

3

Epistemic Closure*Freddie and Betty at the Zoo*

In a logically valid inference, the truth of the premises absolutely guarantees the truth of the conclusion. Suppose you know proposition P is true. And you know P entails Q. From these, you draw the conclusion that Q is true. This seems like a foolproof way of coming to know that Q is true.

This general idea has come to be called “deductive closure” for knowledge. It can be stated this way:

Closure: For any subject S and any propositions P and Q: If S knows that P is true and knows that P entails Q, then, evidentially speaking, this is enough for S to know that Q is true.¹

In other words, if you know something and you know that thing entails something else then you are in a position to know the other thing too. Just make the inference.

Two prominent epistemologists who deny Closure are Fred Dretske and Robert Nozick. We will focus on Dretske. Dretske’s reason for rejecting Closure stems from a deeper idea about what is required for knowledge. According to him, you can know something only if you have a “conclusive reason” to believe

¹ There is a good bit of debate over how to state various epistemic closure principles. For an overview see John Collins, “Epistemic Closure Principles” in *The Internet Encyclopedia of Philosophy* <http://www.iep.utm.edu/epis-clo/>. This particular statement of closure is taken from Fred Dretske’s “The Case against Closure,” in *Contemporary Debates in Epistemology*, Matthias Steup, John Turri, and Ernest Sosa (eds.) (New York: Blackwell, 2005), 13–25.

TELL ME SOMETHING I DON'T KNOW

it. Dretske has his own way of understanding what makes something a conclusive reason. He defines that as follows:

R is a conclusive reason for P =df R would not be true unless P were true (in other words, if P were false R would be false).

Dretske allows both beliefs and perceptual experiences to serve as reasons. (If you find it a bit odd to call a perceptual experience “true” or “false,” you can understand the definition above to say that R “obtains” or “does not obtain” when we are talking about perceptual experiences.)

An example will help make this notion of a conclusive reason clear. Suppose you believe that the used car you just bought is in perfect condition because that’s what the salesman told you. On Dretske’s view, the fact that the salesman told you the car is in perfect condition is a conclusive reason for believing the car is in perfect condition only if:

The salesman would not have said the car is in perfect condition unless it were true that the car is in perfect condition.

In other words:

If the car were not in perfect condition, then the salesman would not have said so.

If you want to know whether the used car salesman has provided you with a conclusive reason, ask yourself whether he would have said it if it weren’t true. (Well, would he?)

Now consider another example: You are at the zoo looking at what appears to be, and in fact is, a zebra. Based on your visual experience—that of a four-legged equine-shaped black and white striped mammal—you believe that there is a zebra in front of you. This visual experience will be a conclusive reason only if the following condition holds:

If there were not a zebra in front of you, then you would not have had that perceptual experience.

How do we tell whether a counterfactual conditional statement like this is true?

A common way to think about counterfactuals is in terms of possible worlds or alternative ways things might have been. There was a zebra in front of you but there might not have been. How would things have looked to you if there

were not? Assuming everything is normal at the zoo on this day—your eyes are working properly and you are not prone to hallucinations or anything like that—then, it seems, if there were not a zebra in front of you, you wouldn't have had that sort of experience. If there were nothing at all in front of you or if there were some other zoo animal there, then you wouldn't have had the visual experience you did. So it looks like we can say that your perceptual experience qualifies as a conclusive reason to believe that's a zebra.

This in essence is Dretske's account of perceptual knowledge. Your perceptual experiences can give you a way to know things about the world around you because they can provide you with conclusive reasons. Now here's where things get interesting. Consider a cleverly painted mule—a mule expertly airbrushed to look just like a zebra. Is your current perceptual experience, an experience as of a four-legged equine-shaped black and white striped mammal, a conclusive reason to believe that the thing before you is not a cleverly painted mule? If the answer is yes, then the following is true:

If there were a cleverly painted mule in front of you, then you would not have had that perceptual experience.

But this, Dretske tells us, is false. If there were a cleverly painted mule in front of you, one that is painted up to look just like a zebra, your perceptual experience would have been the same. You would still have had a visual experience of a four-legged equine-shaped black and white striped mammal. And therefore your perceptual experience does not provide you with a conclusive reason to believe that the thing in front of you is not a cleverly painted mule.

Now consider the following argument:

1. That's a zebra.
 2. If that's a zebra then that's not a cleverly painted mule.
-
3. Therefore, that's not a cleverly painted mule.

This is a deductively valid argument. The second premise is true by definition. Painted or not, anything that is a zebra is not also a mule. You know that. And you also know that the first premise is true. "That's a zebra," we are assuming, is in fact true and, for reasons given above, you have a conclusive reason to believe it. So you know each of the premises and you know that the inference is valid but, Dretske contends, you cannot know the conclusion because you do not have a conclusive reason to believe it. You do not have a conclusive reason to believe that thing in front of you is not a cleverly painted mule. If there were

a cleverly painted mule there, your visual experience would be the same. This is why Dretske rejects Closure. You know the above premises, you know they entail the conclusion but you don't know the conclusion.²

Some contend that this position is absurd. Keith DeRose puts the point by saying that Dretske's view generates "abominable conjunctions."³ For example: "I know that's a zebra but I don't know it is not a cleverly painted mule." This seemingly absurd statement is a consequence of Dretske's view.

Although its consequences strike some as absurd, Closure denial does have its advantages. Imagine you are at the zoo admiring what you believe to be (and in fact is) a zebra. A skeptical philosopher attempts to spoil your fun by asking how you know it isn't a cleverly painted mule. You might think you can't know something like that (after all, a cleverly painted mule would look pretty much the same to you). That might lead you to think you don't really know it is a zebra, and that might ruin your whole afternoon. But, on Dretske's view, you can know it is a zebra even if you cannot know it is not a cleverly painted mule. Closure denial provides grounds for an interesting response to skepticism. (You should be able to see how Dretske-style closure denial can help with the classic problem of the brain in a vat. We will revisit brains in vats in a later chapter.)

In this chapter, Freddie and Betty are at the zoo admiring what they are both convinced is a zebra. Freddie is a Closure denier and Betty is rather critical of that position. The worry about various kinds of abominable conjunctions as well as a few other worries are explored.

Freddie and Betty are at the zoo standing in front of the zebra exhibit.

Freddie: 4:30 already? I guess we missed feeding time.

Betty: I didn't realize you were hungry. Want some of this popcorn?

F: Is there anything on it?

B: "I can't believe it's not butter."

F: You can't *merely* believe it's not butter.

B: One taste and you *know* it's not butter.

2 Astute readers will notice that the above argument has two premises and the Closure principle offered earlier was stated in terms of a single premise. If this is bothersome, note also that the second premise is true by definition and can be dropped out without affecting the overall point. One might also regard proposition P in the Closure principle as the conjunction of 1 and 2 above. Closure principles for very long deductive arguments will be considered in the next chapter.

3 Keith DeRose, "Solving the Skeptical Problem," *Philosophical Review* 104(1): 1-52 (1995).

F: “Fortified with 14 essential vitamins and lubricants.”

B: Well at least the animals here are good. Take this exhibit right here. Now *that* is a zebra.

F: Indeed my friend. Indeed.

B: But we came here to see the mule exhibit. We better move on. I don’t know where the mules are, Freddie, but I know they’re not here. I know that’s zebra.

F: That’s right Betty. You know it’s a zebra but of course, you don’t know it’s not a cleverly painted mule.

**Abominable
Conjunction**

B: A what?

F: A cleverly painted mule.

B: A mule with a paint job.

F: That’s right. And I’m not talking about some \$99.95 Earl Scheib-can’t-believe-it’s-not-a-zebra job here, Betty. I’m talking about a mule that’s been expertly airbrushed to look *exactly* like a zebra. If there were a cleverly painted mule here, you’d never know just by looking at it.

B: Why would anybody paint up a mule and say it’s a zebra?

F: A zoo might do something like that in the event of a zebra shortage. Who’s gonna pay 18 bucks for a zoo ticket if there’s no zebras?

B: That seems pretty far out.

F: Yeah it is far out. Completely implausible even. But just because it’s far out and completely implausible doesn’t mean we know it’s not happening. Have you checked with the zoo authorities? Did you examine the animal closely enough to detect such a fraud?⁴

B: I guess not. But wait. Are you saying I *don’t know* that’s a zebra?

F: Don’t be ridiculous. Of course you know that’s a zebra. But you don’t know it’s not a cleverly painted mule. I don’t either.

B: Well if I don’t know it’s not a cleverly painted mule how do I know it’s a zebra?

F: Well, duh. You’re standing here looking at it.

B: What?

4 Fred Dretske, “Epistemic Operators,” *Journal of Philosophy* 67(24): 1007–23, 1016 (1970).

TELL ME SOMETHING I DON'T KNOW

F: Here's the thing, Betty. Since that's a zebra and you know what zebras look like and there's nothing wrong with your eyesight, you know, by standing here looking at it, that that's a zebra. But since cleverly painted mules look just like zebras, you can't know, by standing here looking at it, that it's not a cleverly painted mule. Same goes for me.

B: Wait. If that is a zebra, then it's not a cleverly painted mule. Right?

F: Correct.

B: And it is a zebra.

F: Yep.

B: And *I know* that it's a zebra? And *I know* that if it's a zebra then it's not a cleverly painted mule?

F: That is my position.

B: I can see how somebody could know that's a zebra, know that if that's a zebra then it's not a cleverly painted mule but fail to know that's not a cleverly painted mule because he never puts the two together and makes the inference.

F: Yeah that's one way it could happen. But even if you make the inference, you still cannot know that's not a cleverly painted mule. At least not by looking at it and reasoning from there.

B: Let's test this out. I'll take my knowledge that that's a zebra and combine it with my knowledge that *if* that's a zebra, then that's not a cleverly painted mule. And then, via modus ponens, I'm gonna arrive at the belief that that's not a cleverly painted mule.

F: Go for it.

Betty closes her eyes and strains for a moment.

B: There. Did it. Now I know that's not a cleverly painted mule.

F: No you don't. You believe it. And you arrived at that belief by a valid inference from premises you know are true. But you are in no position to know you arrived at a true conclusion.

B: Hold on. Take look at the argument I just ran through. Here it is again. Premise One: *That's a zebra*. Premise Two: *If that's a zebra then that's not a cleverly painted mule*. Therefore: *That's not a cleverly painted mule*. That is a valid argument. Right?

F: Yes.

B: So unlike some philosophers, you're not some sort of "deviant logician" who goes around denying the validity of modus ponens.⁵

F: Not at all. My logical tastes run nice and conventional.⁶

B: And my premises are true too, right? I mean that is a zebra and if it's a zebra then it's not a cleverly painted mule.

F: Of course.

B: So we've got here a valid argument with true premises.

F: Yeah that's right.

B: Is the conclusion true?

F: I don't know.

B: You don't know.

F: Yeah that's right. I don't know.

B: But you agree that the premises are true and, since the premises are true and it's valid, there's no possible way that the conclusion can be false!

F: I know what validity is, Betty. The argument is valid. So the only way the conclusion could be false is if a premise is false. But the premises are both true.

B: So is the conclusion true?

F: I don't know.

B: I'm starting to feel like I'm standing in front of the tortoise exhibit.⁷ Let's try this: That's a zebra. If that's a zebra then that's not a cleverly painted mule. Therefore, that's not a cleverly painted mule. Now, is that a *sound* argument?

F: The argument's valid. And the premises are true. Is it sound? I don't know.

5 See Vann McGee, "A Counterexample to Modus Ponens," *Journal of Philosophy* 82(9): 462–71 (1985) and William Lycan, "MPP, RIP," *Philosophical Perspectives* 7: 107–22 (1993).

6 See Dretske, "The Case against Closure," 13.

7 In Lewis Carroll's essay "What the Tortoise Said to Achilles," *Mind* 4(14): 278–80 (1895) the tortoise is presented with an obviously valid argument with true premises. He believes the premises but refuses to believe the conclusion. Achilles insists that if you believe the premises, then you must believe the conclusion. The tortoise responds: that sounds like another premise—better put it in the argument! Achilles obliges. Achilles tells the tortoise if he accepts the premises then he now must certainly accept the conclusion. The tortoise says this too sounds like another premise; better put it in the argument. And so it goes.

TELL ME SOMETHING I DON'T KNOW

B: But you do know that all it takes for an argument to be sound is that its premises are true and it's valid?

F: Of course I know what makes an argument sound. Any valid argument with true premises is sound. And this is a valid argument with true premises. But I don't know whether it's sound. Why is that so hard to understand?

B: Maybe because you are contradicting yourself?

F: Nope. I'd be contradicting myself if I'd said that's a zebra and that's not a zebra or the argument is sound and it isn't or something like that. What I said is I *know* that's a zebra and I *know* that if that's a zebra then it's not a cleverly painted mule. But I'm in no position to know it's not a cleverly painted mule. That's not only consistent but true. I'd be guilty of logical inconsistency only if I also held some general principle to the effect that, for any subject S and any propositions P and Q, if S knows that P and S knows that (P entails Q) then S is in a position to know that Q. But this sort of closure principle is precisely what I reject.

B: Well it sure sounds like you are contradicting yourself.

F: There are things that it sounds absurd to say that aren't really contradictions. It's often true that it's raining but I don't believe it. But it's always absurd for me to say such a thing.

B: Moore's paradox.

F: You got it. Now I'm happy to admit that there is a sense in which it is absurd to *say* that you know that's a zebra but you don't know that's not a cleverly painted mule. My view entails that things like that are true. But so what? Just because something is conversationally inappropriate that doesn't mean it's false.⁸

B: Well if no one should ever say "I know that's a zebra but I don't know that's not a cleverly painted mule" why do you keep doing it?

F: Sometimes I care more about saying something true than saying something appropriate. Philosophers face that choice every day, Betty.

B: I care about truth too. That's why I'm willing to infer the conclusion "that's not a cleverly painted mule" from the premises "that's a zebra" and "if that's a zebra, then that's not a cleverly painted mule." The premises are true, the argument's valid, so the conclusion's gotta be true too! That's my philosophy.

8 Dretske, "The Case against Closure," 17.

F: That's chutzpah, not philosophy.⁹

B: A guy who says—with a straight face, mind you—that a valid argument with true premises might be unsound is gonna say *I've* got chutzpah? Oy vey!

F: Lemme think about that again. *That's a zebra and if it's a zebra, it's not a cleverly painted mule. Therefore it's not a cleverly painted mule.* Yes, I said that this argument is valid and has true premises but I don't know whether it's sound. That sounds pretty silly to me now.

B: So now you think the argument is sound?

F: Yes.

B: And sound arguments have true conclusions?

F: Of course. It's a sound argument and all sound arguments have true conclusions.

B: So is the conclusion true?

F: I don't know.¹⁰

B: Maybe we should start looking for that mule exhibit, Freddie.

Betty and Freddie walk around and look for the mule exhibit.

B: Wait a minute, Freddie. Why are we walking around? After all, maybe that was a mule back there in the exhibit marked "Zebra."

F: That's not a mule. That's a zebra. You and I both know that. Of course, we don't know that's not a cleverly painted mule. But we know it's not a mule.

B: Hold on, Freddie. Even if you are right and I know it's a zebra but don't know it's not a cleverly painted mule, how could I know it's *not a mule* but not know it's not a *cleverly painted* mule? Cleverly painted mules are still mules. If it's not a mule, it's not a cleverly painted mule.

F: Knowledge requires conclusive reasons. R is a conclusive reason for P if and only if R would not be the case unless P were the case. Wasn't that implied in what I said before?

B: I guess. But how does it explain how I can know that's not a mule but fail to know that's not a cleverly painted mule?

⁹ Dretske, "The Case against Closure," 24.

¹⁰ Freddie is here offering what I call a "logical abomination." For more, see Michael Veber, "The Argument from Abomination," *Erkenntnis* 78(5): 1185–96 (2013).

TELL ME SOMETHING I DON'T KNOW

**Conclusive
Reasons**

F: Your reason for thinking that's a zebra is that you had a certain type of visual experience. You would not have had that visual experience if you weren't looking at a zebra. If you were looking at an elephant or that popcorn stand or something else around here, your visual experience would have been much different. But if you were looking at a cleverly painted mule, you would have had an experience just like the one you had when you were looking at that zebra. So your visual experience is a conclusive reason for believing that's a zebra but it's not a conclusive reason for believing that's not a cleverly painted mule. That's why you can know it's a zebra but you cannot know it's not a cleverly painted mule.

B: So conclusive reasons don't need to be entailing reasons.

F: Correct. Your visual experience doesn't logically entail that's a zebra in the sense that there are some possible worlds where you have that sort of visual experience but there is a cleverly painted mule in front of you. But, as we agreed earlier, those sorts of worlds are way far out.

B: Okay and so, in the nearby possible worlds, wherever I have that sort of experience, there is in fact a zebra in front of me.

F: Correct. In other words, the nearest non-zebra world is also a world where you aren't having that sort of experience but a different one. That's enough to make your visual experience a conclusive reason. Notice, however, that your visual experience is not a conclusive reason for thinking that's not a cleverly painted mule. If you were looking at a cleverly painted mule you would have had the same experience.

B: Okay. I get why you think I don't know it's not a cleverly painted mule even though I know it's a zebra. But, if I don't know it's not a cleverly painted mule, how can I know it's not a mule?

F: If you were looking at a mule, you wouldn't have had a visual experience like that. The only way you could've had that sort of visual experience and been looking at a mule is if you were looking at one of those cleverly painted mules. But, as we already agreed, cleverly painted mules are far out. That makes your visual experience a conclusive reason to believe that's not a mule.

B: So I know it's not a *mule* but I don't know it's not a *cleverly painted* mule.

F: Exactly. Since you examined the animal closely enough to know it's a zebra, you examined the animal closely enough to rule out it being a mule. But you can't rule out it being a cleverly painted mule because that would require gathering some additional evidence. Neither of us bothered to do that.

B: Do I also know it's not a member of the set of all mules?

F: Sure. After all, if there had been some member of the set of all mules in front of us, it wouldn't have been one of the cleverly painted mules. It would be one of the ordinary ones and in that case your visual experience would have been much different. So you know it's not a member of the set of all mules.

B: And if something is not a member of a given set, it is not a member of any subsets of that set. I know that's not a member of the set of all mules. And therefore I know it's not a member of any subset of the set of all mules. The set of all cleverly painted mules, of course, is one of the subsets of the set of all mules. But, for all I know, that is a member of the set of all cleverly painted mules.

F: Now you're speaking my language.

B: I also know it's not a mule that somebody put paint on but I don't know it's not a cleverly painted mule.

F: Come again?

B: The American Mule Racing Association often paints numbers on their mules for the races which I regularly attend and lay down a few bets. Had I been standing in front of a mule that somebody put paint on, it would have been one of those. And they're pretty sloppy about it. So, had I been standing in front of a mule that somebody put paint on, it wouldn't have been a *cleverly* painted mule and my visual experience would have been nothing like the one I had looking at that zebra. So I know that's not a mule that somebody put paint on. But I don't know it's not a cleverly painted mule.

F: If you say so.

B: And now that you mention it, I remember reading about how the Gaza City Zoo painted up a mule to look like a zebra.¹¹ I saw pictures and, I have to say, they did a pretty half-assed job.

F: They didn't paint the whole thing?

B: The black stripes weren't dark enough and, unlike zebra stripes, they were all of uniform width. To me, it didn't look anything like a zebra. It looked like a mule in faded prison garb. Even Earl Scheib would be embarrassed. The nearest possible world in which I was just standing in front of a mule that some

11 It was a donkey actually but you get the point. <http://www.telegraph.co.uk/news/worldnews/middleeast/israel/6274874/Gaza-zookeepers-draw-crowds-with-painted-donkeys-after-zebras-die.html>.

TELL ME SOMETHING I DON'T KNOW

zoo officials painted in an effort to make it resemble a zebra is a world where I am at the Gaza City Zoo standing in front of their pathetic fake.

F: Uh, okay.

B: But my experience in that sort of situation would have been much different from the one I had looking at that zebra. Now of course someone cleverer than those guys could paint up a mule so that it looked just like a zebra even to me. In other words, I know that's not a mule that some zoo officials painted in an effort to make it resemble a zebra. But I don't know it's not a cleverly painted mule.

F: Right. Now before you go making any more sport of all this, don't forget that I never said these conjunctions are conversationally appropriate. I just said that they are true. Just like those instances of Moore's paradox.

B: Yeah, about that. I see why moonronic¹² conjunctions like "It's raining but I don't believe it" and "Dogs bark but I don't know it" are absurd things to say. I cannot assert either one because I cannot know either one. But according to you, frediculous conjunctions like "I know that's not a member of the set of all mules but I don't know it's not a cleverly painted mule" are knowable. Why not shout them from the rooftops?

F: Frediculous. That's cute. But let's be careful here. Although frediculous conjunctions are true, it's not clear to me that they are knowable. Consider this one: You know that's a zebra but you don't know it's not a cleverly painted mule. For the first conjunct to be true, you need to have conclusive reasons for believing that's a zebra. And you do. You wouldn't have had a visual experience like that had you not been looking at a zebra. The second conjunct is true because you do not have conclusive reasons for thinking that's not a cleverly painted mule. Had you been looking at a cleverly painted mule, your experience would have been the same.

B: Got it.

F: In order for you to know the conjunction *I know that's a zebra but I don't know it's not a cleverly painted mule*, you would have to know the first conjunct. Knowing the first conjunct amounts to *knowing that you know* that's a zebra. Knowing that one knows is a form of inoculation against skeptical

12 "Moonronic" is Laurence Goldstein's label for a Moore-paradoxical statement. For more, see his *Clear and Queer Thinking: Wittgenstein's Development and His Relevance to Modern Thought* (Lanham, MD: Rowan and Littlefield, 1999).

challenges.¹³ You know that you know that's a zebra only if you know that your reasons for believing that's a zebra are conclusive.

B: Don't I know that my reasons are conclusive?

F: No. If you knew that your reasons were conclusive then you would be in a position to rule out that being a cleverly painted mule. And you aren't. Neither am I.

B: So knowing that P requires having conclusive reasons for P. But knowing that P does not require *knowing* that I have conclusive reasons for P. The reasons just have to be conclusive, whether I know they are or not. *Knowing that I know* that P, on the other hand, does require knowing that my reasons for P are conclusive.

F: Exactly. And thus you can know without being in any position to know that you know. Now, to generalize, Frediculous conjunctions will be true only in cases where subjects lack second-order knowledge. "You know that's a zebra but you don't know that's not a cleverly painted mule" will be true only in cases where you don't know that you know that's a zebra. This is because knowing that you know that's a zebra entails knowing that your reasons are conclusive. So if you know that you know that's a zebra, then you *will* know that's not a cleverly painted mule. But this means that Frediculous conjunctions will be true only when the subjects they concern do not know them to be true. This is why those conjunctions always sound so awful when we assert them.

B: Even though you've been asserting them all day.

F: Like I said before, good philosophy will always involve asserting inappropriate things. Why do you think they killed Socrates?

B: Something else is bothering me here, Freddie. I get that you think S knows that P only if S believes that P on the basis of R and R is a conclusive reason for P. What I'm not seeing is why you think knowing that I know that P requires knowing that R is a conclusive reason for P. Look at it this way. According to you, S knows that P whenever S has a true belief that P grounded upon a conclusive reason. Let S stand for me and let P stand for the proposition *I know that's a zebra*. Not the zoological proposition *that's a zebra* mind you, but the epistemic proposition *I know that's a zebra*. Your definition of knowledge entails that I cannot know the epistemic proposition unless I have conclusive reasons for it. The definition does not entail that I cannot know

Does this commit
Freddy to the
Knowledge Account of
Assertion?

13 Dretske says this in "Conclusive Reasons," *Australasian Journal of Philosophy* 49(1): 1-22, 17 (1971).

TELL ME SOMETHING I DON'T KNOW

the epistemic proposition unless I know that I have conclusive reasons for the zoological proposition.

F: But you will have conclusive reasons for the epistemic proposition only if you know that your reasons for believing the zoological proposition are conclusive.

B: Why would that be? Take my actual situation. Suppose that my visual experience is my reason for believing not only *that's a zebra* but suppose also that, on the basis of that same experience, I believe *that I know* that's a zebra. It is natural to think I'd do this. Had I said "I know that's a zebra" and you asked "What makes you think you know that's a zebra?" I'd say "I'm looking right at it." My perceptual experience is grounds for believing both the zoological proposition and the epistemic one.

F: But your visual experience alone will not be a conclusive reason for believing that *you know* that's a zebra.

B: Why not? Let's apply the test. R is a conclusive reason for P if and only if R would not be the case unless P were the case. So, when we are wondering whether my experience constitutes conclusive reason to believe the proposition *I know that's a zebra*, the relevant question is this: Could I have had that experience had I not *known* that's a zebra? One way for me to fail to know that's a zebra is for it to be one of those cleverly painted mules. And, in that case, I will have the same experience I had when I was looking at that zebra. But of course, that is not the only way I can fail to know that's a zebra. And, as we agree, that is one of the more far out ways for me to fail to know that's a zebra. There are lots of other ways I might have failed to know that's a zebra. I might have failed to know that's a zebra if I hadn't bothered to come over here and look at it or, if the zookeeper had decided to give him the day off and keep him inside, or if I hadn't come to the zoo at all today, or if he had escaped his pen last night and wandered into the lion den and been eaten. All these ways I might have failed to know that's a zebra are much closer possibilities than the cleverly painted mule scenario. And notice that in each of these situations (and in every other plausible way of failing to know that's a zebra that I can think of), I do not have the sort of visual experience I had when I was looking at that zebra. This shows that if I had failed to know that's a zebra, I wouldn't have had that visual experience. I wouldn't have had that experience unless that *was* a zebra and I wouldn't have had it unless I *knew* that was a zebra. My visual experience is therefore a conclusive reason for believing that I know that's a zebra.

F: So you are saying that you can know that you know that's a zebra without knowing that your reasons for believing that's a zebra are conclusive.

B: That's what you say, Freddie. Or at least that's what your own view on the nature of knowledge says.

F: How's that?

B: Look at it this way. Earlier, you were assuming that the only way I could know that I know that's a zebra is to know that I have conclusive reasons for believing that's a zebra. Now that might be *one way* to know that I know but, given your own view on the nature of knowledge there's no reason why it should be the only way. As long as I have some reason R such that R would not be the case unless "I know that's a zebra" is true, then R is a conclusive reason for "I know that's a zebra" and I can on the basis of R know *that I know* that's a zebra.

F: And you think your perceptual experience, the same thing that serves as your reason for believing the zoological proposition, also fits the bill for being a conclusive reason for the epistemic one.

B: Yep.

F: Well I'm still not quite sure that your visual experience passes the test for being a conclusive reason for knowing that you know that's a zebra. Counterfactuals like this are a slippery and shifty business.¹⁴ And some of what you said there kind of sounded like you might be guilty of backtracking.¹⁵ I was thinking the same thing about some of the examples you were using earlier. In other words—Oh Look! There's a sign that says "MULE EXHIBIT THAT WAY" with an arrow pointing left. Let's go.

B: Of course, we don't know that the mule exhibit isn't somewhere else and the sign is wrong but we know the mule exhibit is over that way.

F: That's right, Betty. A signal can provide information but it cannot provide the information that it isn't misrepresenting things.

B: But if you look closely at the bottom of the sign you'll notice it says "AAAOK Zoo Signs, Inc. Your source for accurate zoo signage since 1976."

F: Uh huh.

Is this true? If so, does it mean we can never know that sense perception accurately represents reality?

14 Dretske says this in his "Reply to Hawthorne," in *Contemporary Debates in Epistemology*, 44.

15 A man who hates his boss has a conversation with him where he pretends to like him.

Immediately afterwards, he thinks to himself, "If I had said what I really think of my boss a minute ago, I would have been fired." False! He would never tell his boss what he thinks of him to his face. If he had said what he really thinks of his boss, his boss would not have been in the room. This is an example of backtracking. For a discussion see David Lewis, "Counterfactual Dependence and Time's Arrow," *Noûs* 13(4): 455–76 (1979).

TELL ME SOMETHING I DON'T KNOW

The two begin to walk toward the mule exhibit.

B: Something else just occurred to me, Freddie. Take another look at the belief I arrived at back there at the zebra exhibit. I believed that the animal I was looking at was not a cleverly painted mule. You say I don't know this because my reason for believing it was my visual experience and I would have had that visual experience even if I had been looking at a cleverly painted mule.

F: That's right, Betty. Your reason for thinking that you were not looking at a cleverly painted mule is inconclusive.

B: But what's wrong with this? Suppose I say that my reason for believing that's not a cleverly painted mule is not my visual experience but my belief *that it's a zebra*. It certainly seems like I could believe *that's a zebra* and from there infer *that's not a cleverly painted mule*. In fact that's pretty much what I did back when we first started talking about this stuff. I formed a belief that that's not a cleverly painted mule via an argument whose key premise was *that's a zebra*. Now *that's a zebra* would not be true unless it were not a cleverly painted mule. In other words, if it were a cleverly painted mule, it wouldn't be a zebra. Therefore, *that's a zebra* is a conclusive reason for believing that it's not a cleverly painted mule.

F: Nice try, Betty. But look, your reason for believing *that's a zebra* is your visual experience. Right?

B: Yep.

F: If your visual experience is your reason for believing *that's a zebra* and *that's a zebra* is your reason for believing *that's not a cleverly painted mule*, then your visual experience is your reason for believing *that's not a cleverly painted mule*. And we've already agreed that your visual experience is not a conclusive reason for believing that's not a cleverly painted mule. Therefore, you do not have a conclusive reason to believe that's not a cleverly painted mule.

B: You are assuming that reasons are transitive. If R is a reason for P and P is a reason for Q then R is a reason for Q. Why not reject transitivity?

F: I do reject transitivity. I think your perceptual experience is a conclusive reason for believing that's a zebra. And if it's a zebra then it's not a cleverly painted mule. But your perceptual experience is not a conclusive reason for believing it's not a cleverly painted mule.

B: I'm not talking about the transitivity of conclusiveness. I am talking about the transitivity of reasons simpliciter or, if you prefer, the transitivity of the basing relation. I'm saying my basis for believing that's not a cleverly painted mule is

that it's a zebra. And my basis for believing it's a zebra is my perceptual experience. But it follows that my basis for believing it's not a cleverly painted mule is my perceptual experience only if we assume that the basing relation is transitive. So why not reject that? Why not say that my reason for believing that's not a cleverly painted mule is just that it's a zebra and stop there? And look at what we gain if we do. We get a way to refute skepticism. Not only is *that's a zebra* a conclusive reason for *that's not a cleverly painted mule* but *I have hands* is a conclusive reason for *I am not a handless brain in a vat*. It would not be true that I have hands unless it were true that I am not a handless brain in a vat.

F: The Principle of the Transitivity of Reasons is an obvious epistemic principle, Betty. You can't go around denying obvious epistemic principles just because it gives you a nifty way to defeat skepticism.

B: Is that right?

F: That's right, Betty. And I know what you're smirking about but I don't deny closure just because it gives me a nifty way to defeat skepticism. I deny closure because of general facts about the nature of knowledge.¹⁶

B: But one of these "general facts about the nature of knowledge" is supposed to be the idea that knowledge requires conclusive reasons. This fact about knowledge does not get you a denial of closure. The most it shows is that either closure fails or reasons are not transitive.

F: Well I suppose that if you could give me some independent argument against the transitivity principle,¹⁷ I might—Wow! There it is. Look at that mule. Beautiful.

B: It sure is, Freddie. It sure is.

16 Dretske, "Reply to Hawthorne," 43.

17 For such an argument, see John Post's "Infinite Regresses of Justification and Explanation," *Philosophical Studies* 38(1): 31–52 (1980).