

Epistemic Possibility

1. Desiderata for an Analysis of Epistemic Possibility

Though one of the least discussed species of possibility among philosophers, epistemic possibility is perhaps the kind of possibility most often invoked in ordinary life – as when, having lost my wallet, I ask myself, “Where could it be?” I am not wondering, say, where it is logically or nomologically possible that my wallet should be (in which case “on the moon” and “in Czechoslovakia” would be among the correct answers to the question). Rather, it seems that the “could” in my question is epistemic; it has *something* to do with the wallet-locations that are consistent with what I now know – though as we shall see, there are problems with a *simple* account of epistemic possibility along those lines. In fact, the notion is more puzzling and resistant to analysis than appears at first glance. In the following, I criticize several accounts of epistemic possibility and offer my own account designed to solve the problems on which earlier accounts founder.

First, some semantic assumptions: I assume that sentences of the form “It might be that P” (not to be confused with “It might *have been* that P”) or “It may be that P” – as in “The wallet might/may be in the car” – typically ascribe epistemic possibility. I assume that “It is impossible that P” is the negation of “It is possible that P.” And I take the “definitely” operator as the dual of the “may” operator – for example, “The wallet is definitely not in the movie theater” is equivalent to “It is not the case that the wallet may be in the movie theater.” In other words, “Definitely, P” ascribes *epistemic necessity* to P. These semantic assumptions function merely to help us interpret examples in which I would say epistemic possibility, impossibility, or necessity is being ascribed using ordinary language.

Second, consider some more substantive proposed axioms about epistemic possibility:

- i) If S knows that P, then it is not epistemically possible for S that $\sim P$.

- ii) If it is impossible that P , then $\sim P$.
- iii) If it is possible that P , and P entails Q , then it is possible that Q .

I take (i) and (ii) to be uncontroversial. (i) articulates perhaps the most central platitude specifically about *epistemic* possibility as opposed to other types of possibility.¹ (ii) is a platitude that ought to hold for any sort of “impossibility” worthy of the name. For example, [The wallet is definitely not in the movie theater] entails [The wallet is not in the movie theater].²

Principle (iii) is more problematic. (iii) may appear an obvious truth applicable to any sort of possibility. It is applicable to any species of possibility that can be treated in the usual way in terms of possible worlds, that is, any sort of possibility that can be defined in terms of truth in at least one of the members of a certain class of possible worlds; for if P is true in all the possible worlds in some class C , and Q is true in all the possible worlds in which P is true, then Q is true in all the worlds in C . Both logical possibility and nomological possibility (the latter of which can be defined as truth in at least one of the possible worlds in which all the actual laws of nature hold) satisfy principle (iii). If epistemic possibility could be construed as truth in at least one “epistemically possible world” (regardless of what the set of epistemically possible worlds is taken to be), then (iii) would likewise be true of epistemic possibility. Nevertheless, we shall see below that epistemic possibility does *not* satisfy this principle, plausible as the principle might seem. Intuitively, it would be desirable to find an account of epistemic possibility that satisfies (iii), but this turns out to conflict with a more important desideratum, that of accommodating clear cases of correct epistemic possibility ascriptions.

A further desideratum for an account of epistemic possibility is that such an account should illuminate the uses of epistemic possibility talk—it should be intelligible in terms of the account why, for example, it is appropriate for one who

¹DeRose (1991, pp. 596-601) makes a strong case for this principle.

²Throughout, square brackets are used to indicate that an expression denotes a proposition. Thus, “[The wallet is in the car]” refers to the proposition that the wallet is in the car. For aesthetic reasons, however, I omit the brackets in cases of variables or logical formulae; e.g., the proposition that either P or Q is the case will be denoted simply “ $P \vee Q$ ”.

has recently lost his wallet to be concerned about the epistemically possible locations of the wallet, as opposed, say, to its logically or nomologically possible locations. The account should likewise put us in a position to understand why the judgement “The wallet might be in the car” results in a trip out to search the car, while the judgement “The wallet is definitely not at the movie theater” forestalls a similar trip to the theater.

In short, we should like an account of epistemic possibility that (insofar as this is possible) upholds the intuitive principles (i)-(iii) above; explains the connection between rational action and judgements of epistemic modality; and, of course, coheres with our linguistic intuitions about particular cases.

2. Examples of Epistemic Modals

Before turning to specific proposed analyses of epistemic possibility, let us consider some cases. I take it that the following are examples of intuitively correct and appropriate ascriptions of epistemic possibility or impossibility:

Lost Wallet: Having lost my wallet, I ask myself: “Where could it be? Could it be at the movie theater?” “No,” I reason, “I remember buying gas after that, and I had to have my wallet to do that. Hm, but it might be out in the car . . .” I then go out to search the car.

Mathematical Proof:

Sam has just completed a (correct) five-page mathematical proof, concluding with the proposition M. Sam and Mary then have the following exchange.

Sam: M.

Mary: That was a difficult proof, and you have occasionally made mistakes before. Could you be wrong about that?

Sam: Yes. I could be wrong in believing M.

Airport:

Mary is picking up Sam from the airport, but she’s a little late, so she calls Sam

on his cell phone.

Mary: Where are you?

Sam: I'm on the ground; we've just landed.

Mary: Is it possible that you're still in the air?

Sam: No, it isn't. I can see out the window right now, and we're on the ground. I know what the ground looks like! Sheesh.

Among other things, these examples illustrate commonly-accepted ways of reasoning about epistemic possibility. In the first example, the proposition [My wallet is at the movie theater] is judged epistemically impossible, based on my memory of buying gas after leaving the theater and my knowledge that I had to have my wallet to do that, while [My wallet is in the car] is judged epistemically possible (presumably on the grounds of my inability to think of any similar reason why it may not be the case). In the second example, the difficulty of Sam's proof and Sam's history of mistakes provide grounds for judging $\sim M$ epistemically possible. In the third case, [Sam is still in the air] is judged epistemically impossible, on the grounds that Sam can see that he is on the ground. In these cases, the facts cited in support of the epistemic possibility or impossibility claim – that I remember buying gas, that the proof was difficult, and so on – are not part of the content of the epistemic possibility claim, but rather function, I take it, as *justification* for it. In each case, of course, the inference is likely enthymematic; but, given normal background assumptions, one's seeing that one is on the ground, for example, constitutes adequate grounds for the claim that it is not epistemically possible (for oneself) that one is in the air.

The following, by contrast, are examples of intuitively incorrect ascriptions of epistemic possibility or impossibility:

Rigel 7:

Rigel 7 is the seventh planet in the Rigel star system.³ Sam, however, knows

³For simplicity, assume that "Rigel 7" may be treated as a definite description equivalent to "the seventh planet in the Rigel star system."

nothing of Rigel and consequently has no thoughts about Rigel or any of its planets. Sam looks at his couch in normal conditions and sees nothing on it. Mary (who happens to know of Rigel 7) says: "For all Sam knows, Rigel 7 might be on the couch."

Unconscious Sam:

Sam is currently unconscious. Mary says: "It is epistemically possible for Sam that he is now conscious. In fact, it is epistemically *impossible* for him that he is *unconscious*."

Epistemology Class:

Jon is teaching a class about philosophical skepticism. After reviewing Descartes' skeptical scenarios in the *First Meditation*, Jon gestures at the table at the front of the room and asks the class: "So, is it possible that there's really no table here?" A student replies: "No, it isn't. We can see the table in front of us right now. We know what a table looks like! Sheesh."

The first two of these epistemic modality claims are sufficiently bizarre that it is difficult to imagine their actually being made by a normal English speaker. To see the bizarreness of the claim in the Rigel 7 case, it is important to note that Mary's claim is not to be understood as a metalinguistic one. Mary is not asserting that for all Sam knows, the *sentence* "Rigel 7 is on the couch" might express a truth. Rather, Mary is to be understood as using (not mentioning) "Rigel 7" in the way she (having a correct understanding of the term) would normally use it, and asserting that a certain proposition, one that would be true only if the seventh planet orbiting the star Rigel were on Sam's couch, is among the epistemic possibilities open to Sam.

The epistemic impossibility claim in Epistemology Class is less bizarre. Rather than assuming that the speaker does not understand the concept of epistemic possibility, we would likely suppose in that case that the student did not hear or understand the skeptical arguments, or that the student is presupposing some controversial philosophical theory. Notice that we would not make such

assumptions about Sam in Airport. The grounds cited for the epistemic impossibility claim in Epistemology Class are, of course, parallel to those cited in Airport – the student has perceptual evidence that there is a table, just as Sam has perceptual evidence that he is on the ground. Yet in Epistemology Class, the claim seems inappropriate and probably false, while in Airport the claim seems appropriate and true. This strongly suggests that the meaning of “possible” is context-sensitive in the way that (many epistemologists argue) the meaning of “know” is context-sensitive. Jon’s discussion of skepticism in Epistemology Class serves to establish a very different context for his epistemic possibility query from the context of Mary’s epistemic possibility query in Airport. Of course, if Mary were to start discussing skeptical scenarios with Sam on the phone, she too might shift the context to one in which Sam would correctly admit to the “possibility” that he is still in the air (perhaps he is a brain in a vat in an airplane). But, in the context as described, assuming no skeptical scenarios are under discussion and things are otherwise normal, Mary’s question is bizarre and Sam’s response seems correct.

Perhaps some other, non-contextualist account of Airplane and Epistemology Class exists. For now, what is important is to recognize the desiderata for an account of epistemic possibility that it should license the epistemic modals displayed in Lost Wallet, Mathematical Proof, and Airport, and that it should disallow the epistemic modals displayed in Rigel 7, Unconscious Sam, and Epistemology Class.

3. Failed Attempts at an Analysis

I shall consider seven proposed analyses of epistemic possibility:⁴

⁴DeRose (1991, pp. 584-6) argues that epistemic possibility is relative to a person or group, so that what is possible in the relevant sense for one person or group of people may not be possible for another, and that what the relevant person or group is varies with conversational context. For the sake of simplicity, I shall confine my attention to epistemic possibility claims that are made relative to explicitly specified individuals, as in “P is possible for S.”

In each analysis, the “=” sign indicates that the expression on the right hand side gives the truth conditions for the expression on the left.

- D₁ P is epistemically possible for S = S does not know that \sim P.
- D₂ P is epistemically possible for S = P is consistent with everything S knows.
- D₃ P is epistemically possible for S = If S were to consider whether P, S would not know that \sim P.
- D₄ P is epistemically possible for S = S does not know that \sim P, nor would any practicable investigations by S establish that \sim P.
- D₅ P is epistemically possible for S = The probability of P, on S's evidence, is greater than 0.
- D₆ P is epistemically possible for S = It is metaphysically possible for someone to have the kind and degree of justification that S has for P, for a false proposition.
- D₇ P is epistemically possible for S = S does not know that \sim P, and there is no relevant way by which S can come to know that \sim P.

Most of these proposals, while perhaps *prima facie* plausible, violate in a serious way one or more of the conditions for a satisfactory analysis of epistemic possibility.

Begin with D₁:

- D₁ P is epistemically possible for S = S does not know that \sim P.⁵

This account mishandles the Rigel 7 case. Since Sam has no concept of Rigel 7, he does not believe any propositions about Rigel 7, including the proposition [Rigel 7 is not on the couch]. Therefore, he does not know [Rigel 7 is not on the couch]. So, according to D₁, it is epistemically possible for Sam that Rigel 7 *is* on the couch. More generally, D₁ implies that if a person does not actually believe \sim P, perhaps due to his having failed to consider it or lacking the concepts required to entertain it, then P is thereby guaranteed to be epistemically possible. Granted, in some cases, a person may count as *dispositionally* knowing \sim P despite failure to *occurrently* consider whether P; I take it, however, that we ought not to extend the notion of

⁵Gendler and Hawthorne (2002, p. 3) mention both D₁ and D₂ as possible analyses. Hacking (1967, pp. 148-9) criticizes D₁ persuasively.

belief so far that the Rigel 7 case counts as one of Sam's believing (even dispositionally) that Rigel 7 is not on the couch.

It seems, rather, that the reason [Rigel 7 is on the couch] is epistemically impossible for Sam is that, although Sam does not actually know that proposition to be false, he does know things that are *incompatible* with it. Thus, we should replace D1 with D2:

D₂ P is epistemically possible for S = P is consistent with everything S knows.⁶

But this account of epistemic possibility does not handle the Mathematical Proof case: $\sim M$ is not consistent with everything Sam knows, since Sam knows the axioms of mathematics that his proof is based on, and – though Sam himself does not know this for certain – the proof is in fact valid. Worse, D₂ implies that before Sam even discovered the proof, it was epistemically impossible for him that $\sim M$, since $\sim M$ was, in actual fact, inconsistent with the mathematical axioms that Sam knows.

D₃ is designed to remedy the defects of both D₁ and D₂:

D₃ P is epistemically possible for S = If S were to consider whether P, S would not know that $\sim P$.⁷

If Sam were to consider whether Rigel 7 was on the couch (which happens only in possible situations in which Sam has a concept of Rigel 7), he would immediately know that Rigel 7 was not on the couch; hence, Rigel 7's being on the couch is epistemically impossible for Sam. But $\sim M$ remains epistemically possible in Mathematical Proof, since if Sam were to consider whether M is true (as in fact he has), he would not know that M (as in fact he does not). But this analysis

⁶Fetzer endorses D₂, though he goes on, inconsistently, to recognize that epistemic possibility is "both broader and narrower in range" than logical possibility (1974, pp. 332-3).

⁷D₃ was suggested to me by Christian Lee (in conversation) and Clayton Littlejohn (*Certain Doubts*).

mishandles the case of Unconscious Sam. Though Sam is in fact unconscious, if he were to consider whether he was conscious, he would be conscious, and he would presumably then know that he was conscious. It is therefore true that if Sam were to consider whether he was conscious, he would not know that he was not conscious. So according to D_3 , it is possible for Sam that he is conscious. Furthermore, it is *false* that if Sam were to consider whether he was unconscious, he would not know that he wasn't unconscious (again, if he were to consider whether he was unconscious, he would have to not be unconscious, and he would then presumably know he was not unconscious). So according to D_3 , it is *not* possible for Sam that he is unconscious – even though he actually is unconscious. D_3 thus violates our condition (ii), the condition that epistemic impossibility should entail falsity.

Hacking suggests an analysis along the lines of:

D_4 P is epistemically possible for S = S does not know that $\sim P$, nor would any practicable investigations by S establish that $\sim P$.⁸

This seems wrong. In the Mathematical Proof case, suppose that checking the proof over a few times carefully would suffice for Sam to know that M. Surely this would count as a practicable investigation. Nevertheless, Sam has not done this as of yet, so he is correct to say that he might be wrong. The possibility of later establishing M does nothing to render M epistemically necessary as of now.

An obvious suggestion for analyzing epistemic possibility is:

D_5 P is epistemically possible for S = The probability of P, on S's evidence, is greater than 0.

The relevant interpretation of probability is presumably epistemic rather than, say, a frequency or propensity interpretation. Even so, there are two problems with D_5 . First, on standard interpretations of probability, probabilities are *at least* things that

⁸Hacking 1967, p. 153 (with slight modification by me).

satisfy the axioms of probability.⁹ It follows from these axioms that the probability of a logical falsehood is 0. Therefore, in the Mathematical Proof case, the probability of $\sim M$ is 0. Given D_5 , this implies that $\sim M$ is epistemically impossible.

This consequence can be avoided only by maintaining that epistemic probability violates one or more of the axioms of probability theory, most likely the axiom that the probability of a tautology is 1. A result of this is that intuitive platitudes about probabilities, such as that the probability of $\sim a$ equals 1 minus the probability of a , or that the probabilities of an exhaustive set of possibilities should add up to 1, may no longer hold, at which point it seems to me that epistemic probability's claim to the name "probability" is called into question. One might resolve this doubt by admitting to a nonstandard or stipulative use of "probability," but in that case it becomes unclear what illumination D_5 sheds on the concept of epistemic possibility. If the notion of "probability" in D_5 is a non-standard one to which our ordinary intuitions about probability do not apply, then D_5 is informative only insofar as one defines this sense of "probability." I suspect that any such definition of "probability" will make D_5 into a notational variant on D_6 , which we will come to presently.

A second difficulty with D_5 is that it licenses the claims of the external world skeptic, given that in fact, the probability of all or nearly all of what I believe about the external world is less than 1, even if only ever so slightly less. But as this problem affects D_6 as well, let us first look at D_6 .

D_6 P is epistemically possible for S = It is metaphysically possible for someone to have the kind and degree of justification that S has for P, for a false proposition.¹⁰

⁹I have in mind the Kolmogorov axioms: (i) $P(a) \geq 0$ for any proposition a ; (ii) $P(t) = 1$ if t is a tautology; (iii) $P(a \vee b) = P(a) + P(b)$ if a and b are mutually exclusive; and (iv) $P(a \& b) = P(a) \times P(b | a)$ for any propositions a, b .

¹⁰Kripke (1972, pp. 103-4) introduces something like this to explain the sense in which, at one time, it was epistemically possible for us that Hesperus was not Phosphorus, even though it is metaphysically necessary that Hesperus is Phosphorus. A similar proposal appears in Bealer 2002, pp. 79-80.

D_6 has the virtue of handling the Mathematical Proof case naturally. Sam's justification for believing M consists in his having an apparent proof for M ; but it is possible to have an apparent proof, which seems just as compelling as the present one does and when one has exercised the amount of care that Sam in fact has, and yet for one to be mistaken.

But D_6 has the vice of supporting external world skepticism. It is metaphysically possible for someone to have the kind and degree of justification for a proposition that I presently have for the claim that I have two hands, and yet for that proposition to be false. This is amply shown by brain-in-a-vat and similar scenarios. The skeptic could therefore argue as follows:

1. If you know that P , then it is not epistemically possible for you that $\sim P$. (Our principle (i).)
2. It is epistemically possible for you that you do not have two hands. (As argued above.)
3. Therefore, you do not know that you have two hands. (From 1, 2.)

If the problem were only the counter-intuitiveness of conclusion (3), we might remedy the problem by dropping the assumption that one's knowing P entails the epistemic impossibility of $\sim P$. But D_6 's verdicts about epistemic possibility themselves are counter-intuitive. D_6 entails that the epistemic impossibility claims in Lost Wallet and Airport are incorrect – despite my memory of having bought gas after leaving the theater, according to D_6 , I should still say the wallet may be at the theater. Indeed, the wallet may be on the sun. Similarly, in Airport, rather than observing peevishly that he knows what the ground looks like, Sam might have more appropriately responded to Mary's question with, "Well, *of course* it's possible that I'm still in the air. Isn't that obvious?" After all, it should be obvious to Mary and Sam that Sam lacks the sort of justification for the claim that he is no longer in the air that would be metaphysically incompatible with his being in the air (it is doubtful that anyone could ever have such justification).

As a consequence of its counter-intuitive treatment of these cases, D_6 also renders the point of epistemic modality judgements obscure, as nearly everything

whose epistemic possibility people commonly inquire about is trivially epistemically possible, while this epistemic possibility has no practical implications. For instance, if D_6 is correct, it is obscure why upon realizing that one's wallet is lost (or at any other time), one should ask where the wallet might be. The correct answer to the question would be that the wallet might be literally anywhere, regardless of your circumstances, what you remember, and so on, but this observation seems to be of no relevance for one looking for a wallet.

A skeptic might counter that the reason people consider questions of epistemic possibility interesting is that they *mistakenly think* many things are epistemically impossible that in fact are possible. The objection to D_6 , however, is that we do not talk and think about epistemic possibility as if we were talking about the property that D_6 describes. If we can articulate an account of epistemic possibility on which the way we commonly deploy epistemic possibility claims makes sense, this account will be *prima facie* superior to D_6 .

Finally, consider DeRose's flexible, contextualist proposal:

D_7 P is epistemically possible for S = S does not know that $\sim P$, and there is no relevant way by which S can come to know that $\sim P$.¹¹

According to DeRose, the relevant ways of coming to know vary with conversational context, though he does not specify rules describing how they vary. Suppose we accept that in the Mathematical Proof case, Sam does not know M. Then D_7 can accommodate the case by maintaining that Sam has no relevant ways of coming to know M. Of course, Sam *has* ways of coming to know M; for instance, he could check the proof over carefully, or he could have other experts check the proof over. The defender of D_7 will simply say that those ways are not "relevant" in the context.

But in the absence of an account of why these ways are not relevant, this treatment of the case is too easy to be interesting. After all, the second clause of D_7 could be reconciled with any assertion of epistemic possibility or impossibility by

¹¹This is based on DeRose 1991, pp. 593-4.

simply adjusting one's view of the relevant ways of coming to know. In the Rigel 7 case, we think it epistemically impossible for Sam that Rigel 7 is on the couch. The defender of D_7 can accommodate this intuition by proposing that in this case, "having someone else tell him about Rigel 7 so that he forms a concept of it, proceeding to consider whether Rigel 7 is on the couch, and thereupon seeing that Rigel 7 is not on the couch" counts as one of the relevant ways of coming to know. Why is *that* a relevant way of coming to know in the Rigel 7 case, whereas "checking over his proof again" is *not* a relevant way of coming to know in the Mathematical Proof case? My complaint here is not that D_7 is false, but merely that it is insufficiently informative to be satisfying.

4. Epistemic Possibility and Justified Dismissals

I propose to define epistemic possibility in terms of epistemic impossibility:

D_8 P is epistemically possible for S = It is not the case that P is epistemically impossible for S.

P is epistemically impossible for S =

- a. P is false;
- b. S has justification for $\sim P$ adequate for dismissing P; and
- c. S's justification for $\sim P$ is Gettier-proof.

A number of aspects of this account call for comment. Condition (c) incorporates whatever condition or conditions must be included in an analysis of knowledge to deal with Gettier cases, that is, whatever is the elusive "fourth condition" (beyond justified, true belief) required for knowledge.¹² What (c) amounts to thus depends on what the correct analysis of knowledge is. I shall not here advance any particular such analysis. For illustrative purposes, however, consider two proposed accounts of the fourth condition. First, consider Clark's account in terms of fully-

¹²See Gettier 1963.

groundedness: on this account, condition (c) amounts to the condition that the justification S has for $\sim P$ should not depend upon any false propositions.¹³ Second, consider the defeasibility analysis: on this account, condition (c) would amount to the condition that S's justification for $\sim P$ should be undefeated, that is, that there should be no true propositions that if added to S's beliefs would neutralize S's justification for believing P.¹⁴

Definition D_s does not presuppose any particular account of knowledge. It does, however, presuppose that whatever the correct account of knowledge is, it can be phrased in such a way as to include both a justification condition and some further, Gettier-proofing constraint on the character of one's justification. I do not presuppose any particular account of what justification consists in; for example, I do not rule out a reliabilist or other externalist approach to justification. I do, however, assume that justification should be understood in such a way that one may separate the having of justification *for* believing something from one's actually believing that thing. Intuitively, a person may have evidence for some conclusion even if he does not in fact believe that conclusion. If this evidence is sufficiently strong, it constitutes adequate justification for believing the conclusion. That is at least one way one may have justification for something one does not believe. There may be others; for example, perhaps having a perceptual experience that non-conceptually represents a given state of affairs provides one with justification for believing a proposition describing that state of affairs, independently of whether one believes that proposition.

What one has justification for depends upon one's cognitive capacities. Thus, one individual might have justification for P by virtue of his awareness of the fact that Q, while another lacks justification for P despite being aware of the same fact. This would occur if the first individual had greater cognitive abilities, such that he would be able to see that Q entails or renders probable P, while the second individual would not be able to see this. What matters here is one's cognitive

¹³Clark 1963. I am not suggesting this as an adequate analysis.

¹⁴Klein 1971. A qualification is required to deal with the problem of "misleading defeaters," which I omit for ease of exposition.

abilities, that is, what one would be able to see, rather than what one actually sees, since one need not have actually considered P to have evidence for P. Thus, a rough account of one way of having justification for believing that P is this: S has justification for believing that P, if S justifiedly believes some proposition Q, such that if S were to consider, using S's actual logical abilities, whether Q supports P, it would be obvious to S that Q supports P. The *degree* of justification S thereby has for P depends upon at least two factors: the strength of the apparent support relation between Q and P, and the degree of obviousness, to S, of that relation. Thus, the more probable Q would appear to render P (the limiting case being that in which Q would appear to entail P), the stronger is S's justification for P. Likewise, the more obvious it would be to S that Q supports P, the stronger is S's justification for P. The notion of obviousness here is psychological – a thing is obvious if it *seems* to one to be the case upon consideration, and some things more strongly seem to be the case than others. Lastly, note that obviousness is taken to be a matter of how things would appear upon relatively *direct* consideration, that is, a consideration that does not require significant further reasoning. The purpose of this last stipulation is to block theorems requiring complex proofs that S has not in fact discovered from counting as epistemically necessary. Such theorems become epistemically necessary only when S goes through the proof (or more precisely, only when S gets to the step immediately before the conclusion). At that point, S's knowledge of this penultimate step constitutes justification for believing the conclusion. The strength of that justification depends on the complexity and difficulty for S of the proof.

The suppositions of the preceding paragraph characterize (part of) my own view of justification. They are not built into D_8 ; other views of the conditions for having justification for a proposition could be combined with D_8 . The views of the previous paragraph, however, provide a basis for resolving some of the puzzles concerning epistemic justification that we have encountered in section 3. We shall come to this matter in the following section.

Finally, the most interesting aspect of D_8 is the appeal to the concept of

dismissing a proposition. Dismissal, as I use the term, is a strong form of disbelief.¹⁵ One who dismisses a proposition is disposed to refuse to take it seriously. Consequently, he does not take it as relevant in practical deliberation, nor does he bother to collect evidence about whether it is the case. When S dismisses P, S will typically not consider or discuss P unless a third party brings it up, at which point S will reject it, typically, as not merely mistaken but *irrelevant* to the discussion. Dismissal is often dispositional: a person need have no *occurrent* attitudes toward P to dismiss P, and usually will in fact have no such attitudes. But one's dismissal of P is *compatible* with one's entertaining and discussing P – though not with an eye towards finding out whether P is true, but for other purposes. For instance, a professor in a philosophy of science class might discuss the Ptolemaic system of astronomy, for its historical interest and to illustrate points about scientific methodology – but not, in all likelihood, for purposes of suggesting it as a serious possibility today. Or an epistemologist might discuss the proposition that the Plutonians are going to launch a lethal nuclear strike against Earth in the year 2100, purely for purposes of illustrating a point about epistemically irrational beliefs.¹⁶

Though dismissal typically has practical consequences, it is in itself a purely *cognitive* state; dismissing P is simply disbelieving P in a particularly strong way, disbelieving P and regarding the question as *settled*. The reluctance, for example, to gather further evidence about whether P is simply a natural consequence of this attitude; if the question has been settled against P, then there is no epistemic value in collecting further evidence. Similarly, the *justification* for dismissing P that D_s refers to is epistemic, rather than, say, pragmatic or moral. Thus, P cannot become epistemically impossible by virtue of one's having sufficiently strong prudential or moral reasons to dismiss P. This point remains even if the contextual factors that I have suggested affect epistemic possibility include such pragmatic factors as the importance of being right about whether P: these pragmatic factors might affect how much justification one needs in order to count as being justified in dismissing

¹⁵My conception of dismissal owes much to Owens' (2000, pp. 142-5) conception of the nature of belief in general.

¹⁶Fumerton (1985, p. 39) is a case in point.

P, but the type of justification one needs is still epistemic. Furthermore, it is plausible to think that there is some minimum level of justification one must always have to count as being justified in dismissing a proposition, whatever the context—just as it is plausible for contextualists about knowledge to maintain that there is some minimum level of justification one must always have to count as knowing a proposition. Knowledge contextualists will presumably deny that there is any context, however casual, in which one may “know” a proposition that is only 80% probable. Contextualists about epistemic possibility should take a similar line. We need not attempt here to identify the minimum level of justification, though it could be investigated by canvassing linguistic intuitions about cases. It is worth noting also that there may be qualitative as well as quantitative requirements on justification—for example, it might be that one’s justification must not only render $\sim P$ at least 99% probable, but that justification also must not rely on any set of premises that themselves are known to support a false proposition.¹⁷

Whatever the exact requirements for justified dismissal are, it seems that some actual dismissals are justified, even when the proposition dismissed is “possible” by Cartesian standards (e.g., by the standards of D_5 or D_6). Thus, in the Airplane Case, it seems appropriate for Sam to report that he is on the ground at the airport, without even considering the possibility that he is a brain in a vat on Mars, or that he is having a massive hallucination causing him to mistakenly think that the airplane has landed. While Sam’s existing justification for believing that he is on the ground does not logically rule out such alternatives, and thus does not logically guarantee that he is on the ground, it seems reasonable in the context for Sam to ignore those sorts of possibilities and to plan his actions on the basis of the assumption that he is indeed on the ground. It is reasonable for Sam to regard his visual experience through the window of the airplane as settling whether he is on the ground, and to end the inquiry into his present location with that. Sam may be well aware of the arguments that no one can really know whether there are such things as airports, but he may view all that as irrelevant in the present context.

¹⁷I am not advancing either of these as actual conditions on either knowledge or epistemic possibility. However, the latter, qualitative condition is the sort of condition one might appeal to to explain intuitions about lottery cases.

How does this notion of dismissal relate to the concept of knowledge? Knowledge requires belief. But it does not require just any sort of belief – one must have an especially strong or confident belief to count as knowing a proposition. It is plausible to suppose that this strong form of belief corresponds to the strong form of disbelief involved in dismissing – that is, that dismissing a proposition is just believing its negation in the particularly strong way that would be required for *knowing* its negation.¹⁸ If so, then presumably having justification for $\sim P$ adequate for dismissing P is simply equivalent to having the sort of justification for $\sim P$ that would be required for knowing that $\sim P$. If so, then D_8 is more closely tied to the analysis of knowledge than is apparent at first glance: D_8 makes the conditions for the epistemic impossibility of P simply coincide with the conditions for *knowing* $\sim P$, minus the belief condition.

In fact, I believe this is the case; I believe that “it is epistemically impossible for S that P ” may be understood simply as asserting that S satisfies the conditions, other than the belief condition, for knowing that $\sim P$. Stating the point from the other direction, we might say that knowledge that P requires the sort of justification for P in the light of which the negation of P is impossible.

5. The Virtues of D_8

We are now at last in a position to resolve the puzzles regarding epistemic possibility that bedeviled earlier accounts. D_8 , given the views about justification stated in the previous section, can accommodate nearly all the desiderata for an account of epistemic possibility. It correctly handles all the cases presented in section 2, it upholds principles (i) and (ii) from section 1, and it accounts for the practical import of epistemic possibility judgements.

In *Lost Wallet*, my memory of buying gas after leaving the movie theater and my knowledge that I must have had my wallet to do that, constitute sufficient justification for dismissing the hypothesis that the wallet is still at the movie theater, given normal background assumptions. Likewise, in *Airport*, Sam’s visual

¹⁸I discuss this strong conception of belief further in my (forthcoming).

experience through the airplane window gives him sufficient justification for dismissing the hypothesis that he is still in the air. In neither case does the subject's evidence logically *entail* the falsity of the hypothesis in question, but it is reasonable in the context for the subject to take the matter as settled by these considerations, ignoring the skeptical scenarios in which he is mistaken. But in Mathematical Proof, Sam's justification for believing M is not sufficient to dismiss the possibility that he has made a mistake in the proof and that M is false. Sam knows mathematical axioms that in fact entail M. The entailment, however, is non-obvious; that is why Sam needs the five-page proof. Although this proof provides him with justification for M, the possibility of making a mistake in such a proof is too great for him to dismiss the possibility that $\sim M$.

In Rigel 7, however, Sam has justification for dismissing the possibility that Rigel 7 is on the couch, despite Sam's lack of beliefs about Rigel 7. This justification may consist in Sam's knowledge that there is nothing (of a visible size) on the couch, and/or simply Sam's visual experience of the couch. Though Sam cannot in his present state (lacking the concept of Rigel 7) consider whether these things support the conclusion that Rigel 7 is not on the couch, it is nevertheless true that if Sam *were* to consider whether they support [Rigel 7 is not on the couch], it would be immediately obvious to him that they do. Thus, if someone were to ask Sam, "Does your present visual experience indicate to you that the seventh planet orbiting the star Rigel is not on the couch?" he would have no difficulty answering "Yes," without the need either for any significant reasoning process or for any augmentation of his logical faculties.

In Unconscious Sam, Sam lacks justification for the claim that he is conscious, because (among other things) he does not presently have any evidence such that, if he were to consider whether it supports the claim that he is conscious, it would be obvious to him that it does. Sam has no evidence at all that he is conscious, nor does he appear to have any other sort of justification available for the claim that he is conscious. This is despite the fact that if he were to consider whether he was conscious, he would then know that he was – the reason he would then know that he was conscious is that he would have different evidence (introspective evidence) than he presently has.

In Epistemology Class, the student arguably lacks justification adequate for dismissing the skeptical scenarios, because in that context, the skeptical scenarios are particularly salient and relevant to the conversation, unlike in the Airport case. As this case illustrates, D_s is able to explain the context-sensitivity of epistemic possibility attributions, as due to shifts in implicit standards of justification. Though the student in Epistemology Class has the same level of justification for asserting that there is a table in the room as Sam has in Airport for asserting that he is on the ground, how much justification one must have to dismiss a claim shifts between the two contexts. In the Epistemology Class, the course of the conversation makes fairly clear that one is meant to dismiss a possibility only if one's evidence entails its negation.

Next, consider our proposed "axioms of epistemic possibility":

- i)* If S knows that P, then it is not epistemically possible for S that $\sim P$.
- ii)* If it is impossible that P, then $\sim P$.
- iii)* If it is possible that P, and P entails Q, then it is possible that Q.

According to D_s , the epistemic impossibility of $\sim P$ requires (a) that $\sim P$ be false, (b) that S have justification for P adequate for dismissing $\sim P$, and (c) that S's justification for P be Gettier-proof.¹⁹ (a) and (c) are trivially satisfied whenever S knows that P. It is plausible that (b) must also be satisfied, since if S knows that P, then this knowledge constitutes particularly strong justification for denying $\sim P$. As suggested at the end of section 4, it is plausible that the level of justification required to know that P would be sufficient to justify dismissal of $\sim P$.

Principle (*ii*) is trivially upheld by D_s , since condition (a) simply stipulates that (*ii*) should hold.

But condition (*iii*) is violated. The easiest way to see this is to consider the Mathematical Proof case. In this case, I assume that Sam knows the axioms of mathematics on which his proof is founded. Indeed, we may stipulate that Sam

¹⁹Since we are here discussing the impossibility of $\sim P$, rather than of P, I have restated the conditions in D_s with "P" and " $\sim P$ " interchanged.

knows that *all* of the axioms are true. Nevertheless, $\sim M$ is epistemically possible for Sam. Since the axioms entail M , $\sim M$ entails that at least one of the axioms is false. So if principle (iii) holds, it is epistemically possible for Sam that at least one of the axioms is false. Given principle (i), this means that Sam does not know that all the axioms are true, contrary to our initial stipulation.

In short, our choice is between principle (iii) and the intuitive assessment of the Mathematical Proof case. D_8 holds with the intuitive assessment of Mathematical Proof, against principle (iii). It seems clear that Sam may be justified in dismissing the hypothesis that one of the axioms is false, without being justified in dismissing the hypothesis that M is false.

Finally, D_8 explains the practical import of epistemic modality judgements. When I have lost a wallet, the locations where the wallet might be according to D_8 are of practical interest for me, because those are the only locations I might want to search – that is, I might want to gather further evidence regarding whether the wallet is at one of those locations. I should ignore the locations where the wallet cannot be – in particular, I should conduct no further inquiry regarding those possible locations – since my present evidence settles that it is not in one of them. Even D_8 's exclusion of a belief condition on epistemic impossibility has a role to play: if I don't know whether the wallet is in the neighbor's freezer merely because I lack a belief about that hypothesis, this has no bearing on whether I should check the neighbor's freezer. As long as I *have justification for* dismissing the neighbor's freezer, I should not check there.

What of the condition that one's evidence must rule out an alternative *directly* and in a way that would be obvious to one, for the alternative to be rendered epistemically impossible? Again, this bears on the need for further inquiry: if my evidence rules out P , but in a way that is somewhat obscure to me or that requires intermediate inferences that I have not made, then further inquiry on my part is appropriate, albeit perhaps only inquiry in the form of further thinking.

These observations lead us towards the speculation that the typical function of epistemic modality judgements is to assess which propositions call for further inquiry in view of our epistemic situation. The class of epistemically impossible propositions is of interest to us because when P is epistemically impossible, it makes

sense not only to act on the assumption that $\sim P$, but to conduct no further investigation. Of course, there are non-epistemic reasons why a proposition might call for no further inquiry, such as that the proposition is boring or that we have more important things to do, and these do not bear on epistemic possibility. An epistemic impossibility is unworthy of inquiry for epistemic reasons, because of the character of the evidence against it. This is not to say that “epistemically impossible” means something like “unworthy of inquiry.” Rather, the suggestion is that “epistemically impossible” denotes a certain feature of a proposition (the feature described by D_8) that in fact renders propositions with that feature unworthy of inquiry, and it is because it renders propositions with it unworthy of inquiry that the feature is of interest to us.

6. The Relativist Challenge

I turn finally to a recent challenge that applies to all of D_1 - D_8 . Some philosophers have proposed a *relativist* treatment of epistemic possibility claims, according to which the truth of an epistemic possibility claim depends not merely on the context in which it is made, but also on the context in which it is *assessed* – that is, the same *token* utterance may be true when one person assesses it at one time, and false when another person assesses it at another time.²⁰ I do not propose to discuss this proposal in detail here. Rather, I aim only to respond to a problem relativists may pose for non-relativist accounts. Consider:

Professor Granger:

Professor Granger is somewhere in the South Pacific when she sees a reporter on television discussing her whereabouts. “We don’t know where Professor Granger is,” says the reporter. “She might be in Prague.” Professor Granger then says to herself:

“That reporter is wrong, for:

1. When he says ‘She might be in Prague,’ the reporter is saying that I might

²⁰MacFarlane 2003; Egan, Weatherson, and Hawthorne 2004.

be in Prague.

2. I know I'm not in Prague.
3. If I know I'm not in Prague, then it is not the case that I might be in Prague.
4. If it is not the case that I might be in Prague, and the reporter says that I might be in Prague, then the reporter is wrong."²¹

It seems that the reporter made a reasonable and correct statement when he said "She might be in Prague." But each of (1)-(4) in Professor Granger's argument also seems correct. The relativist solution is to say that the reporter's utterance is true for the reporter (true in the reporter's context of assessment), but false for Professor Granger (false in Professor Granger's context of assessment).

My account of epistemic possibility, D_s , is non-relativist. It would treat the reporter's statement as true and Professor Granger's contrary conclusion as simply false. How, then, might a proponent of D_s respond to Granger's above argument? Briefly, I think that the argument commits a fallacy akin to equivocation. To understand this fallacy, first consider how someone might accuse Granger's argument of equivocation. One might say:

In premise (1), "I might be in Prague" means that it is epistemically possible *for the reporter* that Professor Granger is in Prague. But in premise (3), "I might be in Prague" means that it is epistemically possible *for Professor Granger* that she is in Prague.

Now, my view of the argument is slightly different from this. I believe that epistemic possibility claims are made relative to a person or group—when one says something is possible in the epistemic sense, one means that it is possible *for some person or group*.²² This person or group can be specified explicitly, as when one says, "For me, it was possible that she was in Prague" or "For all John knows, she

²¹The example is from Egan, Weatherson, and Hawthorne (2004, p. 3).

²²See DeRose (1991, pp. 584-6).

may be in Prague.” Other times, the person or group is not specified explicitly, in which case the conversational context is left to determine who the relevant person or group is. But in some cases, there is not enough information in the context to single out one person or group as relevant; in such a case, it may be indeterminate or unknown to observers what the relevant person or group is.²³

In Professor Granger’s argument, I think, the meaning of “might” is either indeterminate or not knowable on the basis of the information specified in the example. The context does not enable us to determine relative to whom epistemic possibility is being referred to in each of the four “might”’s occurring in premises (1), (3), and (4). But the argument is fallacious, because there is no way of interpreting all four “might”’s such that all the premises are true and the intended conclusion follows from them. Here are the plausible interpretations of the premises:

<u>Premise</u>	<u>Interpretations</u>
1. When he says ‘She might be in Prague,’ the reporter is saying that I might be in Prague.	1a. When he says ‘She might be in Prague,’ the reporter is saying that for all he knows, I might be in Prague. 1b. When he says ‘She might be in Prague,’ the reporter is saying that for all I know, I might be in Prague.
3. If I know I’m not in Prague, then it’s not the case that I might be in Prague.	3a. If I know I’m not in Prague, then it’s not the case that for all the reporter knows, I might be in Prague. 3b. If I know I’m not in Prague, then it’s not the case that for all I know, I might be in Prague.

²³I leave open the possibility that there may be further facts, such as speakers’ intentions and dispositions, that determine the relevant persons or groups even in these cases.

4. If it is not the case that I might be in Prague, and the reporter said that I might be in Prague, then the reporter is wrong.
- 4a. If it is not the case that for all the reporter knows, I might be in Prague, and the reporter said that for all he knows, I might be in Prague, then the reporter is wrong.
- 4b. If it is not the case that for all I know, I might be in Prague, and the reporter said that for all I know, I might be in Prague, then the reporter is wrong.
- 4c. If it is not the case that for all the reporter knows, I might be in Prague, and the reporter said that for all I know, I might be in Prague, then the reporter is wrong.
- 4d. If it is not the case that for all I know, I might be in Prague, and the reporter said that for all he knows, I might be in Prague, then the reporter is wrong.

For premise (1) to be true, it must be interpreted as (1a); the reporter obviously did not mean that for all Professor Granger knows, she might be in Prague. Similarly, for (3) to be true, it must be interpreted as (3b). And for the argument to be valid given those interpretations of (1) and (3), (4) must be interpreted as (4d). But (4d) is false; it may very well be the case – as in fact it is in the example – that for all one person knows, P may be true, while given what another person knows, P is definitely not true. Thus, there appears to be no way of disambiguating (1), (3), and (4) that makes the argument sound. Consequently, the case of Professor Granger causes no difficulty for D_s , nor for any other non-relativist account of epistemic possibility.

References

[Author], "Moore's Paradox and the Norm of Belief" (... forthcoming).

- Bealer, George. "Modal Epistemology and the Rationalist Renaissance" in *Conceivability and Possibility*, ed. Tamar Gendler and John Hawthorne (Oxford: Oxford University Press, 2002), pp. 71-125.
- Certain Doubts (weblog). URL: <http://bengal-ng.missouri.edu/~kvanvigj/certain_doubts/index.php?p=148>.
- Clark, Michael. "Knowledge and Grounds: A Comment on Mr. Gettier's Paper," *Analysis* 24 (1963): 46-8.
- DeRose, Keith. "Epistemic Possibilities," *Philosophical Review* 100 (1991): 581-605.
- Egan, Andy, Brian Weatherson, and John Hawthorne. "Epistemic Modals in Context," forthcoming in *Contextualism in Philosophy*, ed. Gerhard Preyer and Georg Peter (Oxford University Press, 2005).
- Fetzer, James. "On 'Epistemic Possibility,'" *Philosophia* 4 (1974): 327-38.
- Fumerton, Richard. *Metaphysical and Epistemological Problems of Perception* (Lincoln, Nebr.: University of Nebraska Press, 1985).
- Gendler, Tamar and John Hawthorne. "Introduction: Conceivability and Possibility," in *Conceivability and Possibility*, ed. Tamar Gendler and John Hawthorne (Oxford: Oxford University Press, 2002), pp. 1-70.
- Gettier, Edmund. "Is Justified True Belief Knowledge?" *Analysis* 23 (1963): 121-123.
- Hacking, Ian. "Possibility," *Philosophical Review* 76 (1967): 143-68.
- Klein, Peter. "A Proposed Definition of Propositional Knowledge," *Journal of Philosophy* 67 (1971): 471-482.
- Kripke, Saul. *Naming and Necessity* (Cambridge, Mass.: Harvard University Press, 1972).
- MacFarlane, John. "Epistemic Modalities and Relative Truth," unpublished ms. (2003).
- Owens, David. *Reason without Freedom: The Problem of Epistemic Normativity* (London: Routledge, 2000).