



Belief Worlds and Epistemic Possibilities

Hylarie Kochiras
University of North Carolina-Chapel Hill

ABSTRACT: This paper develops an individualistic, belief-based account for a limited class of epistemic possibility statements. Section I establishes the need for such an account by reviewing a recent version of the majority view (the "Relevant Community Account") and contesting two key assumptions. I argue that some epistemic possibilities are belief-based-*contra* the assumption that all are knowledge-based. Against the assumption that all epistemic possibility statements are analyzable in terms of the speaker's "relevant community," I contend that the truth value of some statements is a function of the speaker's epistemic states alone. Section II develops an alternative account designed to capture those internal, individual statements. Modeling belief sets as "belief worlds," I explain our epistemic processes in terms of an ability to shift attention among our various belief worlds.

Ever since G. E. Moore set out to explain the phrase "it's not certain that" and ended up offering a definition of epistemic possibility, ordinary uses of epistemic modal sentences have commanded more attention. A number of more recent accounts-which I'll call the "Moorean accounts"-follow Moore by analyzing all statements having the form "It's possible that p" uniformly. Their analysis is uniform in that they make two assumptions: first that some community is relevant in every case, and second that every statement is analyzable in terms of knowledge.

I think that both assumptions are false, so one thing I'll do today is to suggest truth conditions for statements failing of both assumptions-"individual, doxastic possibility statements." I will have time to contest only the first assumption, though, and will therefore help myself to my conclusion regarding the second; I will assume that the locution "It's possible that p" is not restricted to expressions of epistemic, that is, knowledge-based possibility, but may be used to express doxastic or belief-based possibility as well.

The first assumption-that some community is relevant to every use of "It's possible that p"-will be my main concern, then. To make the assumption in its stronger form, as Ian Hacking and Paul Teller do, is to treat all uses of "It's possible that p" as statements to be translated as "For all *we* know, p." But the locution "It's possible that p" is used not only for statements of that sort-for, in my terminology, "community statements"-but also for what I'll call "individual statements"-statements properly translated as "For all *I* know, p."

The assumption that some community is relevant to every use of "It's possible that p" is also made by Keith DeRose, but in a weaker form. DeRose does allow for individual statements, however he treats them as degenerate community statements. It is because he

doesn't treat individual statements as a class in their own right, I think, that he overlooks the purely individual case of the following problem he raises: a constant epistemic situation may generate several true but non-equivalent epistemic possibility statements. In what follows, I'll show how this problem about non-equivalent statements arises even when all statements involved are individual, and I'll then develop the 'belief worlds account' to explain the individualist case.

I

By modifying a predecessor's account, DeRose arrives at the following truth conditions for statements taking the form "It's possible that p":

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. (2)

These truth conditions are designed to resolve the problem about non-equivalent statements. DeRose raises the problem with his 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. (3)

At the first round of tests, then, only a negative result would be conclusive. The first case DeRose constructs from these facts is Cancer Test Case 1A:

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 . . . 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 . . . until tomorrow." (4)

Jane's statement,

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

is a fairly simple one.

The puzzle arises with Cancer Test Case 1B (CTC-1B). Here, 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." (5)

It is interesting that Jane now claims ignorance with (b), since this quite reasonable statement seems to conflict with statement (a) of CTC-1A. *Yet she is epistemically no worse off now than before.* So how could a single epistemic state be productive of non-equivalent, even seemingly contradictory statements?

To explain this oddity, DeRose invokes his flexibility requirement. He suggests that although Jane knows no more in the one case than in the other, the cases still differ in that the two statements are made relative to different communities. In CTC-1A, Jane is expressing her own epistemic position relative to John's family, a community no more knowledgeable than herself. Relative to this equally ignorant community, she makes the stronger statement: "It's possible that John has cancer." In the second case, however, Jane expresses her epistemic position relative to a community whose sum knowledge exceeds her own because it includes the doctor. Relative to that community, Jane's epistemic position is comparatively weak. She consequently makes a weaker claim, saying she doesn't know whether it's possible that John has cancer. Jane's two statements, though they don't mean the same thing, are both true because they are uttered relative to different communities.

II

I think we can accept DeRose's explanation for a good many pairs of true, non-equivalent statements. However, since the explanation turns on a shift in the relevant community, it can apply to only those cases in which at least one of the statements is a community statement. There are also cases in which there is no community to shift, namely cases in which both statements express only the individual's epistemic (or doxastic) position. To see that non-equivalent statements may arise from a constant epistemic position even when the speaker is expressing only her individual position, consider the following (now outdated) case.

In this case, Anne muses, "It's possible that Benazir Bhutto is in power in Pakistan." Since Anne may well think she is the only person ignorant about Pakistan's leadership, it would be odd to read her as trying to express the epistemic position of any community. Her statement is clearly individual. It is also, I think, an expression of doxastic rather than epistemic possibility. She isn't being so circumspect as to consider what she actually knows for certain, and what's compatible with that; she's just speaking from the default perspective of expressing her beliefs, and what is consistent with those beliefs. I suggest the following truth conditions for individual doxastic possibility statements such as Anne's:

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

These truth conditions draw upon Jaakko Hintikka's discussion of consistency, in which he suggests that the epistemic possibility of a proposition turns on its being consistent with a set of known propositions. Since my concern is doxastic possibility, however, my truth conditions appeal to a set of beliefs rather than a set of known propositions. And to keep matters fully internal, my truth conditions require only that the subject perceive the proposition to be consistent with her belief set, not that it actually be consistent.

I now want to resolve the individual case-specifically, the doxastic subcase of the individual case-of the non-equivalency problem. So, we need Anne to make a second, true statement which isn't equivalent to her earlier statement. Let's call the earlier, simple case, in which Anne expresses her uncertainty but has no expectation of finding out about Bhutto, "Bhutto Case A." We now construct a second case, call it "Bhutto Case B," by supposing that Anne now plans to attend a lecture tomorrow that will answer her questions about Pakistan. With the lecture in mind she now utters a second individual statement, "It's possible that I will come to believe that Bhutto's in power." We now have two non-equivalent, individual doxastic possibility statements, uttered by a speaker who has neither learned nor forgotten anything. Since neither statement is uttered relative to any community, we need an explanation of how both could be true. To explain this case, we may model Anne's sets of beliefs as 'belief worlds'.

To develop the belief worlds model, recall that Anne was thinking "It's possible that Benazir Bhutto is in power in Pakistan." Let "p" stand for the embedded proposition, "Benazir Bhutto is in power in Pakistan." As Anne wonders about p, the following beliefs come to mind:

q: Bhutto was forced out two years after coming to power.

r: Bhutto returned to power a few years ago.

s: Pakistan is politically unstable.

The set could tolerate the addition of either p or \sim p without producing a contradiction, allowing Anne to say truly, "It's possible that Bhutto is in power." This latter statement expresses Anne's second-order belief that her relation to p is a state of uncertainty; aware that she believes neither p nor \sim p, she takes her own uncertainty as an object of belief.

By taking Anne's second-order belief, Mp, together with the belief set comprising all the beliefs that led her to believe Mp, we generate Anne's total belief set in relation to the proposition p, at a given time, t0. This belief set is, in precise formulation, Anne's *belief world in relation to p*, or for short, her *belief world*. The total state that leads Anne to say "It's possible that Bhutto's in power," may be represented schematically as the following belief world:

Anne's belief world, represented by the circle, includes both the first-order beliefs which make up her relevant belief set, and the second-order belief which they lead her to adopt, Mp. The dotted lines indicate that either p or \sim p could be added to the set, but not both.

Since the belief world is simply a means of speaking about what really interests us, Anne's relation to p, it will be helpful to be able to pick out whichever proposition of the belief world expresses that relation. I will call any such proposition the *salient proposition* of that world. Here, the salient proposition is "Mp," for as I have described the case, Anne does not anticipate finding out in the near future whether Bhutto is in power. She expects that "It's possible that p" will continue to describe her state regarding p, thus her uncertain state is in equilibrium, and the case is a fairly simple one.

Now recall that we produced Bhutto Case B by complicating Anne's situation. This time, remember, Anne plans (at t0) to attend a lecture tomorrow (t1), and she expects the lecturer to tell her who Pakistan's prime minister is. Since at t0 Anne is uncertain about whether p is true, she may again say, "It's possible that Bhutto's in power." But now Anne anticipates that her present state of uncertainty will be replaced tomorrow by a belief about Bhutto. This added element, anticipation, is simply one or more second-order beliefs that her present state will change.

The complication in Anne's situation calls for a refinement of our notion of a belief world. Anne's present state is her *actual belief world* (ABW). In anticipating that tomorrow's lecture will lead her to believe either p or \sim p, Anne frames two *possible belief worlds* (PBW's)-worlds which she thinks are candidates for becoming her actual belief world. Because Anne projects these potential belief sets in the present moment, they are represented diagrammatically as branching from her ABW. In the case at hand, the possible belief worlds which Anne projects occupy separate branches, since as the two potential fulfillments of a current state of uncertainty, they contain mutually incompatible propositions.

Now should Anne in fact attend the lecture at t1, and hear that Bhutto is power, then what was at t0 PBW1 would be *actualized* for her, becoming her ABW. To be clear, actualization

occurs not when the PBW does come to describe the actual world, but when the subject comes to believe that it describes the actual world.

The introduction of branching, possible belief worlds allows us to represent the various factors that complicate Anne's relation to p : she believes that her relation to p may soon change, and she envisions specific potential outcomes, namely a belief that $\sim p$, and a belief that p . At this point, it is necessary to revise the earlier definition of a salient proposition, so as to incorporate potential as well as actual epistemic positions regarding p :

For any belief world, whether actual or possible, its salient proposition is that proposition which either does express the subject's actual position regarding the proposition in question, p , or which would express her position regarding p , were that world to become her actual belief world.

The salient proposition of Anne's ABW is Mp , just as it was in Case A, but her ABW is now characterized by a projection of two possible belief worlds. These worlds are projected for t_1 , and have p and $\sim p$ as their salient propositions. So, p is the salient proposition of PBW1, because were PBW1 to be actualized, Anne's stance toward p would just be the belief that p .

The psychological complication introduced by Bhutto Case B is that some of our subject's states are self-reflective, taking other epistemic states (here 'epistemic' should be understood broadly) as their objects. Schematically, we use branching worlds to represent self-reflective states, so that from Anne's overall, complex doxastic situation, we may distinguish various belief subsets, some representing outward-directed states, and some representing self-reflective states. The linguistic expression of these self-reflective states will require statements which are more complex than, and not equivalent to, the one Anne made in Bhutto Case A, "It's possible that Bhutto's in power."

We can now explain how Anne, who is fully as uncertain in Case B as she is in Case A, can be justified in making the additional, non-equivalent possibility statement she utters in Case B. A subject's justifiable assertion of non-equivalent statements derives from the ability to shift her attention among her various belief worlds. While it is sometimes convenient to speak of Anne's total state, we can distinguish within it many constitutive elements. We can distinguish first and second-order beliefs, but to explain the various possibility statements which Anne makes, we must draw our distinctions as Anne would draw them. When Anne is thinking about p itself, the salient component of her belief state is the proposition that currently describes her relation to p . When she self-reflectively thinks about her own relation to p , the propositions describing potential future relations to p are salient. Looking at her total doxastic situation, then, its salient components are all the salient propositions of her various belief worlds-and it is because she can focus on one or another of her belief worlds that she can make a number of possibility statements besides Mp .

Take Bhutto Case B. Sometimes Anne focuses her attention on just that part of her state which is held in common with Case A, namely the upper tier world of the belief world tree, her ABW. Her thoughts are now directed outward, in that she concentrates on the proposition p . She expresses her uncertainty about p by saying just what she said in Case A, "It's possible that Bhutto's in power." This is hardly surprising, given that she has heard no new information. Anne can shift her attention, however, away from the proposition p itself, and instead self-reflectively consider her potential belief states. Suppose she is thinking that it is possible for her now that she will come to believe p tomorrow. As she looks ahead to believing that p , the object of her thoughts is not p itself, but the potential state of believing that p . This is why, in addition to what she said in Case A, Anne may now also say "It's possible that I will come to believe that p ."

Employing the belief worlds metaphor, we can say that Anne is focusing her attention on the belief world whose salient proposition is *p*, namely PBW1. Her belief that she may come to think that *p* may then be described as the belief that PBW1 may become her ABW. It is because she is looking ahead to PBW1 that she says, "It's possible that I will come to believe that *p*." Notice that in looking ahead, Anne still speaks of possibility; even though she is thinking about herself believing that *p*, she is not fully imagining herself believing *p*, but is rather aware of PBW1 *qua* possible belief world. There are additional statements Anne could make, and she is led to make one rather than another depending upon which component of her total situation-which belief world-she happens to be thinking of at the time. It seems, then, that some uses of "It's possible that *p*" should be understood as relative to the various subsets of the speaker's beliefs, and not to different communities or ways of coming to know.