This brings us to the most famous puzzle cases in epistemology, possibly the most famous puzzle cases in the whole of philosophy: the Gettier cases. They are fantastically well-known, probably because they're so clear and decisive against the justified true belief analysis of knowledge, or at least they seem to be so.

Let's suppose that P is something that I'm justified in believing, and let's suppose that P clearly implies Q. There's no doubt whatsoever that P implies Q. So, I believe that P. I have a justified belief that P (we're assuming that) and I then infer Q where Q obviously follows from P.

Does it follow that I'm justified in believing Q? Would you agree that's how it sounds, very plausible? If I'm justified in believing P and Q obviously follows from P, surely I must be justified in believing Q. That's how you get the Gettier cases.

So, here's one example. Suppose I'm in the desert, I see what is in fact a mirage. I think, and I'm justified in believing, let us suppose, that I'm seeing an oasis. I infer that there's no oasis over there and so I turn to my companion and I point and I say, "There's an oasis over there." Now, in fact, there is an oasis over there, but it's hidden behind a sand dune. What I'm seeing is a mirage, but coincidentally there is an oasis in that direction. So, I believe that there's no oasis over there. It's true that there's no oasis over there. I'm justified in believing that I see an oasis, and I've inferred from the supposition that I see an oasis - there actually being an oasis over there - it looks like I've got a justified true belief. But we don't want to say it's a case of knowledge.

Another example, which I think is a slightly more plausible case of the kind of example that Gettier himself gives, is this: suppose we have a load of applicants for a job. Just two of them are men. One of the men is very well-qualified, one of the men is very poorly qualified. All the rest are women, and they're much, much better qualified than the badly qualified man. Now suppose I have it on good authority that the well-qualified man is going to get the job. So, I believe that a man is going to get the job, and I've got a justified belief that a man is going to get the job. Unbeknownst to me, however, there's some funny business going on (I don't know, maybe involving the Masons or bribery or blackmail or what-have-you), and actually the other man gets the job. In this situation, I have a justified belief that a man would get the job. It's true that a man did get the job, but I certainly didn't know that a man would get the job because there's a kind of accident. My justification led me to the belief that the well-qualified man would get the job, but in fact, that a man would get the job has turned out to be true by some quite different route.

Now, these sorts of cases do seem to refute the standard justified true belief account of knowledge. So, it's tempting to add a fourth condition. Maybe we should say that S knows that P if and only if P is true, S believes that P, and S is justified in believing that P in a way that doesn't depend on any falsehood. That's a way of trying to evade the Gettier counterexamples where somebody infers a truth from a falsehood and thus achieves a justified belief that isn't knowledge.

Now, back in the days when I was a student, this sort of thing, trying to patch up the traditional analysis of knowledge to avoid the Gettier counterexamples, was quite a major industry. Lots of

papers came out with people trying to invent conditions that would get around the counterexamples. It's a rather sad history because all this effort failed to produce any convincing resolution of the problem.

Here's an example of the sort of problem we face. Suppose I'm organizing some event and I want to know how many people were there. Maybe the reason I want to know how many people were there is because I want to make a judgment as to which room to use for a future event of that kind. So, somebody comes to me and they say, "Oh, there were exactly 78 people there." What interests me is whether there were more than 40. If there were more than 40, then I have to use a big room rather than a small room next time. So, they tell me there were exactly 78 people there. I infer that there were more than 40. Arguably, I know that there were more than 40. But let's suppose the reporter actually got it wrong. There weren't 78, there were 77. I'm still inclined to say I know that there are more than 40. So, I've inferred a truth from a falsehood that hasn't undermined its claim to being knowledge.

Now, you might want to get around that. You might want to say, "No, what you did was infer that there were more than 40 people from the fact that he believes that there were 78, and you have the knowledge that people can go wrong in little ways. But there's an implicit probability judgment there. You're making the judgment that if somebody says there were 78 and says it sincerely and they've got good faculties and that's this kind of thing, it's overwhelmingly probable that there were more than 40. It's not overwhelmingly probable that there were exactly 78. So, maybe that gets you out of it. Let's add a fifth condition."

But there's a general problem with heading in the direction of probabilities - the so-called lottery paradox. It's very tempting to try to get around these kinds of problems by saying, in order to have knowledge, you've got to have a sufficiently high probability. That's the key thing. It's not whether you've inferred something from a falsehood, it's having a sufficiently high probability of truth. But suppose we've got a billion tickets in a lottery. Well, I believe that the ticket with zeros in every place won't win. Indeed, I'm pretty sure it won't win. There's only a one in a billion chance that it's going to win. I believe that the next ticket after that won't win either. In fact, I believe apparently quite reasonably of every single ticket, but it won't win. Nevertheless, one of them will win. And because of this, we're reluctant to call my belief about any of these tickets knowledge. So, if I go and buy a ticket in a billion-ticket lottery and then I say, "Well, I know it won't win," I think you'd probably say, "No, you don't know. You've got a very probable belief, an extremely probable belief that it won't win, but you don't know that it won't win because there's a chance that it will, even if that chance is tiny." Now, if that's right, then the lottery paradox involves real problems for any attempt to explain knowledge in terms of sufficiently high probabilities because however high the probability is, you can make it a trillion-ticket lottery or whatever, you can always get a lottery in which your belief that this ticket won't win will have as high or higher a probability than any belief which isn't 100 percent certain. And yet, we're not going to want to call it knowledge. okay?

So, maybe we want to say it's not exactly a matter of probability, it's a matter of ruling out accidents. What we don't want is to allow us knowledge a belief that accidentally happens to be true. I thought I knew that one person was going to get the job and it's a kind of accident as far as my knowledge was concerned that somebody else got it. That's what rules it out. It's a kind of accident that there's an oasis over there in the same direction as the mirage, so forth. Actually, it's extremely difficult to pin this down.

Suppose I have a car whose speedometer gradually corrodes. Let's say that at a particular time, it just happens to be accurate enough to ensure that I'm complying with the law. Let's say I always drive at what I think is 38 miles an hour along a road with a speed limit of 40 miles an hour. Fortunately, the corrosion of my speedometer is such that it just keeps me within plus or minus two miles an hour of the actual speed, so I'm safe. Well, is it an accident that I'm safe? In a sense, yes, in a sense, no. Given that the speedometer has corroded, I'm very lucky that it's keeping me within that margin. But given that it's keeping me within that margin, it's no accident that I'm safe. Again, if I occasionally hallucinate things, does that mean it's just a matter of chance that my current belief isn't a hallucination? It's very difficult to pin these things down in a way that will give us a satisfactory account of knowledge.

Another problem is known as contextualism. Suppose you want to get a train up north, and I say, "I know that the train is scheduled to leave at 17:36." It's a train I regularly take, so I can assure you it's scheduled to leave at 17:36. But maybe you've got a really important appointment, and you're not content with my saying that, so you say, "Do you really know that it leaves then? I absolutely need to make that appointment." Okay, I say, "I'll check on the web." And you can imagine this going even further: "I know you're familiar with the timetable and you've checked the web, but do you really, really know?" "I really, absolutely have to be there." "Okay, I'll ring up the station." And you can imagine a sequence of checks, each more stringent than the last, which suggests that the threshold we require to count something as knowledge can be variable. We put a higher and higher hurdle, depending on the importance of the task. And that suggests that maybe knowledge isn't an absolute category. Maybe it's dependent on our particular purposes.

Let's also consider the context in which we use the word "knowledge" within ordinary life. So, consider this contrast: "Does she know that her husband is cheating on her?" Imagine that being said in a soap opera or something like that. Now, that probably means something like, "Does she believe he's cheating on her, like we all do?" You could imagine it being said in a context where there's some uncertainty as to whether he's cheating. Maybe there's all sorts of circumstantial evidence, and that question, "Does she know," is not really a question about her epistemological state, it's a question about her belief. Contrast that with the following case: somebody alleges that her husband is cheating, and I say, "Yes, but do you know her husband is cheating?" In which case, I'm asking about what is the case rather than about the belief. Or, you can imagine a case where some train accident is reported and my son was on the train, and I say, 'Do you know that my son is alright?" And I don't actually give a damn about your epistemological state, the only thing I'm interested in is, "Is he okay?" So, it's arguable that when we talk about knowledge in a practical situation, normally we're interested either in somebody's state of belief, or we're interested in the actual facts. It's very, very unusual for us to ask in ordinary life whether something is a case of knowledge when we already know that somebody believes something and we already know that it's true. It's the kind of question only philosophers ask. And you might wonder in that context whether we're actually likely to get any single consistent account of knowledge. Why should we assume that if the word "knowledge" has these different roles in language, that there really must be some single unitary essence of what knowledge is which will give the answer to all of these questions? Maybe we'll find that the concept of knowledge, as we use it in ordinary language, varies depending on our purposes.

Now, this is the kind of message that's very much associated with the later work of Wittgenstein in his book "Philosophical Investigations". We shouldn't just assume that because we've got a

word "knowledge" which seems to be a noun for a certain kind of state that there really is some sort of essential state that it picks out. Well, of course, however we choose to use the word "knowledge," we can still ask in any particular case whether the P is true. And this brings us back to G.E. Moore and his hands.