{25} WEDNESDAY

WEIROB: You can serve yourselves some fresh mushroom soup from the stove. Mushrooms are supposed to be good for clear thinking. Or is that avocados? Whatever. I baked some cookies, too. Dig in.

MILLER: I'm glad I'm not a zombie, so I'll enjoy the taste of your lunch.

WEIROB: Huh?

COHEN: Sam and I have been talking about zombies.

WEIROB: Zombies? Good Lord. Do you mean like zombies and vampires? We're going to do the philosophy of zombies and vampires? Perhaps discuss the fine conceptual points that determine whether something is a zombie or a vampire? I think zombies don't drink blood and vampires do. Is there more to be said? Give me a break!

MILLER: No, not *that* sort of zombies. We've been talking about *philosophical* zombies, an idea Dave and I came up with while discussing Ewing's experiment. We came up with two arguments for property dualism based on Ewing that we call the "the zombie argument" and "the knowledge argument."

{26} WEIROB: Zombie argument! That's a bit eerie. And intriguing. Let's look at it first. Tell me more.

COHEN: Zombies are beings exactly like us—that is, exactly like us in every *physical* respect—but without experiences, without consciousness. If there *could* be such beings, if they are possible, if their existence does not involve a contradiction, that would show that experience and consciousness are not physical, wouldn't it? We would have to be property dualists and believe that key features of our minds are not just physical.

WEIROB: But there aren't any zombies. Your philosophical zombies are no more real than ordinary zombies. You are giving me science fiction. How can you make a philosophical point with science fiction?

COHEN: That's not quite fair, Gretchen. Philosophical views have consequences for what's *possible.* If a philosophical theory says something is possible that isn't possible, or something isn't possible that is possible, that's a good argument against it.

You used this kind of argument once in class. You were talking about a Humean theory, that something is morally wrong if normal, well-brought-up people who think of someone else doing that thing have a feeling of aversion. You said that we can conceive of a world in which normal, well-brought-up people think torturing babies is just fine. But then it's possible that torturing babies is not morally wrong. But, you said, it is not *possible* that torturing babies is

anything but morally wrong. So we must reject Hume's theory.

{27} WEIROB: Did I say that was a definitive argument?

COHEN: You never say that anything is a definitive argument for anything. But you did say something like this. Historical accounts have to be rejected if they entail that things didn't happen that did happen. Scientific accounts have to be rejected if they say that things wouldn't happen in the actual world under particular circumstances, but they would. And philosophical accounts have to be rejected if they say that something is *not* logically and conceptually possible when it *is* logically and conceptually possible. That's just the principle we are using.

WEIROB: Okay. I admit the form of argument works. According to my view, physicalism, it is impossible for there to be beings that are physically just like us but have no experiences —what you are calling "zombies." But you are saying that it is possible for there to be such beings. So we must reject my theory.

I admit, that's a good form of philosophical argument. If I accept your premises, I must accept your conclusion. So I'll do my best not to accept your premises.

MILLER: Dave thinks there must be something wrong with the argument, but I think it's a great argument. So we thought we would find out

what you think—although I can already make a pretty good guess.

{28}
WEIROB: I think it's a fine argument, because it's an interesting argument. But I don't think it can be right. On that, I'm with Dave. In the end, I don't think we will find that zombies should convince us that our minds are not simply physical.

MILLER: Are you ready to explain why?

WEIROB: First of all, let's get a bit clearer by what we mean by "consciousness."

MILLER: Here we go again. I would've thought that consciousness is consciousness and that's that. But I suppose that's too simple.

WEIROB: We need to make sure we agree on what Ewing thought the materialist *would* have when she grasped the red-hot poker—and what your zombies *wouldn't* have, even if they grabbed it. Ewing is thinking about what Herbert Feigl calls "raw feels." Some philosophers call them "qualia." But I don't know whether it's pronounced "qu-AH-lia" or "qu-AI-lia." So I prefer "raw feels."

MILLER: I think "raw feels" sounds kind of sleazy. Why don't we just call them experiences?

WEIROB: Fine with me. The main point is, as Thomas Nagel* puts it, it's *like* something to have a pain, or see a tree, or think about whether to get out of bed. That's the raw feel—I mean the experience. Nagel says that bats must have

experiences when they fly around guided by sonar. But that's a kind of experience humans don't have, and so we {29} can't know what it's like to be a bat. Even if we dissect a bat, we won't know. But we know what it's *like* to feel pain and how that is different than what it's *like* to taste a chocolate chip cookie. It doesn't *seem* like this is just a difference in which neurons are firing in my brain. People know a lot about the experiences they are having, while knowing next to nothing about what the neurons in their brains are doing.

MILLER: Okay, so far so good. I have experiences, and you're at least admitting that having an experience doesn't seem like just having various chemical and electrical things and subatomic stuff happening in one's brain. So what else might we mean by "consciousness"?

WEIROB: We naturally take humans to be the paradigms of consciousness. But there is a lot more to a human mind than just experiences. We have reason, or at least most of us have some reasoning abilities. We have self-consciouness. As Locke says, a person is aware of itself, as itself, in different times and places. And we have some kinds of higher-order consciousness. When Ewing's materialist grabs the red-hot poker, she will feel a pain, but she will also be aware she is feeling the pain. And if she is in a philosophical mood, she will have the pain, be aware of having the pain, and also be aware of being aware of having the pain. And so on. That's what human consciousness involves.

MILLER: Of course. What's the point?

{30}
WEIROB: I assume you will admit that dogs and frogs and perhaps even worms have pains and pleasures, and that it is like something for a dog to see or smell or hear things.

MILLER: Sure.

WEIROB: Well, I think so, too. But it's not obvious that they have much if any of the other things characteristic of human consciousness.

COHEN: If dogs aren't self-consciousness, does that mean your dog doesn't know who she is? She certainly seems to. When she's hungry, she doesn't seem to have any doubts about who is hungry. When she buries a bone, in order to dig it up and chew on it the next day when it is a little more interesting, isn't she making plans for her own future—sort of?

WEIROB: Hey look, I love my dog Penelope. I think she's great. She deals with life better than most people I know. And I'm sure that she knows who she is, in *some* sense. When she feels hungry, she makes sure to get the food in her own mouth. But I doubt very much it's anywhere near as complicated and layered as human self-consciousness. And I would be even more certain of the point if I had a pet frog or a pet worm. We each have a conception of ourselves as one human among others, with a long past and hopefully a long future. We think about ourselves, criticize ourselves, try to improve ourselves, and the like.

MILLER: Penelope is a very cute dog, but why are we discussing dogs?

{31}
WEIROB: I'm just trying to nail down the point that it's experiences we are worried about. Consciousness, in the sense of having experiences, may be *necessary* for all the complicated things involved in human consciousness. But it doesn't seem *sufficient.*

COHEN: I think I get your point. In the sense of relevant sense of "consciousness" a person who is asleep and dreaming is conscious. A senile person who no longer seems to be conscious of who they are or where they are or much of anything else is still conscious in this sense. Anything that has experiences is conscious.

WEIROB: Here's another way of making the point. Computers can do a lot of things, and every year they are able to do more and more things. We can imagine robots made to look like humans that we could talk to in the way we talk to other humans. It seems they can do, or soon will be able to do, lots of things that we consider part of human consciousness. They have sensors to determine what the world around them is like. They can compute, analyze information, and make inferences—the things we regard as thinking in humans.

COHEN: Well, some philosophers, like John Searle*, would say that robots can't even do that. They can't compute, analyze, make inferences, or do anything that requires understanding. They do things that *simulate* those human activities,

and that makes it convenient to use the same words we use for humans to describe what they are {32} doing. People in artificial intelligence want to say they have made computers that understand language, think, and reason. But they are exaggerating their accomplishments, impressive as they are. Philosophically speaking at least, we shouldn't use those words.

WEIROB: Good point, Dave. Maybe we can talk about Searle and artificial intelligence on another occasion. But for now, let's go along with the AI folk and say robots can think, reason, and understand, and can do all of those things in some reasonable sense. My point is that, even so, there still seems to be a big question: Would such robots, with computers for brains, really *experience* anything? When we put this as, "Would intelligent robots be *conscious*?" we're using "consciousness" in the way I'm getting at.

COHEN: So our zombies are sort of like robots, if we assume that robots *don't* have experiences. Now when you look inside a robot, you don't find anything like a human brain, but instead a computer or computers. They are silicon-based, not carbon-based. Perhaps that's a good reason to at least doubt that they have experiences. When you look inside a zombie, however, you find a human brain. Even so, Sam and I are suggesting that they don't have experiences. One might want to say that they compute and reason and plan and all of those things—we say those things about computers.

But they don’t have experiences, and so, in the sense we are using the word, they are not conscious.

{33} MILLER: Great, we understand, so let’s get on with it. If the pain Ewing’s materialist feels when she grabs the red-hot poker involves nonphysical properties, I win. If there could be zombies with brains just like ours but no experiences, I win. If there could be zombie dogs just like Penelope, or zombie frogs just like the ones I hear croaking at night,
I win.

COHEN: And I guess if we have a robot that has sensors and a great computer brain, seems to be aware of its surroundings, makes complicated plans, and writes a moving robot autobiography, it still wouldn’t be *conscious* unless, in addition to all of that, it had experiences.

WEIROB: Right.

MILLER: Yes, right, right, agreed. Are we done with distinctions and terminology? Can we start discussing the zombie argument?

WEIROB: Okay. Let’s look at the the argument. Explain it to me again.

MILLER: I’ll let you do the honors, Dave.

COHEN: It’s pretty simple. Consider an imaginary but possible world that is physically indiscernible from ours. “Indiscernible” here doesn’t merely mean *looks the same,* but *is* exactly similar.

But in this world, the zombie counterparts of humans have no experiences. It's not *like anything* to be a zombie. Their minds are blank. They don't have any *raw feels*. Excuse me, *experiences*. We say {34} that computers "think" and robots "perceive" things. So we could say the same about zombies, if we wanted to. But, for them, it's not *like anything* to "think" and "perceive." No experiences are involved.

WEIROB: Suppose I succeed in imagining such a world, and admit that it is possible. What follows from that?

COHEN: Every physical event that happens in our world happens in the zombie world, too. But there are no experiences.

WEIROB: And . . . ?

COHEN: Since all the physical events are there, but the experiences aren't, experiences are not physical events. If the physical events and the physical properties of the brains are all there, but the experiences aren't, then experiences must not be physical properties. They may be properties of brains, but they have to be nonphysical properties, or they couldn't be missing in the zombie world. That's the argument.

WEIROB: Well, as I said, that is an interesting argument. An intriguing argument, even.

MILLER: How about a *convincing* argument.

WEIROB: I think I see a problem. Since your zombie world is physically indiscernible from the real one, there are three of these zombies sitting around eating lunch in a house just like this one. Words {35} come out of their mouths—utterances are surely physical events—the same words that are coming out of our mouths as ours. But they aren't thinking. Right?

COHEN: They may be thinking the way computers think. But there is nothing going on inside the way there is when we think. When we think, it's *like* something. What it's like to think is much different than what it's like to eat a chocolate chip cookie, or have a toothache or Ewing's throb of intense pain. As experiences go, it's a bit subtle. Hume would say that what happens when we think and imagine and the like is less "lively and vivacious" than what happens when we perceive. But still it's like something to think.

MILLER: Well put, Dave. Ready to concede and become a dualist, Gretchen?

WEIROB: Well, maybe not quite yet. I need to know a little bit more about what it means for a world to be physically indiscernible from ours.

MILLER: How is that a problem? For every physical event that occurs in our world, the actual world, one exactly like it occurs in the zombie world. But that's all the events there are in the zombie world. So, as far as brains go and all other physical things, and all their physical properties are concerned, it's exactly similar. The only difference between the zombie world

and the real one is the missing experiences. Is that clear enough for you?

{36} WEIROB: Well, actually, no. I'm still not quite clear about what is going on in the zombie world. Let's try a couple of examples, to make sure I understand.

MILLER: Fine.

WEIROB: Dave, you just took a spoonful of the mushroom soup I made. How was it?

COHEN: Delicious, Gretchen. You are a great cook. Thanks for making lunch.

WEIROB: Now you said it was delicious. You made the sound, "delicious." That was a physical event, right? And a different physical event than if you had said, "Not the worst I've ever tasted," or "Completely foul."

COHEN: Quite right.

WEIROB: Now, assuming you are being truthful, is it fair to say that your response was due, at least in part, to how the soup tasted to you? That is, to what *experience* you had when you ate it?

COHEN: I was being truthful. I said it was delicious because of its taste—that is, because of the experience I had when I put it in my mouth and swished it around and swallowed it.

WEIROB: So you are saying that the cause, or at least part of the cause, for the physical event of saying "delicious" was the experience you had when you ate the soup?

{37} COHEN: I think I see where you are going. But yes, part of the cause of my saying “delicious” is how it tasted, the experience I had.

WEIROB: Of course there was a lot more involved. You were responding to a question I asked. Your tongue and mouth muscles were all working. You are disposed to tell the truth. All sorts of things. But we can say, can’t we, that the experience was a *necessary part* of the totality of conditions that were sufficient for you to say “delicious.”

MILLER: Slow down, you’re losing me.

COHEN: A sufficient condition for something to occur is a condition—usually a complex condition or set of conditions—that is enough to bring about the occurrence.

MILLER: Yes, I understand that.

COHEN: A necessary condition for something to occur is a condition without which the event wouldn’t happen.

MILLER: I understand that, too. But when Gretchen talks about a necessary part of a sufficient condition, I get a little lost.

COHEN: She’s getting at the point that what we usually call the “cause” of an event isn’t usually, as one might think, *all* of the conditons that are necessary and sufficient for the event to happen.

{38} Let’s take another example. Suppose a

wheel fell off a boxcar while the heavily loaded train it was part of was going around a curve. The boxcar leaves the rails, and then the whole train goes off the rails.

MILLER: It happens from time to time, even on the Rock Island Line.

COHEN: We'd say the wheel falling off the boxcar was the cause of the derailment. But the wheel falling off wasn't a *sufficient* condition for the derailment. If the train hadn't been moving it wouldn't have derailed. Perhaps if it hadn't been going around a curve, it wouldn't have derailed. Perhaps if it had been more lightly loaded, or going a lot slower, it wouldn't have derailed.

MILLER: I see. The wheel falling off all by itself was not sufficient for the derailment. It was the wheel falling off, *together* with the train going a fair speed, and going around a curve, and being fully loaded, that was sufficient.

COHEN: Right. Now notice that the wheel falling off isn't a *necessary* condition for the derailment either.

MILLER: Well, it was as I understood the example. I assumed that the train could have made it around the curve without derailing had the wheel not fallen off.

COHEN: Quite right. But the train could have derailed due to a different set of conditions. Perhaps someone had removed a rail. Or set off a bomb just as {39} the train passed over it.

The sufficient condition for the derailment would be different, and not involve the wheel falling off.

MILLER: Okay, it's not necessary, in an unconditional sense. But *given* the way the derailment happened, if you leave everything the same, and the wheel doesn't fall off, it doesn't derail. So it's necessary in that way.

COHEN: Exactly. And that's what Gretchen means by "a necessary part of a sufficient condition." The sufficient condition is the whole combination of things—going around a curve, heavily loaded, moving at a fair rate of speed. If you take away one of them, no derailment. So, they are each necessary parts of the sufficient condition that actually led to the derailment. At least, that's what I think she meant.

WEIROB: You've got it just right, Dave. Well explained. By the way, we philosophers have a term for this. We call things like the wheel falling off an "INUS condition."

MILLER: Huh?

WEIROB: That's short for an *i*nsufficient but *n*ecessary part of an *u*nnecessary but *s*ufficient condition. The actual sufficient condition—speed, curve, and missing wheel—isn't the only way the derailment could have occurred. So it's not a necessary condition of the derailment. It's an *u*nnecessary but *s*ufficient condition.

{40} MILLER: I get it. The wheel falling off, as Dave pointed out, is *insufficient* by itself for the derailment—

the train might have been going slow on a straight stretch of track. But it is a *necessary* part of the actual sufficient condition. So it's an INUS condition. Great, a new addition to my philosophical vocabulary.

COHEN: The idea comes from Mackie*, right?

WEIROB: Exactly. Brilliant fellow, Mackie.

COHEN: Mackie's point was that in a case like the train derailment, we would ordinarily say the wheel falling off *caused* the derailment. What we ordinarily call "causes" of an event are usually not *total* causes—not the whole complex of conditions that suffices for the event, but INUS conditions.

MILLER: This is all quite fascinating, I guess. But can we get back to the zombie argument now? Does anyone remember where we were?

WEIROB: I was getting Dave to admit that the experience he had, when he slurped up his soup, was an INUS condition for the physical event of his producing the sound "delicious."

COHEN: I never denied it. And I wasn't slurping. But, aside from that, that's where we were.

WEIROB: So then the experience Dave had, the pleasant experience of tasting perfectly concocted mushroom soup, was a cause, in the INUS sense, of his saying "delicious."

{41} MILLER: Okay, okay. We grant you that. Get on with it.

WEIROB: Patience, Sam. My point is that in the zombie world, that pleasant experience doesn't occur. But, since the zombie world is physically indiscernible from ours, in that world the physical event of zombie-Dave producing the sound "delicious" occurs, just as it does in the real world.

MILLER: Of course.

WEIROB: But *why* does it occur? If we remove the cause, or an essential part of the cause, don't we have to remove the effect, too?

MILLER: Well, it just happens.

WEIROB: But if he didn't have the pleasant experience, what accounts for his saying "delicious" rather than "foul" or something else? Or nothing?

MILLER: I don't see a contradiction there. You haven't shown that the zombie world isn't possible.

WEIROB: But if *nothing* causes him to say one thing rather than another, doesn't that mean that the zombie world isn't physically indiscernible? If it's physically indiscernible, shouldn't the *same* physical events have the *same* causes in the zombie world that they do in the real world?

MILLER: Maybe the cause is the same, and it is purely physical in both worlds. It just *seems* to us that experiences cause physical events. Maybe the {42} zombie world shows that Dave's experience wasn't just an INUS condition even in the real world—it just seemed like it was.

WEIROB: So you are willing to embrace *epiphenomenalism*?

MILLER: Geez, more terminology. Embrace epi-whatism?

COHEN: Epiphenomenalism. In general, an epiphenomenon is an event that is caused, but that doesn't have any further effects, or at least none that are of interest. In philosophy, epiphenomenalism is the doctrine that experiences are just epiphenomena. That is, they have physical causes, but no physical effects.

WEIROB: Epiphenomenalism is a way of accommodating three things that seem plausible, at least to some people. First, property dualism. Experiences are nonphysical brain events. Second, physical events in our brains determine which experiences we have. Third, the physical world is a closed system. That means that physical events have only physical causes. That doesn't prevent them from having nonphysical effects. That's an assumption that has guided a lot of science for centuries.

COHEN: The epiphenomenalist picture goes like this. The soup being in my mouth, and its chemistry and the way my tastebuds work, caused a chain of physical events, including my saying "delicious." That physical process *also* caused me to have an experience. But, contrary to how it seemed to me, that experience *wasn't* the cause, or even part of {43} the cause, of my saying "delicious." Experiences don't

cause physical events, so the assumption that the physical world is “closed” is safe, even if physical events cause experiences. The experience wasn’t an INUS condition, as it seemed, but merely an epiphenomenon. So, when we leave experiences out of the zombie world, we aren’t leaving out any causes of physical events. So we don’t have to leave out any effects.

MILLER: Well, if I understand what you two are saying, given epiphenomenalism, the zombie world *can* be physically indiscernible from the actual world. So the zombie world makes sense. I don’t see how Gretchen can claim that the failure of physical events to cause nonphysical events constitutes a *physical* difference.

WEIROB: I’ll agree to that. I agree that *if* we accept epiphenomenalism, then the zombie world can be physically indiscernible from the actual world.

MILLER: So, don’t I win?

WEIROB: Not if you’re trying to convince me of dualism. You’d first have to convince me of epiphenomenalism. The zombie argument works only if we add epiphenomenalism as an extra premise. And I don’t accept the extra premise. Epiphenomenalism seems crazy to me.

MILLER: Crazy? That’s kind of harsh.

WEIROB: Okay, I retract “crazy.” I’ll just say extremely implausible and unattractive.

{44} MILLER: Right this minute it seems plausible and attractive to me.

COHEN: Look, Sam, do you really want to accept that the pains and joys we experience, the experiences we have when we see or feel or taste something, have no effect on what we do? No effect on the physical actions we perform—movements of our limbs of our tongues and vocal chords, as when I said “delicious”? Isn’t that a high price to pay for accepting the zombie argument?

MILLER: When you put it that way, I must admit I’m not too enthusiastic about accepting epiphenomenalism. But isn’t there some evidence for it?

WEIROB: Such as?

MILLER: Suppose after Gretchen took the soup off the electric burner, while the coil was still red hot, I put my hand on it.

COHEN: Back to Ewing?

MILLER: But I’m making a different point. I would feel pain, and I would remove my hand from the burner very quickly. But is it so clear that the pain *causes* my hand to move? Suppose some physiologist does a careful study and shows that the movement of my hand actually begins a nanosecond or two *before* I feel that pain. I think I’ve read about studies along these lines. The physiologist concludes that the pain doesn’t cause the movement, but that the effect of the heat on my hand causes both the

pain {45} and the movement, in such a way that it *seems* that the pain caused the movement. I don't see anything to show that he wouldn't be right. My pain would be an epiphenomenon, in your terminology.

WEIROB: Okay, I'll grant you that. But suppose that when you touched the burner you were a small child. You didn't know why the burner was red. You put your hand on it. You felt pain and removed your hand. We grant that pain doesn't cause the removal of your hand. But won't it have *other* physical effects? Won't you be much less likely to put your hand on a red burner in the future? Suppose when Dave ate the soup, the taste he experienced was foul. Wouldn't he be much less likely to take another spoonful of the soup? Even if we grant that experiences don't have all the physical effects we think they do, don't we have to suppose that they have *some* effects on our physical brain that affects our future physical behavior?

MILLER: That's a good point, I admit. It seems like having one experience rather than another sometimes affects future behavior. It's hard to see why that would happen if the experience doesn't affect the brain somehow. Hard, but not impossible. The physical events that cause the pain could also cause a physical state that prevents us from doing the same thing again.

COHEN: But think about it, Sam. Insofar as having a profound effect on what we think being a human being amounts to, wouldn't the truth of

{46} epiphenomenalism be an even *bigger* blow than the truth of physicalism? We see ourselves as beings that do what we do because of the experiences we have—not everything we do, but a lot of it. If epiphenomenalism is true, this is just a big illusion. I'm not finishing this soup because it tasted so good. Gretchen isn't talking to us because she enjoys talking philosophy. You don't go to the beach on summer days because of how good it feels to lie in the sun. And so on. It seems to me a very dismal view of human life.

MILLER: Well, I guess I can't deny that. So where does that leave us?

WEIROB: I think it leaves us here. If epiphenomenalism is true, then the zombie world is possible. The absence of experiences doesn't imply that there is a physical difference between it and the real world. But if you don't accept epiphenomenalism, removing experiences removes some of the causes of physical events. So removing the experiences from the world will make it physically different. All the things for which experiences are INUS conditions won't happen.

So maybe the zombie argument should persuade epiphenomenalists to become dualists. But it won't persuade me, because I'm not an epiphenomenalist. So I don't think you can remove experiences without removing their effects, which means that your zombie

world will *not* be physically indiscernible from the real world.

{47} MILLER: You've given me something to think about. But I'm not convinced. It's getting late. I need to meet with some parishioners over at the church. We need to continue with epiphenomenalism when we meet Friday. By then I hope to have figured out some sort of response. Or maybe Dave will.

COHEN: And remember, Sam, we've still got our knowledge argument to discuss. So I'll see you both Friday? Can we count on you for lunch again, Gretchen?

WEIROB: I think so. I'm in all the right mental states to make another nice meal. But if epiphenomenalism is true, who knows what will happen?

MILLER: Ha, ha. So long.