

BELIEF CONTEXTS AND EPISTEMIC POSSIBILITY

HYLARIE KOCHIRAS

University of North-Carolina at Chapel Hill

Abstract

Although epistemic possibility figures in several debates, those debates have had relatively little contact with one another. G. E. Moore focused squarely upon analyzing epistemic uses of the phrase, 'It's possible that p', and in doing so he made two fundamental assumptions. First, he assumed that epistemic possibility statements always express the epistemic position of a community, as opposed to that of an individual speaker. Second, he assumed that all epistemic uses of 'It's possible that p' are analyzable in terms of knowledge, not belief. A number of later theorists, including Keith DeRose, provide alternative accounts of epistemic possibility, while retaining Moore's two assumptions. Neither assumption has been explicitly challenged, but Jaakko Hintikka's analysis provides a basis for doing so. Drawing upon Hintikka's analysis, I argue that some epistemic possibility statements express only the speaker's individual epistemic state, and that contra DeRose, they are not degenerate community statements but a class in their own right. I further argue that some linguistic contexts are belief- rather than knowledge-based, and in such contexts, what is possible for a speaker depends not upon what she knows, but upon what she believes.

Epistemic possibility statements fill our everyday speech, and our philosophical debates. The means for recognizing epistemic uses of the phrase, 'It's possible that p', are widely agreed upon: epistemic uses are those that express uncertainty by taking the indicative rather than the subjunctive mood. There has been less agreement about what epistemic possibility is, however, and little attempt to reach any consensus. For the most part, the debates concerning epistemic possibility proceed in relative isolation from one another.

The debate begun by G. E. Moore investigates the concept directly. Moore set out to explain the phrase 'it's not certain that' and ended up offering a definition of epistemic possibility. Ian Hacking, Paul Teller, and Keith DeRose, who continued the discussion, may be seen as his successors. These theorists give all statements that both take the indicative mood, and have the form 'It's possible that p', a uniform analysis by retaining Moore's two central assumptions. According to what I shall call the *community assumption*, the truth-value of ev-

ery epistemic use of ‘It’s possible that p ’ depends upon the epistemic states of some community, rather than the speaker alone. Their second assumption — the *knowledge assumption* — holds that every epistemic use of ‘It’s possible that p ’ is analyzable in terms of knowledge rather than belief.

Outside the Moorean tradition, however, neither of these two assumptions has been consistently maintained. When Descartes suggests the possibility that no external world exists, he is expressing only his own epistemic position. Similarly, in the recent conceivability literature that Descartes’ investigations spawned, we sometimes find Stephen Yablo explaining the epistemic possibility of p as finding “that it is true for all you know”, rather than looking to any community (Yablo 1993, p. 7). There have also been some departures from Moore’s knowledge assumption. In a theological debate, for example, Thomas Morris argues that Christ’s sinning was not a metaphysical possibility but an epistemic one; and he then explains epistemic possibility in terms of belief-sets.¹

One theorist who appears to diverge from both the community and the knowledge assumptions is Jaakko Hintikka. Hintikka explains epistemic possibility in terms of an individual speaker’s knowledge, and he also seems to suggest that a similar account may be developed in terms of belief. He does not, however, explicitly challenge or examine the Moorean assumptions, so that is what I propose to do in this paper.

I challenge the two assumptions by examining the Moorean account developed by Keith DeRose, and I draw upon Hintikka’s analysis to develop an alternative. While the truth conditions of what I will call *community statements* do depend upon persons other than the speaker, I argue that there is another class, *individual statements*, whose truth conditions involve the epistemic states of the speaker alone. It will not do to treat these individual statements as degenerate community statements, as DeRose has done. DeRose draws our attention to the following problem: a constant epistemic situation may generate several true but non-equivalent epistemic possibility statements. But while his explanation assumes that the problem arises only for community statements, it in fact arises for individual statements as well. We can account for that case only by rejecting the Mooreans’ first, community assumption and instead treating individual statements as a class in their own right. Against the Mooreans’ second assumption, I shall argue that while most statements of the form ‘It’s possible that p ’ do express knowledge-based possibilities, some others express belief-based ones. In other words, the locution ‘It’s possible that p ’ is not reserved for epistemic possibility, but may be used to express doxastic possibility as well.

The first section of the paper traces the evolution of the Moorean view. The second section is devoted to countering the Moorean assumption about relevant communities. There, I demonstrate that the problem about non-equivalency arises for individual statements. Modeling a complex epistemic state in terms of actual and possible epistemic worlds, I show that the speaker's non-equivalent statements may both be true because each is relative to a different one of these worlds. In the paper's third section, I argue that some uses of 'It's possible that p ' express doxastic possibility, which is to say that they are relative not to knowledge but to belief.

I

Moore offers the following definition for the epistemic use of 'It's possible that p ':

"It's possible that p " = "It's not certain that $\sim p$."²

It is significant that he arrives at his definition of epistemic possibility while trying to make sense of statements having the form 'It's certain that p .' Because they employ the dummy subject 'it', such statements provide no indication of viewpoint when separated from context. Indeed, knowing the truth value of the definiens, 'It's not certain that $\sim p$,' requires being able to answer the question, 'Not certain for whom?' Because the answer to this question will vary across statements and contexts, Moore concludes that epistemic possibility has at least two senses:

- (1) No one *in this group* knows that $\sim p$, where "this group" is understood to include the speaker and whichever human beings the speaker addresses.
- (2) No human being knows that $\sim p$.³

For any statement, we determine which sense of the epistemic 'possible' is at work by determining whether the group whose members could falsify the statement is universal or restricted in scope. I will therefore supply Moore with the term *restricted* for his first sense, and *universal* for his second.

A statement contains the restricted sense of epistemic possibility if it is falsifiable only by either the speaker or a member of her audience. Adapting an example of Moore's, a speaker's remark to a circle of friends, 'It's possible that they're not yet married,' may be true even though someone — the minister for instance — knows that the marriage has taken place.⁴ Only a member of the

circle can falsify the statement, saying, ‘No it isn’t, I was present at the ceremony yesterday.’ Whenever a statement is falsifiable by the knowledge of someone outside the group comprising speaker and audience, the universal sense of epistemic possibility is at work. Modifying another of Moore’s examples, I may say ‘It’s possible that theorem M is false’ to a crowd as ignorant as I. All that is needed for my statement to be false, claims Moore, is for some mathematician to have proved M. Whether I have ever heard of that mathematician is beside the point.

Ian Hacking retains Moore’s central idea about the variability of community membership,⁵ but modifies the account to respond to a worry he raises with the following example:

Imagine a salvage crew searching for a ship that sank a long time ago. The mate of the salvage ship works from an old log, makes some mistakes in his calculations, and concludes that the wreck may be in a certain bay. It is possible, he says, that the hulk is in these waters. No one knows anything to the contrary. But in fact, as it turns out later, it simply was not possible for the vessel to be in that bay; more careful examination of the log shows that the boat must have gone down at least thirty miles further south. The mate said something false when he said, “It is possible that we shall find the treasure here,” but the falsehood did not arise from what anyone actually knew at the time. (Hacking 1967, p. 148.)

According to Hacking, since the mate could have used the log to discover that the ship was not in the bay, his statement is false. More generally, a proposition may fail to be epistemically possible, even when no one knows it to be false. Hacking suggests this “working hypothesis” to accommodate the case:

A state of affairs is possible if it is not known not to obtain, and no practicable investigations would establish that it does not obtain. (Ibid., p. 149.)

As its first condition, the hypothesis retains Moore’s condition about knowledge states, but it adds a second condition appealing to whatever practicable investigations could improve those states. If Hacking is right, then the mate indeed speaks falsely, for a look at the log will establish that the embedded proposition does not obtain. His hypothesis will falter, however, given a case in which there exist practicable investigations that would reveal a proposition to be false, yet we nonetheless think that the proposition is epistemically possible for the speaker.

Paul Teller objects to Hacking’s second condition by introducing such a case. An expectant father, Teller writes, truly may say both ‘It is possible that my child

will be a boy' and 'It is possible that my child will be a girl,' while knowing full well that the fact of the matter could be established via a practicable investigation, e.g. by amniocentesis. It would be hard to gainsay Teller here — both statements seem clearly true. To accommodate the case, Teller suggests a definition that follows both his forerunners in its first condition, but replaces Hacking's second condition:

It is possible that p if and only if

- a) p is not known to be false by any member of community C , nor
- b) is there a member, t , of community C , such that if t were to know all the propositions known to community C , then he could, on the strength of his knowledge of those propositions as basis, data, or evidence, come to know that p is false. (Teller 1972, p. 311.)

Now, by replacing Hacking's second condition, Teller can say that the father speaks truly. But notice that the first condition alone is sufficient to get this result. In fact, given the heritage of that first condition, Moore's own analysis should be sufficient to get the right result in the Expectant Father Case. Moore requires only that no person in the relevant group know that $\sim p$, and this requirement is clearly met, since no one in the father's audience knows either that the child will not be a girl, or that it will not be a boy. What then, does Teller expect his second condition to accomplish? He does not mean the condition to resolve the case at hand, but rather the Salvage Ship Case, for he shares Hacking's intuition there. His definition is designed to imply that the mate speaks falsely, yet without giving the same result for the expectant father's statements.

What this means is that if we find reason to reject the claim that the mate speaks falsely, then we should not side with Teller over Moore, whose simpler account resolves the expectant father case just as well. The question at hand is, then, does Hacking's Salvage Ship Case really call for a revision of Moore's account? I think it does not, for Hacking reaches the conclusion he does only because he conflates epistemic possibility with physical possibility. Given the truth conditions that Moore and Hacking have each provided, a minimal notion of an epistemic possibility upon which they could agree might be: a proposition that is consistent with what the subject knows. What is ruled out by the facts recorded in the log does not fall under this kind of possibility, however, for if it did, Hacking would be talking about what is possible for a person to know. Instead, he talks about what is possible for the vessel:

It simply was not possible for the vessel to be in that bay ... the log

shows that the boat must have gone down at least thirty miles further south. (Ibid., p. 148.)

Presumably, what the log contains is information about physical conditions antecedent to the time that the ship went down: how quickly it was capable of traveling, and where it was located at various points in time. Given that the mate's method for discovering whether it was possible for the vessel to have gone down in the bay is to consider factors such as its location prior to sinking, it seems that what the mate is doing is considering antecedent conditions in the light of physical laws. The mate is investigating whether, given the circumstances described in the log, reaching the bay was physically possible for the ship.

Hacking reasons that since it was physically impossible for the ship to be in the bay, the mate must have spoken falsely. But epistemic possibility does not depend upon physical possibility, and so there is no contradiction in saying that the ship's being in the bay is physically impossible, but nonetheless epistemically possible for the mate. So we have no reason to agree with Hacking when he says that the mate speaks falsely.⁶ We therefore have no reason to accept either Hacking's or Teller's account over Moore's.

Still, Moore's account will not help us resolve an interesting problem raised by Keith DeRose. DeRose introduces the problem via a series of cancer cases, all of which

involve John, who has some symptoms indicative of cancer, and a "filtering" test which John's doctor decides to run and which has two possible results: If the results are "negative," then cancer is conclusively ruled out; if the results are "positive," then John might, but also might not, have cancer: further tests will have to be run. (DeRose 1991, p. 582.)

At the first round of tests, then, only a negative result would be conclusive. In the first case which DeRose constructs from these facts, Cancer Test Case 1A, John's wife, Jane, and the doctor make contradictory statements.

CTC 1-A: John's doctor has received the results of the test, which are negative, but has not told anyone else what the results are. ... John's wife, Jane, has received the call, so she knows that the doctor knows the results of the test, but she does not know what the results are. John's estranged brother, Bill ... who has heard a rumor that John has cancer, calls Jane and says, "I've heard that John has cancer. Is it true?" Here, it seems, Jane might well say to Bill, "It's possible that John has cancer.

... They've run a test on him which may rule cancer out, but they won't tell us the results of the test until tomorrow." However, at the very same time, John's doctor might well say to another doctor, "It's impossible that John has cancer, so we should start planning tests for other diseases." (Ibid., p. 582.)

This first case presents no trouble for Moore, since his first sense of epistemic possibility easily explains why both statements are true: an epistemic possibility statement of the restricted sense is true just in case neither the speaker nor any member of her audience knows the negation of the embedded statement. So, two speakers may truly assert apparently contradictory epistemic possibility statements, as long as the two groups to which they belong have appropriately different epistemic positions. Let 'p' stand for the embedded statement, here, 'John has cancer.' Jane speaks truly by saying,

- (a) 'It's possible that John has cancer,'

since neither the doctor nor anyone else who knows that $\sim p$ is part of her audience. This statement could not truly be uttered by the doctor, knowing as he does that John doesn't have cancer. So, the doctor speaks the truth when he asserts a statement contrary to Jane's.

What Moore's analysis cannot explain is DeRose's Cancer Test Case 1B (CTC-1B). In this case, the facts are just as they were in CTC-1A, except that instead of asking Jane whether it's *true* that John has cancer, Bill asks whether it's *possible*. Jane now responds,

- (b) 'I don't know whether it's possible that John has cancer; only the doctors know.' (Ibid., pp. 584–5.)

Does Jane mean to retract what she said earlier? In CTC-1A she said that it was possible that John had cancer, but now she denies knowing whether it is possible. This is puzzling, because she is epistemically no worse off now than she was before. Moore has said nothing that will help us answer the question DeRose is raising: how can a constant epistemic situation generate non-equivalent epistemic possibility statements?

The answer, DeRose suggests, is that while Jane's epistemic state is unchanged, the factor to which her statement is relative is not constant between the two cases:

The explanation seems to be that in the first case Jane is (truly) saying that it is possible that John has cancer relative to the epistemic situation of a fairly small relevant community — perhaps John’s family — while in the second case Jane, it seems, is professing to be ignorant as to whether it’s possible that John has cancer relative to what is known to a somewhat larger group of people that includes John’s doctor. For all she knows, the doctor may now know that John does not have cancer, and she is assuming that if he does . . . then it is not possible that John has cancer. Thus, it seems that in CTC-1B, John’s doctor, though he is *not* one of Jane’s listeners, is a member of the relevant community to which her use of “possible” is relative. I will not address the issue of how we can determine what the relevant community is for a given utterance. (Ibid., pp. 585–6.)

So, variation in the composition of the relevant community constitutes a variation in the epistemic factor to which a statement is relative. Jane says that John’s having cancer is epistemically possible relative to one community, but relative to a more knowledgeable community, she retreats to a weaker claim. By adding a caveat about variability to Hacking’s proposal, DeRose produces the following truth conditions:

S’s assertion, “It is possible that *P*” is true if and only if (1) no member of the relevant community knows that *P* is false, and (2) there is no relevant way by which members of the relevant community can come to know that *P* is false, where it is remembered that there is a good deal of flexibility in what the relevant community is and what is to count as a relevant way of coming to know: that these matters will vary according to the features of the context in which “It is possible that *P*” is uttered. (Ibid., p. 594.)

Once we realize that Jane’s two statements are not in fact relative to the same community, the phenomenon of non-equivalent statements is explained. Whenever a speaker who has neither learned nor forgotten anything diverges from an earlier statement, she does so because she is speaking relative to a different community or way of coming to know.

II

I think we can accept DeRose’s explanation for a good many cases of non-equivalent statements. It does not apply universally, however, because some uses

of 'It's possible that p ' do not express the epistemic state of a community. In this next case, which I have constructed from DeRose's basic facts and which I will call 'Cancer Test Case 3A', taking the speaker to be expressing anyone's epistemic state other than her own leads to a counterintuitive result.

Again, John has undergone a "filtering" test, in which only a negative result is conclusive. Jane knows that the doctors have the results by now, but she and John will not find them out until tomorrow. John's estranged brother, Bill, telephones, asking, "I've heard that John has cancer. Is it true?" Jane replies, "It's possible that he has cancer. He has some symptoms that have us worried." What Jane doesn't know is that Bill is hiding something from her, something he thinks John should hear first. Bill is actually one of the team of doctors. Moreover, he already knows the result of the test — it's conclusively negative.

There is a strong intuition that in saying, 'It's possible that John has cancer', Jane speaks truly. Yet we would have to say her statement is false if we interpreted her as expressing the epistemic position of a community and applied DeRose's truth conditions. The first of DeRose's conditions requires that no member of the relevant community know p to be false. But Bill knows that p is false, and as Jane's audience, he would seem to belong to the relevant community if anyone does.

We can obtain the result that Jane speaks truly only if we take her to be expressing her individual epistemic position. The proper translation of her statement is not, then, 'For all *we* know, John has cancer' but rather 'For all *I* know, John has cancer.' Does DeRose's account include any provisions for interpreting Jane's statement this way? In principle, yes, for unlike his predecessors, he allows for "solitary uses" of epistemic possibility statements. He considers such statements to be degenerate community statements, for which the relevant community is a "community of one". By interpreting Jane's statement as a solitary use, DeRose could obtain the result that Jane speaks truly in CTC-3A by excluding Bill from the relevant community. But Bill is again Jane's audience, just as he was in the earlier cases, and the absent clear reasons to the contrary, the audience belongs to the relevant community. It would therefore be ad hoc to exclude Bill from the community in only this case.

Even supposing, however, that DeRose could find substantive grounds for classifying Jane's statement in CTC-3A as solitary, another problem arises. DeRose explains the phenomenon of true, non-equivalent possibility statements

by saying that each statement is relative to a different community. This means that he cannot explain the phenomenon when both statements are solitary, for when both statements are relative only to the speaker's own knowledge, there is no shifting community. Such pairs of solitary statements do arise, as we may see by introducing Cancer Test Case 3B. Here the facts are just as in CTC-3A, except that this time, instead of asking whether it's true that John has cancer, Bill asks Jane whether it's possible. Jane now replies, "I don't know whether it's possible. Tomorrow's tests might not be conclusive."

We now have two cases, analogous to DeRose's earlier cases. In the first case, CTC-3A, Jane says,

(c) 'It's possible that John has cancer'

and in the second case, CTC-3B, she says,

(d) 'I don't know whether it's possible that John has cancer.'

Once again, she has gained no new information, yet though in (d) she denies knowing whether (c) is true, she seems to speak truly in both cases. Additionally, both are solitary, or as I prefer to call them, individual statements. In CTC-3A, her statement was true precisely because she was expressing only her own epistemic position. Similarly, in CTC-3B, she gives no indication of trying to express what Bill may know, or what the doctors may know. She is simply expressing her own uncertainty. Because the filtering tests may not yield a conclusive result, however, Jane is not simply uncertain about whether a proposition is true or false. Her state of uncertainty is more complex than that, because Bill has asked her whether John's having cancer is *possible* rather than *true*. The complexity, then, is that she knows her relation to p may soon change. In expectation of this change, she frames different possible outcomes — different propositions which she may come to know. Her total epistemic state is made up of various sets of propositions, some of which she knows, and some of which she may or may not come to know. I will model these sets as epistemic worlds. This model has no ontological implications, but is simply a device for explaining how Jane's individual statements can both be true.

Before developing that model, however, we need truth conditions for individual epistemic possibility statements. The conditions I shall suggest derive from Jaakko Hintikka's analysis. Hintikka explains the epistemic possibility statement, "It is possible, for all that a knows, that p "⁷ as follows.

I shall take [this statement] to mean the same as “It does not follow from what a knows that not- p .” In other words, I shall take [it] to be (very nearly) synonymous with “It is compatible with everything a knows that p .” (Hintikka 1962, p. 5.)

Here Hintikka clearly gives the statement an individualist reading, explaining epistemic possibility in terms of consistency with the speaker’s own knowledge. (He hints that a similar account may be developed in terms of belief, but I shall postpone that suggestion until Section III.) He subsequently suggests that consistency with everything the speaker knows may be too stringent a condition, and in agreement with that, I suggest the following truth conditions.

$\diamond p$ is true for a person a iff p is consistent with the relevant set of propositions known by a .

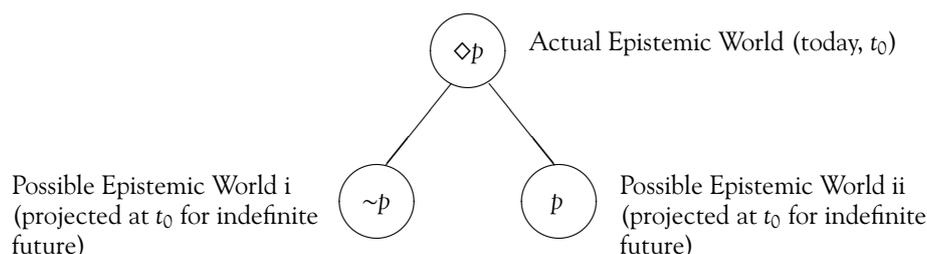
The notion of a relevant set simply acknowledges that epistemic possibilities shift according to changes in epistemic states. It includes the propositions known by the speaker at the time-slice to which her statement is indexed. To develop the model of epistemic worlds, let us restrict our focus to the following propositions.

- q : John’s grandfather had cancer.
- r : Many people with symptoms like John’s have cancer.
- s : Some people with symptoms like John’s do not have cancer.

Letting ‘ p ’ stand for the embedded proposition, ‘John has cancer.’ Jane speaks truly in saying, ‘It’s possible that p ’, since either p or $\sim p$ could be added to that set without producing a contradiction.

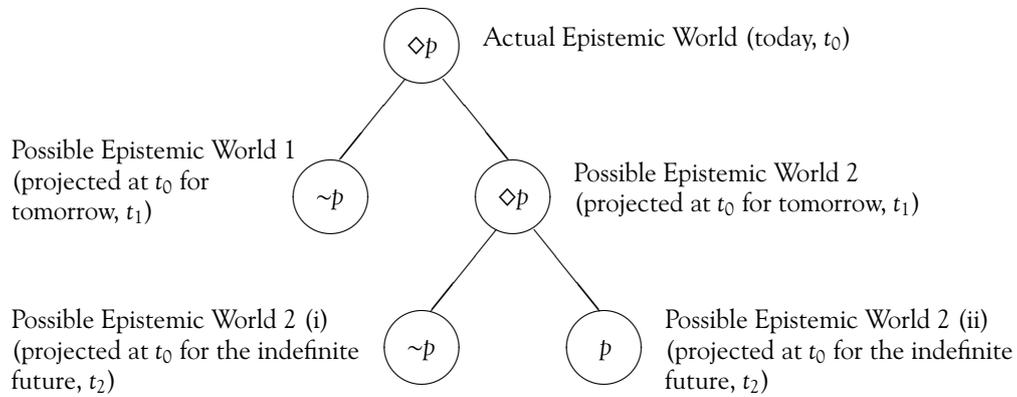
These two first-order propositions, p and $\sim p$, are potential outcomes Jane considers, so neither expresses her relation to p today, t_0 . Her actual relation to p at t_0 is a state of uncertainty, as expressed by her second-order statement, $\diamond p$. The propositions relevant to p that she currently knows make up her *actual epistemic world*. Thus her actual epistemic world at t_0 comprises the relevant set noted above, q , r , s , together with the proposition that expresses her relation to p , $\diamond p$. The potential outcomes she frames regarding p , namely p and $\sim p$, are her *possible epistemic worlds*, and she regards them as such. She says, ‘It’s possible that p ’ precisely because she frames these two possible epistemic worlds. The proposition that interests us in each world is the one expressing the speaker’s relation, whether actual or potential, to p . So let us call that the world’s *salient proposition*. The salient proposition of Jane’s actual epistemic world is $\diamond p$, for

this is her actual relation to p at t_0 . The salient propositions of her two possible epistemic worlds are, respectively, $\sim p$ and p . In the following diagram, I omit q , r , and s , for simplicity's sake, and show only the worlds' salient propositions.



If Jane did not anticipate learning anything new, or if she expected tomorrow's test to be conclusive, then the above diagram would completely represent her epistemic state. In the filtering test, recall, only a negative result would be conclusive, whereas a positive result would indicate the need for more tests. Jane's total epistemic state therefore includes the expectation that her current uncertainty may be succeeded by a further stage of uncertainty. To represent this complication in her total epistemic position, we need to add an intermediate tier of possible worlds to our diagram. One of the possible epistemic outcomes Jane frames is knowledge that $\sim p$, should the test have a negative result. This is P.E.W.1. Since a positive test result would be inconclusive, giving rise to a second stage of uncertainty, P.E.W.2, contains $\diamond p$ as its salient proposition. Having introduced this intermediate tier, we see that P.E.W.(i) and P.E.W.(ii) belong in a lower tier, branching from P.E.W.2. Now that we have placed those worlds in the third tier, let us rename them 'P.E.W.2(i)' and 'P.E.W.2(ii)'. This yields a diagram that represents Jane's full epistemic state (see next page).

Using this diagram, we may easily explain the phenomenon of true but non-equivalent statements when there is no community to shift: each of those individual statements is made relative to a different one of the speaker's epistemic worlds. In CTC-3A, Bill asked whether it was *true* that John had cancer. This direct inquiry about p caused Jane to think about her actual relation to that first order proposition. Using our worlds-metaphor, we might say that it caused her to focus upon her actual epistemic world. She therefore replied by stating the salient proposition of her actual epistemic world, $\diamond p$, that is, statement (c), 'It's possible that John has cancer.' In CTC-3B, Jane knows neither more nor less, but expresses greater uncertainty, saying, (d) 'I don't know whether it's possible that



John has cancer.’ This time, however, she is replying to the question of whether John’s having cancer is possible, and she therefore gives a more detailed account of her epistemic state. Since she realizes that tomorrow’s test may be inconclusive, she focuses on PE.W.2, telling Bill that she does not know whether p is possible. Her statement, ‘I don’t know whether it’s possible that p ’, is the iterated statement, $\diamond\diamond p$, may be translated as follows: ‘For all I know today (t_0), it will be epistemically possible for me tomorrow (t_1) that p .’ This iterated statement, $\diamond\diamond p$, Jane expresses her current uncertainty about whether she will subsequently enter a new state of uncertainty once the results of the filtering test have been received.

Given that Jane may make a variety of statements, different propositions may become the salient proposition of her actual epistemic world. Which proposition is salient at any given moment depends upon which epistemic world she is considering at that moment. When her attention is upon the lower tier of worlds, the salient proposition of her actual epistemic world is $\diamond p$. When she turns her attention to the middle tier worlds, the salient proposition is the iterated statement, $\diamond\diamond p$. So, Jane may truly say both $\diamond p$ and $\diamond\diamond p$, and whether she utters one statement or the other will depend upon whether she is thinking of her actual, current relation to p , or of some potential, future relation to p . Thus, individual uses of ‘It’s possible that p ’ should be understood as relative to the various components of the speaker’s total epistemic situation, not to different communities or ways of coming to know. It therefore makes sense to consider individual uses of ‘It’s possible that p ’ not as degenerate community statements, but as a class in their own right.

III

In addition to assuming that all uses of ‘It’s possible that p ’ express the epistemic state of a community, the Mooreans also assumed that all these possibilities are relative to knowledge. Sometimes, however, the speaker is expressing what is possible relative to a set of beliefs. Examples of such statements can be found with speakers whose beliefs are motivationally biased. Some of the most disturbing examples of motivationally biased beliefs come from people who deny the Holocaust. However, a belief need not be obviously false in order to count as motivationally biased. All that matters is that the speaker holds it because of some need or desire rather than because of evidence. Thus believing in God because one fears the thought of a purposeless universe is quite different from accepting the argument from design, or was in pre-Darwinean times. The example I shall discuss comes from a tragic case, reported recently, in which all five of a young couple’s children died in infancy. For each of Tim and Waneta Hoyt’s children, the cause of death was ruled to be Sudden Infant Death Syndrome, though many suspected Waneta of murder. When she was charged with the murders, nearly two decades later, Tim refused to believe that she might have done it. He continued to insist upon her innocence even after she confessed, claiming that the confession had been coerced. Let us attribute the following statement to Tim: ‘It’s not possible that she killed the babies’.⁸ If we let ‘ p ’ stand for the embedded proposition, ‘she killed the babies’, Tim is asserting ‘ $\sim\Diamond p$.’ Waneta did kill their babies, however, as was subsequently proved at trial. So, $\sim p$ is false, and though Tim is convinced that $\sim p$, he does not know it. Is it epistemically possible for him that his wife is a murderer? Yes, for though p contradicts what he believes, it does not contradict what he knows. So far, I am in agreement with the Mooreans; that his wife murdered their children is an epistemic possibility for Tim, because it is true for all he knows.

But is that the sort of possibility that Tim’s statement expresses? In vehemently denying that his wife could have murdered their children, what is he trying to say? Does he mean to say that her committing such a crime is not compatible with what he really *knows* about her — and does he therefore speak falsely? The Mooreans, who explain all uses of ‘It’s possible that p ’ in terms of knowledge, must say yes. Yet that response does not take account of the purpose our speaker’s statement serves. Has he carefully weighed the evidence? No, he has deliberately ignored it. Is he attempting to represent the world accurately? No, this is a case of motivationally biased belief, and Tim’s aim is not to deter-

mine what actually happened. He has a strong desire to preserve his beliefs about his wife and their past together, a set of beliefs that crucially includes the belief that their children died of natural causes. I think it is the speaker's purpose that determines which sort of possibility his statement employs. Since his statement does not attempt to express facts but rather serves the purpose of preserving cherished beliefs despite the facts, it is best understood as expressing what is possible relative to his beliefs.

If this suggestion is correct, then there are uses of the locution 'It's possible that p ' that express the speaker's internal relation to p . In such uses, p is a *doxastic possibility* — that is, p is possible not relative to what the speaker knows, but to what she believes. In this sense of possibility, a subject's failure to know that a given proposition is true does not make that proposition possible for her. A proposition that the speaker firmly believes to be false may be epistemically possible for her, but it is not doxastically possible, because relative to her beliefs about the world, that proposition is ruled out. I suggest the following truth conditions for doxastic uses of 'It's possible that p '.

$\diamond p$ is true for a person a iff a does not perceive any inconsistency between p and her relevant set of beliefs.⁹

These truth conditions indicate that the speaker is expressing her internal relation to the embedded proposition, p , for they relativize the proposition to a set comprising beliefs rather than known propositions. Additionally, they require perceived rather than actual consistency between p and the set of beliefs, for perceived consistency is an internal relation between the speaker and the proposition. So, when the speaker denies that the proposition, p , is possible, she asserts that that proposition is not consistent with her beliefs. Applying this to our example, Tim Hoyt speaks truly, for his beliefs about his wife's character rule out the proposition that she committed murder.

The Moorean theorists will object, of course, to my claim that Tim's statement is true. They will argue that I have simply conflated truth with sincerity, for however sincere Tim is in denying that his wife could be guilty, his statement is nonetheless false. Indeed, they will note, we may easily suppose that someday, Tim acknowledges the evidence and then concedes the possibility that his wife is guilty. Should he eventually concede that possibility, the Mooreans will say, he would thereby admit that his earlier statement was false. And if that earlier statement was false, it was so because it expressed epistemic possibility; it expressed what was possible relative to what he knew, not to what he believed.

My reply to this objection draws upon David Lewis' work on shifting contexts. Lewis argues that the boundaries of permissible conduct in a language game may shift in accordance with shifting presuppositions. In the case of modal claims, participants in a conversation may tacitly agree that certain possibilities shall be ignored. Certain statements may be true, so long as those possibilities go unspoken, but once those possibilities are mentioned, the context changes. A statement that was true in the earlier context may not be true in the new one. Lewis provides the following example:

Suppose I am talking with some elected official about the ways he might deal with an embarrassment. So far, we have been ignoring those possibilities that would be political suicide for him. He says: "You see, I must either destroy the evidence, or claim that I did it to stop Communism. What else can I do?" I rudely reply: "There is one other possibility — you can put the public interest first for once!" That would be false if the boundary between relevant and ignored possibilities remained stationary. But it is not false in its context, for hitherto ignored possibilities come into consideration and make it true. And the boundary, once shifted outward, stays shifted. If he protests, "I can't do that", he is mistaken. (Lewis 1979, pp. 354–5.)

Thus in Lewis' example, the politician's claim, 'I can't do that', would have been false only in the second context, not in the first.

Similarly, if we suppose that Tim someday admits the possibility of his wife's guilt, this does not commit us to saying that his earlier statement, which denied that possibility, is false. The truth-value of the statement depends upon the sort of context in which it was uttered. His earlier statement was true, despite his eventual admission, if it was said in a belief-based context. To determine whether the context was belief-based, we can apply a retraction test. We can ask whether Tim would view his second statement as an admission that his first statement was false. I doubt he would. Rather than saying that he was wrong about what was possible, he is much more likely to speak in terms of what was 'true for him'; he is likely to say, 'That is what was *true for me* at the time'. Much as philosophers might like to edit this locution out of existence, we might do better to ask what people use it to mean. I think they use it to express how things seemed, and how things seem to a person is a function of belief. For Tim, it was his beliefs that defined his experience, and so defined the context. In that belief-based context, his wife's guilt was not a possibility for him. Only later, when he begins to

acknowledge the evidence he had hitherto ignored, will he shift to a knowledge-based context.

There is more to be said about how contexts are determined. The example just considered arose from motivationally biased beliefs, but I am not claiming that such beliefs will ensure that a belief-based linguistic context. To see why, we may consider the thoughts of another participant in the above-mentioned tragedy, Dr. Al Steinschneider, whose study of SIDS had included two of the Hoyt babies. Dr. Steinschneider built his career on the apnea theory he derived from the study, and also had significant financial incentives for persisting in his theoretical beliefs. Thus it is plausible, though not proven, to suppose that by the time he was called as a witness at Waneta's murder trial, many of his beliefs were motivationally biased. The prosecution observed that some of the data that Dr. Steinschneider reported in his publications were at odds with the nurses' observations, and asked whether it was possible that he had fabricated any of the data. Steinschneider replied that it was not. His assertion amounts to this: 'It's not possible that I fabricated the data.' Despite the similarities to the earlier case of Tim, I think this statement expresses epistemic rather than doxastic possibility. Even if we suppose that Steinschneider simply believed what he wanted to believe — just as Tim did earlier — the contexts are crucially different. It is not simply the speaker's interests and purposes that determine the context, but the presence or absence of other interests as well. In Tim's case, what he believed had no great effect upon others, so his own purpose was sufficient to determine the context. Since his interests lay in preserving his beliefs, that was sufficient to generate a belief-based context. An identical purpose on Steinschneider's part, however, would not produce the same result because his interests were not the only ones involved. From the beginning, Steinschneider had, in Lewis' terms, 'shifted the boundary outward'. He had done so by embarking upon scientific research, subjecting his claims to public scrutiny, and ensuring that a great many people had a stake in what he believed. Since his position as a scientific figure, even a discredited one, ensured that his claims would continue to be of public interest, his claims belong to a knowledge-based context.

IV

My concluding remarks shall be brief. The Moorean theorists were right about some epistemic possibility statements, but there are other statements to which their assumptions do not apply. There are individual uses of the phrase 'It's
Principia, 10(1) (2006), pp. 1–20.

possible that p ', and those uses are not degenerative community statements, as DeRose suggested. Only if we recognize them as a distinct class can we explain how a speaker's non-equivalent possibility statements can all be true, even though her epistemic position remains unchanged. There are also uses of the phrase 'It's possible that p ' that express what is possible relative to belief rather than knowledge. These statements express doxastic possibility because they occur in belief-based contexts. It is the speaker's intentions that determine whether a context is knowledge-based or belief-based. This does not mean, however, that the speaker determines the context by fiat. As our final case demonstrated, having belief-preservation as one's current goal is not sufficient to ensure a belief-based context. The context will shift to a knowledge-based one if the speaker's original intentions invited the scrutiny of the public or, we might say, of a community.¹⁰

References

- Bealer, G. 2002. "Modal Epistemology and the Rationalist Renaissance." In T. Gendler & J. Hawthorne (eds.), *Conceivability and Possibility*. Oxford: Clarendon Press, pp. 71–125.
- DeRose, K. 1991. "Epistemic Possibilities." *The Philosophical Review* 100: 581–605.
- Firstman, R. & Talan, J. 1997. *The Death of Innocents: A True Story of Murder, Medicine, and High Stakes Science*. Bantam Books.
- Hacking, I. 1967. "Possibility." *The Philosophical Review* 76: 143–68.
- Hintikka, J. 1962. *Knowledge and Belief: an Introduction to the Logic of the Two Notions*. Ithaca: Cornell University Press.
- Lewis, D. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8: 339–59.
- Moore, G. E. 1962. *Commonplace Book: 1919–1953*. Casimir Lewy (ed). London: George Allen & Unwin Ltd.
- Morris, T. 1986. *The Logic of God Incarnate*. Ithaca: Cornell University Press.
- Teller, P. 1972. Epistemic Possibility, *Philosophia* 2: 303–20.
- Yablo, S. 1993. "Is Conceivability a Guide to Possibility?" *Philosophy and Phenomenological Research* 53(1): 1–42.

Keywords

Belief context, doxastic possibility, epistemic possibility, DeRose, Hacking, Hintikka, Lewis, Moore, Teller.

Hylarie Kochiras
Department of Philosophy
University of North Carolina at Chapel Hill
Chapel Hill, North Carolina, 27599-3125
USA
kochiras@email.unc.edu

Resumo

Embora a possibilidade epistêmica apareça em vários debates, tais debates têm tido relativamente pouco contato entre si. G. E. Moore concentrou-se diretamente na análise de usos epistêmicos da expressão ‘É possível que p’, e nisso ele fez duas suposições fundamentais. Primeiro, pressupôs que os enunciados de possibilidade epistêmica sempre expressam a posição epistêmica de uma comunidade, em vez da posição de um falante individual. Segundo, pressupôs que todos os usos epistêmicos de ‘É possível que p’ sejam analisáveis em termos de conhecimento, não de crença. Alguns autores mais recentes, inclusive Keith DeRose, apresentam explicações alternativas da possibilidade epistêmica, ao mesmo tempo em que conservam as duas suposições de Moore. Nenhuma dessas pressuposições foi explicitamente contestada, mas a análise de Jaakko Hintikka fornece uma base para tal. Baseando-me na análise de Hintikka, argumento que alguns enunciados de possibilidade epistêmica expressam somente o estado epistêmico individual do falante, e que, contra DeRose, não são enunciados comunitários degenerados mas uma classe existente por si mesma. Afirmando ainda que alguns contextos lingüísticos são antes baseados em crença do que em conhecimento e, em tais contextos, o que é possível para um falante não depende do que ele sabe, mas do que ele acredita.

Palavras-chave

Contexto de crença, possibilidade doxástica, possibilidade epistêmica, DeRose, Hacking, Hintikka, Lewis, Moore, Teller.

Notes

¹ See Morris 1986, p. 148: “Jesus could be tempted to sin in case it was epistemically possible for him that he sin. If at the times of his reported temptations, the full accessible belief-set of his earthly mind did not rule out the possibility of his sinning, he could be

genuinely tempted.” It should be noted, however, that Morris’ discussion of belief-sets is spare, and could therefore be simply a terminological convenience.

² Moore 1962, p. 279. Moore does not use the term ‘epistemic possibility’, but speaks instead of the senses of ‘possible’ which are “different from that sense of ‘possible’ in which ‘possible’ = not self contradictory.” While this remark alone may not be determinative of the sort of possibility at issue, it is clear enough from the context that he means epistemic possibility.

³ As Moore actually states the condition for each sense, it requires that no one in the group know either $\sim p$, or anything from which $\sim p$ follows. I have omitted the clause about inferences, since it is not central to the discussion at hand, and since Moore indicates that he is not committed to it. See his *Commonplace Book*, p. 279.

⁴ I have adapted both of these examples to suit my purposes. Moore actually uses them to show the conditions under which ‘It’s not certain’ is false. Since he uses this phrase to define ‘possible,’ I doubt that he would quibble with my adaptation. See his *Commonplace Book*, p. 278.

⁵ Although Hacking does not expressly include Moore’s two senses of epistemic possibility, we may assume that he accepts the distinction, since he states his disagreement with Moore to be over the matter of practicable investigations.

⁶ This is not to say that the mate’s statement does not permit a reading in which ‘possible’ is interpreted as physical possibility. However Hacking clearly intends the statement to involve an epistemic use of ‘possible’.

⁷ Hintikka 1962, p. 3. This statement seems intended as a straightforward expression of epistemic possibility, not as an iterated possibility statement. “It’s possible, for all I know” is one of our idioms for simple epistemic possibility, and absent any contextual reason to think the speaker intends a more complex statement, it should be taken as such.

⁸ In this section, I am modifying actual statements of the participants, or attributing statements to them that are consistent with their actual statements, as quoted in *The Death of Innocents: A True Story of Murder, Medicine, and High Stakes Science*, by Richard Firstman and Jamie Talan (Bantam Books, 1997). See pp. 439 and 526.

⁹ Thomas Morris’s suggestion, mentioned in the introductory section of this paper, is similar. He writes: “Very roughly, but sufficient for our purposes here, we can say that some proposition P is epistemically possible for some subject S at a time t just in case it is epistemically possible relative to a full accessible belief-set B of S at t , where that relation consists in something like the following: B neither contains nor self-evidently entails the denial of P , nor does B contain or self-evidently entail propositions which seem to S to show P to be either false or impossible.” (Morris 1986, p.148.)

¹⁰ **Acknowledgements** For comments and discussion, I would like to express my gratitude to an anonymous referee, Dorit Bar-On, Keith DeRose, Douglas Long, William Lycan, Ram Neta, and especially, Keith Simmons.