

Moore's Paradox and the Norm of Belief

ABSTRACT: Reflection on Moore's Paradox leads us to a general norm governing belief: fully believing that p commits one to the view that one knows that p . I sketch conceptions of both the nature of belief and the nature of knowledge that account for this norm. The norm for belief revealed by Moore's Paradox leads us to an appreciation of the deep philosophical significance of the concept of knowledge.

1. Moore's Paradox

G.E. Moore discovered what at first seems a minor puzzle concerning such statements as "It is raining, but I don't believe it". This statement strikes us as something akin to a contradiction, but in fact the statement is consistent, since it is logically possible that it should be raining at a time when the speaker does not believe that it is. As Moore also observed, it is possible for another person to say of me, "It is raining but *he* doesn't believe it", or for me to say, at a later time, "It *was* raining, but I didn't believe it", and in either of these cases the speaker would apparently assert the same proposition as in the original "It is raining but I don't believe it" (said by me), yet there is no air of inconsistency or absurdity about these statements.¹ I shall refer to sentences that are absurd in the same way (whatever that is) as "It is raining but I don't believe it" and "It is raining but I believe it isn't" as "Moore-paradoxical sentences".² Moore's Paradox is the puzzle of explaining why Moore-paradoxical sentences are absurd.³

In my view, the epistemological lessons of Moore's Paradox are rich. Reflection on the paradox assists us in the analysis of the concept of knowledge, it leads us to an account of the deep philosophical significance of that concept, and it suggests an

¹Moore 1993.

²Here I follow Shoemaker's (1996, p. 74) terminology.

³The term "Moore's Paradox" is from Wittgenstein, who appears to use it to refer to this puzzle, rather than to the Moore-paradoxical sentences themselves (1968, p. 190).

important general constraint on rational belief.

2. Linguistic Accounts

Most who have weighed in on Moore's Paradox have offered purely linguistic solutions. These are solutions that seek to explain the absurdity of Moore-paradoxical sentences in terms, broadly speaking, of rules of language. This includes G.E. Moore's own solution, according to which, in asserting that p , one *implies*, even though one does not actually say, that one believes it. I take it that the sense of "implies" here is something like the sense in which, in asking, "Have you stopped beating your wife?" one implies that the addressee has beat his wife. Moore observed that not only is it absurd to say, "It is raining but I don't believe it", but it is also absurd to say, "It is raining but I don't know that it is". He went on, therefore, to propose that in asserting that p , one generally implies that one *knows* that p . This explains the air of contradiction: "It is raining" implies that I know that it is raining, but "I don't know that it is" denies that I know this; so when I say, "It is raining but I don't know that it is", something that I *say* contradicts something that I *imply*.

Wittgenstein offers another linguistic solution. He contends that although in most contexts " x believes y " means that the person denoted by " x " believes the proposition denoted by " y ", sentences of the form "I believe that p " are a special case: rather than meaning that the speaker believes that p , they function as a kind of hesitant assertion of p . Thus, "It is raining, but I believe that it isn't" is actually contradictory: the first part asserts that it is raining, while the second part denies, albeit hesitantly, that it is raining.⁴ Perhaps the best way to see the futility of the Wittgensteinian approach is to introduce a new term, say, "schmelieve", with the stipulation that "I schmelieve that p "—perhaps unlike "I believe that p "—is to mean only that the speaker believes that p . We could then ask why "It is raining, but I schmelieve that it isn't" is an absurd thing to say, even though it is not contradictory. And since this question would be just as puzzling as our original question involving "believe"—indeed, "It is raining but I schmelieve that it isn't" seems to be absurd in

⁴Wittgenstein 1968, pp. 190-2; 1980, pp. 472-8, 501. Malcolm (1995) defends Wittgenstein's view at greater length.

exactly the way we originally thought that “It is raining but I believe that it isn’t” was absurd—nothing is accomplished by positing the sort of ambiguity in the word “believe” that Wittgenstein posits.

A third linguistic solution is Timothy Williamson’s, according to which the activity of assertion is governed by a *constitutive norm* (roughly, a norm whose general recognition is a precondition on there being such a thing as assertion). This constitutive norm is the rule that one ought only to assert what one knows to be the case. Thus, in saying, “It is raining but I don’t know it is”, I necessarily violate a norm essential to the practice of assertion, since I cannot know that it is raining and that I do not know that it is raining.⁵

One or more linguistic solutions may well be correct, as far as they go. I suspect that Moore is right to say that an assertion of *p* implies that the speaker knows that *p*, and this fact may in turn be explained by Williamson’s thesis that the norm governing assertion is “assert only what you know”. But all purely linguistic solutions are incomplete. Moore’s Paradox cannot be fully resolved by appeal to rules governing solely the use of language, because it is easy to construct non-linguistic versions of the paradox. It would be absurd to think to oneself, even without giving any overt expression to the thought, that it is raining but that one does not believe this. And presumably if we can explain why this *thought* would be absurd, we would also thereby understand why the corresponding assertion expressing such a thought is absurd. So it seems that we should rather focus on the question of why what we may call “Moore-paradoxical beliefs” are absurd.⁶

3. The Self-Intimation Account

Perhaps Moore-paradoxical beliefs are absurd because (typical) beliefs are *self-intimating*, in the sense that if one believes that *p* (or if one *explicitly* believes that *p*, or *consciously* believes that *p*), then, if one considers whether one believes that *p*, one must believe that one believes that *p*. The Moore-paradoxical thinker believes

⁵Williamson 1996, especially pp. 506-7.

⁶Shoemaker (1996, pp. 75-6) and de Almeida (2001, p. 33) stress this point.

something like (1) or (2) below:

- (1) It is raining, but I do not believe that it is.
- (2) It is raining, but I believe that it is not.

This (plausibly) entails that he believes that it is raining.⁷ This, given the self-intimation thesis (and suitable background conditions), implies that he believes that he believes that it is raining. If he also accepts the second half of (1), then the Moore-paradoxical thinker has contradictory beliefs: he believes both that he believes it is raining, and that he does not believe it is raining. And if he accepts the second half of (2), then the Moore-paradoxical thinker must *take himself* to have contradictory beliefs: he believes both that he believes that it is raining and that he believes that it is not raining. In either case, it is plausible that the Moore-paradoxical thinker must be irrational.⁸

This sort of account seems plausible when applied to (1) and (2). But what about the epistemic version of the paradox, in which a thinker accepts something of the form of (3)?

- (3) It is raining, but I do not know that it is.

Can something along the lines of the self-intimation account explain the absurdity of believing (3)? Only if there is a plausible self-intimation thesis that involves *knowledge* in the appropriate way. Some epistemologists find it plausible that, if one knows that *p*, and one considers whether one knows that *p*, one can always know by reflection that one knows that *p*.⁹ It is important to realize that *this* is not the self-intimation thesis required to explain the absurdity of (3). Rather, to explain the absurdity of

⁷I do not assume that belief is closed under entailment in general. But the idea that believing a conjunction implies believing the first conjunct seems an especially plausible form of doxastic closure.

⁸Shoemaker (1996) defends essentially this account of Moore's Paradox, albeit with more qualifications and complications.

⁹Sosa 1997, pp. 232-3.

believing (3), one would have to adopt the stronger assumption that if one *believes* that *p*, and one considers whether one *knows* that *p*, then one must believe that one knows that *p*. From there, we could infer that the thinker who accepts the first conjunct of (3) accepts (if he considers the question) that he knows that it is raining. If he also accepts the second conjunct of (3), this thinker must have contradictory beliefs.

But in fact it is not true that whenever one believes that *p* and considers whether one knows it, one believes that one knows that *p*. To take an extreme case, Peter Unger undoubtedly has many beliefs. But when he considers whether he knows the things he believes, he does not come to believe that he knows them; instead, he comes to believe that he does not know any of them.¹⁰ Doubtless Unger is *misguided* to think this way; but as long as he does, he is a counter-example to the thesis that if one believes a thing and one considers whether one knows it, one will take oneself to know it.

4. Metacoherence and the Knowledge Norm

While it is undoubtedly false that anyone who believes *p* will on reflection believe that they know *p*, I think there is a principle in this general vicinity that is very likely correct: if one believes that *p*, one is thereby *rationally committed* to taking one's belief to be knowledge.¹¹ I shall call this principle the *Knowledge Norm (for Belief)*. If correct, the principle provides a satisfying resolution of Moore's Paradox. A person who believes (3), "It is raining, but I do not know that it is", must thereby believe that it is raining. According to the Knowledge Norm, this commits the believer to the view that he knows that it is raining. In denying, in the second half of (3), that he knows that it is raining, the believer thus contradicts something to which he is rationally committed. This seems like a satisfying account of our sense that Moore-paradoxical beliefs are something akin to contradictions. The absurdity of Moore-paradoxical statements is also easily explained, in terms of the fact that such statements express

¹⁰Unger (1975) makes the case for skepticism at length.

¹¹This principle is advanced by David Owens (2000, pp. 37-41) and criticized by Klein (2004).

Moore-paradoxical beliefs; a statement expressing something akin to a contradictory belief is itself something akin to a contradictory statement.

There is more to be said for this account of Moore's Paradox. The Knowledge Norm account predicts that we should find all statements of the form " p but $\sim\Phi(p)$ ", where $\Phi(p)$ is a condition on the speaker's knowing that p , absurd in much the same way as the original examples of Moore-paradoxical statements. Thus, consider the following proposed examples of Moore-paradoxical statements:

- (1) It is raining, but I do not believe that it is.
- (2) It is raining, but I believe that it is not.
- (3) It is raining, but I do not know that it is.
- (4) It is raining, but that isn't true.
- (5) It is raining, but I have no justification for thinking so.
- (6) It is raining, but my reason for thinking so is false.
- (7) It is raining, but there are (non-misleading) facts that neutralize my reasons for believing that.
- (8) It is raining, but my belief that it is was formed in an unreliable way.
- (9) It is raining, but I would believe that even if it were false.
- (10) It is raining, but I am not sure that it is.
- (11) It is raining, but it is not certain that it is.

Each of these is based on a different putative requirement on knowledge.¹² And it is worth noting that each of (2) through (11) seems an irrational thing to assert or believe, something akin to a contradiction, just as (1) does. With respect to each of (10) and (11), note that the statement is not to be read as one in which the speaker asserts that it is raining and then, a second later, thinks better of his assertion and adds a qualifier; rather, the speaker means to assert that it is raining, and to continue to stand by that assertion while also admitting that he himself is not sure it is raining or

¹²(6) is based on Clark's (1963) fully-groundedness condition. (7) is based on the defeasibility analysis (Klein 1971). (8) is based on reliabilism (Goldman 1992, pp. 105-26). (9) is based on Nozick's (1981, pp. 172-8) tracking condition.

that it is not certain that it is raining.¹³ A speaker might very well say, as one often does, something of the form of “It is *probably* raining, but I am not sure (or, it is not certain) that it is” or “I *think* it is raining, but I am not sure (or, it is not certain) that it is”. But again, we are to imagine a speaker who asserts that it *is* raining, and not merely that it is probably raining or that he thinks it is. That said, it should be fairly clear that (10) and (11) are Moore-paradoxical. (Compare the similar statement, “It is raining, but it may not be”, whose unassertability is perhaps more patent.) Furthermore, there do not seem to be any *other* kinds of Moore-paradoxical propositions, of the form “*p* but $\sim\Phi(p)$ ” where $\Phi(p)$ is *not* plausibly regarded as a condition on the speaker’s knowing that *p*. So it seems that the Knowledge Norm for Belief can be used to explain the absurdity of all Moore-paradoxical statements and beliefs.

One might object that, while the second conjunct in each of (1)-(11) contradicts something that is a condition on one’s knowing that *p* according to *some* proposed analysis of knowledge, there is no proposed analysis of knowledge on which *all* of these second conjuncts contradict a requirement on knowledge; no epistemologist has yet proposed, for example, that knowledge is justified, true, fully-grounded, certain belief that tracks the truth, formed by a reliable method, with no non-misleading defeaters. I think, however, that it is plausible that knowledge really does have all or nearly all of these requirements. This does not mean that the best analysis of knowledge ought to be a very complicated one, containing all these miscellaneous conditions. Rather, it is plausible that some of these conditions entail the others, making the latter redundant. For instance, if a belief is formed in an unreliable way, that fact, plausibly, constitutes a defeater for the belief’s justification; hence, the no-defeater condition entails the reliability condition. Be that as it may, I think that if one can really *know* that it is raining while the second conjunct in one or more of (1)-(11) holds, then it is not absurd to assert that conjunct while affirming that it is raining. In fact, I think this is true of proposition (9)—I think one may know that *p* even when one’s belief does not track the truth—but I shall discuss this point

¹³The difference between (10) and (11) lies in the fact that while (10) refers merely to the speaker’s actual attitude toward *p*, (11) refers to the state of the evidence supporting *p*.

further below.

The thesis that it is irrational to believe any of (1)-(11), if true, renders the Knowledge Norm for Belief very plausible. For if, in believing that it is raining, one were *not* rationally committed to the belief that one knows it is raining, then why could one not rationally go on to accept something that implies that one does *not* know this?

We have not yet satisfactorily resolved Moore's Paradox, because we have not yet explained *why* the Knowledge Norm should hold. And this is something that is initially puzzling. In believing that *p*, I am, obviously, rationally committed to accepting that *p* is true (I take it that this is trivial on a suitable construal of "rational commitment"); and if belief is self-intimating in some sense, or under suitable conditions, then I will (under suitable conditions) also be committed to accepting that I believe that *p*. But my *knowing* that *p* is a much stronger condition than my having a true belief that *p*. So the question is why I should be rationally committed to thinking that this very strong condition holds.

The Knowledge Norm, in my view, is best explained by two further principles, the Principle of Metacoherence, and the Endorsement Theory of Knowledge:

Metacoherence: Consciously believing that *p* commits one, on reflection, to comprehensively, epistemically endorsing one's own belief that *p*.

Endorsement Theory of Knowledge: Knowledge attribution is the most comprehensive epistemic endorsement.

To explain what these claims mean: First, what is it to *endorse* a belief? It is to hold a meta-belief that positively evaluates the first belief—more specifically, it is to believe that one ought to hold the first belief, or at least that it is not the case that one ought not to hold it. An *epistemic* endorsement is, roughly, an endorsement according to one of the criteria or goals peculiarly applicable to belief, or according to the criteria of epistemic rationality, as opposed, say, to a moral or prudential endorsement. "Belief *B* is true" is an epistemic endorsement; "Jon's faith in God is very comforting to him" is a prudential endorsement.

Second, what is it to be rationally committed to Φ ing? "Believing that *p*

commits one to Φ ing” does not entail that, if one believes that p , one rationally ought to Φ ; it entails only that, if one believes that p , then one rationally ought either to Φ or to withdraw one’s belief that p . For example, suppose that, on reflection, it becomes clear to me that my belief in the afterlife is purely a product of wishful thinking. In that case, it is false that I ought rationally to endorse my own belief in the afterlife; rather, I ought to surrender the belief. Indeed, I ought to surrender it *because* I cannot rationally endorse it epistemically (this is, roughly, because I know that beliefs based on wishful thinking are unlikely to be true). This is in accord with the principle of Metacoherence, which simply dictates in this case that I cannot rationally continue to believe in the afterlife while reflectively refusing to epistemically endorse that belief.

Third, why the qualifier, “on reflection”? A subject believing that p might simply fail to consider whether his belief that p is epistemically acceptable; this would not render the subject irrational. What the Metacoherence principle requires is that, if the subject comes to reflect on whether his belief is epistemically acceptable in some respect, then the subject is committed to the view that his belief *is* acceptable in that respect; that is, the subject should believe that his belief that p is acceptable or, if he cannot rationally do that, withdraw the belief that p . This does not require a subject to endorse his belief that p with respect to every epistemic dimension, but only with respect to those dimensions that the subject considers. Similarly, the qualifier “consciously” is included because a subject who merely unconsciously believes that p might on reflection rationally believe that he ought not to believe that p ; perhaps if he is unaware of his actual belief that p , its persistence under such conditions would not mark him as irrational. In any case, I wish only to consider the case of conscious beliefs.

Fourth, what is meant by the qualifier “comprehensively”? The modifier “comprehensively” indicates that a subject is committed to epistemically endorsing his belief across the board, that is, with respect to every dimension of epistemic evaluation that he considers, as opposed, say, to endorsing it on some epistemic criteria but not others. For instance, a subject who considers both whether his belief that p is true and whether it is formed by a reliable method is committed to taking the belief *both* to be true *and* to have been formed by a reliable method; if he doubts either

of these things but continues to believe that p , then he is to some extent irrational.

Fifth, what does it mean for knowledge to be the most comprehensive epistemic endorsement? The idea here is that to say that a person knows that p is to endorse his belief that p on all criteria, or along all dimensions, of epistemic evaluation; it is to say that there is no sufficient (epistemic, as opposed to prudential, moral, etc.) reason why that person should not believe that p . This is not meant as an illuminating *analysis* of the concept of knowledge, partly because I have given at most a sketchy account of the concept of epistemic endorsement, and because it may be impossible satisfactorily to explain that concept without employing the concept of knowledge. Nevertheless, the Endorsement Theory of Knowledge is non-trivial and interesting—it is non-trivial and interesting to say that there exists a distinctive type of evaluation applicable to beliefs, such that “ S knows that p ” offers the most comprehensive positive evaluation, of that kind, of S ’s belief that p .

The principle of Metacoherence and the Endorsement Theory of Knowledge, I suggest, offer a satisfying account of why the Knowledge Norm for Belief should hold, and hence, in turn, of why Moore-paradoxical propositions are unfit for belief or assertion. It is not hard to see why Metacoherence should be true. It seems reasonable that if I cannot endorse my belief that p , even in the minimal sense of holding it to be epistemically *acceptable*, then I ought, epistemically, to withdraw that belief. Consider some practical analogies. If, on reflection, one cannot rationally hold a course of action to be morally permissible, then one ought, morally, to refrain from that course of action. If on reflection, one cannot rationally hold a course of action to be an acceptable way of pursuing one’s goals, then one ought, prudentially, to refrain from that course of action. There ought, in short, to be a sort of coherence between what one does (believes) and one’s attitudes about what one ought to do (believe).

The Endorsement Theory of Knowledge is more controversial. In my view, one’s failure to satisfy one of the conditions (other than the belief condition) on knowing p implies that, in some sense, one ought not to believe that p . For instance, that p is false implies that, in some sense, one ought not to believe that p . Consider one objection to this claim: suppose that p is adequately *justified* for me, even though p is false. It seems to follow from this that I *should* believe that p . A similar objection applies, arguably, to every condition on knowledge that goes beyond justified belief.

Thus, suppose that p is adequately justified for me, although I lack a reliable belief-forming method whose output would be p . Still, given just the fact that p is justified, it seems to follow that I ought to believe that p . Some would question, of course, whether p 's being justified is compatible with my lacking any available reliable belief-forming method whose output would be p .¹⁴ Nevertheless, it is generally agreed that knowledge requires more than justified belief. So whatever further conditions may be required for knowledge, suppose that one of these conditions does not hold, yet one is still justified in believing that p ; in such a case, it seems, one should believe that p , even though one does not know that p . This suggests at least that knowledge-ascription is not *merely* epistemic endorsement of a belief in my sense.

But in fact, while my having justification for believing that p may entail in some sense that I ought to believe that p , there is still a sense in which, if I do not satisfy the other conditions (apart from justified belief) on knowledge with respect to p , I ought not to believe that p . Consider a practical analogy. In driving to the airport, I must decide whether to take Interstate 25, or 104th Avenue. I-25 is normally about ten minutes shorter, although the scenery is less attractive. I decide that, because I do not care much about the scenery, I shall take I-25. Assume that this is a rational decision. But suppose that, after getting on the freeway, I discover that an accident, which I could not have anticipated, has caused a traffic jam. As I sit there stuck in traffic on I-25, I say to myself: "Damn, I should have taken 104th Avenue". This is a reasonable thing to say. But I would not say that my choice was irrational or unjustified given my information at the time. The "should" here is one that takes account of the actual, external facts, whether they were available to me at the time of decision or not, that bear on the achievement of my goals. In *this* sense of "should", that a decision was justified is compatible with the observation that one should have chosen differently. Similarly, there is a sense in which a person epistemically should not believe that p —even if p is epistemically justified—if p is in fact false, or his grounds for believing that p are false, or his belief-forming method is in fact unreliable, etc.

But, whether or not you are convinced that this external sense of "should" to

¹⁴Goldman 1992, pp. 122-3.

which I refer is legitimate, it is clear in any case that if I *take* my belief that *p* to be false, my grounds for the belief to be false, my method of forming the belief to be unreliable, or the like, then I should not believe that *p*. Since, pursuant to the Metacoherence principle, I am committed by my belief that *p* to the view that it is not the case that I should not believe that *p*, I must *not* take my belief that *p* to have any of those undesirable, knowledge-canceling features. And more than this, it seems plausible that I am committed, positively, to taking my belief to have the desirable features of truth, reliability, and so on. If, on reflection, I find that I can at best remain neutral concerning whether my belief-forming method is reliable, then it seems that I am not justified in relying on that method, and thus that I should not believe that *p*.

5. The Problem of Certainty

Perhaps the most puzzling aspect of Moore's Paradox is that it seems almost contradictory to think, "It is raining, but it may not be". Why should this be? I take it that the "may" here indicates epistemic possibility, so it would seem that "it may not be raining" would be true as long as one's evidence that it is raining is less than conclusive. The statement that it *may* not be raining is of course compatible with the possibility that it is actually raining. Moreover, surely one might have adequate justification for believing that it is raining, while realizing that this justification is less than 100% conclusive. In such a situation, would one not be both justified in believing that it is raining, and justified in believing that it *may* not be raining? If so, what is so odd about the belief that it is raining but it may not be? The fact that "It is raining, but it may not be" strikes us as absurd suggests that the belief that it is raining commits one to its being *certain* that it is raining. The problem is that it seems that the standards for justified belief ought to be lower than the standards for certainty.

To add to the puzzle, consider that it does *not* strike us as odd to say or think, "I *believe* it is raining, but it may not be". This suggests that while one cannot rationally believe both that it is raining and that it may not be, one *can* rationally believe both *that one believes* it is raining and that it may not be. But this would be very odd.

What I think is going on here is this. First, we need to draw a distinction between *high degree of belief*, and *outright belief* or *categorical belief*.¹⁵ Degree of belief, or subjective probability, is the sort of thing that explains the betting odds one would be willing to accept. High degree of belief is necessary but not sufficient for outright belief. Two individuals may agree in regarding p as 99% probable, and yet one may consider this good enough to accept p , while the other continues to withhold judgment. There would be a qualitative difference in these two individuals' attitudes: the first individual, for example, can be expected to cease to gather evidence or conduct inquiry about whether p . In his reasoning and practical deliberations, where the truth of p is relevant, he will proceed on the assumption that p , without taking account of p 's probability. He will not treat alternatives to p as relevant in his decision-making or theoretical reasonings about further facts. However, the second individual, the one who has a high degree of belief but not outright belief, will enter the *probability* of p (at least in some rough sense) into his practical and theoretical reasonings. He will consider it appropriate to gather evidence and conduct further inquiry into whether p , assuming the costs of doing so are manageable and the question of whether p is of some interest to him; he will not, at any rate, reject the idea of conducting further inquiry on the grounds that it has already been settled whether p . And while the first individual might agree to conduct inquiry into whether p for some ulterior purpose—for instance, to humor more skeptical colleagues—he can not do so for the purpose of answering the question as to whether p , since, again, he regards that question as already settled in favor of p . This attitude of categorical belief need not be irrevocable—one's categorically believing p *now* is compatible with one's later losing that belief, perhaps as a result of evidence of the kind that one did not consider worth looking for, or perhaps even as a result of pragmatic considerations (if the costs of error should suddenly rise, one may rethink one's earlier view of the amount of evidence required to justify outright belief in a given proposition).

Next, we need to understand the concepts of certainty and epistemic

¹⁵My conception of outright belief here is largely based on Owens' (2000, pp. 142-5) conception of belief, though Owens and I differ over whether one can have reflective control over belief and over whether inconclusive reasons can motivate belief.

possibility. Certainty, at least in one important sense of the word, is a matter of one's being justified in taking the attitude of categorical belief as I have just described it. In other words, it is certain that p when the available evidence makes it reasonable to close the inquiry into whether p , regarding the matter as having been settled in favor of p ; to conduct future deliberation and reasoning on the assumption that p ; and to dismiss alternative possibilities as irrelevant. Epistemic possibility I take to be the dual of certainty: it is epistemically possible that p (it may be that p) if and only if it is not certain that $\sim p$. In this sense, many propositions are certain for me now. For example, notwithstanding skeptical scenarios, it is certain that other people exist, and that they have conscious experiences. I do not need to conduct any further inquiry into that. When deciding whether to rob the local liquor store, I do not need to take into account the possibility that the store owner is a mindless automaton or the product of an elaborate hallucination.

Stronger conceptions of certainty exist. One could propose that it is certain that p only if one's justification for p (or one's having the justification one has) is logically incompatible with the falsity of p . Or one might propose that p is certain only if it is impossible that anything be more justified than p is. I reject these notions because they do not appear to match the ordinary usage of such words as "certain" and "possible". Thus, imagine the following dialogue:

- A: Do you know where your shoes are?
B: Yes, they're by the door. I remember leaving them there just ten minutes ago.
A: Might they be on the moon instead?

In any normal context, B's only appropriate response is "No" (or perhaps, "What on Earth are you talking about?"). But on either of the stronger conceptions of certainty just mentioned, a more apt response would be something like this: "Well, *of course* they might be on the moon. Isn't that obvious?" For after all, it is obvious that one's justification for believing that one's shoes are by the door, when one remembers leaving them there ten minutes ago, is logically consistent with their now being on the moon.

We now have an explanation for how certainty can be a requirement on rational belief: certainty just *is* justification for outright belief in a proposition. When we think that it would be absurd to believe “It is raining but it may not be”, we are imagining the believer *categorically believing* that it is raining but that it may not be; on my view, this amounts to his categorically believing that it is raining and that he is not justified in categorically believing that. If we explicitly imagine a weaker doxastic attitude, Moore’s Paradox does not arise. Thus, imagine someone regarding it as highly probable that it is raining, but also believing that it *may* not be. Or imagine (what is perhaps slightly different) one’s having a high degree of belief that it is raining (being willing to offer strong odds on the proposition, and so on) while believing that it may not be. In these cases, I think it is clear that there is nothing self-defeating or contradiction-like in the subject’s attitude. To find something contradiction-like, we must ascribe to the believer some quite robust sort of belief.

However, in those contexts where one says something of the form, “*I think that p, but I may be wrong*” or “*I believe that p, but perhaps ~p*”, “think” and “believe” refer merely to having a high degree of belief, not outright belief. This is why these statements escape Moore-paradoxicalness.

6. Lessons for the Analysis of Knowledge

Moore’s Paradox can aid us in the analysis of knowledge. The fact that it is in some manner self-defeating to declare, “*p, but I do not know that p*”, suggests a test for proposed conditions on knowledge: $\Phi(p)$ is a genuine requirement on knowing p only if believing (or asserting) that p rationally commits one to accepting that $\Phi(p)$.

Some alleged conditions on knowledge fail this test. Robert Nozick’s proposed counter-factual tracking condition—if p were false, S would not believe that p —is a case in point. The same arguments that show tracking to be unnecessary for knowledge also show that one is not rationally committed to believing that one’s beliefs track the truth. First, consider the failures of closure.¹⁶ Though Nozick was happy to reject the closure principle, *some* failures of closure seem particularly counter-

¹⁶Here, I shall take the closure principle as the principle that if one knows that p and that p entails q , then one knows that q .

intuitive, such as cases in which one knows that $(p \ \& \ q)$ without knowing p . On Nozick's account, for example, one may know that one is not a brain in a vat *and* one has two hands, but one cannot know that one is not a brain in a vat.¹⁷ This seems an unacceptable result.

The Knowledge Norm for Belief leads us to a similar conclusion; indeed, the Knowledge Norm generates a general argument for closure. Suppose that $\Phi(p)$ is a condition on knowledge that violates the closure principle; then it is possible that I know p and I know that p entails q , but because $\Phi(q)$ does not hold, I do not know q . Suppose in addition that, as presumably might also be the case, I know myself to be in such a situation. In this case, I am rational in believing p , but—given the Knowledge Norm for belief—I cannot rationally accept q since I know that I do not know q . This seems unacceptable, because it also seems that my acceptance of p , together with my knowledge that p entails q , rationally *commits* me to accepting q . So my only rational recourse would be to withhold both p and q . One problem with this is that it seems that my knowing p should give me license to (continue to) believe p . Another problem is that, at least on Nozick's view, failures of closure are so pervasive that in the end I would be forced to withhold judgment about nearly everything—for nearly every proposition p that I believe, there exists a proposition q such that q does not satisfy the tracking condition and I can see that p entails q . This can be seen from the fact that most things I believe entail that I am not in one sort of skeptical scenario or another, and as Nozick notes, we do not satisfy the tracking condition with respect to the negations of skeptical scenarios.¹⁸ The point is made even clearer by the example immediately below.

The second problem with the tracking condition is that it entails, falsely, that it is impossible ever to know that one is not mistaken about something that one

¹⁷The nearest possible world in which it is false that (I am not a brain in a vat and I have two hands) is a world in which I do not have two hands, perhaps through some horrible accident a few years ago, but I am still not a brain in a vat. In this world, I do not believe (I am not a brain in a vat and I have two hands). So I track the truth with respect to the proposition (I am not a brain in a vat and I have two hands). But I do not track the truth with respect to the proposition (I am a brain in a vat), since in the nearest world in which I am a brain in a vat, I falsely think I am not one.

¹⁸Nozick 1981, pp. 200-204.

believes. For, if one were mistaken about whether p , one would still believe that one was not mistaken.¹⁹ This example also shows that one is not in general committed to believing that one's beliefs satisfy tracking. For although it is obvious that the belief that I am not mistaken about p fails to satisfy the tracking condition, surely it is false that I therefore ought to withdraw the belief. I can recognize that if I were mistaken, I would think I was not mistaken, but still reasonably go on to say that in fact I am not mistaken. (Compare this case: a murder suspect asserts, "If I *had* committed the murder, I would probably lie to you and say that I did not do it; however, in fact I am not lying and I did not do it". Though the hearer may well distrust the subject's word, the assertion makes perfect sense and is not Moore-paradoxical.)

The tracking condition, then, is one example of a spurious condition on knowledge. Doubtless there are others; the writings of philosophical skeptics are the most likely place to find them. Thus, suppose a skeptic maintains that one knows p only if one's justification for p is logically inconsistent with the negation of p . This putative condition on knowledge would make it mysterious why "It is raining, but I do not know that it is" should be Moore-paradoxical. For it hardly seems that, in believing that it is raining, one is committed to believing that one's evidence for the claim that it is raining logically entails that it is raining—notice that "It is raining, though my evidence for that claim does not logically entail that it is raining" does not seem Moore-paradoxical. Therefore, if the skeptic's condition were a genuine condition on knowledge, there should also be nothing paradoxical about saying, "It is raining, though I do not know that".

7. The Significance of Knowledge

Finally, Moore's Paradox illuminates what is so philosophically important and interesting about the concept of knowledge. Initially, it may seem puzzling that this

¹⁹Nozick (1981, pp. 179-80) later adds a qualification to the tracking condition, resulting in roughly the following: S knows that p only if, if p were false and S used the same method to form a belief about whether p as S actually used, then S would not believe that p . My objection can be posed also to this principle: on Nozick's view, one could never know that one's actual belief-forming method did not deliver a false result.

particular concept should give rise to an entire major branch of philosophy. Why is the concept of knowledge more philosophically significant than, say, the concept of belief, or of matter, or of happiness? Our puzzlement is likely, if anything, to increase once we see the sort of complex and seemingly gerrymandered accounts of knowledge that philosophers in this field have devised.

The Knowledge Norm and the Endorsement Theory of Knowledge alleviate this puzzlement: they suggest that knowledge is the strongest condition that one must take to hold whenever one believes a thing. All areas of human inquiry, both inside and outside of philosophy, are centrally concerned with the attempt to form and transmit beliefs about their subject. Physicists form and transmit beliefs about the physical world; historians form and transmit beliefs about the past; and so on. The most comprehensive intellectual commitments of any inquirer are therefore of the form “I know that p ”. The nature and extent of knowledge are thus of central importance for all intellectual endeavor. In a sense, epistemologists study the conditions for the legitimacy of all claims.

Now in spite of what I have said, some philosophers will doubt my claims. Some may feel that the Knowledge Norm is too strong, that there are conditions on knowledge that one need not take to obtain when one believes a thing. This is a particularly likely position for those who take knowledge to require certainty in some very strong sense. Nevertheless, the concept of *the most comprehensive rational commitment of belief*, or the strongest condition that one is committed to taking to obtain when one consciously believes that p , is an important and interesting one, whether or not that condition is knowledge. Furthermore, I should think it clear that this concept is at least strikingly close to that of knowledge. Clearly, if I consciously believe that p , I am rationally committed to taking my belief to be both true and justified; I am also committed to taking it to be fully grounded (not based directly or indirectly on any false beliefs). This is already quite close to knowledge. If “knowledge” in ordinary English actually requires some further, very stringent condition, such as a demanding form of certainty, such that one is not committed to taking this condition to hold when one believes a thing, then I should think the concept of *the most comprehensive rational commitment of belief* more interesting than that of *knowledge*. In any event, I consider the former concept worthy of the attention

of epistemologists, and it is to that concept that Moore's Paradox directs our attention.

References

- Clark, Michael. "Knowledge and Grounds: A Comment on Mr. Gettier's Paper", *Analysis* 24 (1963): 46-8.
- de Almeida, Claudio. "What Moore's Paradox Is About", *Philosophy and Phenomenological Research* 62 (2001): 33-58.
- Goldman, Alvin I. *Liasons: Philosophy Meets the Cognitive and Social Sciences* (Cambridge, Mass.: MIT Press, 1992).
- Klein, Peter. "A Proposed Definition of Propositional Knowledge", *Journal of Philosophy* 68 (1971): 471-82.
- Klein, Peter. "Skepticism: Ascent and Assent?" in *Ernest Sosa and His Critics*, ed. John Greco (Malden, Mass.: Blackwell, 2004), pp. 112-25.
- Malcolm, Norman. "Disentangling Moore's Paradox" in *Wittgensteinian Themes: Essays 1978-1989*, ed. G. H. von Wright (Ithaca: Cornell Univ. Press, 1995), pp. 195-206.
- Moore, G. E. "Moore's Paradox" in *G. E. Moore: Selected Writings*, ed. Thomas Baldwin (New York: Routledge, 1993), pp. 207-12.
- Nozick, Robert. *Philosophical Explanations*. Cambridge: Cambridge University Press, 1981.
- Owens, David. *Reason without Freedom: The Problem of Epistemic Normativity* (London: Routledge, 2000).
- Shoemaker, Sydney. "Moore's Paradox and Self-Knowledge" in *The First-Person Perspective and Other Essays* (New York: Cambridge University Press, 1996), pp. 74-93.
- Sosa, Ernest. "How to Resolve the Pyrrhonian Problematic: A Lesson from Descartes", *Philosophical Studies* 85 (1997): 229-49.
- Unger, Peter. *Ignorance: A Case for Scepticism* (Oxford: Clarendon, 1975).
- Williamson, Timothy. "Knowing and Asserting", *Philosophical Review* 105 (1996): 489-523.
- Wittgenstein, Ludwig. *Philosophical Investigations*, 3rd ed., tr. G.E.M. Anscombe (New

York: Macmillan, 1968).

Wittgenstein, Ludwig, *Remarks on the Philosophy of Psychology*, vol. 1, ed. G. E. M. Anscombe and G. H. von Wright, tr. G. E. M. Anscombe (Oxford: Basil Blackwell, 1980).