Rationally Speaking #159: Colin Allen on, “Do fish feel pain?”

Julia Galef: Welcome to Rationally Speaking, the podcast where we explore the borderlands between reason and nonsense. I'm your host, Julia Galef, and with me is today's guest, professor Colin Allen.

He is a provost professor at the University of Indiana, Bloomington, where he has joint appointments in the department of history and philosophy of science and medicine, and the program in cognitive science, so as you can guess, his work spans the intersection of many fields. He has published many books and edited many books, and is also the associate editor of the Stanford Encyclopedia of Philosophy.

One of the primary areas of focus of his work is philosophy of animal behavior and cognition, and that's what we're going to be focusing on in today's episode. Professor Allen, welcome to the show.

Colin Allen: Thanks, Julia. Pleased to be here, and call me Colin.

Julia Galef: Colin, great. The specific question that motivated me to reach out to you and invite you to the show was a much narrower slice of animal behavior and cognition. I was specifically interested in fish, and in whether fish, A, are conscious, and B, can feel pain or suffer, which are distinct but related questions.

I guess to kick things off, I'll just explain why I found this question interesting and important, and then I'd love to hear how you got interested in this topic in the first place, and why you think it's important.

Colin Allen: Sure.

Julia Galef: For me it was two things basically. First, I think fish occupy an interesting position in the sort of loose spectrum of non-human animals, where on the one end we have, say, great apes, that most people would agree that great apes have a kind of moral status by virtue of having some kind of consciousness and the ability to suffer. Whereas on the other end of the spectrum you have something like insects, where there's not too many people who are worried about insects suffering, for example.

Colin Allen: But maybe they ought to be.

Julia Galef: Yeah, so there's a spectrum of controversy where ... All parts of the spectrum have at least some controversy, but I think that fish occupy a position towards the middle of the spectrum where there's genuine uncertainty and debate about fish consciousness, ability to suffer, and their moral status.

That brings me to the second part of why I think this is an important question, which is: As I've thought over the years about the ethics of eating animal products, or eating animals, it has occurred to me that something I think is really underappreciated when people think about these questions is the issue of, how

many meals does one animal produce? If you kill a cow, cows are huge, especially the ones that we breed to be huge. If the cow suffers in its experience on the farm, or its experience of being killed, but it produces meals for dozens or a hundred people, then that's very different, I think, than the experience of, say, a chicken or a fish, which only produces a meal for maybe one person. Do you know what I mean?

Colin Allen: There's a lot of things to pick up on there. Let me start with a thought ...

Julia Galef: I just wanted to add one caveat, which is that the question of the ethics of eating fish is a separate question with lots of other related topics in it. I want to focus mainly in this episode on the empirical question, empirical-slash-philosophical question, of whether fish can suffer, and are they conscious. Then there's ethical questions that can be tacked onto that, but I want to focus mainly on the former.

Colin Allen: I always joke that ethics is above my pay grade anyway, so I'm happy to do that, although I can't help but point out that there are some fish that are every bit as massive as a cow. Part of the problem-

Julia Galef: That we eat?

Colin Allen: Not that we eat, but that we could potentially-

Julia Galef: That we could, sure.

Colin Allen: Some of the sharks are pretty darn large for a start, and we eat sharks.

Julia Galef: Although people generally just take the fin and throw the shark away, right?

Colin Allen: Generally no. Generally shark fin soup, it particularly goes after the particular fins of particular species, but there are lots and lots of species of sharks.

That's part of the problem I want to get to here. If we talk about fish generically, as a category, first, fish constitute 60% of all vertebrate species. It's a hugely diverse taxonomic group. Asking what it's like to be a fish, or are fish conscious, or do they suffer or feel pain, is even a broader question than asking, what is it like to be a primate, for instance. Primates are a relatively small taxonomic group compared to fish.

I think that it's first necessary to be careful to say, which fish are we talking about? As I said, it's a really broad taxonomic group that includes a bunch of things, that include the sharks and rays, include in fact an even more primitive family of fish, namely the jawless fish, which include the lampreys. Then also divides in more recently evolved groups of species, the ray-finned fish and the lobe-finned fish. You and I are technically osteichthyes, so we are actually descended from that fish group that includes the lobe-finned fish.

The ray-finned fish lead to what are called the teleosts, in particular the upper teleosts, which are probably most of the fish you're thinking about when you're thinking about what's swimming around on a coral reef or in an aquarium, largely with the exception of the sharks and what's on your dinner plate. It's also important to realize that, among the upper teleosts, the evolutionary trajectory of those is roughly the same as the mammals, so roughly contemporaneous with mammals. They're relatively recently evolved, and as I indicated, incredibly diverse. They fill all sorts of ecological niches, and we can get into this later, but that means their behavior and their cognitive capacities might be rather different depending on which niches they occupy.

Julia Galef: Makes sense. To the question of what got you interested in this topic of the conscious experience of animals, or to what extent are non-human animals conscious, how did that happen?

Colin Allen: I made the joke that ethics was above my pay grade, I used to make the joke that consciousness was above my pay grade.

Julia Galef: You've gotten a promotion.

Colin Allen: Exactly. I started off just being interested in cognition generally, and because of some interest in biology, thinking about that in a broad biological context. I was lucky enough as a Ph.D student to fall into the lap, so to speak, of people who were working on vervet monkey alarm calls. Robert Seyfarth and Dorothy Cheney were at UCLA at the time when I was a Ph.D student, and the now somewhat notorious Ph.D student Mark Hauser was also a student at UCLA at the same time.

We got talking, and I got very interested in this issue of animal communication. Because I had gone to UCLA, as most of us did at that time, and still to some extent, because it was a center for philosophy of language. That had been my undergraduate orientation, but I always wanted to do something that had a little bit broader application than some of the traditional questions in analytical philosophy.

I was interested, as I said, in mind and cognition broadly speaking. This provided a perfect opportunity to think about things like Gricean levels of intentionality and so on, in the context of an animal and evolutionary context that wasn't normally being discussed by philosophers. I was sort of in the right place at the right time for that.

Later on, after I got my Ph.D and started my first job, and got in touch with more people who were working with ethology, which is the evolutionary study of animal behavior, in particular Mark Bekoff, he of course already had this interest in ethics and consciousness, and we decided to write a book together that was really going to focus on the cognition side of things. We realized that it would be impossible to get away with writing a book that just deferred the consciousness question. You can imagine what the reviewers would have said. That's where I started getting into the consciousness issues in a more academic way, rather than just having a

few opinions of my own perhaps.

Then of course, as you've indicated, this has implications for the ethics, although I didn't think I did the ethics myself, but it was very much something the people who were doing the ethics were interested in. I also, because of a colleague at Texas A&M, got very interested in some of the neural mechanisms involved in nociception, which is the capacity to detect noxious stimuli, which is in the sensory system before it gets up into the brain and becomes experienced pain. That led to some thinking about how learning and pain systems might be part of a single system, and that might give us an empirical route into studying these things.

Julia Galef: A question that I think often confuses people about consciousness, among the many questions, is, to what extent is this an empirical scientific question, versus a philosophical question?

Colin Allen: I think if we could answer that question, we'd have the need for no more philosophy or science. The point is that there are some ways of framing the question which makes it seem very much like no amount of empirical evidence could ever be relevant. There are other ways of framing the questions in which it seems like empirical evidence would be extremely relevant. I don't think there's any way to determine which is the right way of framing those issues, and whether any of the ways that we currently frame the issues will stand the test of time.

My other orientation is that, if you look at the history of philosophy and science, and you look at the relationship, the back-and-forth between those two, that naturalistic approaches to problems that philosophers have thought intractable have generally, over the years, gained a lot of traction. Even if you think that no kind of empirical or natural approach is possible, really there's no a priori argument that's going to establish that, because such an argument would pretend to now more than it could possibly know. Even if you're anti-naturalistic, the only way to defend that is to keep going with more naturalistic methods and see whether they pay off. At some point maybe you throw up your hands and say, "Okay, we've tried everything and nothing is going to work," but I don't think we're anywhere near that point.

Julia Galef: I would agree with that. Maybe the important question to ask at this point is to break down the concept of consciousness. Because I think people do mean very different things by it, and some of those things are abstract enough that experiments couldn't possibly bear on the question. It's almost not even clearly a meaningful question. Then on the other hand we have things like, when I'm asleep, and if I sleepwalk, it makes sense to say that I'm not conscious when I'm sleepwalking. We understand what that means, there's this clear distinction between someone behaving and being conscious, versus someone behaving and being unconscious.

Colin Allen: Sure.

Julia Galef: I'm sure there are other conceptions of consciousness as well, like self-consciousness, like awareness, being able to think about your thought processes, right?

Colin Allen: That's right, yes. There are lots of ways in which people carve up the space, and there are ways of doing it which, for some of them, will just make it a matter of definition that it's not something that's empirically studyable. As you pointed out, there are also ordinary senses of the term, such as the difference between being asleep and awake, or even awake and sleepwalking, in which we would ordinarily say there's clear evidence of consciousness.

People have generally thought that those issues, for animals, are not so challenging, so it's easy to tell when an animal is sleeping or not sleeping. There has been a little bit of discussion, coming back to our starting point of fish, about whether fish engage in the particular kind of sleeping that is called REM sleep, which stands for rapid eye movement sleep. With that often, in people, being correlated with dreaming, and therefore thought to involve some kind of conscious experience, even though it's not going as far as motivating the organism to engage in any coherent activity such as in sleepwalking. That's sort of controversial.

Some people have taken the absence of REM sleep to be a fact, and claim for instance that it doesn't occur in fish or in reptiles, and have then wanted to identify the emergence of consciousness later in the evolutionary time than the origins of those phylogenetic groups. But there's evidence coming back in the other direction, that in fact there is perhaps some form of REM sleep in reptiles, and also sleeping fish.

Julia Galef: Just to clarify, the reason I brought up the example of sleepwalking versus regular walking was not exactly to bear on whether we can detect if fish are sleeping or awake, but more to bear on the question of, when it seems like an animal is behaving consciously, could it just be similar to sleepwalking? Maybe it's not that they are asleep, but they're executing behaviors that don't have a kind of internal subjective experience going along with them, the same way we are when we're sleepwalking. Could that just be what it's like, there is nothing that it's like to be a fish? There's no light on in the inside?

Colin Allen: Depends what you think is going to be the kind of evidence you can get from sleepwalking, how similar you think it is to what you see in animals behaving normally. Sleepwalkers are generally actually not very responsive to features of their environment. They don't learn from those experiences. There's all sorts of ways in which whatever happens during sleepwalking is not integrated with the rest of their activities. I would actually disagree with the claim that it's behaviorally similar to what you see going on in specific cases of animals engaged in their day-to-day living. I think it's easy to squint at any two things and say, "Yeah, they're similar," but when we get into the details it's not so clear.

Your point about experience, if we're going to make this tractable then we also, I

think, have to get away from the idea that the right way to frame the question is, "Is the light on". Because the light is sort of a metaphor which suggests it's just epiphenomenal, it's a nice rider that goes along with whatever else is going on, but it's not actually doing any work. I think that kind of framing of the issue is endemic in a fair amount of philosophy. There's not enough serious thought about what the biological functions of being conscious versus not conscious are.

I would point to attempts among scientists to think of consciousness in terms of information integration. There are some philosophers who are also following this kind of line as well, where consciousness constitutes something like a workspace, in which information from various sources is combined and used to provide global control of the organism. In Bernie Baars' global workspace theory, in Giulio Tononi this is what he calls hi-phi, for the Greek letter phi, which is a kind of integration information model. When we look at animals we can ask, to what extent do they actually integrate the experience that they're getting over different timescales in order to adjust their behavior in an adaptive way?

Fish are interesting here too, because they do seem to have a kind of Jamesian window of influence, so things that just happened can influence what they're about to do. James's idea was that consciousness is this specious present, it's not an instant in time, it's spread out somehow, with what's just happened shaping one's expectation of what's about to happen. I think we can actually tie that to certain kinds of learning capacity that other animals besides humans may have, and begin to get a story which gets empirical evidence and ties it into a functional, theoretical story about what consciousness is good for. That's how we move more in the direction of empirical questions.

Julia Galef: The issue that I think that I have with some of those lines of reasoning is that it seems like they could apply to computers or computer programs as well. A computer can learn, their expectations will be shaped by things that recently happened ... I forget if you said this one, but some people will point to the reactions, like a fish's reaction to a noxious stimulus, as evidence of the fish consciously experiencing that stimulus as pain for example, but a computer will also take actions to avoid potentially damaging stimuli like overheating. It seems like it's still pretty hard to get at ... There is something that is consciousness that goes beyond these things, since these things can also be found in computers.

Colin Allen: I think if you describe things thinly enough, then yes, you can say that computers do whatever animals do, but in fact it's just not the case. We don't know how to make a robot that has the adaptive survival capacities of a trout. There simply isn't a robotic trout that can do everything that a trout can do, that can learn in the same way. Yes, there's been great advances in machine learning recently, but it looks like those kinds of advances are actually quite limited to certain domains, and require lots of repetitive experience.

Animals, and humans included, have that kind of repetitive learning capability, but some of what goes on in us is also much faster than anything we know how to do in

a computer. Yeah, you can program a computer so that one keystroke will change its behavior, but you have to set up the computer in a specific way with a program to have that kind of response. It's all built in by the programmer in a way that I think is not the case for animal behavior more generally. Animals are much more adaptive, and capable of taking what is the relevant piece of a very complex set of inputs and adjusting their behavior accordingly.

I think just to follow up on that a little bit, and then give you a chance to ask a question again, you say "learning"; in fact, there are many, many different kinds of learning. I don't want to argue that every kind of learning is good evidence for consciousness. In fact some of the spinal cord work in pain system nociception shows that learning can occur in the spinal cord, and I don't think that there's any consciousness involved there.

I think that there are higher forms of learning, which actually don't take place in the spinal cord, which depend on a kind of openness to experience, which I can be more specific about, that in fact, in humans, correlate extremely well with conscious awareness of the conditions that are being learned about. That we can get evidence for those kinds of things in other animals, as well as fish, makes me think, "Here's the right place where we should be pushing empirically."

Julia Galef: I was asking about consciousness in the last few minutes because I think I was implicitly assuming that it's necessary for the experience of pain, but it might be a good point here to stop and ask you what you think the relationship is between consciousness and pain. Is consciousness necessary for pain, but also, if we take as a given that an animal is conscious, and it reacts against noxious stimuli, then is that enough to say that it's experiencing pain from those stimuli, et cetera?

Colin Allen: I think that it's useful and pretty standard to separate two dimensions of pain. One is its localization to some area of potential injury or damage, and the other is the affective, bothersome side of the pain. It's hard to think of a pain experience that is affectively bothersome without it having some conscious component to it. That's the part that we really care about, we don't want to be bothered by pain. We know from human cases that using morphine will actually lead people to say, "The pain is still there, but it's not bothering me." They're aware of something, but it's lost that affective dimension.

We also know that if you take, for instance, rats, and give them moderate doses of morphine, they will still show withdrawal response to a noxious stimulus such as an electric shock, but they won't do anything to get themselves out of the situation. You might say by analogy they're sensitive to it, in the sense that they withdraw from it, but they're not going to get out of the part of the chamber where they're getting shocked. You set things up where they could be in one of two places in the experiment, and they can freely go to the other place, but if you've got a moderate dose of morphine, they're not going to learn to go to the other place.

You can also tie this in to the neurophysiology, so there are things that you can do

to the rat's brain which will produce roughly similar kinds outcomes, so those are the areas that we think are particularly important for this affective component of pain.

There's a bit of a terminological issue here, so is the morphine patient who says, "I'm still in pain but it doesn't bother me," using "pain" correctly? I don't think there's any good right or wrong answer to that question. Normally in our ordinary experience those two things don't come apart. You've got some experience which is not bothering you, but in every other component is what you would be experiencing if you had jammed your finger or something. It's natural that language would just allow you to use the word "pain" in that context. I don't think much hangs on that terminology.

With these kinds of approaches we're beginning to break down these two different components of the pain experience, and show that the basic neuroanatomy and physiology is identical across all the mammals, at least, and may have good analogous counterparts in birds. Then when we get to animals that are more distantly related to us, reptiles, amphibians, and eventually fish, the neuroanatomy gets a lot less well understood.

The behavior is still a possible route. When we can show that animals, certain fish, will choose one location over another, I think you're right, the reflex alone doesn't show conscious pain. That that's integrated with a learning system that allows the animal to learn, "Stay out of this area," then we get back to a functional characterization of what the affective component of pain is doing.

Julia Galef: There are other experiments that I've read about, that looked at additional behaviors beyond whether the animal learns to stay out of an area or tries to avoid the stimulus, that felt more compelling to me, and I'm curious what you think about them. For example, I know there are some experiments that looked at whether fish would continue rubbing their ... If you applied some acid to a part of their body, would they continue rubbing their body against things to, presumably, try to get the acid off or to dull the pain?

If you give a fish morphine, which I suppose blocks the ... Actually I don't know how morphine works, I won't speculate on that, but in humans at least it dulls the pain. Then does the animal stop avoiding the stimulus, or stop trying to nurse its injured parts, et cetera? Are there any experiments like that, that you think give us more evidence than the basic behavioral avoidance?

Colin Allen: I'm a little curious why you think that the lip-rubbing is more compelling, because after all that would be a relatively trivial thing to implement in a robot, to be honest. To go back to your computer argument, that seems much less compelling to me because we can explain why it might be functional. Maybe it's a way of dispersing the, it was bee venom they actually used in some of those experiments, dispersing the bee venom mechanically from the affected part. You can imagine a feedback system which is reinforcing that behavior.

The other kinds of behaviors, the use of morphine again; that, I think, ties into exactly what I was saying before. It turns out that if you give high doses of morphine, animals, including humans, become just sort of inert, and won't even show a withdrawal reflex. It's these intermediate doses that are interesting, where you get a continued reflexive response to a noxious stimulus, but you get a change in the behavior of removing oneself from the location where that stimulus is being presented or delivered. I don't know whether anybody's done anything specifically on fish though, to see whether morphine actually reduces the amount of lip rubbing behavior and that kind of thing, the amount of protective behavior.

One thing, also, to bring into this conversation, is that there are differences between acute pain and chronic pain. There are differences between particular stimuli that will produce a reflexive withdrawal response, and stimuli that will produce more sorts of protective responses. When you're injured you may or may not rub the affected spot. You may rather just prefer to hold your arm, if it's broken or something, in a particular position. Rubbing it or touching it might hurt.

I think that we need to be more sophisticated about the range of behaviors, and the range of interactions, between drugs and other treatments, and how that gives us a more complete picture of the kind of system that's underlying this, and what the temporal requirements of what that system are. Nursing an injured part is something that one wants to do over hours or days; removing oneself from a source of noxious stimuli is something one typically wants to do within seconds or minutes.

Morphine affects parts of the brain with a particular temporal dynamic that's going to have different effects on those behaviors. The more we fill out that picture, the more I think we find that, as a matter of fact, the systems that we have in humans and other mammals are evolutionarily conserved, or rather similar to the systems that we find in more remotely distant species.

Julia Galef: I can try to explain why I thought the rubbing of the lip that had venom injected into it was stronger evidence. It's something like, I'm trying to separate between responses to noxious stimuli that seem designed just to reduce risk to the creature, or reduce damage to the creature, versus responses to those noxious stimuli that seem designed to reduce an unpleasant sensation, independent of risk or damage introduced. Moving away from a noxious stimulus is the kind of response that could totally occur even if there was no subjective experience of pain, but just an algorithm designed to reduce risk posed by a noxious stimulus.

But rubbing a lip, or even vocalizations, like a yelp that we would want to call a yelp of pain, or a grimace or something like that, those kinds of responses don't seem like they're serving any kind of evolutionary purpose, in terms of reducing harm to the animal. It's hard for me to understand why those would exist if they weren't responses to some subjective experience of pain. Does that make sense?

Colin Allen: It makes sense, but I think, again, we need to think a little bit harder about possible

functions here. Take the yelping. One obvious function would be to alert conspecifics, members of your own group, to the presence of some potential injury source, so there's a social function it can play. In fact there's other dynamics here at work too. Not just the social aspect of alerting other individuals, but also showing whether or not you're injured to others can be useful in soliciting caregiving responses from them.

It can also be a disadvantage. I saw a study some time ago that I'm afraid I can't give you an exact reference for, where they looked at dogs post-operatively in a vet school. The vets had always been surprised by dogs, because you could give them really major surgery, go visit them in the recovery room an hour later, and they're up on their feet wagging their tails at you, and it looks like any person would be curled up and groaning at this point. They took the simple step of installing cameras in the recovery room, and found that this behavior was only being shown when people walked into the room. There was the social drive to be up and interacting with the human, plus, layered on top of that, what we know about dogs, they're pretty vicious about attacking any other individual who's perceived to be weaker, and there's additional evidence for that. It could be really important for a dog not to show that you're in pain.

I think when we understand the social dynamics of signaling, either showing some signal isn't necessarily about revealing some inner state of your own; it could be just a warning to others, or solicitation of certain kinds of caring behavior from others. Conversely, not showing that kind of thing doesn't show that there's no pain felt. I find that cuts both ways, and what we need is a much richer understanding of the organism in its social and ecological niche.

That's what's missing, to my mind, from this discussion among applied ethicists and philosophers. They just are not tuned in to the actual empirical details that are really necessary to understand these issues thoroughly.

Julia Galef: While we're talking about evolutionary arguments, it seems to me that we could get some suggestive evidence about the likelihood that, say, fish feel pain, by asking whether it would be adaptive -- not to show pain, but in this case -- to feel it. Is there some reason that it would be better for an organism's survival to have the subjective experience of pain, as opposed to just be programmed with reflexes to avoid stimuli that correspond with harm or risk?

Colin Allen: Good. Here's where I think the learning story comes in. I think there's a lot of learning that can occur without conscious awareness. I think there's even some protectively-organized learning that can occur without conscious awareness. But if you're going to be able to pair, generally, any condition in the world that might be a cue for something that is potentially harmful, then you need something that can actually do, back to information integration, a kind of openness to experience that will require you to hold, in some sort of temporally extended buffer, what the conditions were when some process started, and what the outcome was some seconds later.

Let me try to make that more concrete, because that was very abstract. A lot of cases, learning can take place completely unconsciously when there are two stimuli that are occurring at the same time, and that have some particular consequence for the organism.

In fact you can even get the rat's spinal cord to learn to hold a leg up when that leg, if lowered, gets electric shock, even when, and this is not going to please animal rights people that this kind of experiment's done, but it's an essential source of knowledge I think, even when the spinal cord of the rat has been cut such that no signals are getting from the neck to the brain. You get this kind of instrumental learning. The rat's leg position determines whether or not a shock is received, and that rat will learn to hold that leg up. I say, "that rat"; the conscious part of the rat, which I take to be in the brain, is not involved at all in this. It's completely spinal.

What you cannot get that rat to do is arbitrarily associate any stimulus with that electric shock. It's got a very specific system in the spinal cord that allows it to correlate leg position with noxious stimuli, and learn from that. In fact you can even induce learned helplessness under these conditions, which is an interesting phenomenon where, under certain conditions, a rat that's getting shocked won't know what to do, and will basically give up. Again, when I say "rat" here, I just mean the spinal cord on downwards.

What goes on in general learning is that we and other animals can arbitrarily pair, not in a completely open way, but in a very open, flexible way, various stimuli, lights flashing, sounds going off, whatever, with various outcomes: getting an electric shock, getting a bad taste in your mouth. All of these things are pairable.

There are limits, there are biological limits. It's not like anything can be paired with anything else, but a whole lot of things can be paired with a whole lot of other things, and in a way that's not predictable by evolution. You can learn new associations of things that are not part of your evolutionary history. That's what's invoked, the system to do exactly that. That's much more adaptive than having something that's tied into things that have only been experienced lots and lots of times in the past lineage.

This comes out most clearly in a distinction within classical Pavlovian conditioning, between what's called delay conditioning and trace conditioning. Everybody's familiar with the story of Pavlov's dog, and the bell ringing, and the dog salivating. The idea is that the sound starts, while the sound is going on… the dog acquires the response of salivating in response to the sound, the bell in the story. The key there is that the unconditioned stimulus, which is the food, is delivered while the conditioned stimulus, which is the sound, is ongoing. They might be a little offset, the sound comes on and a second or so later the food arrives, but they're experienced simultaneously.

If you vary the protocol slightly, and you produce a sound and then remove it, and then produce the other stimulus, what's going to be the conditioned stimulus, then

not all individuals will acquire the response of producing the behavior to that training stimulus.

This has been studied extensively, for instance, in eye-blink conditioning. In eye-blink conditioning, if I send a puff of air toward your eye, you will blink. If I precede it by a tone, and the puff of air comes while the tone, the sound, is on, you will learn to blink in response to the tone. 100% of people will acquire this response regardless of whether or not they're aware that's what's happening to them. Consciousness seems not to be necessary. If we turn on the tone, and then turn it off, and then deliver the puff of air, only about half of people acquired the response.

If you take people who've been through this protocol and you ask them after the fact, "Do you know what was going on there," you ask them some questions to get at what they were aware of, in the standard delay conditioning protocol, the standard classical Pavlovian conditioning protocol, about half the people will be able to tell you something about the tone and the puff of air, and about half won't. In the trace conditioning, where there was this delay ... It's confusing, because there's a gap between the tone and the puff of air. Only those who actually acquired the response can answer the questions about what they experienced.

There's this really tight correlation between the two, which gets at my point about there being a temporal window here. Where if you think about the functional demands here, in order to pair an arbitrary tone with something that happens later, you have to buffer that. You have to keep that active in some way. It's the keeping it active which I think provides the building block for certain aspects of conscious experience.

It turns out, if we do this eye-blink conditioning experiment with rabbits, 100% of them will acquire the response in the classical delay Pavlovian conditioning experiment, and only about half of them acquire the response when it comes to the trace conditioning, which is really suggestive. The mechanism is the same, we can't ask them, but they're also having to buffer it. Depending what else they're doing, they might not always be buffering things that would be relevant to certain kinds of learning, but it is absolutely central for other kinds of learning.

Julia Galef: Is there any reason to think that there would be some other way to buffer that learning without having the subjective experience of pain, or is that just sort of what the subjective experience of pain is?

Colin Allen: I think of the buffer as much more general than the subjective experience. I think that there's a lot more going on, also, in pain than simply this buffering. I think of the buffering as a way to think about consciousness generally. Pain, and particularly the suffering, or affective component of it, I think actually taps into another part of the system.

One of the things that's going on is, when you've got certain noxious stimuli, it may

or may not be easy for you to get that to go away. I'll use an example. I have a mild incipient case of arthritis in one of my hands, and occasionally my thumb will get in a position where it's like, "Ow," but I instantly know how to move it out of that position. I would say there's pain there, but there's not much suffering, because I have this ability to instantly make that acute pain go away. If it's chronic arthritis, nothing you do will make that pain go away. I sometimes have days where it just aches, and it's really hard to find a position. Those are the days where I would say I'm suffering more from this pain.

Interestingly, the very systems that are invoked in that are the parts of the anterior cingulate cortex which are also where the opioids, morphine, have their major effect, and that we know also are in an interaction with your prefrontal cortex, which is where action planning is also being done. You have a kind of system which is trying to find ways to remove some source of pain, and getting a signal back from ... In this interaction there's a signal going backwards and forwards, in which a plan is being assessed by the ACC, "Yeah, that'll work, no that won't work."

When you find people who are in chronic pain, have lots and lots of ongoing pain that they can't get rid of, there are actually changes to that system. There's, again, a temporal dynamic here. It's one in which it's trying to get rid of a certain kind of potentially harmful stimulus, or find a resting position in which recovery can happen, which ties into the systems that we know that are important, neurologically and pharmacologically, for treating the affective dimension of pain, and which are also important, the ACC turns out to be really important, in this kind of place preference conditioning I was talking about earlier. It's when you disable, pharmacologically, the ACC, then you get animals that will show a withdrawal reflex, that's all going on mostly spinally I think, but won't do anything to get out of the situation.

Julia Galef: So I realized that I have never in this conversation directly asked you: “Do you think that fish feel pain? Do you think that fish can suffer?” You've sort of implicitly hinted at maybe some of your views on the topic throughout our conversation.

But we have only a few minutes left. So what I want to do now is zoom out, and look at some of these different kinds of arguments for fish pain that we've touched on in the episode. Like the behavioral arguments, the arguments about brain structures, the evolutionary arguments, and ask you -- which of those kinds of arguments do you think is most compelling? Or least compelling? How does it all add up to give you an overall impression of the likelihood that fish can consciously suffer?

Colin Allen: That's a big question.

Julia Galef: Sorry!

Colin Allen: First, are there fish that feel pain? Again, I want to come back to my point, 60% of all vertebrate species, it depends which fish we're talking about. We do have

evidence now of trace conditioning in at least three species of fish, and I find that intriguing. That gets at the general issue of consciousness. I think when you pair that with the more specific kinds of things that you were talking about, the lip-rubbing and everything else, I'm reasonably confident that some fish, particularly among the upper teleosts, are likely going to have conscious pain. However, that's very controversial.

There are some fish biologists arguing that they just don't have the brain structures that are necessary, that they claim we know are necessary, for pain. I think we only know those are the ones that are actually doing it in mammals and birds; there's no reason why you couldn't have similar functionality in differently-organized brains. That's where the neurobiology is not where we need it to be. Fish brains in general are highly diverse, but they also develop in a different way from avian brains, and that makes it really hard to do a mapping between the two.

What kind of arguments do I find most persuasive? I think I've tried to indicate that I find the arguments most persuasive that tell a whole systems view. I want to know, not just superficially, is this behavior analogous to that. Because I think any argument by analogy of that kind is vulnerable to somebody pointing out a disanalogy, or claiming that the analogy is superficial, which is where your computer objection came from. When we start thinking hard about the functions of consciousness in general, and pain in particular, we think about what organisms are doing with those experiences. Not just a light bulb coming on, grooving to it, but that actually plays a real functional role in certain kinds of learning.

Furthermore, by virtue of having to play that role, then you get some of the other things that people have theorized about consciousness, such as, come back to the Jamesian specious present point. Also the idea that consciousness seems like it's forward-looking in some way, there's an anticipatory aspect to it. I think all of those things, to me, when we can find evidence of similar things in non-human animals, adds up to a more compelling case than simply a checklist that says, "Item A yes, item B yes, item C yes, humans and whatever other species has exactly that." I realize what I'm doing is more of the same, but it's putting it also within a kind of functional-theoretical framework within which these are not just items on a list, but items on a list that connect to each other in a way that constructs a whole systems-level theory.

Julia Galef: Makes sense. Excellent. We are just over time at this point, so I'm going to wrap up this part of the podcast and we'll move on to the Rationally Speaking Pick.

[interlude]

Julia Galef: Welcome back. Every episode, we invite our guest on Rationally Speaking to introduce the Rationally Speaking Pick of the episode. That's a book, or website, or blog, or article that influenced their thinking in some interesting way. Colin, what's your pick for this episode?

Colin Allen: I'm going to put in a plug for Kim Sterelny's The Evolved Apprentice, which has been out a couple years, now. I recently read it, and in many ways he has anticipated some of the things that I'm doing in a project we've got going on the evolution of human cognition. What I really like about the book ... A couple of things.

One is, he's very explicit about doing a kind of philosophy of nature project. I think in many ways this gets philosophy back to its roots. Philosophy has been in a discipline, or a silo, for some time now, having rightly tried to professionalize itself in a way that it perhaps wasn't so professionalized in the 19th century. The 20th century really went in this sort of direction. If you look back historically, philosophers, very important philosophers, have always done their best to try to make sense of the latest science or their own investigations of nature. That spans Aristotle and Descartes and Kant. I think a return to that kind of philosophizing is really good, and I think Sterelny does an excellent job of illustrating how that might work.

Then the other thing I really like about it is, he is very critical of single-factor explanations. That's something I resonate with a lot, I think he's done a great job of explaining how we should think about human cognitive evolution as the product of lots of influences and feedback loops among those influences. Certain changes mean longer development periods, which means more capacity to learn, which means more capacity to learn in certain kinds of ways, and the whole thing snowballs in a certain way.

The drawback of the book is that, while the story itself is very compelling, the specific details need to be worked out. That's what makes it so challenging and interesting, because I think that there are really interesting things to say and do, and for philosophers to be engaged with here. Because these are cutting-edge problems in epistemology and scientific epistemology. How do we infer things from the fossil record? It's just a very rich book that will lead to lots of potentially interesting work, it seems to me.

Julia Galef: Excellent. The warning against single-factor explanations seems like a very generalizable warning to me.

Colin Allen: Yeah, I think too often, and maybe it's a good thing in science and in philosophy, it's "my theory versus your theory" and only one left standing. I'm a much more pluralist sort of attitude, and I think that's coming back in certain quarters too, but the hard work is to try to figure out how these different perspectives on things might actually be compatible when you reframe them, perhaps in some higher-dimensional understanding. What looks like competitive theories can actually be parts of the same reality. I think of it as taking two-dimensional slices of three-dimensional figures. We could argue about whether that's a circle or a rectangle, and not realize that in fact they're just different slices of a cylinder. That kind of metaphor, I think, applies to a lot of what goes in science and philosophy.

Julia Galef: Nice. Excellent, we will link to that pick on the podcast website. Colin, thank you so much for joining us, it's been a pleasure having you on the show.

Colin Allen: My pleasure too, thanks very much.

Julia Galef: This concludes another episode of Rationally Speaking. Join us next time for more explorations on the borderlands between reason and nonsense.