OTHER MINDS

THE OCTOPUS, THE SEA, AND THE DEEP ORIGINS OF CONSCIOUSNESS

PETER GODFREY-SMITH

4

FROM WHITE NOISE TO CONSCIOUSNESS

What It's Like

What does it feel like to be an octopus? To be a jellyfish? Does it feel like anything at all? Which were the *first* animals whose lives felt like something to them?

At the start of the book I quoted William James's plea for "continuity" in our understanding of the mind. The elaborate forms of experience found in us derived from simpler forms in other organisms. Consciousness surely did not, James said, suddenly *irrupt* into the universe fully formed. The history of life is a history of intermediates, shadings-off, and gray areas. Much about the mind lends itself to a treatment in those terms. Perception, action, memory—all those things creep into existence from precursors and partial cases. Suppose someone asks: Do bacteria *really* perceive their environment? Do bees really *remember* what has happened? These are not questions that have good yes-or-no answers. There's a smooth transition from minimal kinds of sensitivity to the world to more elaborate kinds, and no reason to think in terms of sharp divides.

For memory, perception, and so on, this gradualist attitude makes a lot of sense. But the other side of the coin is subjective experience, the feel of our lives. Many years ago Thomas Nagel used the phrase *what it's like* in an attempt to point us toward the mystery posed by subjective experience. He asked: What is it like to be a bat? It's probably like *something*, but very different from what it's like to be a human. The term "like" is misleading here, as it suggests that the problem hinges on issues of comparison and similarity —*this* feeling is like *that* feeling. Similarity is not the issue. Rather, there *is a feel* to much of what goes on in human life. Waking up, watching the sky, eating—these things all have a feel to them. That's what has to be understood. But when we take an evolutionary and gradualist perspective, this takes us to strange places. How can the fact of life feeling like something slowly creep into being? How can an animal be halfway to having it feel like something to be that animal?

~ Evolution of Experience

I aim to make progress on those problems here. I don't claim to solve them entirely, but to take us closer to the goal James laid down. I'll set the topic up as follows. *Subjective experience* is the most basic phenomenon that needs explaining, the fact that life feels like something to us. People sometimes now refer to this as explaining *consciousness*; they take subjective experience and consciousness to be the same thing. Instead, I see consciousness as one form of subjective experience, not the only form. For an example that motivates this distinction, take the case of pain. I wonder whether squid feel pain, and whether lobsters and bees do. I take this question to mean: Does damage feel like anything to a squid? Does it feel *bad* to them? This question would often now be expressed by asking whether squid are conscious. That always sounds misleading to me, as if it's asking too much of the squid. To use an older term, if it feels like something to be a squid or octopus, then these are *sentient* beings. Sentience comes before consciousness. Where does sentience come from?

It's not a soul-like substance that is somehow added to the physical world, as *dualists* think. Nor is it something that pervades all of nature, as *panpsychists* believe. Sentience is brought into being somehow from the evolution of sensing and acting; it involves being a living system with a point of view on the world around it. If we take that approach, though, a perplexity we run into immediately is the fact that those capacities are so widespread—they are found far outside the organisms that are usually thought to have experience of some kind. Even bacteria sense the world and act, as we saw in chapter 2. A case can be made that responses to stimuli, and the controlled flow of chemicals across boundaries, are an elementary part of life itself. Unless we conclude that all living things have a modicum of subjective experience—a view I don't regard as insane, but surely one that would need a lot of defense—there must be something about the way *animals* deal with the world that makes a crucial difference.

One way to approach this question is just to talk about the complexity of different kinds of organisms, and the complexity of their dealings with the world. But there are many kinds of complexity, and we want something more specific. I'll look now at one such factor—something that I'm sure is part of the story, though it is not easy to see exactly where it fits in. In animal evolution, along with the sheer elaboration of sensing and acting, there's the evolution of new kinds of connection between these activities, especially connections that loop, that involve feedback.

For an organism like you or me, here are some familiar facts. What you'll do next is affected by what you're now sensing; and also, what you'll *sense* next is affected by what you now *do*. You read, and turn the page, and turning the page will affect what you see. Sensing and acting each affect the other. We know this explicitly and can talk about it, but their intertwining also affects in more fundamental ways how things feel, in quite a raw sense of "feel."

Consider the case of *tactile vision substitution systems* (TVSS), a technology for the blind. A video camera is attached to a pad that sits on the blind person's skin (for example, on their back). Optical images picked up by the camera are transformed into a form of energy (vibrations, or electrical stimulation) that can be felt on the skin. After some training with this device, the wearers start to report that the camera gives them an experience of *objects located in space*, not just a pattern of touches on their skin. If you are wearing such a system and a dog walks past, for example, the video system will make a moving pattern of presses or vibrations on your back, but under some circumstances this will not be experienced as vibrations on your back; instead you'll experience an object moving out in front of you. This happens, though, only when the wearer is able to *control the camera*, to act and influence the incoming stream of stimulation. The user of the device has to be able to move the camera closer, change its angle, and so on. The simple way to do this is to attach the camera to the person's body. Then the wearer can make objects loom, and come and go from the visual field. Subjective experience is intimately tied here to the interaction between behavior and sensory input. The moment-to-moment feedback between sensing and acting affects how sensory input itself feels.

Though the idea that our actions affect what we perceive seems routine and familiar, philosophers through many centuries did not treat it as especially important. In philosophy, this is the territory of unorthodoxies, of works beside, rather than within, the main development of ideas. That is true even in recent years. Instead, a huge amount of work has looked at a small *piece* of the total picture; it has looked at the link between what comes in through the senses and the thoughts or beliefs that result. Little was usually said about the link to action, and even less about the way action affects what you sense next.

Some philosophers have always disliked this obsession with sensory input, with receptivity, seen in theories of the mind. But their response was to reject the importance of input in a wholesale way, and to try to tell a story about the self-determining organism, about the subject as *source*, imposing itself on the world. This is a kind of overcompensation, as if philosophers are only able to concentrate on one side at a time. It is a big thing, apparently not easily achieved, to accept that there is *traffic*, a to-and-fro.

In everyday experience there are two causal arcs. There is a sensory-motor arc, linking our senses to our actions, and a *motor-to-sensory* arc as well. Why turn the page? Because doing so will influence what you'll see next. The second arc is not as tightly controlled as the first, because it extends into external, public space, rather than remaining inside the skin. Perhaps as you turn the page, someone grabs the book, or grabs you. The sense-to-motor and the motor-to-sense pathways are not on a par. But the neglected junior partner, the effect of action on what we sense next, is surely important. This, after all, is *why* we do much of what we do: to control what our senses will encounter.

Philosophers often use the metaphor of a *stream* of experience. Experience, they say, is something like a river in which we are immersed. This image is quite misleading, though, as the flow in a river is almost entirely outside our control. We might change our location within the river, swimming to one spot or another, and that gives us some control over what we'll encounter. But in real life we can usually do much more than that; we can reshape the things themselves that we interact with. Rivers are rather resistant to such efforts when we're alone in the middle of them.

What you sense next has two sources: what you just did, and what the larger world beyond you is up to. The overall shape of the cause-effect relations looks like this.

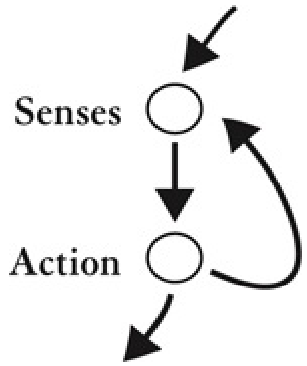

Two arrows lead into the senses. They have different roles in different contexts and sometimes one is more important than the other, but they are both almost always there.

Loops that link actions back to the senses are not only seen in us. They are present also in very simple forms of life. But they become more marked in animals, especially because animals can *do* more. The evolution of muscle, derived from tiny fiber-like elements inside cells, created a new means by which life impresses itself on the world. All living things affect their environment by making and transforming chemicals, and also by growing and sometimes by moving, but it is muscle that gives rise to rapid, coherent action on large spatial scales. It makes possible the *manipulation* of objects, the deliberate and rapid transformation of what is around us.

The evolution of animals is affected by these looping causal paths in a number of ways. Often these loops lead to a *problem*, as an animal attempts to work out what's going on around it. Some fish, for example, send out electric pulses for communication with other fish, and also electrically sense other things going on around them. The self-produced pulses will affect their own senses, though, and it may be difficult for a fish to distinguish the pulses it has made from electrical disturbances that are due to external things. To deal with this problem, whenever a fish emits a pulse it also sends a *copy* of the command around to the sensing system, enabling that system to counteract the effect of the pulse it has produced. The fish is tracking and registering the distinction between "self" and "other," between the effects on its senses of its own actions and the effects of events going on around it.

An animal need not send out electrical pulses to encounter this problem. As the Swedish neuroscientist Björn Merker notes, it results just from being able to move. An earthworm withdraws when something touches it—the touch might be a threat. But every time the worm crawls forward, it causes part of its body to be touched in just the same way. If it withdrew at every touch, it could never move at all. The worm succeeds in moving forward by canceling the effects of those self-produced touches.

For all organisms there is a distinction between self and the external world, even if only onlookers can see it. All organisms also affect the world outside them, whether they register that fact or not. Many animals, though, acquire their own glimpse, their own registration, of these facts, because action would be so difficult otherwise. Plants, in contrast, have quite rich senses but don't move. Bacteria move, but their simple senses don't threaten to tangle them, in the way seen in Merker's worm.

This interaction between perception and action is also seen in what psychologists call *perceptual constancies.* For us, an object can remain recognizable as the same object while our viewpoint on it changes. As you move closer to or further from a chair, it does not usually seem to grow, shrink, or move, because you tacitly compensate for the changes in appearance that are due to your actions, along with some changes that are not due to you, such as shifts in lighting conditions, and so on. Perceptual constancies are seen in a fairly wide range of animals, including octopuses and some spiders, as well as vertebrates. This ability has probably evolved independently in several different groups.

Another path in the evolution of experience leads to *integration*. As streams of information come in from different senses, they are brought together into a single picture. This is vivid in our own case; we experience the world in a way that ties together what we see with what we hear and touch. Our experience is usually of a unified scene.

This might seen inevitable, a consequence of having eyes and ears attached to the same brain, but it is not. It's one way of being wired up, and some animals do not integrate their experience nearly as much as we do. For example, in many animals the eyes are on each side of the head, not the front. The eyes then have separate visual fields, either largely or entirely, and each connects just to one side of the brain. In such an animal it is easy for scientists to control what each side is exposed to, by masking one eye. Then we can ask a question that might seem to have an obvious answer: If we show something to only one side of the brain, does the other side get the information too? We are not looking at damaged or altered animals here, so the two sides of the brain have all their natural connections. One would think that the information should get across. Why should evolution set things up so that only half the animal knows what has been seen? But when this question was studied in pigeons, it turned out that the information was *not* being passed across. The pigeons were trained to do a simple task with one eye masked, then each pigeon was tested on the same task while being forced to use the other eye. In a study using nine birds, eight of them did not show any "inter-ocular transfer" at all. What seemed to be a skill learned by the whole bird was in fact available to only half the bird; the other half had no idea.

These experiments have also been done on octopuses. An octopus trained on a visual task using just one eye initially remembered the task only when tested with the same eye. With extended training, they could perform the task using the other eye. The octopuses were unlike the pigeons in that some information did get across; they were unlike us in that it did not get across easily. In more recent years, animal researchers such as Giorgio Vallortigara of the University of Trieste have uncovered a number of similar “fissures” in information processing related to the separation between the two halves of the brain. A variety of species appear to be more reactive to predators seen on the left side of their visual field. Several kinds of fish, and even tadpoles, seem to prefer to position themselves so that an image of another individual of the same species is on their left side. On the other hand, various animals are better at perceiving what’s on their right side when the aim is to detect food.

This specialization seems to have clear disadvantages, leaving the animal vulnerable to attack on one side, or less able to find food on one side. Vallortigara and others think it might make good sense, though. If different tasks require different kinds of processing, it may be best to have a brain with specialized sides that deal with each task, and not tie them too closely together.

These findings are reminiscent of experiments on “split brain” humans. In cases of severe epilepsy, it sometimes helps the patient to cut the *corpus callosum*, the connector between the left and right hemispheres that make up the upper part of the human brain. People tend to behave fairly normally after these operations, and it took a while for researchers to realize that anything unusual was going on. But if the different halves of such a patient’s brain are exposed to different stimuli, quite dramatic disunity often emerges. The operation seems to have given rise to two intelligent selves, with different experiences and skills, in a single skull. The left side of the brain usually controls language (though not always), and when you talk to a split-brain patient, it is the left side who speaks back. Though the right side cannot usually speak, it can control the left hand. So it can choose objects by touch, and draw pictures. In various experiments, different images are provided to each side of the brain. If the person is asked what they have seen, their verbal response will follow what was shown to the left side of the brain, but the right side—controlling the left hand—may disagree. The special kind of mental fragmentation seen in split-brain humans seems to be a routine part of many animals’ life.

Animals seem to have a range of ways of dealing with this situation. In the case of birds, incoming visual information can be even more fragmented than it was in those eye-masking experiments I described above. In birds like pigeons, each retina has two different “fields,” the yellow field and the red field. The red field sees a small area in front of the bird where there is binocular vision, and the yellow field sees a larger area that the other eye cannot access. Pigeons not only failed to transfer information between eyes; they also did quite badly at transfer between different regions of the *same* eye. This might explain some distinctive bird behaviors. Marian Dawkins did a simple experiment showing a novel object (a red toy hammer) to hens, who were allowed to approach and inspect it. She found that hens approached such an object in a weaving way that seemed designed to give the different parts of each eye access to it. That, apparently, is the way the whole brain gets access to the object. The weaving gaze of a bird is a technique designed to slosh the incoming information around.

To some degree, unity is inevitable in a living agent: an animal is a whole, a physical object keeping itself alive. But in other ways, unity is optional, an achievement, an invention. Bringing experience together—even the deliverances of the two eyes—is something that evolution may or may not do.

~ Latecomer versus Transformation

The story I am working toward is one of gradual change: as sensing, acting, and remembering became more elaborate, the feel of experience became more complex along the way. Our own case shows us that subjective experience is not an all-or-nothing matter. We know half-conscious states of various kinds, such as waking from sleep. Evolution includes an awakening on a different time scale.

But perhaps all that is a mistake. A gradual development of subjectivity from simple and early forms is one option, but perhaps the best evidence we have tells against it, evidence that comes from our own brains.

One path to this view begins with an accident, a case of carbon monoxide poisoning from a shower's bad water heater in 1988, which led to a case of brain damage in a woman known only as "DF." As a result of the accident, DF felt almost blind. She lost all experience of the shapes and layout of objects in her visual field. Only vague patches of color remained. Despite this, it turned out that she could still *act* quite effectively toward the objects in space around her. For example, she could post letters through a slot that was placed at various different angles. But she could not describe the angle of the slot, or indicate it by pointing. As far as subjective experience goes, she couldn't see the slot at all, but the letter reliably went in.

DF has been studied extensively by the vision scientists David Milner and Melvyn Goodale. By linking her case to other kinds of brain damage and to earlier work in anatomy, Milner and Goodale put together a theory of what's going on—in us, as well as in special cases like DF's. They argue that there are two "streams" by which visual information moves through the brain. The *ventral stream*, which takes a lower path through the brain, is concerned with categorization, recognition, and description of objects. The *dorsal stream*, which runs above it, closer to the top of the head, is concerned with real-time navigation through space (avoiding obstacles as you walk, getting the letter through the slot). Milner and Goodale argue that our subjective experience of vision, the feel of the visual world, comes only from the ventral stream. The dorsal stream does its work unconsciously, both in DF and in ourselves. After her accident, DF lost her ventral stream, and hence felt almost blind—even as she walked around the obstacles before her.

A simple interpretation of these cases holds that you need the ventral stream to have any experience of what's coming in through your eyes at all. That is probably too simple. It's likely that dorsal stream vision feels like something, though it doesn't feel much like seeing. The details of these two “streams” are less important than the larger surprise coming out of this work. That is the fact that quite complicated processing of visual information—processing that runs all the way from eyes, through brain, to legs or hands—can take place without the subject experiencing any of this *as seeing*. Milner and Goodale link this discovery to what I described a bit earlier as the integration of sensory information. They think that the activity in our brains that leads to visual experience is the building of a coherent “inner model” of the world. It’s certainly reasonable to think that building an integrated model of this kind has effects on subjective experience. But perhaps without such a model there is no subjective experience at all?

Milner and Goodale discuss various animals whose perception of the world is less integrated than ours. In the 1960s, David Ingle rewired the nervous systems of some frogs by means of surgery (he was aided by the fact that frog nervous systems regenerate unusually well). By crossing some wires in the brain, he was able to produce a frog that snapped at prey to its left when the prey was really to its right, and vice versa. It saw prey in a left-right reversed way. But this rewiring of part of the visual system did not affect all of the frog’s visual behavior. The frogs behaved normally when they were using vision to get around a barrier. They behaved as if some parts of the visual world were reversed, and other parts were normal. Here is Milner and Goodale’s comment:

> So what did these rewired frogs “see”? There is no sensible answer to this. The question only makes sense if you believe that the brain has a single visual representation of the outside world that governs all of an animal’s behavior. Ingle’s experiments reveal that this cannot possibly be true.

Once you accept that a frog does not have a unified representation of the world, and instead has a number of separate streams that handle different kinds of sensing, there is no need to ask what the frog sees: in Milner and Goodale’s words, “the puzzle disappears.”

Perhaps one puzzle disappears, but another is raised. What does it feel like to be a frog perceiving the world in this situation? I think that Milner and Goodale are suggesting that it feels like *nothing*. There is no experience here because the machinery of vision in frogs is not doing the sorts of things it does in us that give rise to subjective experience.

Milner and Goodale’s comments illustrate, in one form, an idea that quite a few people who work in this area now accept. The senses can do their basic work, and actions can be produced, with all this happening “in silence” as far as the organism’s experience is concerned. Then, at some stage in evolution, extra capacities appear that do give rise to subjective experience: the sensory streams are brought together, an “internal model” of the world arises, and there’s a recognition of time and self.

What we experience, in this view, *is* the internal model of the world that complex activities in us produce and sustain. Feeling starts *there*—or, at least, it begins to creep into existence when these capacities creep into existence—in the brains of monkeys and apes, dolphins, perhaps other mammals, and some birds. When we think of simpler animals as having subjective experience, according to this view, we're projecting onto them a fainter version of our *own* kind of experience. This is a mistake because our experience relies on features they just don't possess.

A view like this has been defended also by the neuroscientist Stanislas Dehaene, whose laboratory near Paris has done some of the most penetrating work on this topic over the last twenty years or so. Dehaene and his coworkers have spent years looking at perception right at the edge of consciousness—images that come and go just a bit too quickly for subjects to know they are seeing them, or that are presented while attention is diverted, and that nevertheless affect what the subject thinks and does. It turns out that we often process this unexperienced information in quite sophisticated ways. For example, sequences of words can be flashed so quickly that a person has no idea they were shown at all. But sequences with incongruous meanings—such as "very happy war"—are registered by the brain differently from combinations with more reasonable meanings—"not happy war." One might think that conscious thought is necessary to distinguish such meanings, but this is not so.

We can do a lot without consciousness, Dehaene thinks, but some things we can't do. We can't unconsciously perform a task that is novel, rather than routine, and requires a series of acts, step-by-step. We can unconsciously learn associations between experiences—learn to expect A when you see B—but only if B and A come close together. Once there is a significant gap between them, we can learn the association only if we are aware of it. You can learn to blink when you see a light if the light is followed by an irritating puff of air, but only if the light and puff are very close together. Once the light and puff are separated by a second or so, the association can't be learned unconsciously. What we've learned over the last thirty years or so, Dehaene thinks, is that there's a particular *style* of processing—one that we use to deal especially with *time*, *sequences*, and *novelty*—that brings with it conscious awareness, while a lot of other quite complex activities do not.

Back in the 1980s, in one of the first modern attempts to explain consciousness, the neuroscientist Bernard Baars introduced the *global workspace* theory. Baars suggested that we are conscious of the information that has been brought into a centralized "workspace" in the brain. Dehaene adopted and developed this view. A related family of theories claim that we are conscious of whatever information is being fed into *working memory*, a special kind of memory which holds an immediate store of images, words, and sounds that we can reason with and bring to bear on problems. My colleague at the City University of New York, Jesse Prinz, has defended a view of this kind. If you think that a global workspace is needed for subjective experience, or a special kind of memory, or some other mechanism along these lines, you'll hold that only complex brains that are fairly similar to ours can give rise to experiences that feel like something. These brains will probably be found outside of people, but perhaps only in mammals and birds. The result is what I'll call *latecomer* views about subjective experience. These views don't hold that the lights went on in a sudden flash, but they do hold that the "waking up" came late in the history of life and was due to features that are clearly seen only in animals like us.

When I described a number of these theories just above, including the theories of Baars, Dehaene, and Prinz, I said they were theories of *consciousness*. I used that word because that's the word they use. It's sometimes hard to work out how these theories relate to my own target here: subjective experience in a very broad sense. I treat subjective experience as a broad category and consciousness as a narrower category within it—not everything that an animal might *feel* has to be *conscious*. A person might then say that a "global workspace" is necessary for consciousness without it being needed for the most basic kind of subjective experience. Not only is this possible, but I think it's approximately right. In the literature I'm describing here, it's often hard to work out what people think about this. But some of these people think there's no distinction between consciousness and subjective experience; they say that they're giving us a theory of when a mental activity *feels like something*.

The work that inspires latecomer views of experience has led to a great deal of progress. People like Dehaene have found a way *in* to the study of human consciousness, a path of a kind that seemed a fantasy not many years ago. We should not hang on to an alternative picture just because it's more generous or feels right. But I do think that arguments can be given against the latecomer view, and certainly there's an alternative to consider. I'll call this the *transformation* view. It holds that a form of subjective experience preceded late-arising things like working memory, workspaces, the integration of the senses, and so on. These complexities, when they came along, transformed what it feels like to be an animal. Experience has been reshaped by these features, but it was not brought into being by them.

The best argument I can offer for this alternative view is based on the role in our lives of what seem like old forms of subjective experience that appear as *intrusions* into more organized and complex mental processes. Consider the intrusion of sudden pain, or of what the physiologist Derek Denton calls the *primordial emotions*—feelings which register important bodily states and deficiencies, such as thirst or the feeling of not having enough air. As Denton says, these feelings have an "imperious" role when they are present: they press themselves into experience and can't easily be ignored. Do you think that those things (pain, shortness of breath, etc.) *only feel like something* because of sophisticated cognitive processing in mammals that has arisen late in evolution? I doubt it. Instead, it seems plausible that an animal might feel pain or thirst without having an "inner model" of the world, or sophisticated forms of memory.

Let's look at the case of pain. One might initially say it's obvious that even simple animals respond to pain in a way that indicates they feel it, squirming and wriggling in distress. But things are not so straightforward. Many responses to bodily damage that seem to involve pain probably do not. For example, rats with a severed spinal cord, and hence no channel from the site of body damage to the brain, can exhibit some of what looks like “pain behavior,” and can even show a form of learning that responds to the damage. Various reflex responses in animals might look to us like pain, because we empathize with them. We need to go past these mere appearances.

Fortunately, we can. The most telling evidence is based on pain-related behaviors that are too flexible to be dismissed as reflexes, even though the animals in question have quite different brains from ours, and are not likely to meet the requirements of the “latecomer” views. Here is an example from fish. Zebrafish were tested first to see which of two environments they preferred. They were then injected with a chemical suspected to cause pain, and in some cases, the less preferred environment had a painkiller dissolved in it. The fish now preferred this environment, but only when it contained dissolved painkiller. They made a choice they’d not normally make, and they made it in a situation where the idea of a more painful or less painful *environment* would be quite novel to them: evolution could not have set them up with a reflexive reaction to this situation.

Similarly, in a study in chickens, birds with damaged legs chose a food that would usually be less preferred, provided that it contained painkillers. Robert Elwood has done experiments of a similar kind in hermit crabs, the small crabs who live in shells made by various mollusks. Hermit crabs are arthropods, relatives of insects. Elwood gave the crabs small electric shocks, and found they could be induced to leave their shell by a shock. But not always: they were more reluctant to leave a higher-quality shell than a low-quality one—they had to be shocked more. They were also more likely to put up with the shock when the odor of a predator was around and the shell was more valuable for protection.

Tests of this kind don’t suggest that *all* animals feel pain. Insects are in the same large animal group (arthropods) as crabs. Insects appear to behave normally, to the extent that they physically can, even after quite severe injuries. They don’t groom or protect injured parts of their body, but keep doing whatever they were doing. Crabs and some shrimp, in contrast, will groom injured areas. You can still doubt that these animals feel anything, yes. But you can doubt that about your next-door neighbor. Skepticism is always possible, but a case is being built here. These results do provide support for a view of pain as a basic and widespread form of subjective experience, one present in animals with very different brains from ours.

In this picture, there are early and simple forms of subjective experience that are then transformed as evolution makes nervous systems more complicated. With this transformation, new capacities—such as sophisticated kinds of memory—are added which have a subjective side, while other things that once contributed to experience might be pushed into the background. How might we imagine the earlier forms? This is perhaps impossible, as our imaginations are tied to our present-day, complicated minds. But let’s try.

The title of this chapter borrows a phrase from a paper by Simona

Ginsburg and Eva Jablonka. Two Israeli scientists working in different fields within biology, they wrote a paper a while ago trying to sketch the evolutionary origins of subjective experience. At one point in their paper, they offer a stab at a description of experience in a simpler and distant animal: *white noise*. Imagine a poorly differentiated buzz as the beginning of it all.

I keep coming back to that metaphor when I'm trying to get my mind around this topic. It *is* a metaphor—very much so. It's a metaphor of sound applied to organisms that, at least in most cases, probably could not hear at all. I'm not sure why the image stays so consistently with me. Somehow it seems to point in the right direction, with its evocation of a crackle of metabolic electricity, and the *shape* of the story suggested. That shape is one in which experience starts in an inchoate buzz, and becomes more organized.

In our own case, looking inside, we find that subjective experience has a close association with perception and control—with using what we sense to work out what to do. Why should this be? Why shouldn't subjective experience be associated with other things? Why isn't it brimful of basic bodily rhythms, the division of cells, life itself? Some people might say it *is* full of those things—more than we realize, anyway. I don't think so, and suspect there's a clue here. Subjective experience does not arise from the mere running of the system, but from the modulation of its state, from registering things that matter. These need not be external events; they might arise internally. But they are tracked because they matter and require a response. Sentience has some *point* to it. It's not just a bathing in living activity.

Ginsberg and Jablonka imagined their "white noise" as the first form of subjective experience. Perhaps, though, white noise corresponds to the *absence* of experience; it's what was present before subjective experience arose. Maybe that distinction takes the metaphor too far. Anyway, from some such state, the older forms of subjective experience arose—forms linked to the primordial emotions, pain and pleasure, feelings that must be acted on.

If this is right, some tentative conclusions might be drawn about the first animals with nervous systems, discussed in chapter 2. Suppose it's true that much of the work done by very early nervous systems was just pulling the animal together and making coordinated action possible. The patterned contraction of a swimming medusa is a present-day illustration, and the self-possessed lives that might have been lived by Ediacaran animals are also in this category. Here a nervous system acts mostly to generate and maintain an activity, and modulation of that activity is a more minor player. If so, perhaps this is a form of animal life that feels like nothing at all. The Cambrian—with all its richer forms of engagement with the world—would then be where simple experience begins.

This beginning would not have been a single event, or even a single extended process occurring on one evolutionary path. Instead, there would have been several such processes, taking place in parallel. By the time of the Cambrian, many of the different kinds of animals I've been discussing in this chapter had already branched off from one another—the branchings probably occurred in the Ediacaran, when all was quieter. By the Cambrian, the vertebrates were already on their own path (or their own collection of paths), while arthropods and mollusks were on others. Suppose it's right that crabs, octopuses, and cats all have subjective experience of some kind. Then there were at least three separate origins for this trait, and perhaps many more than three.

Later, as the machinery described by Dehaene, Baars, Milner, and Goodale comes on line, an integrated perspective on the world arises and a more definite sense of self. We then reach something closer to *consciousness*. I don't see that as a single definite step. Instead, I see "consciousness" as a mixed-up and overused but useful term for forms of subjective experience that are unified and coherent in various ways. Here, too, it is likely that experience of this kind arose several times on different evolutionary paths: from white noise, through old and simple forms of experience, to consciousness.

~ *The Case of the Octopus*

Let's now return to the octopus, our unusual and historically important animal. How does it fit in? What might *its* experience be like?

An octopus is, first, an organism with a large nervous system and a complex active body. It has rich sensory capacities and extraordinary capacities for behavior. If there is a form of subjective experience that comes along with sensing and acting in a living system, an octopus has plenty of that. But that's not all. In elusive and alien form, the octopus has some of the sophistications, some of the steps beyond the basics, described in this chapter.

Octopuses, of at least some species, have an opportunistic, exploratory style of interaction with the world. They are curious, embracing novelty, protean in behavior as well as in body. These features are reminiscent of what Stanislas Dehaene associates with consciousness in human mental life. As he says, the demands of novelty jolt us from unconscious routine into conscious reflection. An octopus's explorations are sometimes mixed with caution and sometimes with a puzzling recklessness. I noted in the previous chapter how my collaborator Matt Lawrence, diving near the Octopolis site, came across an octopus who seized his hand and led him over the sea floor, pulling him along behind it. We have no idea why it might have done this. In contrast, once while scuba diving at another site, I was hovering off the sea floor, anchored with a few fingers of one hand, as I photographed tiny sea slugs. I became aware of something below, and saw that a single slender octopus arm was slowly extending toward my fingers on the bottom, from a clump of seaweed next to me. The octopus was curled up in a ball in the weed, with most of its body hidden but one eye visible through a hole, sending one arm out cautiously as it watched. This was an act of exploration, accompanied by what seemed to be very close attention, keeping me in view as the arm went out. I was a novel object of uncertain importance. The seaweed provided both cover and a viewing hole. From this shelter, one arm was sent out to inspect, perhaps to taste.

Earlier I discussed *perceptual constancies*. These are abilities an animal has to re-identify objects despite changes in viewing conditions—distance, lighting, and so on. The animal must factor out the contribution of its own location and perspective to identify the object itself. Psychologists and philosophers often associate this ability with sophisticated, as opposed to rudimentary, forms of perception. Perceptual constancies show that an animal is perceiving external objects *as* external objects—as objects that can stay the same while the animal's vantage point changes. In an old 1956 experiment some octopuses were taught to approach particular shapes and avoid others. In some of the experiments, the relevant difference was between small and large square shapes. The octopus sat in a tank, a square was introduced at the other end, and the octopus had to approach some squares (for reward) and not approach others (or be punished, with electric shock). That was the routine, and the octopuses were able to do this. The researchers then say, almost in passing, that on "several" occasions, the octopus was presented with the small square at half the usual distance from its body. The small square would then initially look larger—or at least, the size of the squarish shape on its retina would be larger. In every such trial, the experimenters say, the octopus performed the right action for the square's real size. The octopus was able to factor out the change in distance.

A surprising thing about this report is that it is quite an important observation, and yet it's little more than a brief aside in the paper. No numbers are given for the cases that tested for perceptual constancy, and no one seems to have followed the idea up. If the finding is accepted, it shows that octopuses do have at least some perceptual constancies. So, apparently, do some other invertebrates, honeybees and some spiders; this is not one of the octopus's unique invertebrate achievements.

Octopuses are also good at navigation. Whenever I see an octopus wander from its den, I follow it if I can, and I've been taken on a lot of tours. If I don't get too close as they roam and explore, the octopuses often pay me no attention at all. The octopuses are usually foraging for food, and this takes them on long, rambling paths that eventually return to their dens. I am often surprised at how well they do this, as the ramblings can take a good fifteen minutes or so, traveling through quite murky water. If they head off from their den in one direction, they may well return to it from another. The tour has the form of a loop, not an out-and-back path. Some years ago, Jennifer Mather did a careful study of this kind of behavior, watching an octopus in the Caribbean as it went on hunting trips, and she charted looping paths of this kind. It's not known how octopuses do this—what sorts of landmarks, guides, and memories they make use of. But some octopus species are certainly good navigators.

Remember again that our most recent common ancestor—a worm-like creature in the Ediacaran—almost certainly had none of these skills. It seems that once an animal starts to lead an active and mobile life, full of controlled, goal-directed, and rapid movement, there are ways of seeing and handling the world that make more sense than others. Different animals have independently evolved perceptual constancies. Though in some ways they must see the world very differently from us, octopuses seem to deal with the world by identifying and re-identifying objects, and to have some grip on a distinction between self and other. When you are around an octopus, it's impossible not to think they can also direct considerable *attention* on objects, especially new ones.

In the previous section I discussed some work on pain behavior in fish, chickens, and crabs. Trying to work out how octopuses relate to pain is not easy. At Octopolis, our site in Australia, we once got a lot of video of a large male octopus engaged in a series of aggressive interactions, roaming and wrestling with others at the site. He often "stood tall," up on stretched legs, and sometimes raised his rear end high in the water, above his head. We think he was doing this to make himself look as large as possible; these displays were often precursors to an attack on another octopus. Once when he had his body positioned in this way, a small but vicious fish (a leatherjacket) darted in and bit him right on the rear. Here's the bite as it happened, with the fish at top center:

The octopus responded in a way that was very human-like, with a startled jump, arms going everywhere.

He then went straight back to beating up other octopuses. The bite was fortunate for us, as it left a noticeable mark, which we could use to identify this individual from some distance for the rest of that trip.

As we've seen, some animals tend and protect a wounded spot on their body. Our octopus bitten on the rear did not do that. His initial response suggested that he certainly felt the bite, but the aftereffects were not noticeable. We suspect this was because it was a minor injury and he was busy with ongoing pugilistic activities. A recent paper written by Jean Alupay and her colleagues looked carefully at pain-related behaviors, including wound tending, in another species of octopus. There was some reason to expect oddities, because some octopuses, including Alupay's species, nip off their own arms, when necessary, to escape predators. The study found that octopuses who'd had their arms crushed (not *too* crushed) in an experiment amputated them in some cases, not in all, and they all tended and guarded the injured site for some time. That tending and guarding is, as I discussed, usually seen as an indicator of pain.

In the case of the octopus, everything about experience is made more complicated by the unusual relations between brain and body. Let's assume the octopus has a kind of mixed control over what its arms do, an interpretation supported by the behavioral experiments discussed in chapter 3. Octopuses, as they evolved their complex behavioral abilities, opted for a partial delegation of autonomy to their arms. As a result, those arms are brimming with neurons and seem able to control some actions locally. Given that, what might octopus experience be like?

The octopus may be in a sort of hybrid situation. For an octopus, its arms are partly *self*—they can be directed and used to manipulate things. But from the central brain's perspective, they are partly *non-self* too, partly agents of their own.

Let's consider some analogies with our case, beginning with acts like blinking and breathing. These are activities that normally happen involuntarily, but through attention you can assert control over them. An octopus's arm movements have something like this combination. The analogy is imperfect, as breathing, while normally involuntary, can be subjected to very fine-grained control when you do intervene to breathe voluntarily. Attention is used to take over what is normally an automatic process. In the octopus, if the mixed-control interpretation is right, central guidance of the movements is never complete, and the peripheral system always has its say. To put it too anthropomorphically: you would send an arm out deliberately and *hope* the local fine-tuning goes right.

Action by an octopus, then, would mix elements that are usually distinct, or at least seem that way, in animals like us. When we act, the border between self and environment is usually fairly clear. If you move your arm, for example, you control the arm both on its general path and also in many fine details of its motions. Various other objects in the environment are not under your direct control at all, though they can be moved indirectly by manipulating them with your limbs. Uncontrolled movements by an object around you are usually a sign that it is not part of you (with partial exceptions for knee-jerk reflexes and the like). If you were an octopus, these distinctions would be blurred. To some extent you would guide your arms, and to some extent you would just watch them go.

To tell the story this way is to tell it from the vantage point of the "central octopus." That might be an error. In addition, I might be assuming too simple a contrast with the human case. When a person becomes able to play a musical instrument well, various actions—including adjustments—become too rapid to be controlled consciously. Bence Nanay, a philosopher based in Antwerp, also sent me a quite different interpretation of the octopus/human comparison. Bence thinks that some relationships that look weird and novel in the octopus case are present in our case too, if we look hard enough. They are usually invisible to us, but they are there. Suppose you are reaching for an object with your hand. If the location or size of the target you are reaching for suddenly changes, your reaching movement changes extremely quickly—in less than a tenth of a second. This is so fast that it is unconscious. Subjects in experiments don't notice the change—they don't notice that *they've* changed their own movement, and don't notice the change in the target object. When I say "subjects" in the experiment, I mean that if you ask the person if there was a change, they will say no. The person does not notice the change, but their arm alters its path.

As in the octopus, there's a top-down decision to extend the hand, but also a fine-tuning which is fast and unconscious. In the octopus case, the fine-tuning is greater—it's more than just fine-tuning—and it does not only happen quickly. The octopus might watch some of the arm's wandering as if it is a spectator. In us, these adjustments are too fast to see.

In the case of humans, these rapid adjustments of the arm come from the brain, and they are visually guided. In the octopus, the motions are guided by the arm's own chemical and tactile senses, not by vision (though I'll qualify that statement in the next chapter; the issue is not quite so clear). In any case, Nanay's interpretation is that the octopus displays an extreme case of something that is also present in human action, though in a milder, less noticeable form. In the human case, there's a top-down command and then the addition of whatever fine-tuning is needed. In the octopus case, there's probably an ongoing interaction between commands from the center and decisions at the periphery. The arm is sent out, it wanders, and the octopus might respond by adjusting—perhaps by using attention, exerting some octopus willpower—to redirect the arm and keep it on track.

In the paper on "embodied cognition" I quoted earlier, Hillel Chiel and Randy Beer contrast an old and a new view of how action works. The old view has it that the nervous system is the "conductor of the body, choosing the program for the players and directing exactly how they play." Instead of this, they say, "the nervous system is one of a group of players engaged in jazz improvisation, and the final result emerges from the continued give and take between them." I'm not convinced by that as a general claim; I think it understates the role of the nervous system in most animals, to see it as one player among many. But in the case of the octopus, such a metaphor may well apply. The contrast now is not between the nervous system and the body, but between the central brain and the rest of the organism, which has its own nervous organization.

In the octopus's case there is a conductor, the central brain. But the players it conducts are jazz players, inclined to improvisation, who will accept only so much direction. Or perhaps they are players who receive only rough, general instructions from the conductor, who trusts them to play something that works.

Farrar, Straus and Giroux
18 West 18th Street, New York 10011

First edition, 2016

An excerpt from *Other Minds* originally appeared, in slightly different form, in *Scientific American*.

All photographs were taken by the author except for the following, which are frames from video taken by unmanned cameras (collaboration by Peter Godfrey-Smith, David Scheel, Matt Lawrence, and Stefan Linquist): frontispiece photo; the second photo in the “Visiting Octopolis” section in chapter 3 (showing the octopus on the left leaping toward the fleeing octopus on the right, on a shell bed); both photos in chapter 4; and the first two photos in chapter 8. Drawn figures are by the author with the following exceptions: the first illustration in chapter 3, of a cephalopod, figure by Ainsley Seago; the “Evolution of the Cephalopods” diagram in chapter 3, figure by Eliza Jewett; and the drawing showing differences in pattern between two octopuses in chapter 8, figure by Eliza Jewett.

Library of Congress Cataloging-in-Publication Data
Names: Godfrey-Smith, Peter.
Title: Other minds: the octopus, the sea, and the deep origins of consciousness / Peter Godfrey-Smith.
Description: First edition. | New York: Farrar, Straus and Giroux, 2016. | Includes bibliographical references and index.
Identifiers: LCCN 2016016696 | ISBN 9780374227760 (cloth) | ISBN 9780374712808 (e-book)
Subjects: LCSH: Nervous system—Evolution. | Consciousness. | Cephalopoda—Behavior. | Cephalopoda—Psychology.
Classification: LCC QM451 .G58 2016 | DDC 612.8—dc23
LC record available at https://lccn.loc.gov/2016016696

Our e-books may be purchased in bulk for promotional, educational, or business use. Please contact the Macmillan Corporate and Premium Sales Department at 1-800-221-7945, extension 5442, or by e-mail at MacmillanSpecialMarkets@macmillan.com.

www.fsgbooks.com
www.twitter.com/fsgbooks • www.facebook.com/fsgbooks