CATHERINE Z. ELGIN

THE EPISTEMIC EFFICACY OF STUPIDITY*

ABSTRACT. I show that it follows from both externalist and internalist theories that stupid people may be in a better position to know than smart ones. This untoward consequence results from taking our epistemic goal to be accepting as many truths as possible and rejecting as many falsehoods as possible, combined with a recognition that the standard for acceptability cannot be set too high, else scepticism will prevail. After showing how causal, reliabilist, and coherentist theories devalue intelligence, I suggest that knowledge, as contemporary theories construe it, is not a particularly valuable cognitive achievement, and that we would do well to reopen epistemology to the study of cognitive excellences of all sorts.

Socrates maintained that he was the wisest of men in that he alone knew that he knew nothing. Although his avowal is typically taken to be ironic, he may have been telling the truth. For currently popular theories of knowledge have the surprising consequence that stupidity can enhance, and intelligence diminish, one's prospects for knowledge. So if any of these theories is correct, Socrates may have known less than others precisely because he was wiser than they.

I will show that an unwitting bias in favor of stupidity is characteristic of both internalist and externalist theories of knowledge. It derives from the shared convictions that (a) our epistemic goal is to accept (or believe) a sentence if it is true and reject (or disbelieve) it if it is false, and (b) the standard for acceptability cannot be set too high, else scepticism will prevail. The epistemic inutility of intelligence that follows is not the sceptic's fatalistic conclusion that since no one knows anything, dullards are no worse off than the rest. It is the more disconcerting result that since qualities of mind like sensitivity, breadth, and logical acumen often interfere with the satisfaction of the requirements for knowledge, individuals deficient in such qualities have an epistemic edge. The quest for knowledge may then be furthered by the cultivation of obtuseness.

It would be tedious to demonstrate that this follows from all current theories of knowledge. So I have chosen to focus on four – two externalist and two internalist. They represent dominant strains in contemporary epistemological theorizing. And the difficulties I find

Synthese **74** (1988) 297–311.

are not difficulties in detail. So if all four find cognitive deficiencies conducive to knowledge, there will be reason to suspect that a commitment to the epistemic efficacy of stupidity is endemic to current epistemology. Toward the end of my paper I consider what to make of this finding.

1.

Contemporary epistemologists agree on this much at least: however good one's grounds for *p*, one cannot know that *p* if *p* is false; knowledge then requires truth. Moreover, one cannot know that *p* without being cognitively committed to *p*; knowledge also requires belief or acceptance. And one cannot know that *p* if one's true belief that *p* is accidental; so knowledge requires a tether.

Internalists take the tether to be epistemic. Knowledge, they maintain, is tied down by justification that is epistemically accessible to the knowing subject. Disagreements among them concern the criteria for epistemic accessibility and the range of accessible information the subject need take into account. So they differ over, e.g., the epistemic status of unacknowledged implications of things one explicitly knows, or of undermining evidence one does not, but could, possess.

Externalists take the tether to be metaphysical. For a true belief to amount to knowledge, they contend, it must be necessarily connected to the fact that makes it true, or to facts from which its truth follows. They differ over the type of necessity required, but agree that it need not be within the subject's ken. An individual can know that *p* even if he is unaware that his belief that *p* is appropriately related to the facts. Some take the metaphysical tether, known or unknown, to constitute the justification for a belief, thereby conceding that justification may be epistemically inaccessible. Others follow internalism in requiring justification to be epistemically accessible, but deny that justification is integral to or necessary for knowledge. To avoid confusion, I will speak of external tethers, leaving it open whether a belief's tether provides its justification.

2.

Causal theories of knowledge maintain that for a subject to know that *p*, his true belief that *p* must be caused by the fact that *p* or by facts

from which it follows that *p*. Sophisticated versions require that the causal connection be lawlike, so that knowledge cannot result from a fortuitous commingling of circumstances (Goldman 1967). Such theories account for inferential knowledge by claiming that inferential and logical relations may be parts of causal chains.

According to a causal theory then, my true belief that there is a yellow surface before me is caused by a neurophysiological response to the presence of yellow in my visual field. A sequence of optical and neural events linking the surface with a brain state is responsible for the production of my belief. If that sequence instantiates a natural law, I know that the surface is yellow. It is no accident that I believe what I do; for, given the laws of nature and the circumstances in which I find myself, my belief is a necessary consequence of the fact that the surface is yellow.

Such causal connections are common. It is no accident that normal perceivers typically believe objects to be the colors those objects actually are; for their beliefs are normally caused by the law-governed response of the human nervous system to the presence of those colors. That being so, causal theorists contend, normal perceivers generally know the colors of the objects they perceive. Since causal theories do not require epistemically accessible justification, they can recognize that unreflective and unintelligent people are often in a position to know. Watson is as capable as Holmes of knowing that the surface before him is yellow. And this is as it should be. Cognitive virtuosity is hardly required for knowledge of this kind.

But seemingly parallel conclusions are less comfortable. Consider one involving the sense of taste. Holmes, we may suppose, is an oenophile, while Watson is oblivious to all but the most obvious differences among wines. The two share a bottle of Bordeaux, and because it stimulates the appropriate nerve endings and brings about the proper neurological connections, it causes each to believe that he is drinking Bordeaux. (For vividness we can assume that exactly the same neurological events occur in both.) According to causal theories, both Holmes and Watson know that the wine they are drinking is Bordeaux. The fact that Watson cannot tell a Bordeaux from a muscatel does not prevent him from knowing about this wine, for it does not intrude upon the causal chain leading to his current belief. And unless we are prepared to conclude that Holmes lacks knowledge, we cannot dismiss the chain of neurological events as

anomalous. If a causal law is instantiated in the production of Holmes's belief, it is instantiated in the production of Watson's. For exactly the same events occur in each. If Holmes knows what he's drinking, so does Watson.

It follows from causal theories that subjects can 'luck into' knowledge. Given Watson's insensitivity to distinctions among wines, it is accidental that the lawful causal chain eventuates in a true belief. Despite its impeccable breeding, Watson's belief is unreliable.

The conviction that unreliability precludes knowledge leads some externalists to reliabilism – the view that knowledge depends on a belief's relation to truth in counterfactual as well as in actual circumstances. On a reliabilist account, a properly tethered belief is, roughly, one the subject would harbor if it were true and would not harbor, at least on account of that tether, if it were false (Nozick 1981; Dretske 1971). The truth of a properly tethered belief is no accident; for such a belief tracks truth across possible worlds.

Reliabilism concludes – correctly, it seems – that Watson does not know, since he would believe he was drinking Bordeaux even if he were drinking muscatel. The problem is that Holmes apparently fares no better. Although he can tell Bordeaux from muscatel, he cannot infallibly discriminate Bordeaux from all other sources of sensory stimulation. So Holmes, like Watson, fails the subjunctive test; there are non-Bordeaux he would believe to be Bordeaux, and Bordeaux he would believe to be non-Bordeaux.

Indeed, a fullblooded subjunctive requirement seems practically impossible to satisfy. So reliabilists moderate their demands by restricting the scope of the counterfactual condition to relevant alternatives. Holmes's belief needn't track truth through the minefields set out by malevolent demons in order to qualify as knowledge.

His epistemic prospects clearly depend on what alternatives count as relevant. If knowledge is to be possible at all, the sceptic's bogies – the machinations of malevolent demons and manipulative neurosurgeons – must be excluded as irrelevant. If all other actual wines are relevant alternatives for Holmes, then in order to know, he needs the ability to discriminate between Bordeaux and every other wine. The obstacles to knowledge then remain formidable. But the class of relevant alternatives might be narrower still. Perhaps it is limited to wines Holmes is likely to encounter, or even to the wines in his own cellar. Then his powers of discrimination need not be so great. If he

can distinguish Bordeaux from the other members of fairly restricted classes of wines, he is in a position to know what he's drinking.

Reliabilist requirements for knowledge are variable, expanding and contracting with the range of relevant alternatives. Against the background of one set of alternatives, Holmes knows; against the background of another, he does not. Indeed, if the range is sufficiently restricted or gerrymandered, even Watson turns out to know; the possibility that the wine is muscatel can be excluded as irrelevant. Apparently any true belief can be constituted as knowledge by suitably configuring the range of relevant alternatives. The epistemic status of a true belief thus depends on the selection of such a range; and without criteria to guide us, it is hard to avoid the appearance of begging the question in making a selection. Still, Holmes's epistemic situation is better than Watson's in that significantly more austere restrictions are required to constitute Watson's belief as knowledge. In this respect at least, the smarter man has an epistemic advantage.

It is not clear though, that Holmes can sustain his advantage. Watson, we may suppose, reliably classifies wines as *rotgut*, *table wine* and what he calls '*vintage stuff*'; and his beliefs about wine quality result from lawlike causal chains. So according to both reliabilist and causal theories, Watson knows he's drinking rotgut.

Holmes knows nothing of the sort. 'Rotgut' is not part of his conceptual repertoire, so he formulates no beliefs about rotgut. Since belief is required for knowledge, Watson knows something about their shared experience that Holmes does not. Still, Holmes brings to the wine tasting a wealth of refined, delicate distinctions. The first sip convinces him that he's drinking a 1986 Thunderbird, made from a resoundingly inferior grape grown in vacant lots just off the Santa Monica Freeway; a wine aged for a week in a plastic vat previously used to launder sweat socks. Holmes, with his more sensitive perceptual and conceptual categories, seems to be in a position to know a good deal more than Watson. Being able to frame more hypotheses, he has more candidates for knowledge than Watson does.

The problem is this: the more distinctions a system of categories admits, the less difference there is between adjacent categories. As we refine our conceptual schemes, we increase our chances of error. Although Holmes can usually tell the vintage of the wine he's drinking, no more than anyone else is he infallible. The perceptible differences among vintages are often extremely subtle and difficult to

discern. Common conditions – the beginnings of a head cold, a poorly rinsed glass, a moment's inattentiveness, a stuffy room – can throw the most sensitive palate off, leading the taster to confuse a Margaux with a St. Julien. So Holmes's true belief that he's drinking a Margaux does not track truth very far. Were he victim to such contingencies, he would think he was drinking one wine when he was drinking another. The sources of error here are not hyperbolic constructions or remote possibilities, but everyday eventualities. So they cannot legitimately be excluded by circumscribing the range of relevant alternatives. On a reliabilist theory, Holmes does not know; nor does anyone else whose judgments are vulnerable to such contingencies. The more delicate our distinctions, the more easily circumstances conspire to confound judgment. So as we refine our categories, we diminish our prospects for knowledge.

Causal theories seem to do better here, being indifferent to the counterfactuals that confute the reliabilist. If Holmes's belief that he is drinking a Margaux is caused by the fact that he's drinking a Margaux, and if the causal chain that eventuates in that belief instantiates a law of nature, Holmes knows that he's drinking a Margaux. It seems then that the causal theory can accommodate increasing categorial refinement, being concerned solely with the genesis of actual beliefs; for there is no a priori limit to the precision of beliefs that can be lawfully generated.

The problem is that Holmes is no dummy. He is well aware of the circumstances that might mislead – of the availability of wines easily mistaken for a Margaux, and of the physiological and environmental conditions that can affect the palate. And he realizes that he cannot be confident that no such circumstances obtain. This gives him pause. Although he strongly suspects that he is imbibing a Margaux, he can't bring himself fully to believe it. And without belief, there is no knowledge. So Holmes's appreciation of the precariousness of his epistemic situation prevents him from knowing.

Respect for evidence may also inhibit knowledge. Suppose there is such a thing as extrasensory perception, and that the absence of evidence for such a faculty is due to the fact that genuine extrasensory perceptions are extremely hard to distinguish from a variety of unreliable sources of intimation.[1] Watson and Holmes are equally extrasensorily perceptive. But Watson is credulous; Holmes is not. So Watson believes the deliverances of ESP, dismissing the evidence out of hand.

Holmes respects the evidence and the methods of the sciences that produced it. So he does not credit his extrasensory perceptions. Although he cannot prevent himself from experiencing them, he withholds belief, for he can find no legitimate grounds for the suspicions they produce. Holmes then does not know; for his epistemic scruples prevent him from forming the requisite beliefs.

On both causal and reliabilist accounts, Watson does know. Extrasensory perceptions yield true beliefs via lawful, if unrecognized, causal chains. And since ESP is reliable (even though we have no reason to think it is), Watson would believe its deliverances if they were true, and would not believe them via ESP if they were false. So Watson's obliviousness to the evidence serves him well; it enables him to know.

In summary, externalism favors the employment of crude categories; for refinements invite error and unreliability. If our objective is to believe what is true and disbelieve what is false, it is reasonable to restrict opportunities for belief to cases in which truth and falsity are easily distinguished.

Externalism also favors obliviousness to evidence. A subject is affected by evidence if that evidence initiates the causal chain or activates the reliable mechanism responsible for his belief. But there is no epistemic advantage to his being aware of the evidence, for a belief's tether is not strengthened by the subject's cognizance of its constitution. Indeed, knowledge may be lost by his attempt to give evidence its due. For evidence can mislead, inhibiting the adoption of true, tethered beliefs and encouraging the adoption of false, untethered ones. We do best then to let evidence exercise its effects subliminally when it is integral to a belief's tether, and to ignore it when it is not.

Finally, externalism favors unreflectiveness about one's epistemic circumstances. Indeed obliviousness to evidence is but a special case of this. Appreciation of the opportunities for error and of the claims of alternative hypotheses cause reservations, leading the reflective agent to suspend judgment. A heady, if unfounded, confidence, borne of the ability to overlook obstacles, supplies the unreflective subject with a goodly store of beliefs, many of which turn out to be true and tethered. The unreflective subject succeeds or fails depending on the proportion of true, tethered beliefs in his doxastic system. But the reflective subject is bound to fail; for unless he is willing to believe, he

is in no position to know. He neither believes what is true nor disbelieves what is false; lacking sufficient evidence, he suspends judgment.

3.

Internalism maintains that a claim is justified to the extent that it is reasonable in light of what is already known. Justification thus depends on coherence with a system of already accepted claims.[2] Some take the relevant system to be individualist; others take it to be social. I shall consider accounts of both kinds.

Keith Lehrer is an individualist (Lehrer 1986). He holds that the justification for a hypothesis is a matter of its coherence with a system of claims the subject already accepts, where a statement coheres with a system if its acceptance is more reasonable relative to that system than is the acceptance of any competing claim. Epistemic justification does not, of course, demand coherence with everything the subject holds. He may accept statements for purposes other than knowledge, and coherence with such statements confers no epistemic status. If, e.g., he accepts a religious doctrine on the basis of faith and for the purpose of salvation, the coherence of a claim with that doctrine would be epistemologically irrelevant. What is required for epistemic justification, Lehrer contends, is that a hypothesis cohere with the statements the subject accepts for the purpose of knowledge. These statements constitute his personal acceptance system, and he is personally justified in accepting anything that coheres with that system. But personal justification is not enough, for personal acceptance systems typically contain falsehoods. And a statement that coheres with antecedently accepted falsehoods is not on that account a viable candidate for knowledge. Candidacy is restricted to statements that also belong to the subject's verific acceptance system – the system that results when his personal acceptance system is purged of all error. A claim that coheres with both is, Lehrer believes, completely justified for the subject. For its justification does not depend essentially on any false belief; and relative to the truths the subject believes, it is more reasonable than any of its rivals. Indeed, on Lehrer's account, an accepted, completely justified truth is knowledge.

What coheres with a narrow system can fail to cohere with a broader one. So Watson, with his limited purview, knows things that

Holmes, burdened with a more comprehensive one, does not. Upon sighting a bird, Holmes and Watson form the belief that it is a superb starling. Watson's relevant background beliefs are truths about the characteristic markings of superb starlings. He has no beliefs about the bird's habitat; for, although he studiously attends to the pictures in the bird watcher's manual, he ignores the accompanying text. Given the information in his acceptance system, Watson's belief is completely justified. And since the bird, an escapee from the London Zoo, is in fact a superb starling, Watson knows that it is. Holmes, however, does not. Although he too recognizes that the bird in question has the markings of a superb starling, he realizes that such birds, being indigenous to equitorial Africa, are unlikely to be found on Baker Street. So relative to Holmes's acceptance system, it is at least as reasonable to suspect that they've sighted a strangely marked local bird. Watson's ignorance thus enables him to know what Holmes cannot. The fact that prevents Holmes from knowing, being external to Watson's acceptance system, cannot undermine Watson's justification (Ginet 1980).

The point is not that Watson benefits from ignorance of one specific, and in this case misleading, fact. It is rather that relatively sparse systems may be better sources of knowledge than richer systems. Watson can appeal only to markings to determine what kind of bird he's looking at. Still, his resources are sufficient for complete justification and knowledge. Holmes's system includes information about markings and about habitat. So coherence with Holmes's system is harder to achieve. But its achievement gives Holmes no more than Watson already has – viz., complete justification and (often) knowledge. When, as in the present case, beliefs about habitat undermine an identification based on markings, Holmes is not justified in accepting any identification. Watson may then know what the bird is; Holmes surely does not. So if an acceptance system is sufficient to generate knowledge about a subject, additional information about that subject is otiose and potentially detrimental. Its incorporation into the system increases the difficulty of achieving coherence, making it harder to know.

This suggests that Holmes could protect his justification and enhance his epistemic prospects by isolating his system from potential defeaters. So long as he remains ignorant of history, for example, his justification cannot be undermined by unfortunate historical prec-

edents. Lehrer suggests, however, that such self-conscious protectionism would be incompatible with the quest for knowledge. "A person who seeks after truth in a disinterested and impartial manner would not arbitrarily restrict his beliefs in this way" (Lehrer 1974, p. 209). Such protectionism is not obviously arbitrary. Given the goal of knowledge, it seems a reasonable strategy to accept the minimum required to generate completely justified true beliefs. For, as Holmes's predicament shows, to include superfluous information is to ask for trouble. Still, if such intentional restrictions on the scope of his acceptance system are arbitrary, Holmes could not, Lehrer believes, adopt Watson's stance without forsaking the quest for knowledge.

But Watson comes by his limitations naturally. So his motives as a knowledge seeker cannot be impugned because of his failure to incorporate certain information into his acceptance system. Indeed, he may be incapable of doing so. Suppose Holmes's confounding belief derives from a complex statistical generalization correlating the intensity of a bird's coloration with the mean temperature of its habitat – a generalization from which it follows that a brightly colored bird like the superb starling is unlikely to be found in a temperate climate. Watson does not know the generalization; moreover, he could not understand or appreciate its import, were it imparted to him. So neither it nor its denial can enter into his personal acceptance system. As a result, the generalization cannot defeat any of his completely justified beliefs. For epistemically inaccessible truths are, for the internalist, epistemically inert. It is then his stupidity, not just his ignorance, that enables Watson to know what the more intelligent Holmes cannot.

Like externalism, individualist internalism favors the employment of crude categories, where differences are stark and instantiation is easily verified. Reasonably conscientious application of such a system typically produces knowledge. But when category systems admit of subtle distinctions, knowledge is much harder to achieve. It is, e.g., fairly easy to tell whether something is a bird; fairly hard to tell whether it is a tree pipit. So since Watson is given to entertaining hypotheses at the level of

x is a bird

he's likely to generate a good deal of (trivial) knowledge. Since Holmes draws finer distinctions, he has a harder time. Often no single

classification at the species level, for example, is most reasonable; so none coheres with his personal acceptance system. Moreover, if one alternative does prevail, it is apt not to beat its competition by much. So the least inaccuracy in his relevant background beliefs may exclude it from his verific acceptance system. If, e.g., the only discernible difference Holmes recognizes between a tree pipit and a buff meadow pipit is that the former is slightly plumper than the latter, his justification is defeated if he's even slightly wrong about how plump a tree pipit is expected to be. So Watson is likely to come away from a bird watching expedition with a lot more knowledge than Holmes. For Watson will have formed many completely justified beliefs:

x_1 is a bird;
x_2 is a bird;
...
x_n is a bird.

Holmes, having attempted more precise classifications of $x_1, \ldots, x_n$ will have encountered some birds he could not identify, some whose identification he was not personally justified in accepting, some whose identification he was not completely justified in accepting, and some in which the identification he was completely justified in accepting was nonetheless false. Indeed, under the circumstances, Holmes might reasonably refrain from accepting any claims at this level of refinement. Since he desires to disbelieve falsehoods as well as to believe truths, he would be wise to suspend judgment where the prospect of error looms large. Here again, it seems rational to revert to Watson's safer stance. For Watson achieves the goal of believing truths and disbelieving falsehoods far better than Holmes does.

This might be doubted. It might seem that Holmes, having a richer cognitive repertoire, is in a position to form more undefeated justified true beliefs than Watson. If so, he knows more than Watson, even though Watson knows some things he does not. But the premise is false; for Watson can generate undefeated justified true beliefs at least as quickly as Holmes. Of course, Watson's will tend to be trivial, banal, and boring, while Holmes's are often original, interesting, and important. But contemporary epistemology does not have the resources to discriminate between significant and insignificant beliefs. So it has no basis for ruling that Holmes's justified true beliefs are epistemically better than Watson's.

4.

We are reluctant to credit Watson with knowledge because he neglects information that seems plainly relevant to the justifiability of his beliefs. But we can't fault him merely for ignoring information that bears on the topic of his concern; such information is inexhaustible, so some of it is bound to be ignored. If we consider Holmes's beliefs justified, it is because we think he has taken enough into account. He has neglected no important information, and no significant inference.

But if taking enough into account isn't taking everything into account, and isn't taking into account only what the subject considers relevant, how is it determined? Gilbert Harman suggests that 'enough' here is a social matter – that the standards of the epistemic community decide what information is important, what inferences are significant, and how much is required in order to know (Harman 1973). Since knowledge is an ordinary, not an extraordinary, cognitive achievement, the standards in question must be ones that normal members of the community normally meet. Watson's justification is inadequate then when it omits reasons that normal members of the community would normally invoke to justify such a belief.

The socialization of justification has two consequences worth noting. First, justifiedness varies with community standards. Holmes's belief may be justified according to the standards of one epistemic community and fail to be justified according to those of another. And variability in the requirements on justification results in variability in knowledge (Cohen 1986). Holmes knows relative to one set of standards, but does not know relative to another. Moreover, relative to a community with sufficiently lax or peculiar standards, Watson too has justification and knowledge.

Second, epistemic accessibility is construed socially. Epistemic resources count as accessible if they are available to normal members of the community, even if the peculiarities of an individual's situation make them unavailable to him. It follows that information the subject does not possess and inferences he does not draw can defeat his justification, if that information is known to, or those inferences drawn by, normal members of the community. The widely known fact that London is inhospitable to tropical birds can defeat Watson's justification, even though he is unaware of that fact. And the easy inference from the extreme improbability of encountering a tropical

bird on Baker Street to the unreasonableness of believing one has done so can defeat his justification, even though he fails to draw that inference.

Since Holmes more than satisfies community standards, his epistemic success may seem assured. But it is not; for by excelling, he invites trouble.

Being an acute reasoner, Holmes validly infers that recent fluctuations in grain prices discredit the claim that the Prime Minister lied about the prospects for peace. Normal members of the community lack the acumen to recognize the relevance of grain prices to the Prime Minister's statement, and to draw the proper inference from them. So the considerations Holmes adduces are too arcane to undermine the social justification for the claim. As it turns out, the Prime Minister did lie; the economic indicators are misleading. Since normal members of the community are not bound to consider those indicators, their justification is intact. So they know that the Prime Minister lied. What about Holmes? If he need only satisfy the community's standards, 'The Prime Minister lied' is justified for him, since he is privy to all the information that justifies normal members of the community in their belief. But he pays for his justification by sacrificing his belief. Realizing that the evidence of the grain prices discredits the community's justification, he cannot consider himself justified in believing the Prime Minister lied. So, being rational, he does not believe it. But without belief, there is no knowledge.

Because he is smarter than others, Holmes is willy-nilly answerable to more exacting standards. He cannot ignore truths within his ken merely because others are incapable of appreciating their significance. Holmes thus fails to know, although his intellectual inferiors succeed.

Again, if Holmes's categories are more refined than those employed by the community at large, his judgments cannot be sustained by community standards. But without socially shared standards for their justification, those judgments are not candidates for knowledge. Holmes's additional conceptual and perceptual sensitivity does not then enable him to know what normal members of the community cannot. Indeed, the social requirements on justification are such that it is impossible for someone to know what normal members of the community cannot.

Since social internalism takes knowledge to be relative to an epistemic community, it might seem to have the resources to block these

untoward results. Can't we give Holmes his due by evaluating his beliefs in terms of the standards of a more intelligent, sensitive community? The difficulty is that *any* epistemic community will have members whose cognitive abilities exceed the norm. So Holmes's predicament will recur no matter how we adjust the membership conditions on the relevant epistemic community.

Internalism, whether individual or social, thus favors conformity and a sort of cognitive minimalism. A person's epistemic prospects are best if his doxastic system includes no more than is necessary to justify his beliefs. Additional information and greater abilities produce no epistemic advantage; and they have the capacity to undermine the justification the minimal system supplies.

5.

It should now be obvious that Holmes's predicament is endemic to contemporary epistemology. This is no surprise: for it results from features that proponents count as virtues of their theories – features that yield the ability to make do, in one way or another, with less than ideal justification. The very limitations on the requirements for knowledge that make it possible for the Watsons of the world to know make knowledge more difficult for individuals like Holmes.

But this conclusion should not be construed as a counterexample to currently popular theories. In some cases at least, it seems reasonable to believe that Watson knows more than Holmes. The blunt man of solid, uninspired common sense, being untroubled by subtleties, may know what's what, while the more sensitive, finely tuned intelligence is distracted by nuances. Nor is it likely that a further condition on knowledge could redress the balance in favor of intelligence. For to attempt to redress that balance would be to move again in the direction of ideal justification, back into the snares of the sceptic.

What Holmes's predicament shows, I believe, is that knowledge, as contemporary theories conceive it, is not and ought not be our overriding cognitive objective. For to treat it as such is to devalue cognitive excellences like conceptual and perceptual sensitivity, logical acumen, breadth and depth of understanding, and the capacity to distinguish important from trivial truths. Even when Watson knows more than Holmes, he does not appear to be cognitively better off.

This suggests that it is unwise to restrict epistemology to the study

of what contemporary theories count as knowledge. What is wanted is a wide-ranging study of cognitive excellences of all sorts, and of the ways they contribute to or interfere with one another's realization. The fruits of such a study might enable us to understand how Socrates, knowing nothing, could be the wisest of men.

NOTES

* I am grateful to Warren Goldfarb for sharing his knowledge of wines with me, and to Kenneth Winkler for sharing his knowledge of birds.

[1] This example is a variant of one developed in Bonjour (1980).

[2] Some internalists – e.g., Chisholm – recognize basic statements that are supposed to be inherently reasonable. But they acknowledge that most statements are not basic, so justification is mostly a matter of coherence (Chisholm 1982). Moreover, their admission of inherently reasonable statements is problematic. For it would not be reasonable on internalist grounds to accept a putatively basic statement that conflicted with the appropriate background system. For example, I ought not accept the claim that I see something red if I am justifiably convinced that I am color blind.

REFERENCES

Bonjour, Laurence: 1980, 'Externalist Theories of Empirical Knowledge', in Peter A. French, Theodore E. Uehling, Jr., and Howard K. Wettstein (eds.), *Midwest Studies in Philosophy V*, University of Minnesota Press, Minneapolis, pp. 53–73.

Chisholm, Roderick: 1982, *Foundations of Knowing*, University of Minnesota Press, Minneapolis.

Cohen, Stewart: 1986, 'Knowledge and Context', *The Journal of Philosophy* **83**, 574–83.

Dretske, Fred: 1971, 'Conclusive Reasons', *Australasian Journal of Philosophy* **49**, 1–22.

Ginet, Carl: 1980, 'Knowing Less by Knowing More', in Peter A. French, Theodore E. Uehling, Jr., and Howard K. Wettstein (eds.), *Midwest Studies in Philosophy V*, University of Minnesota Press, Minneapolis, pp. 151–62.

Goldman, Alvin: 1967, 'A Causal Theory of Knowing', *The Journal of Philosophy* **64**, 357–72.

Harman, Gilbert: 1973, *Thought*, Princeton University Press, Princeton, New Jersey.

Lehrer, Keith: 1974, *Knowledge*, Oxford University Press, Oxford.

Lehrer, Keith: 1986, 'The Coherence Theory of Knowledge', *Philosophical Topics* **14**, 5–25.

Nozick, Robert: 1981, *Philosophical Explanations*, Harvard University Press, Cambridge, Massachusetts, pp. 172–96.

Department of Philosophy
Wellesley College
Wellesley, Mass. 02181
U.S.A.