

Day 3: Consciousness and rational belief

Setting: M and V sharing a third wonderful vegan meal.

M: Hey V, I think I figured out why human pain matters more than animals’.

V: Do tell.

M: Okay, so there are positive and negative mental states, right? Pleasure, happiness, and other forms of enjoyment are positive. Pain, unhappiness, and other kinds of suffering are negative.

V: Sounds right.

M: But mental states, in general, can be conscious or unconscious. You can have conscious or unconscious desires, beliefs, and even emotions.

V: All this is well known.

M: Here’s the interesting part. Conscious versus unconscious isn’t a binary distinction; it’s a matter of degree. Mental states can be *more* or *less* conscious, not simply conscious or unconscious.

V: So you can have a semi-conscious belief or desire, one that you’re only half aware of?

M: Right. Now, when it comes to negative states, the less conscious they are, the less bad. If there could be a completely unconscious pain, then it wouldn’t be intrinsically bad at all.

V: Maybe. But I’m not sure the idea of an unconscious pain makes sense. Unconscious belief, sure. But unconscious pain?

M: But there can be *more* or *less* conscious pains. Say you have a back pain. I decide to distract you from it by engaging you with a delightful philosophical paradox. You start to pay less attention to the pain, and so you start to become less and less conscious of it. After we’ve been arguing for a half hour, you’ve forgotten about your pain.

V: Yeah, I’ve had experiences like that.

M: And my distracting you would be a good thing, right?

V: Fair enough. So you’re saying that the pain becomes less *bad* as it becomes less *conscious*. Where are you going with this? You’re not going to claim that animal pain is always unconscious, are you?

M: No, but it might be *less* conscious than typical human pain. Animals have much less self-awareness in general than humans. Some people doubt whether animals are self-aware in general. So it’s plausible that all their mental states have only a low level of consciousness. They’re only dimly aware of the things they are aware of. In that case, their pain would be less bad than typical (fully conscious) human pains.

V: Interesting theory. This is the first time I’ve heard someone give an explanation for why animal pain matters less than human pain that makes sense. Usually, you guys pick on arbitrary properties, like IQ.

M: So you don’t think level of consciousness is morally arbitrary?

V: No, that really seems to make a difference to the intrinsic value of an experience. This time, your moral claim is actually intuitive.

M: At last, you’ve conceded that I made a good point!

V: Yes, but let’s explore a few implications of this theory. First, say we have a newborn baby . . .

M: Oh no, it’s back to the infants and retarded people again?

V: Well, they seem to have similar cognitive capacities to nonhuman animals, so it’s good to test our

intuitions on them, to make sure we're not influenced by mere bias against other species.

M: Oh, I'm sure I'm not biased against other species. I just have a rational assessment of their capacities.

V: I'm not convinced that infants or severely retarded people have more self-awareness, or more "consciousness" as you say, than animals. So would it be okay to torture infants and retarded people?

M: That doesn't seem right. Maybe their experiences are all still fully conscious, even though they have much lower general intelligence.

V: Maybe. Or maybe not. Do you want to rest the ethical treatment of these people on that speculation?

M: Hm. Well, since we're not sure of their level of consciousness, I would say it makes sense to err on the side of caution and not inflict needless suffering on them.

V: That sounds completely reasonable. Similarly, since we're not sure of the level of consciousness of nonhuman animals, it makes sense to err on the side of caution and not inflict needless suffering on *them*.

M: I guess I'm saying that animals are less conscious than human infants, or at least less likely to pass any given threshold level of awareness.

V: Is there any evidence for that?

M: Maybe the fact that infants are *going* to be fully conscious later?

V: That's true, but it's also true that they *were* fully non-conscious in the recent past (at an early stage of fetal development). They're in transition from fully unaware to fully aware beings. I don't see why you should assume that, upon emerging from the womb, they're immediately more conscious than an adult animal. Who knows, maybe animals are *more* conscious than infants.

M: I admit, it's hard to say. I can't think right now of a way we could test for degrees of consciousness of someone's mental states.

V: So it looks like we should err on the side of caution and avoid hurting *any* of these beings – animals, infants, or retarded people – unless we have a very good reason.

M: Or maybe we should reason in the opposite direction. Maybe we should say that since we aren't sure of their level of awareness, we should discount the interests of all of these beings. We should give preference to normal, intelligent adults.

V: Maybe. But how much preference? Would you be willing to say that the pain of a normal adult matters a million times more than the pain of an infant?

M: What do you think I am, some kind of crazy extremist?

V: I'll take that as a no. Then you shouldn't take the crazy extremist view about animals either.

M: But do you agree that human pains are more important than animal pains?

V: I don't know, but it doesn't matter. It doesn't matter if a human pain is 50% worse, or twice as bad, or ten times as bad, as a similarly-caused animal pain. The amount of animal suffering we're causing each year is still vastly greater than the benefits we gain. Remember, we're torturing and killing *56 billion* land animals per year.

M: I remember. But why do you keep using that emotionally charged word, "torture"?

V: It's an accurate description. If a human being were confined in a tiny cage all day, forced to sit in his own excrement, forced to breathe ammonia, with a small part of his body having been cut off, you wouldn't hesitate to call it "torture." Of course the word has negative emotional connotations, because

the *phenomenon* that it refers to is awful. That doesn't make it an inaccurate or unfair description. Now, even if we discounted the torture of 56 billion beings by a factor of 1,000 due to their possibly lower level of consciousness, factory farming would still be the worst problem in the world.

M: (*sigh*) Okay, so even if my theory about degrees of consciousness is correct, I'm *still* obligated to give up meat.

V: That's about it.

M: But I don't understand why you're so obsessed with this problem. Shouldn't we first solve the enormous problems our own species faces, before we start worrying about other species?

V: What problems do you mean?

M: You know, like war, poverty, and disease.

V: We can work on all those things while at the same time being vegetarians.

M: Yeah, but you seem to spend a lot more time telling people about the problems with the meat industry than you do talking about those other problems.

V: That's because the problem of factory farming is much larger than all those other problems. It's estimated that the number of people who have *ever* lived on Earth is about 108 billion.³² So in just two years, the meat industry slaughters more animals than the total number of humans who have ever existed. Most of those animals endured great suffering before the slaughter. All this makes it plausible that factory farming over the past few decades has caused more suffering to animals than the total amount of suffering endured by all human beings throughout all of history.³³

M: So, if you could either end factory farming or achieve world peace . . . ?

V: End factory farming, hands down.

M: But war is terrible, you know. World War II, for example, killed *55 million* people.

V: And that's almost one thousandth the number of animals slaughtered on factory farms in one year.

M: I'm sorry, but this is making you sound like a crazy extremist.

V: Do you dispute my statistics?

M: No . . .

V: Is there something wrong with my reasoning?

M: I don't know what's wrong with it. But the idea that a year of animal farming is worse than the most destructive war in history, including history's most notorious genocide, just *sounds* to me so extreme and crazy on its face that it makes me want to say there must be something wrong with your argument.

V: And you think that's enough to reject the argument?

M: I do. I learned that from G.E. Moore: if you have an argument for a conclusion that seems crazy, you should reject it, even if you can't say exactly what's wrong with it.³⁴

³² Live Science, "The Dead Outnumber the Living," Feb. 7, 2012, <https://www.livescience.com/18336-human-population-dead-living-infographic.html>.

³³ Compare Rachels' argument that American factory farms over the last twenty years have caused at least five thousand times more suffering than the Holocaust ("Vegetarianism," p. 897).

³⁴ See, for example, G.E. Moore, "Hume's Theory Examined" in *Some Main Problems of Philosophy*, ed. H.D. Lewis

V: G.E. Moore was responding to philosophical skeptics who argue that no one knows anything about the world outside their own minds.

M: Right. The idea that I don't know, e.g., whether I have hands is so implausible on its face that if I hear an argument for that conclusion, I should infer that *some* step in the argument is wrong, even if I can't say which one or why.

V: And you think I'm like the philosophical skeptic.

M: Well, some of your views also sound crazy at first glance.

V: Okay, let's examine that reaction. Would you agree that *sometimes* we should accept conclusions that initially sound crazy?

M: I don't know. What do you have in mind?

V: Here's an example I once heard. Imagine that you have a very large but very thin sheet of paper, one thousandth of an inch thick. You fold it in half, making it two thousandths of an inch thick. Then fold it in half again, making it four thousandths of an inch thick. And so on. After folding it fifty times, how thick would it be?

M: I don't know. Let me get out my calculator. (*pulls calculator out of backpack*)

V: First just take an intuitive guess.

M: Um . . . ten feet?

V: Sounds reasonable. Most people will agree that the answer is something under a hundred feet. What would you say if I told you that the correct answer is over seventeen *million miles*?

M: That's crazy!

V: Yeah, it's crazy. But it's definitely correct. Enter it on your calculator. 0.001 inches, multiplied by two to the fiftieth power.

M: (*types on calculator, reads result*) 1.12×10^{12} .

V: That's the number of inches. To convert to feet, divide by twelve. Then to convert to miles, divide by 5,280.

M: You're right, it's almost eighteen million. But that's crazy.

V: Do you think your calculator is lying to you?

M: Of course not.

V: Do you think there must be something wrong with the argument because the conclusion is so crazy?

M: (*sigh*) No, it's correct. I'm not unreasonable, you know. It's just very surprising.

V: So sometimes we should accept conclusions that sound crazy.

M: Yeah, but that's a math problem. Ethical judgments are different.

V: Different how?

M: I'm not sure. Maybe because ethical premises are less certain and less reliable than descriptive, factual premises.

V: Wouldn't that mean that your sense of what is "crazy" in ethics is also less reliable?

M: I guess so. But my point is that your ethical argument is less reliable than a mathematical calculation or a scientific claim or an observation of the physical world. That's why the sense of "craziness" could be enough to defeat an ethical argument, even though it wouldn't defeat a mathematical calculation, scientific theory, or physical observation.

V: Perhaps. But before we conclude that, let's first try to figure out where the craziness comes from.

M: What do you mean by "where it comes from"?

V: Sometimes, we can identify the particular point in an argument where things become surprising. Take the example of the folded paper. First I say that the thickness of the paper after fifty folds is 0.001 inches times 2^{50} . That statement isn't weird or surprising or controversial. What's surprising is just how enormous two to the fiftieth power turns out to be. *That's* where the "craziness" of the final answer comes from.

M: Yeah, I was pretty surprised by that. I guess I'm not so good with large numbers. But I trust the calculator.

V: That's why it's not reasonable to conclude that there must be something wrong with the argument.

M: Okay, but how does this apply to your argument? You say factory farming is worse than the Holocaust, which sounds crazy. Where does the seeming craziness of *that* conclusion come from?

V: Let's review the major premises in my reasoning. Some of them are *moral*, and some are empirical, factual premises. First, I have the moral premise that suffering is bad. Anything surprising there?

M: No, that seems obvious enough. But I think it's surprising that animal suffering matters just as much as human suffering.

V: But I don't need to assume that. I can just say that animal suffering is at least one one-thousandth as bad as qualitatively similar human suffering. That's enough for my argument. Is that surprising?

M: I'm surprised that an animal welfare nut would admit that humans might matter a thousand times more than animals.

V: Right, *that's* surprising. But it's not surprising that animal welfare matters at least one thousandth as much as human welfare, is it?

M: Not particularly.

V: Well, those are all of my moral premises. The next step in my argument is just a factual, empirical premise: that life on factory farms is extremely unpleasant. Is that surprising?

M: You know, after we talked last time, I watched one of those PETA videos, "Meet Your Meat," about conditions on factory farms.³⁵ I had no idea how cruel they were.

V: So that part *is* surprising.

M: I guess so.

V: Here's my other factual premise: the number of animals killed in two years of factory farming is greater than the total number of humans who have ever existed. Were you expecting that?

M: Okay, that's surprising. I was kind of shocked to hear that.

³⁵ People for the Ethical Treatment of Animals, "Meet your Meat," Nov. 22, 2010, <https://www.youtube.com/watch?v=32IDVdgmzKA>.

V: That's where the "craziness" comes from. My moral claims aren't surprising; it's the empirical facts that are surprising. It's shocking that factory farming could be worse than the Holocaust, not because it's shocking that animal suffering might matter, but because the quantity of animal suffering is shockingly large.

M: So if I find your conclusion incredible, I should question the empirical claims about the quantity of suffering.

V: Do you think that would be the rational thing to do?

M: Somehow, it doesn't seem rational. I'm not sure why, though.

V: You can't use a moral assessment of some case to figure out what the empirical facts of the case are. That's because a moral assessment isn't reliable unless it is based on independently-known empirical facts to begin with. For instance, your moral assessment of meat-eating isn't reliable unless it takes into account the empirical facts about the effects of meat-eating. Therefore, you can't figure out what those empirical facts are based on your initial sense that meat-eating isn't extremely wrong.

M: I guess you're right. But then maybe I should deny one of your moral premises.

V: I don't think that would make sense either. That's why I made the point about how the surprisingness of my conclusions is due to the empirical facts, not my moral premises. You generally shouldn't reject an obvious moral principle based on a moral assessment of a particular case that didn't take the empirical facts into account.

M: That's a bit abstract for me.

V: Okay, an illustration. Let's say that you're initially extremely confident that Alice is a good person. You also believe that a good person wouldn't murder someone for money. Now suppose you learn, to your great surprise, that Alice has murdered someone for money. There's compelling video evidence, and so on. What should you conclude: (a) that Alice isn't a good person after all, or (b) that murdering for money is consistent with being good?

M: Obviously (a).

V: Good. That's like our case. You're initially convinced that meat-eating is okay, or at least not awful. You also believe that causing enormous suffering for trivial reasons is awful. Then you learn that meat-eating causes enormous suffering for trivial reasons. What should you conclude: (a) that meat-eating is awful, or (b) that causing enormous suffering for trivial reasons isn't awful?

M: I see your analogy. But why does your conclusion still sound hard to believe to me?

V: I can think of several plausible explanations for that.

M: Start with the biggest one.

V: Okay. Number one: status quo bias.

M: What's that?

V: It's a bias in favor of the current practices of your own society. We often form moral beliefs by looking at how other people behave, and the reactions of others to that behavior, and assuming the common behavior and reactions are appropriate. That's why people from different cultures with radically different practices tend to all think that their own culture is the best. And why a proposal to radically alter the practices may strike us as "crazy."

M: Well, maybe this tendency is a good thing. It's how we preserve our culture and traditions.

V: Perhaps it's a good thing in most cases. It stops you from stealing, driving on the wrong side of the street, and so on. But it can also lead to mass atrocities. In the nineteenth century and earlier, it led many people to accept slavery, to treat slave masters with respect and runaway slaves as criminals.

M: So you think factory farming is like the slavery of our day.

V: I do. Our society has always had flaws and moral blind spots, which later generations look back at and shake their heads at. Slavery was one of them. It would be surprising, wouldn't it, if today was the first time in history when there weren't any major moral flaws in our society?

M: Sure, I guess. But that doesn't mean that factory farming is one of those flaws.

V: That's what we're trying to figure out. But on the face of it, the movement for animal welfare seems to fit the pattern of past moral progress. Much of the progress of the past was about overcoming prejudices against non-dominant groups – prejudice based on race, based on sex, based on religion, based on disability.

M: And then prejudice based on species?

V: That's the next one.

M: But that's different. Sex and race differences are obviously morally irrelevant. But species differences are obviously relevant.³⁶

V: Two centuries ago, people would have said sex and race differences are obviously morally relevant. Then, the abolitionists and the advocates for women's suffrage were the "crazy extremists."

M: When I introspect, it doesn't *seem* to me that I'm just accepting meat-eating because other people are doing it.

V: People are often mistaken about what accounts for their own intuitions or beliefs. We're unaware of a lot of the things that influence how things seem to us. But status quo bias is extremely widespread and well established, so you are probably subject to it too.

M: But there are other cases in which I criticize the status quo. For instance, I oppose the current President and many of his policies.

V: True, but meat-eating is something you actually see done in front of you on a daily basis, by almost everyone. Government policies are just things that you hear about on the news.

M: So the status quo bias mostly applies to actions that you observe in your day-to-day life?

V: I think so. Also, your political views have more "social proof."³⁷

M: What do you mean, "social proof"?

V: It's similar to status quo bias. People have a bias toward believing what other people believe – or at least things that are *close* to what a *good number of their peers* believe. A view that is too far out of the mainstream of opinion will tend to strike us as "crazy."

M: Well, you're definitely out of the mainstream. I note that this explanation only works if there are already other reasons why your view is unpopular; then this factor might contribute to making it even less popular.

V: Right. Another influence is self-interest. People tend to be biased in favor of beliefs that serve their

³⁶ M again follows Richard Posner's views from his debate with Singer.

³⁷ See Robert B. Cialdini, *Influence: The Psychology of Persuasion* (New York: William Morrow & Co., 1993), ch. 4.

own interests. For example, people in the slavery era who owned slaves had an interest in believing that slavery was okay. Otherwise, they'd have to give them up, at great financial cost. Plus, they'd have to believe unflattering things about themselves.

M: Well, of course it's in my interests to keep eating meat, and I'd prefer to think that it was okay while I'm doing it. But it doesn't seem to me that I'm being influenced by that.

V: It probably wouldn't. Most people who are influenced by a bias can't themselves detect the bias. You have this sense that my conclusion is "crazy," and you don't know why it seems that way. So the seeming is probably caused by some unconscious factors like these.

M: I see. Is that all?

V: Not yet. Another factor is something called the "affect heuristic."³⁸ It's the tendency to evaluate how good or bad something is by reference to the strength of the emotional reaction we feel when we contemplate the thing.

M: That doesn't sound so wrong. Usually, the worse something is, the worse I feel about it; the better it is, the better I feel.

V: Usually, yes. But there are at least two reasons why we might go astray in the case of animal ethics. One is that our capacity for empathy with other species is limited. We find it harder to empathize with other species than with other humans. So we have diminished affective reactions when we think about animal suffering, compared to human suffering.

M: Okay, but maybe the explanation goes the other way: maybe I have diminished empathic responses to animal suffering *because* I know that animal suffering is less important than human suffering.

V: I think that's unlikely; I have a better explanation: human beings evolved as social animals. The capacity for empathy probably evolved to enable us to get along better with other humans in our social group, not to get along with other species. That's why we don't empathize as readily with other species.

M: That sounds speculative.

V: True. I'm just trying to offer plausible explanations for your intuitive reactions.

M: Okay. What's the other problem with the affect heuristic?

V: We have a well-known problem dealing with large quantities. We can't intuitively grasp them. Moreover, as we imagine larger quantities of something that's good or bad, our affective response doesn't increase proportionately.

M: I suppose that's a good thing. Otherwise, we'd be in constant emotional agony from listening to the news.

V: Right. If you hear about a disaster that killed five people, you feel sorry about that. If you hear about one that killed five *thousand* people, you don't feel a thousand times as sorry.

M: Do we just feel the maximum degree of sorriness then?

V: Not even that. If there is a maximum intensity of negative emotion, we don't necessarily feel it even in

³⁸ For discussion of the affect heuristic as used in decision-making, see Paul Slovic, Melissa Finucane, Ellen Peters, and Donald G. MacGregor, "The Affect Heuristic" in *Heuristics and Biases: The Psychology of Intuitive Judgment*, ed. Thomas Gilovich, Dale Griffin, and Daniel Kahneman (Cambridge: Cambridge University Press, 2002), pp. 397-420. On the application of the affect heuristic to moral philosophy, see Michael Huemer, "Transitivity, Comparative Value, and the Methods of Ethics," *Ethics* 123 (2013): 318-45, at pp. 328-30.

response to colossal evils. A vivid description of one death, by a sympathetic person, might make us feel worse than a report of a million deaths.³⁹ Many factors affect our emotional response. It clearly isn't simply proportional to the size of the good or bad event – not even close.

M: So then this leads us astray when we try to assess the badness of large evils.

V: Yeah, like when we talk about something happening to *billions* of creatures. Our minds can't really appreciate, or respond proportionately, to such quantities. Harming a billion creatures is a thousand times worse than harming a million creatures – but we don't *feel* that way.

M: But this doesn't seem to lead us astray much when we are thinking about human harms. If I ask someone, "How much worse is it to kill a billion people than to kill one person?", I bet most people would get the correct answer: one billion times worse.

V: Probably. But that's because they don't have to rely on an independent moral intuition to make the comparison – they just look at the numbers. It's when you compare bads belonging to different categories that you deploy the affect heuristic.

M: Different categories? So like, if I'm asked to compare a broken promise to a sprained ankle?

V: Right. Or animal suffering to human suffering.

M: But according to you animal advocates, those *are* in the same category.

V: But *most* people *think* of them as belonging to different categories. Most people think you have to deploy an independent moral intuition to compare animal and human harms. So they do – and that intuition is affect-driven.

M: So you're against relying on affect in moral evaluation?

V: Not necessarily, not across the board. It's just that in some cases we can predict that it would be unreliable.

M: But in order to say that the affect heuristic is unreliable in this case, don't you have to already know what the correct moral judgment is? If it's leading us away from the truth, it's unreliable; but if it's giving us the correct answer, then it's reliable.

V: No, when I say it's unreliable, what I mean is this: we shouldn't have any *independent expectation* that it would get us the truth. The affect heuristic would lead us to judge human interests more important than animal interests, whether or not that was true. So, if you start out with no opinion about whether human interests matter more, you can't trust your emotional responses to tell you the answer.

M: Okay, so couldn't the large-numbers problem be avoided by just asking people to think about individual cases? Like, imagine a single cow suffering on a factory farm for a day, then imagine a single person enjoying the pleasure of a hamburger. We should be able to compare those two, right?

V: You're right, that would avoid the large-numbers problem, though it's still subject to the other biases I mentioned.

M: Well, when I think about the cow on the factory farm, it doesn't seem very bad to me.

V: Really? When I think about it, it seems very bad to me – clearly much worse than someone being deprived of the pleasure of a hamburger.

M: I guess we have a basic clash of intuitions. I wonder why we have such different reactions.

³⁹ Joseph Stalin is often quoted as saying, "A single death is a tragedy; a million deaths is a statistic."

V: People vary in their capacity for empathizing with other species.

M: That's true, I find it pretty hard to empathize with a cow. But why should I trust your intuitions, rather than my own?

V: Remember all the biases we were just talking about?

M: Sure. But your intuitions are also biased.

V: How do you figure that?

M: You just admitted it: you empathize with nonhuman animals. That's biasing your moral judgment.⁴⁰

V: I said that was explaining the difference in our reactions. I didn't say it was a bias on my part.

M: You don't think empathy can function as a bias?

V: I don't see any reason to think it's a bias in this case. Compare another case: the case of psychopaths. Psychopaths lack the capacity for empathy in general. Does that mean that they make the most objective, unbiased moral judgments?

M: I'm not sure they make moral judgments at all.⁴¹

V: Right, their lack of empathy prevents them from taking others' experiences into account. It doesn't make them objective; it makes them ethically blind.⁴²

M: Okay, obviously a complete lack of empathy is a problem. But too much empathy can also be a problem. I know someone who has too much empathy, and it messes up her life. She winds up feeling anguish a lot of the time because of other people's problems. She's even gone into serious debt trying to help others.

V: Yeah, that sounds like a practical problem. But I'm not sure it's relevant to the point here.

M: Why not? You were talking about how important empathy is.

V: Yeah, but I'm not saying empathy is good in all ways and in all contexts. Of course, it's not necessarily good *from the standpoint of self-interest*. What I'm saying is that empathy helps us to perceive morally relevant factors that depend on the interests of others; without it, we just think about our own interests. So your overly empathic friend is doing poorly with respect to promoting her own interests, but she's probably doing quite well with respect to appreciating the moral relevance of others' interests.

M: Point taken. But too much empathy can also lead to *moral* errors. For instance, we might give money to charities that help people in a visible way – like those ones where you sponsor a child and they send you pictures of the child and stuff – instead of giving to charities that don't send you pictures but that are actually more cost effective. That's because of empathy.

V: Yeah, that's true too. So empathy isn't *sufficient* for making good moral choices. But it might be *necessary*. You need it in order to be moved to take account of interests other than your own. You still need to use reason to decide what to do about those other interests. But if you lack the capacity for empathy, you'll just ignore others' interests. Like how psychopaths just ignore other people's interests,

⁴⁰ On the biasing effects of empathy, see Paul Bloom, *Against Empathy: The Case for Rational Compassion* (New York: Ecco, 2016).

⁴¹ For discussion, see Walter Sinnott-Armstrong, "Do Psychopaths Refute Internalism?" in *Being Amoral: Psychopathy and Moral Incapacity*, ed. Thomas Schramme (Cambridge, Mass.: MIT Press), pp. 187-207.

⁴² Iskra Fileva ("Reflection Without Empathy," ms.) argues that psychopaths are unable to reason morally due to their incapacity for affective empathy.

and factory workers just ignore the interests of animals.

M: But I don't lack the capacity for empathy; I just have more trouble empathizing with other species than with my own – which is perfectly normal.

V: That's normal, true. But it's also true that it prevents you from fully taking other creatures' experiences into account. That doesn't make you more objective; it makes you less aware.

M: Well, you've made some interesting points that I'll have to think about, V. But I still don't think I can accept that the meat industry is worse than *the Holocaust*, for God's sake.

V: I understand. It would take a lot to overcome your initial intuitive reaction. But remember that that doesn't mean eating meat is perfectly okay. A lot of things are less bad than the Holocaust but still wrong.

M: (*laughs*) Fair enough. But I'm not even sure that it's wrong. Maybe your arguments are pieces of sophistry that I'm just not clever enough to see through.

V: The G.E. Moore shift again? I thought we already discussed why that isn't a rational response.

M: I know, but maybe the arguments you gave to show why that isn't a rational response were *themselves* just pieces of clever sophistry.

V: If you're going to say stuff like that, there's no way I could ever convince you. No matter what I say, you can always just say that maybe I'm wrong for some reason you can't identify. That's called being dogmatic.

M: No, I'm not being dogmatic. I'm not saying you're definitely wrong. I'm just saying I'm not fully convinced. And I'm not saying I'll never be convinced; I just need to think about it more.

V: Okay, so for now, you're not sure whether it's wrong to eat meat. Do you think it might be *obligatory* to eat meat?

M: Don't be silly. I just mean that I think it might be morally okay.

V: Alright, it might be wrong, or it might be just okay. In that case, I would suggest that, until you figure out which it is, maybe you should stop doing it. If there's even a fair chance that it's extremely wrong, better stop until you're more sure. You want to be on the safe side, right?

M: In general, yeah. But I can't avoid every action that *might* be wrong. I mean there's *some* chance that just about anything I do might be wrong. But I can't be avoiding everything.

V: Fair enough. I'm not asking you to avoid every action that merely has some nonzero probability of being wrong. I'm saying: avoid an action if it has a *pretty good* chance of being *very* wrong, where you have no moral reasons *to* do it, and where you can avoid it without unreasonable personal cost.

M: Well, that's hard to object to. But until I've finished thinking through all the arguments, I'm not sure if I should even say there's a "pretty good chance" that you're right.

V: I think you know enough to say there's at least a pretty good chance. You know that the issue turns on a moral intuition about the badness of animal suffering. This intuition is held by many people who appear to be in general reasonable, smart, and morally sensitive. Many of them consider it extremely obvious. The great majority of the literature in ethics on the topic also agrees that meat-eating in our society is generally wrong. Many of these experts consider the case decisive.⁴³

M: But most people in our society seem to think eating meat is fine. And even most philosophers seem to be okay with it.

⁴³ See Rachels, "Vegetarianism," pp. 884, 898.

V: Right, so there's a divergence between ethicists who work on the topic, and lay people or philosophers who work in other areas. Now all of this that I just said – this is all stuff that you can know, independently of your direct evaluation of the arguments. I mean, you don't have to first figure out what you think of the arguments, to know that most ethicists who work on the topic think meat-eating is wrong.

M: Okay, but we shouldn't just defer to the experts on a controversial topic like this.

V: Yeah, I'm not saying we should just defer to the experts. I'm saying that the opinion of these experts, together with the prima facie plausibility of the arguments we've been discussing, is enough for you to say that there's at least a pretty good chance that I'm right – at least unless and until you can come up with a good argument against my views.

M: Maybe you're right. But going vegetarian is going to make my life so much worse. I can't commit to such a big lifestyle change.

V: How about you try being vegetarian just for the next week, and then we'll talk about it here at the same time next week?

M: *(sigh)* Oh, alright. I hope you appreciate the big sacrifice I'm making for you, V.