

**“A Puzzle About Doubt,” by Gary Ebbs, published by MIT Press, appears here by permission of the author and publisher.**

**This material is protected by copyright. All rights reserved. Please contact the publisher for permission to copy, distribute or reprint.**

4. I have had it suggested to me that the on/off light found on some measuring instruments has the function of indicating that the instrument is on—thus carrying the information that the instrument is in a representational mode. When the instrument registers “0” with the light off, it does not represent *Q* to be 0. When it registers “0” with the light on, it represents (possibly misrepresenting) *Q* to be 0. So the light carries the information that the instrument is representing (possibly misrepresenting) *Q*. I agree that the light indicates—at least it has the function of indicating—when the instrument is turned on, but I don’t agree that it indicates that the instrument is representing *Q*. Swampmeter also has a light that goes on when a switch is closed, but swampmeter *never* represents anything. The light doesn’t tell us—it has no way of telling us—that it is installed in a genuine *Q* meter rather than, say, a swampmeter.

5. Though I confess to being uncertain in this area. See Dretske, forthcoming.

6. See Evans 1982: 205–233 and Shoemaker 1988 for (what I regard as) attractive accounts of how we might learn to self-ascribe thoughts and experiences to ourselves without being aware (in a perceptual way) of the thoughts and experiences themselves.

## A Puzzle about Doubt

Gary Ebbs

### 1 What Can an Anti-individualist Know A Priori?

My central goal in this paper is to identify and dissolve a puzzle that lies behind a vexing debate about what an anti-individualist can know a priori. In this opening section, I will review common assumptions and contested points of the debate, and briefly explain my misgivings about one of the common assumptions. In later sections I will identify and dissolve the puzzle that lies behind the debate.

Anti-individualism is the view that what a person believes and thinks is not settled by his linguistic dispositions, internal physical states, or phenomenal experiences, described independently of his social and physical environment. One central question about anti-individualism is whether it is compatible with *minimal self-knowledge*—the familiar fact that (in a sense yet to be clarified) we each know without empirical investigation what thoughts our own utterances express. If what we know without empirical investigation is what we know a priori, then the question of whether anti-individualism is compatible with minimal self-knowledge is linked to the question of what an anti-individualist can know a priori.

This question has been much discussed recently in the literature about anti-individualism. Although several answers to the question have been proposed, a single debate now dominates the discussion. On one side are those who argue that an anti-individualist who assumes she has minimal self-knowledge is committed to the unacceptable conclusion that she has a priori knowledge of some truths that in fact she *cannot* know a priori.<sup>1</sup> On the other side are those who argue that anti-individualists are not committed to this unacceptable conclusion.<sup>2</sup> Philosophers on both sides of the debate assume that minimal self-knowledge is second-order, in the sense that statements that express

such knowledge have the logical form 'I am thinking that  $p$ ', where ' $p$ ' is replaced by a declarative sentence.<sup>3</sup> They also assume that both reason and introspection can be sources of a priori knowledge.<sup>4</sup>

Much of literature about the debate has focused on an argument by Michael McKinsey (1991a: 9). The argument presupposes that a priori knowledge is "knowledge obtained independently of empirical investigation." Suppose that I utter the sentence 'Water is a liquid at room temperature', thereby expressing the thought that water is a liquid at room temperature. Suppose also that I have minimal self-knowledge, so I know without empirical investigation that I am thinking that water is a liquid at room temperature. By assumption, a priori knowledge is knowledge obtained without empirical investigation, so I know a priori that (1):

(1) I am thinking that water is a liquid at room temperature.

If anti-individualism is correct and I accept it, McKinsey assumes, I am in a position to know by reasoning alone, hence a priori, a *conceptual* truth of the form 'If I am thinking that water is a liquid at room temperature, then  $E$ ', where  $E$  is a statement that most philosophers would say I cannot know a priori.<sup>5</sup> I am in a position to know a priori, for instance, that (2):

(2) If I am thinking that water is a liquid at room temperature, then either I or members of my linguistic community have seen or touched water.

And if I can know a priori that (1) and (2), then by *modus ponens* I can deduce, and thereby know a priori, that (3):

(3) Either I or members of my linguistic community have seen or touched water.

But it seems that no one can know a priori that either he or members of his linguistic community have seen or touched water. Hence anti-individualism apparently implies that a person can have a priori knowledge of some statements that in fact he cannot know a priori. I will call this McKinsey's argument.

The orthodox reply to McKinsey's argument (implicit in Burge 1982a, explicit in Brueckner 1992a and McLaughlin and Tye 1998) is that despite appearances to the contrary, an anti-individualist who assumes he has minimal self-knowledge has no reason to think he can know a priori *any* statement that most philosophers would say he cannot know

a priori. According to this reply, to know what one is thinking when one utters a particular sentence, one need not know or presuppose any empirical statements. Premise (2) of McKinsey's argument may be true, but even if it is true, it is not a *conceptual* truth, as McKinsey assumes, and so we cannot know it a priori.

The initial plausibility of this reply masks a deep problem with both sides of the debate. The problem can be traced back to a widely accepted but unexamined assumption about how an anti-individualist should analyze epistemic possibility. Twin Earth thought experiments suggest that for each person we can describe *subjectively equivalent worlds* in which everything that is relevant to the person's subjective assessment of her situation seems the same to her as it does in the actual world, but her social or physical environments are different from her social or physical environment in the actual world. According to the standard analysis of epistemic possibility, no one can know by reasoning or introspecting—without empirical investigation—which of her subjectively equivalent worlds she is actually in.<sup>6</sup>

The problem, I will argue, is that this standard analysis of epistemic possibility conflicts with the truism that *to express a thought, one must have some idea of what that thought is*. To defend and clarify the truism, I will argue that contrary to the standard analysis of epistemic possibility, if we accept anti-individualism, there are some apparently empirical statements that we cannot make sense of doubting.

## 2 Anti-individualism and Self-Knowledge

To see the conflict between the truism and the standard analysis of epistemic possibility, it helps to reflect first on the methodology behind Putnam's (1975) Twin Earth thought experiments, which have persuaded many to accept anti-individualism. In my view, these thought experiments are persuasive because they are based in our practice of taking fellow English speakers' words at face value.<sup>7</sup>

Recall the thought experiment involving Oscar, an ordinary English speaker who is competent in the use of the English word 'water' but does not accept (or reject) the sentence 'Water is H<sub>2</sub>O'. Suppose that Oscar utters the sentence 'Water is a liquid at room temperature'.<sup>8</sup> Since Oscar is a competent English speaker, other English speakers take his words at face value—they take him to have said that water is a liquid at room temperature. If they think his utterance is sincere, they also take him to believe this.

Now suppose there is a planet called Twin Earth, which is just like Earth except that wherever there is water on Earth, there is twin water, a liquid with an underlying chemical structure very different from the chemical structure of water, on Twin Earth. On Twin Earth there lives a physical, phenomenological, and behavioral twin of Oscar, Twin Oscar, who is a normal speaker of Twin English, the Twin Earth counterpart of English. When Twin Oscar utters the sentence 'Water is a liquid at room temperature', his fellow Twin English speakers take his words at face value—they take him to have said (translated into English) that twin water is a liquid at room temperature. If they think his utterance is sincere, they also take him to believe this.

Together with our trust in our practice of taking other speakers' words at face value, these observations show that what a person believes and thinks is not settled solely by his linguistic dispositions, internal physical states, or phenomenal experiences, described independently of his social and physical environment. This negative thesis is what I call anti-individualism.<sup>9</sup>

Putnam (1975) also argued that even in 1750, *before* scientists on Earth and Twin Earth discovered the chemical properties of water and twin water, respectively, a competent English speaker who uttered the sentence 'Water is a liquid at room temperature' thereby expressed the thought that water is a liquid at room temperature, while his twin on Twin Earth expressed the thought (translated into English) that twin water is a liquid at room temperature. In my view, it is because we endorse our practice of taking each other's words at face value across time, from moment to moment, and even for centuries, that we can see that in both linguistic communities, looking backward from today, the referents of the words for water and twin water did not change when the chemical properties of the liquids to which they apply were discovered.<sup>10</sup>

If we endorse our practice of taking each other's words at face value, we can also see that minimal competence in the use of a word requires more than simply writing or uttering sentences in which the word occurs. We wouldn't take a speaker to be using the English word 'apple' competently if she applies it only to points of light visible in the night sky. Yet a child who at first refuses to call a green apple an 'apple' might still be taken to be able to use the word 'apple' to express thoughts about apples, and to believe of the green apple that it is not an apple, provided that she has some other beliefs about apples, including some true beliefs that she expresses by using the sentence 'That's an apple'.<sup>11</sup> Our firmest grip on the requirements for minimal competence is our

practice of taking each other's words at face value in a given context, unless we see some concrete reason in that context for not doing so.<sup>12</sup>

These observations about minimal competence are intimately linked with our judgments about when a speaker has minimal self-knowledge.<sup>13</sup> To credit a speaker of a given natural language with minimal self-knowledge is to take her to be able to use words of her own language to express thoughts, make claims, raise questions, and so on. Any situation in which we are willing to take another's words at face value is thereby also one in which we will credit her with having minimal self-knowledge.<sup>14</sup> Viewed in this way, minimal self-knowledge is a practical aspect of ordinary competence in the use of language, not a kind of second-order propositional knowledge, as many philosophers assume. Unlike second-order propositional knowledge of what one is thinking, minimal self-knowledge is as widespread as the everyday use of language to express thoughts, evaluate beliefs, raise questions, and so on.<sup>15</sup> In taking other speakers' words at face value, we thereby also take them to know what they are talking about in a minimal sense that goes with competence.<sup>16</sup>

These observations clarify the truism that *to express a thought, one must have some idea of what that thought is*. To take someone to express a thought by using a given word is also to take him to have some beliefs that he expresses by using that word. These beliefs may be false or misleading, but not just any utterances of sentences containing a word suffice for minimal competence in the use of the word, as the 'apple' example shows.<sup>17</sup>

### 3 Apriority and Epistemic Possibility

The only assumption that McKinsey explicitly makes about a priori knowledge is that it is "knowledge obtained independently of empirical investigation" (1991a: 9). Most philosophers involved in the debate sketched above simply repeat McKinsey's characterization of a priori knowledge and agree with him about which beliefs the person can know a priori. It is widely agreed, for instance, that no one can know a priori that members of her linguistic community have seen or touched water.<sup>18</sup>

But McKinsey's characterization of a priori knowledge does not explain this agreement. To see why, consider my belief that physicists in my linguistic community have detected quarks. I have read this in authoritative books, but have never undertaken any empirical investigation into

whether it is true. By ordinary standards, I am epistemically *entitled* to believe that physicists in my linguistic community have detected quarks; if this is true, then I know it independently of empirical investigation. Similarly, an unusually sheltered person who is *told* that members of her linguistic community have seen or touched water might be epistemically *entitled* to believe this without undertaking any empirical investigation into whether the person who told her this is trustworthy or whether it is true; if it is true, then by ordinary standards she knows it independently of empirical investigation. It therefore seems that according to McKinsey's characterization of a priori knowledge, she knows a priori that members of her linguistic community have seen or touched water. But it is supposed to be obvious that no one knows a priori that members of her linguistic community have seen or touched water.

One might think that in both of these cases the knowledge gained by testimony is not independent of empirical investigation, because it can be traced back to empirical observations made by others. But minimal self-knowledge cannot be independent of empirical observation in this sense, since it requires minimal competence, which typically depends on accepting testimony from others.<sup>19</sup> For a large number of words, the testimony that we accept when we acquire competence in the use of those words can be traced back through chains of similar testimony to speakers who have made empirical observations that support it. In the context of anti-individualism, then, we cannot assume that a given person's minimal self-knowledge is independent of everyone *else's* empirical observations. The most we can say is that to have minimal self-knowledge is to know what thoughts one's utterances express without going through any empirical investigation of one's own.

For this reason, McKinsey's characterization of a priori knowledge does not explain why so many philosophers agree with him about which statements can be known a priori. What does explain this? The answer, I believe, is that most philosophers presuppose a tempting but misguided analysis of epistemic possibility that looks like an immediate consequence of the Twin Earth thought experiments themselves. As I noted earlier, the Twin Earth thought experiments suggest that for each individual, we can describe *subjectively equivalent worlds* in which her physical and phenomenal states, described independently of her environment, are the same, but her environments are different.<sup>20</sup> Most philosophers assume that no one can distinguish between any of her subjectively equivalent worlds just by reasoning or introspecting—that all of a person's subjectively equivalent worlds are epistemically possible for her.

They find this analysis of epistemic possibility appealing on its own terms, for two reasons that I will soon discuss. In addition, I believe, they find it attractive because it is like the analysis of epistemic possibility that Saul Kripke introduced in *Naming and Necessity* to solve a puzzle about his view of reference and necessity. The puzzle is that in Kripke's view, if Hesperus is identical to Phosphorus, then there is no possible world in which Hesperus is not identical with Phosphorus. Prior to our discovery that Hesperus is identical to Phosphorus, we assumed that 'Hesperus is identical with Phosphorus' may actually be false. Hesperus is in fact identical to Phosphorus, however, and so, by Kripke's theory, Hesperus is *necessarily* identical with Phosphorus: we can't express our prior assumption by saying that it could have turned out that Hesperus is not identical to Phosphorus. Kripke therefore had to provide a new analysis of our previous assumption that 'Hesperus is identical with Phosphorus' may actually be false. He stipulated that "given the evidence that someone has antecedent to his empirical investigation, he can be placed in a sense in exactly the same situation, that is a qualitatively identical epistemic situation, and call two heavenly bodies 'Hesperus' and 'Phosphorus', without their being identical" (1972/1980: 104).

Kripke suggested that since prior to our empirical investigation, we could not discriminate between these worlds on the basis of our evidence, we may actually have been in a world in which 'Hesperus is not identical to Phosphorus' is true. He explicitly connected this analysis with the traditional idea of a priori knowledge: "Two things are true: first, that we do not know a priori that Hesperus is Phosphorus; and are in no position to find out the answer except empirically. Second, that this is so because we could have evidence qualitatively indistinguishable from the evidence we have and determine the reference of the two names by the positions of two planets in the sky, without the planets being the same" (Kripke 1972/1980: 104). I suggest that this characterization of a priori knowledge is what lies behind the agreement that a person cannot know a priori, for instance, that members of her linguistic community have seen or touched water. The idea is that she can't know without empirical investigation that she is not in any one of her subjectively equivalent worlds in which no member of her linguistic community has seen or touched water.

Most philosophers writing about what an anti-individualist can know a priori take this analysis of epistemic possibility for granted, without giving any reasons why we should accept it.<sup>21</sup> Kripke himself does not argue for his analysis of epistemic possibility; he presents it as obvious

and beyond question.<sup>22</sup> To evaluate it, however, we need a better idea of why so many philosophers accept it.

I think that there are two main reasons. The first is that they want to make sense of Descartes' radical skeptical hypotheses about the nature and existence of an "external" world, including the hypothesis that all my experiences, from the beginning to the end of my life, are parts of an elaborate dream, and the hypothesis that I am massively deceived by an evil demon. To find these hypotheses compelling is to picture a vast gulf between how things *seem* and how they *are*. In the context of anti-individualism, a tempting way of picturing this supposed gulf is by holding our subjective experiences constant and specifying different external environments compatible with all those subjective experiences. Many philosophers assume that to entertain the thought that they are in one of these worlds, it is enough for them to *picture* the subjective experiences they would have in these worlds (experiences that are by definition the same as the subjective experiences they have in the actual world) and then add a caption that describes an "external" world that is compatible with the picture. In what follows, I'll say that to combine one's subjective experiences with a caption in this way is to *picture oneself* in a specified subjectively equivalent world. Most philosophers assume that for a person to entertain the thought that she is actually in one of her subjectively equivalent worlds, she need only *picture herself* in it.<sup>23</sup>

A second reason why so many accept Kripke's analysis of epistemic possibility is that it provides a natural interpretation of the traditional view that a priori knowledge is based on reasoning or introspecting, independent of any evidence from the senses, and that neither reasoning nor introspecting, by themselves or in combination, can tell us which possible world we are in.

#### 4 The Puzzle

It follows from the standard analysis of epistemic possibility just described that if I restrict myself to what I can know without empirical investigation, I must accept (4):

- (4) I may actually be in any of my subjectively equivalent worlds.

Yet I assume that without empirical investigation, I know what thoughts my utterances express. For instance, I assume without empirical investigation that I am epistemically entitled to accept (5):

- (5) My utterances of 'Water is a liquid at room temperature' express the thought that water is a liquid at room temperature.

If I am justified in accepting (4) and (5) without empirical investigation, then I am justified in concluding (6) without empirical investigation:

- (6) In all of my subjectively equivalent worlds, my utterances of 'Water is a liquid at room temperature' express the thought that water is a liquid at room temperature.

The trouble is that the normal procedure for conducting thought experiments that support anti-individualism implies the negation of (6):

- (7) In some of my subjectively equivalent worlds, my utterances of 'Water is a liquid at room temperature' do not express the thought that water is a liquid at room temperature.

Thus I seem committed to accepting the conjunction of (6) and (7)—a contradiction.

One strategy for trying to avoid this contradiction is to question whether we can know without empirical investigation all the premises that generate it.<sup>24</sup> With this in mind, I constructed this puzzle so that each of its claims, (4) to (7), appears to be independent of empirical investigation. The key premise, (4), apparently follows from our understanding of the phrase "subjectively equivalent world." And we can't give up (5) without abandoning the presumption that anti-individualism is compatible with minimal self-knowledge. Given (4) and (5), we cannot deny (6). It seems that the only claim that *may* require empirical justification is (7).

To construct a thought experiment that supports (7), we need to understand the possible worlds we are describing well enough to see that in those worlds our utterances of sentences would express different thoughts from the ones we take them to express. (This is just another illustration of the truism that to express a thought one must have some idea of what that thought is.) For instance, to support (7) by constructing an anti-individualistic thought experiment involving my word 'water', I must presuppose that the subjectively equivalent world that I *take* to be different from the one that I am actually in is *in fact* different from it. But if all my subjectively equivalent worlds are epistemically possible for me, I can't know without empirical investigation which of my subjectively equivalent worlds I am in. One might therefore think that I cannot support (7) without empirical investigation.

Let's say that a *substantive* statement for a given person is any statement of hers that according to the standard analysis of epistemic possibility she cannot know without empirical investigation. A substantive statement for a person is true in some of her subjectively equivalent worlds and false in some of her subjectively equivalent worlds.<sup>25</sup> For instance, my statements that water is a liquid at room temperature and that I am not in the subjectively equivalent world in which I was born, raised, and now live on Twin Earth are substantive, because according to the standard analysis of epistemic possibility, I cannot know these statements without empirical investigation.

Let's also say that if a person affirms a substantive statement that *p*, then she holds a *substantive belief* that *p*, and that if she *suspends* this belief, then she does not affirm or deny that *p*. Then the reasoning presented two paragraphs above presupposes that I can have minimal self-knowledge of what thoughts I express by using a given group of terms even if I suspend all the substantive beliefs I express by using those terms. On this view, I know what thoughts I express by using my sentence 'Water is a liquid at room temperature', for instance, even if I suspend any substantive beliefs that I could express by using the terms 'water', 'liquid', or 'temperature'.

The trouble is that for most terms I use, I can't have minimal self-knowledge of what thoughts I express by using those terms if I suspend all the substantive beliefs I express by using those terms. To have minimal self-knowledge, one must be able to use one's own words to make claims, raise questions, express thoughts, and so forth. Moreover, as I've already noted several times, it is a truism that to make claims, raise questions, and express thoughts, one must have some idea of what those claims, questions, or thoughts are. For most terms of English, including such terms as 'water', 'liquid', or 'temperature', for instance, a person who suspends all substantive beliefs that she would express by using one of those terms is incompetent in its use and does not count as expressing any thoughts by using it.

One can appreciate this aspect of the puzzle without endorsing my view that anti-individualism implies that to make claims, raise questions, and express thoughts, one must have some idea of what those claims, questions, or thoughts are. It is enough simply to *assume* that we have minimal self-knowledge, and to accept the truism. But the account of anti-individualism that I sketched earlier deepens and consolidates this aspect of the puzzle, by explaining why anti-individualism requires that we have minimal self-knowledge, and why we can't have minimal self-knowledge unless we have some idea of what thoughts our utterances

express. The key point is that in a large number of ordinary cases, to take a person to have expressed a particular thought by uttering a given sentence is to take her to be minimally competent in the use of the terms that make up the sentence. This requires that she have some substantive beliefs—affirm some substantive statements—that she expresses by *using* those terms. As I argued above, a speaker is incompetent in the use of a word—whether it be a widely shared word of a public language, or a word that only a few idiosyncratic speakers share—if she refuses to affirm *any* substantive statements in which it occurs.

The puzzle, then, is this. It seems that without empirical investigation, we are each epistemically entitled to accept (4), because it follows from the standard analysis of epistemic possibility; (5), because we have minimal self-knowledge; and (7), because it follows from the Twin Earth thought experiments. But (4) and (5) together entail (6), which is the negation of (7). If we reject (7) for this reason, we must suspend all our substantive beliefs, so we can't be credited with having minimal self-knowledge, and hence we must reject (5). Yet (5), an instance of minimal self-knowledge, is in fact a consequence of anti-individualism, as I argued above.<sup>26</sup> We are therefore apparently committed to each of (4) through (7), including the contradictory pair (6) and (7).

### 5 My Strategy for Dissolving the Puzzle

The weakest premise of the puzzle is the one that almost everyone accepts without reflection—premise (4). I will argue that in the same sense of 'know' in which we know without empirical investigation what thoughts our utterances express, we can know without empirical investigation that (4) is false. It is a formidable task to make this seem plausible, however, given the popularity of the standard analysis of epistemic possibility, which is sustained by two almost irresistible assumptions: first, that we can make sense of radical Cartesian doubts, and second, that we can't know by reasoning or introspecting which of our subjectively equivalent worlds we are actually in. It is no answer to these deeply entrenched assumptions simply to *assert* that we can know without empirical investigation that (4) is false.<sup>27</sup>

To challenge the standard reasons for accepting (4), we must distinguish between two questions:

(Q1) Given what I know without empirical investigation, is it epistemically possible for me that I am actually in any one of my subjectively equivalent worlds?

(Q2) Are there worlds  $w_1, \dots, w_n$  such that (a)  $w_1, \dots, w_n$  are among my subjectively equivalent worlds and (b) I can know by reasoning or introspecting, without relying on any of my substantive beliefs, that I am not in any of  $w_1, \dots, w_n$ ?

These questions are not explicitly addressed in the literature about what an anti-individualist can know a priori, but I speculate that philosophers who are attracted to (4) would reason roughly as follows:

"The answer to (Q2) is 'No,' because I can't discriminate between my subjectively equivalent worlds by introspecting, and reasoning can only be a source of knowledge about what is in some sense necessary, but all my subjectively equivalent worlds are possible. Hence I cannot know by reasoning or introspecting, without relying on any of my substantive beliefs, which of my subjectively equivalent worlds I am in. Therefore, for all I know without empirical investigation, I may actually be in any one of my subjectively equivalent worlds. Hence the answer to (Q1) is 'Yes.'"

Against this, I will argue that to solve the puzzle we must see that the answer to both (Q1) and (Q2) is "No." Since I agree with the standard assumption that the answer to (Q2) is "No," I must show why, despite this answer to (Q2), the answer to (Q1) is "No," and, as a consequence, we can know without empirical investigation that (4) is false.

The heart of my argument is that even if the answer to (Q2) is "No" and I restrict myself to what I know without empirical investigation, not all of my subjectively equivalent worlds are epistemically possible for me. I start by assuming that it is epistemically possible that  $p$  for a given person only if she can make sense of its actually being the case that  $p$ . Most philosophers assume that all of a person's subjectively equivalent worlds are epistemically possible for her *because* they assume that each person can make sense of actually being in any one of her subjectively equivalent worlds. I will argue that this is an illusion sustained by the mistaken assumption that for a person to make sense of actually being in any one of her subjectively equivalent worlds, it is enough for her to *picture herself* existing in that world.

## 6 Epistemic Possibility and Doubt

To get clear about what is epistemically possible, we must make a number of distinctions and clarifications. First, at any given time  $t$ , sub-

ject  $A$  will have a number of beliefs and a range of observational evidence; what is epistemically possible for  $A$  (at  $t$ ) depends on  $A$ 's beliefs and observational evidence (at  $t$ ). Second, to say that  $p$  is epistemically possible for  $A$  (at a given time  $t$ ) is to say that  $A$  can make sense of its actually being the case that  $p$  (at  $t$ ). Third,  $A$  can make sense of its actually being the case that  $p$  (at  $t$ ) only if  $A$  can *express* the possibility that  $p$  (at  $t$ ).

Fourth, human fallibility tells us nothing about what is epistemically possible for a person at a given time. To see why, suppose that Alice has just constructed what she regards as a proof of a mathematical theorem  $T$ ; she has checked her work carefully, and shown it to a number of prominent mathematicians, who all find it compelling and correct. It does not follow that her proof is correct, or that  $T$  is true, and she knows this. Nevertheless, it is not epistemically possible for her that not  $T$ . Epistemic possibility requires more than mere fallibility; it requires that we be able to specify a way in which the supposed epistemic possibility may be actual.<sup>28</sup> To specify a way in which it may actually be the case that not  $T$ , Alice would have to be able to specify a way in which one of her axioms may actually be false, or a way in which the logic she used may actually be inconsistent. But this she cannot do, if she has what she regards as a proof of  $T$ .<sup>29</sup>

These preliminary clarifications may be summed up as follows:  $p$  is *epistemically possible* for  $A$  (at  $t$ ) if and only if  $A$  can make sense of its actually being the case that  $p$  (at  $t$ ), in the sense that  $A$  can *specify* a way in which it may actually be the case that  $p$  (at  $t$ ). These clarifications don't by themselves *rule out* the standard analysis of epistemic possibility. Together with anti-individualism, however, they can help us to see why the standard analysis of epistemic possibility is incorrect.

Anti-individualism provides a framework for investigating, for a given speaker  $A$  and a statement  $p$ , whether or not  $A$  can specify a way in which it may actually be the case that not  $p$ . For instance, suppose Alice believes that she is not in the subjectively equivalent world in which she was born, raised, and now lives on Twin Earth, but she has not looked for empirical evidence that might settle this question. She wonders whether that world is epistemically possible for her—whether she can specify a way in which she may actually be in the subjectively equivalent world in which she was born, raised, and now lives on Twin Earth.

According to the standard analysis of epistemic possibility, she can easily specify a way in which she may actually be in that world. All she has to do is describe the world and *picture* herself in it. To do this, she