs, Wayne D. 2002. "Beyond Truth and Falsehood: The Real Value of Knowing That P." *Philosophical Studies* 107 (1): 87–108.
Sartwell, Crispin. 1991. "Knowledge Is True Belief." *American Philosophical Quarterly* 28(2): 157–65.
———. 1992. "Why Knowledge Is Merely True Belief." *Journal of Philosophy* 89(4): 167–80.
Sellars, Wilfrid. 1997. *Empiricism and the Philosophy of Mind*. Cambridge, MA: Harvard University Press. Originally published in Herbert Feigl and Michael Scriven, eds. *Minnesota Studies in the Philosophy of Science*, vol. 1 (1956). Minneapolis: University of Minnesota Press.
Sosa, Ernest. 1997. "Reflective Knowledge in the Best Circles." *Journal of Philosophy* 94: 410–30.
———. 2003. "The Place of Truth in Epistemology." In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, ed. Michael DePaul and Linda Zagzebski, 111–34. Oxford: Clarendon Press.
Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, volume I. Oxford: Clarendon Press.
Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
Unger, Peter. 1975. *Ignorance: A Case for Scepticism*. Oxford: Clarendon Press.
Williams, Bernard. 1978. *Descartes: The Project of Pure Enquiry*. New York: Penguin Books.
Williamson, Timothy. 1996. "Knowing and Asserting." *Philosophical Review* 105(4): 489–523.
———. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
Zagzebski, Linda. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge, England: Cambridge University Press.
Zagzebski, Linda. 2003a. "Intellectual Motivation and the Good of Truth." In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, ed. Michael DePaul and Linda Zagzebski, 135–54. Oxford: Clarendon Press.
———. 2003b. "The Search for the Source of Epistemic Good." *Metaphilosophy* 34(1/2): 12–28.