

Putting concepts under stress (II) – Pains

Sleep and pain tend to inspire poets and philosophers; micturition [urination] and defaecation do not. With psychoanalysts, it is the other way round. – Peter Nathan, *The Nervous System* ([142], 104)

7.1 Case study: Shared out-of-body pains

The constructing of possible-worlds tales to place concepts under stress in an attempt to refine or revise those concepts is no talisman guaranteed to produce success. It is just one method, among several, used by philosophers in their work. And its results, far from being consensual, sometimes provoke disputes as intense as those which prompted the use of the method in the first instance.

To see how the technique might be used, and to see the sorts of objections it might elicit, let us discuss the possibility of public, objective, out-of-body pains. At the outset, it is necessary to say that the question we will be examining – viz. “Might there be public, objective (i.e. shareable), out-of-body pains?” – is not to be regarded, in the first instance at any rate, as an *empirical* question. Only at a second stage of inquiry, can such a question be regarded as amenable to investigation by empirical means. For it is essential, first, to determine whether such a question is even meaningful (intelligible), and only if the answer to this preliminary question is affirmative, can we then proceed to the second. The preliminary question, then, might be stated this way: “Could the concept of *pain* apply to something public, out-of-body, or is its use reserved exclusively for sensations located within a living body and perceivable only to the person whose body it is?” The challenge posed by such a question is to determine how one might go about trying to answer it.

Newcomers to philosophy often immediately turn to the resource they have standardly used to determine the extent or range of applica-

tion of concepts, namely, their dictionaries. But it quickly becomes clear that dictionaries are not about to help with this particular question. There will always be vagueness in our concepts. Although we may from time to time reduce the vagueness of some of our concepts, we will never eliminate it entirely. Dictionaries cannot adjudicate the proper decisions to make when one operates within the penumbra of the meaning of a term; they are incapable of settling boundary disputes. When we ask ourselves whether we should apply a concept to some phenomenon which falls within the penumbral region of a term, dictionaries, which are designed to report standard usages, must fail us. We must take recourse to other means.

What the task comes down to is looking to see whether we can describe a situation in which it would be reasonable to say that something has occurred which is enough like ordinary ‘in-body’ pain to be regarded as pain, and yet, is outside of one’s body and is perceivable by more than one person. We will begin by reviewing current theories of pain and its causes.

In the case of seeing, hearing, and touching, the perceived ‘object’ (for lack of a better name) is almost invariably on or, more likely, external to (i.e. at some distance from) our skin. To be sure, we occasionally do hear the rumblings of our own stomachs and do use our hands to touch and feel various parts of our bodies, but for the most part, what we see, hear, and touch are physical objects external to our bodies. We see houses, buildings, human beings, trees, motor vehicles, dogs, etc.; we hear screeching tires, the voices of human beings, the music of a piano, the high-frequency whistle of the flyback transformer of a television set, etc.; we feel the knife and fork in our hands, the slipperiness of a bar of soap, the grip on the handle of a golf club; etc. In short, most often the ‘objects’ of seeing, hearing, and touching are external physical objects.

The ‘object’ of pain – or, to put it somewhat more perspicuously, what we feel when we feel a pain – is not a physical object. I may accidentally prick my finger with a pin, or cut my thumb with a blade, but the ensuing pain is not my feeling *of the pin* or *of the blade*. The pain begins, to be sure, with the pin’s pricking my finger, but it lasts for some time after the pin is removed. The actual pricking may take only a minute fraction of a second, but the resulting pain may last several minutes, long after the pin is removed from the site of the injury. In feeling pain, I am not feeling the pin, but the *injury* caused by that pin.

Physical injury may occur anywhere within one’s body. But while

no part of our body is immune to injury, only certain parts contain nerves which give rise to pains. Most persons are surprised to learn that their intestines, for example, are insensitive to incision, and that bowel biopsies may be performed painlessly with no anesthetic whatever. Similarly, the human brain lacks pain-generating nerve endings. But many sites throughout the body are sensitive to pain.

How do we know where a pain is? How do you know, for example, that there is a pain in, let us say, the thumb of your left hand? Young children, even those who have acquired a certain facility with language, are often notoriously bad at localizing their pains. Many children, obviously in pain and obviously having a serious ear infection, frequently are unable to pinpoint their pain; they may not even be able to localize the region of their body where it hurts (i.e. are unable to localize the pain as being above their necks). It might seem, then, that we *learn* over a period of time, by trial and error, to locate pains within our bodies. Perhaps at some early stage of our lives, we may, for example, have experienced a pain and, in looking about our bodies, spotted an injury on our hand.¹ In touching the wound, we may have discovered that the pain increased (or decreased) and thus came to believe that the pain was occurring at the site of the injury. Later, we were able to identify pains which originated from that location directly, i.e. without our having to look with our eyes, or probe with our fingers, for an injury. But in spite of its initial plausibility, this conclusion – that we *learn* through trial and error to identify the sites of our pains – may be too strong. For there is some contrary experimental data which suggest that newborns are able to ‘home in’ on the site of (at least some) pains directly, without a learning process: “... a newborn can remove an irritant from his nose with his hand or get rid of an irritant on one leg with his other foot” ([32], 19). What confuses the issue is the fact that it is conceivable that this ability of newborns to localize (some) pains is only temporary, like their ability to reach for visual objects. Perhaps, just as in the case of the latter visual-motor ability ([32], 45), the newborn may lose the innate ability to localize pain directly, and may have to ‘regain’ it in the ensuing

1. This picture is somewhat oversimplified. Some, but by no means all, injuries cause an energetic reflex movement, a flexion withdrawal of the limb, a fraction of a second before the pain impresses itself on one’s consciousness. Given this muscular reflex, one’s subsequent attention is naturally drawn to that limb itself.

months by a learning process. In short, it is unclear how much of our ability to locate pains directly is innate and how much is a learned response. In any event, well before we reach adulthood our ability – whatever its origin – to locate pains directly is firmly in place.

(Some researchers have suggested that we locate pains directly by utilizing an ‘internal map’ [either learned or innate] of our own bodies, on which we place, more or less correctly, the incoming signals from the many nerve pathways according to their points of origin. How much stock is to be put in such a theory? Does postulating an ‘internal representation’ of our bodies help to explain this ability to locate pains directly, or is this merely a metaphorical manner of redescribing the very phenomenon itself? We can well understand why researchers differ in their attitudes toward such models. Some researchers regard this model – of an internal, representational map – as the best, if not the only, way to account for the ability to locate pains and to know directly the disposition of our limbs. Other researchers regard the explanatory content of such a hypothesis as nil. They regard the postulating of such a map as unempirical and, indeed, as an entirely dispensable piece of baggage. The debate is but one instance of a century-old controversy about the utility of models in scientific explanations. The dispute can be expected to continue indefinitely.)

Our ability to locate pains directly, good though it often is, is far from perfect as being an ability to locate the injuries which are the physical causes of those pains. For there are a number of instances in which persons will locate pains in their bodies far from the sites of the injuries causing those pains. The best known of these is the pain caused by a rupture (popularly misnamed a “slip”) of the fifth lumbar disk. The ruptured disk presses on the sciatic nerve within the spinal column. But the ensuing pain is nearly always felt some centimeters, or even nearly a meter, removed from the site of the trauma. The pain is often felt in the hip, or down the back of the leg, or even in the toes. Such so-called referred pains sometimes mislead diagnosticians. Doctors will sometimes misdiagnose spinal injuries as pulled back ligaments or as hip injuries. It is essential, then, to distinguish carefully the site of a pain (i.e. where the pain is felt) from the site of the injury which causes that pain. Usually they coincide; occasionally, they do not.

The most dramatic case of referred pain is the phenomenon, well-documented in medical literature, of the so-called phantom limb. A person who has had a limb amputated may complain of pains which

feel as they would had that limb not been amputated. For instance, a person who has had his left leg amputated may complain of pains of just the sort he would have had if that limb were still attached to his body, i.e. it *feels to him as if he still has a left leg and that there is a pain in that leg*.

From a physiological point of view we can explain the phenomenon in this way. The nerve endings of the nociceptor class of 'Ad' (fast) and 'C' (slow) fibers at the site of the amputation (i.e. on the remaining stump of the limb) are firing and sending impulses along these fibers to the spinal cord, where they interact with a variety of other impulses (e.g. indicating touch or pressure) along with descending signals from the brain. If these impulses are not masked or blocked in the spinal cord, then signals proceed to the thalamus and eventually to the cerebral cortex (see, e.g., [219], 103). These latter signals are so like those which used to originate within his leg that they are 'mistaken' for signals which originate, not on the stump, but from within the (nonexistent) leg itself.

Now while this, or something very like this, physiological explanation is probably true, it does not by any means imply that the pain is not exactly where the person says it is. The crucial thing to recognize is that the location of pain, and the location of the injury causing that pain, need not – and occasionally do not – coincide. The person complaining of pain in his hip, who in fact has no injury whatever to his hip but has, rather, a ruptured fifth lumbar disk, has made no mistake about the location of his *pain*. *The pain really is in his hip*. He and his doctor make a mistake only if they infer from the pain's being in a particular place that the injury causing that pain must also be in that same place.

So too with the case of the phantom limb. The pain, as distinct from the cause of that pain, is just where it is felt, e.g. 20 cm or so below the stump.

It is just at this point, where we say that a sufferer's pain may be 20 cm beyond the surface of his skin, that many persons will, perhaps unwittingly, suddenly switch theories about the criteria for locating pain. In cases where the pains are *within* a sufferer's body, most persons are perfectly content to use the patient's own report as to the location of the pain as being definitive. The orthopedic specialist who asks her patient to place his fingers on his pain, sees the sufferer place his fingers on his hip. The doctor does not correct her patient, saying, "No, you are wrong. The pain is really where your ruptured disk is pressing on your sciatic nerve, just about here" (placing her own

fingers close to the injury) “in the spinal column.” Instead she might say, “That sort of pain in your hip is caused by an injury which is actually several centimeters away from the site of the pain, here in your spinal column.” None of us has any trouble understanding this distinction.

But let the location of the pain be out of one’s body, let it, for example, be felt as being 20 cm below the stump of an amputation, and immediately many persons will abandon the clear knowledge that pains and their causes may be remote from one another and will revert to a radically different account of pain, one which totally blurs the distinction between pains occurring where they are felt to be and pains occurring where the injuries giving rise to them are located. Faced with the report of a patient who says that it feels to him as if he still has a leg and that there is a pain in that leg, these revisionists will argue that the person’s pain simply is not where it is reported to be, but must instead be located at the site of the injury, i.e. on the stump of the amputation.

A person who, under the circumstances described, takes recourse to this latter theory has, I would strongly suggest, taken a retrograde step. To argue in the case of a phantom-limb amputee that he is mistaken as to the location of his pain challenges not just the amputee’s ability to locate pain, but everyone’s, amputee and non-amputee alike. If the amputee’s report of the location of his pain is not to be given credence – if, that is, the actual location of his injury is to be given primacy over his report – then there is no reason not to apply the same criterion for every other report of pain as well, in the case of toothaches, spinal injuries, etc.

The revisionists’ theory – that pains occur, not where they are felt to be, but rather where the injuries giving rise to them occur – thus departs flagrantly from our ordinary concept of pain. It is a proposal which sacrifices much of what we ordinarily say and think about pain for the expediency of not having to attribute pains to locations outside of bodies in the case of phantom limbs. But the price is too high. It is far preferable to allow that pains can and do occasionally occur outside of one’s body (as in the case of phantom limbs) than to subscribe to the theory that in every case of referred pain we have made a mistake as to the location of the pain itself.

An orthopedic specialist still would need to know, as an aid to making her diagnosis, that it *feels* to the patient *as if* there is a pain in his hip. That crucial medical symptom does not disappear in adopting the revisionist’s proposal: it simply makes its description awkward. If

we were to adopt the revisionist's proposal, then where we had earlier spoken simply and directly of the 'location of the pain', we would now have to talk clumsily of 'the impression as to the location of the pain', or 'the place where the patient *reports* or *believes* the pain to be'. Far preferable, it seems to me, is to argue that the revisionist has confused two quite different things: the pain and its cause. Pains are exactly where they are felt to be: often at the site of an injury, but sometimes at another place.

If we resist the revisionist's illicit conflating of the location of pains with the location of their causes, then the phantom-limb phenomenon must count as a genuine case of an out-of-body pain. The only reasonable conclusion, it seems to me, is to insist that not only are out-of-body pains possible, they are in fact actual, i.e. they exist, in the case of phantom limbs.

But having argued that there are in fact out-of-body pains is not yet to have proved what I initially set out to establish, namely, the possibility of the existence of *public, objective* out-of-body pains. While the phantom-limb experience is, I want to urge, best regarded as being a genuine out-of-body experience, it still falls short of demonstrating the possibility of there being publicly objective, i.e. shareable, out-of-body experiences. To have proceeded this far is still only to have taken the first step along the way to the intended goal. To proceed past this point, we must now take recourse to a possible-worlds tale.

So prodigious was his talent, so obvious was his promise, that even as a young teenager, having not yet graduated from high school, Michael Robins had been wooed by all the best music schools of the country. But there really was no choice. He had, since childhood, set his heart on studying the cello with Janos Starker, and so was quick to accept the offer from Indiana University when it came.² To persons unfamiliar with the school at

2. Just as certain novels and romances, e.g. E.L. Doctorow's *Ragtime* ([57]), are fictionalized histories through which ostensibly real persons parade, many possible-worlds tales incorporate characters and places patterned after those in the actual world. These fictive entities, e.g. Indiana University and Maestro Janos Starker in the present case, are said to be "counterparts" of the similarly named entities in the actual world. They are, however, still creatures of fiction.

Bloomington, the Midwest had seemed an odd choice. But that was only because they did not know of the visionary presidents and deans who in a labor of love, in a small town surrounded by corn fields and limestone quarries, had conceived and created one of the world's finest and largest music schools.

Michael flourished under Starker's instruction, and when he graduated, he landed the position, which had just fallen vacant, of first cellist in the recently resurrected NBC Symphony Orchestra.

His career seemed to be virtually assured, except that he had begun to experience, at first slight and then increasingly severe, pains in his right knee. X-rays and a biopsy confirmed the worst: a malignant tumor. Surgery was performed within hours, but his right leg, from mid-thigh, had to be amputated. Being a cellist, there was no question that he would be fitted, as soon as possible after the surgery, with a prosthetic leg so that he could resume playing.

For nearly two weeks after the surgery, Michael was given heavy doses of morphine to kill the pain. But as the drug was gradually withdrawn he began to experience classical phantom-limb pains. Mercifully, this new pain was not at all as debilitating as that prior to the surgery, but it was there, naggingly, nonetheless. It felt to Michael again as if there were a pain in his right knee. Of course he had no right knee, he could see that he had no right knee. But the pain 'knew nothing of that'. The pain *felt* as if it were in his right knee. The doctors tried to assure him that in most cases such phantom pains gradually subside, but they did warn him that he could not be absolutely assured that his pain would. In some few cases, phantom pains had been known to continue for the rest of a person's life.

A few weeks later, Michael was fitted with his first prosthesis, a mechanical affair, with a spring and hydraulic knee and a similarly contrived mechanical ankle. There was no electronic circuitry in the prosthesis and no connections whatever to either the muscles or nerves in the stump of his leg. The apparatus simply strapped onto the healing stump.

The phantom pain continued. But as he grew used to the prosthesis and would occasionally forget that he was wearing the mechanical device, he would absentmindedly lean over to rub the location of his pain in an attempt to assuage it a bit. Each time, as his fingers felt the cold plastic of the artificial leg,

his hand recoiled as he discovered his 'foolish' mistake and he felt chagrined (much as does a university lecturer who, in the agonies of nicotine withdrawal, has been known to 'take a puff' from the piece of chalk in his hand).

After having worn the prosthesis for two years, Michael became quite used to the novel phenomenon of moving about the world with a pain that was 'not quite in his body'. The pain moved along with his body; but instead of being – like most persons' pains – *inside* the body of its sufferer, Michael's pain was outside his body, not terribly far removed to be sure, only 20 cm or so, but definitely outside.

There would be nothing particularly remarkable about this case were it not for what happened on the day of Michael's third annual checkup. His surgeon wanted him to try a new prosthesis. Where the original device had used stainless-steel strengtheners, this new one used carbon fiber. The new leg was attached and seemed to work perfectly. The doctor then asked Michael where the phantom pain was. Without looking, Michael reached down and tapped the knee joint and said "right here". The doctor grinned. "Mike", he said, "look what you've done. You've pointed to your old appliance, not the one you're wearing." Michael looked down, dumbfounded. He moved both mechanical legs about, first the one attached to his stump by contracting (what remained of) his thigh muscles, and then the discarded original prosthetic leg by taking it in his hands and waving it about. "I don't understand", he said. "The pain moves about as I shake the old unattached leg, not as I swing the attached leg. What's going on? The pain seems to be *in* the knee joint of the unattached prosthesis, not 20 cm below the stump of my right leg as it has been for the last three years." Panic was mounting in Michael's voice and the doctor was quick to recognize the symptoms. "Perfectly natural", he lied with all the credibility he could muster. "Here, let me give you a sedative. It will calm you down a bit." Michael did not protest.

The next few days made medical history. Michael's initial report of the incredible transference of the pain to his first artificial leg turned out to be irrefutable. In one test, Michael was blindfolded and the old prosthesis was moved silently about the room. Michael could unfailingly point to the pain and correctly judge how far it was from him.

But still stranger developments were in store. One of the

researchers working with Michael's surgeon wondered what would happen if another person were to wear Michael's first prosthesis, the one in which Michael's pain seemed now to be housed. Another amputee, who had never experienced a phantom pain, was fitted with the leg. Within eight hours, he too was complaining of a phantom pain: a pain, that is, several centimeters below his stump, as if he had a leg and there was a pain in it. The prosthesis was immediately removed, but it was too late. The pain did not subside, it simply moved farther away, as the prosthesis itself was moved about. Now there were two patients feeling the identical out-of-body pain! (The hospital administrators began to have visions of a malpractice suit.)

Obviously, research would have to proceed cautiously. Since the pain was not terribly severe, it was agreed to use paid, informed volunteers, explaining carefully to them that they might be left with a permanent out-of-body pain. Most interviewees were repulsed by the prospect, but a few were so intrigued that, in spite of the dangers, they actually begged to participate in this historical experiment. They gladly signed the necessary waivers absolving the researchers and the hospital of any legal liability. One even made a sizable donation to the hospital.

The experiment lasted for years. Every precaution imaginable was taken. Elaborate measures were instituted to prevent cheating or fraud. In due course, it was discovered that one did not have to be an amputee to experience the effect. Merely strapping the artificial leg to one's own good leg for twelve or more hours would induce the remarkable phenomenon. The subjects in the experiment were thereafter able to sense the pain even when the leg was removed out of sight to other rooms; indeed, the pain could be felt at distances of up to 150 km, well beyond the visual horizon. Ultrasophisticated electrograms revealed that there was spinal-cord and brain activity associated with this pain sensation but there were no particular afferent signals originating in the peripheral nervous system, i.e. the nervous activity giving rise to the sensing of 'distant pain' (as it had come to be called) appeared to originate directly within the spinal column itself. All the standard analgesics – morphine, acupuncture, etc. – remained as effective in alleviating distant pain as they were for ordinary in-body pain.

Many skeptics believed that the phenomenon of distant pain

was a case of mind-reading, mass hysteria, or some such thing; but a few of them volunteered for the experiment and, without exception, every one reported that he felt the pain in the artificial limb and that so far as he could tell he was not reading anybody else's mind. Eventually, some 237 different persons could all together feel the same pain, in the same place.

In fact, it was because of the great number of persons all sharing the same pain that the experiment was eventually stopped. Several of these subjects had increasingly, over the years, grown annoyed at feeling a distant pain which moved about from time to time. As long as the mechanical leg was left in one spot, the pain was so constant as to hardly intrude upon their consciences. But when the leg was moved about, in being fitted to a new subject, etc., the movement proved distracting to many of the previous subjects. Eventually several of them who lived within the critical 150-km radius of the research center protested the continued activity so vigorously that the matter was put to a mail vote. Over two-thirds of the affected subjects responded, and of them, more than 80% requested that the leg be retired. Michael's surgeon then donated the leg to his alma mater, the Harvard Medical School, where it may now be viewed in a sealed display case. No one, it is reassuring to mention, has ever experienced an out-of-body pain from merely looking at the leg.

Controversy over how to explain the phenomenon swirled for a generation. But the account which gradually seemed to win favor was that somehow (no one offered an explanation of this particular part of the phenomenon) Michael Robins had, quite unintentionally, managed to 'project' (or 'displace' [there really never was a very good word to describe such an unparalleled occurrence]) his pain into his first prosthesis. The pain was really 'there' *in the prosthesis*: it could be felt – and indeed would continue to be felt – by anyone who was keen (or fool) enough to strap it tightly onto, or alongside, his own leg for a half-day or more.

As the extraordinary discovery became widely known (*New England Journal of Medicine*, *Lancet*, etc.), was discussed in medical and psychology textbooks, and was the subject of several science programs on television ("Nova", "National Geographic", "The Nature of Things", etc.), a gradual broaden-

ing took place in popular thinking about pain. No one ceased to regard pains as sensations; but what did change was that people now began to talk easily and matter-of-factly about the possibility of pain sensations occurring not only within one's own body, but also in external objects where they might even be experienced (shared) by two or more persons. And when they came across early-twentieth-century philosophy texts which denied the logical possibility of pains being shared or being external to one's body, many of them thought it odd that any writer should have had such a blinkered conception. But, of course, persons who adopted this uncharitable view of their predecessors had not placed the earlier views in their historical context. Would they, themselves, have had a view much different from the earlier one if they had not witnessed for themselves the extraordinary events in Michael Robins's life?

There is, so far as I can tell, nothing logically incoherent in this possible-worlds story. It is, granted, wildly implausible. Indeed, I am certain that nothing like this is likely to occur in this world. But foretelling the future was not the point of the tale. The point of telling the tale was to see whether one could, without logical self-contradiction, describe a case of a public, objective, out-of-body pain. I submit that I have done this. If so, then we may conclude that it is no part of the *concept* of pain that pains *must* be private, internal sensations.

Certain philosophers will sometimes object to exercises such as the one just gone through here, by protesting that the results are 'linguistically deviant'. They might put their objection by saying, "You cannot say such things." By this they mean, not to deny that one can utter certain sentences, but that if one does, then one speaks nonsense. "No meaning has been given in ordinary English", they might say, "to an expression such as 'a publicly sensed pain in a nonliving physical object'." My reply is: The possible-worlds tale, once told, explains how the notion might apply, and thus succeeds – if there had not been an intelligible sense previously – in giving us that sense. If the notion was 'linguistically deviant', then it is no longer; if "public, objective pain" has heretofore lacked a sense, then it has one now. A concept need not apply to anything actual to be intelligible. There may never have been any unicorns, but the concept of *unicorn* is intelligible; there may never be any public, objective pains, but the concept of a *public, objective pain* is perfectly intelligible.

7.2 Case study: Unfelt pains

The sensing of pains (e.g. headaches) is usually regarded as quite unlike the sensing of ‘external’ objects (e.g. tables and chairs). External objects, we usually think, are not dependent upon our experiencing them for their existence: they can – and usually do – exist without our experiencing them. But pains are usually thought to be different sorts of things. Pains are thought to exist only insofar as they are experienced. A pain must be someone or other’s pain; there are no such things as ‘free floating’, unexperienced pains, in the way, for example, a chair may exist without being experienced. To account for this alleged difference, some philosophers adopt different theories of perception for external physical objects (e.g. tables and chairs), on the one hand, and internal ‘feelings’ (e.g. headaches, ennui, anxiety, euphoria, fear), on the other. For the first kind of experience, they will adopt what is called an “act/object” or “relational” theory of experience. The experiencer is regarded as being in a perceptual relationship with a certain kind of ‘object’.³ For the second kind of experience, they will adopt what is called an “adverbial” theory of experience. On this latter account, the experiencer is not aware of an ‘object’, but is having a certain kind of sensation. Pain, for example, on this account is regarded not as the object of an act of sensing, but is regarded as sensing *in a certain way*.

These two theories are not so much competing theories as they are complementary. They are designed to account for different *kinds* of experience. And it is thus possible for a person to subscribe to both of these theories without inconsistency.

But how viable is the alleged distinction which prompts the creating of the relational and the adverbial theories of experience? Is there something about the ‘felt aspect’ (the phenomenology*) of sensing colors, for example, which is different from the ‘felt aspect’ of having a pain, a difference which warrants our creating, and subscribing to, two different theories of experience? If the difference is not to be accounted for by anything inherent in the nature of the sensations themselves, then might it be something we learn, through science, about the nature of this world?

3. We will here ignore the various versions of this first theory. We will ignore, for example, whether the ‘objects’ of perception are physical objects or what some philosophers have called “sense-data” or “sensa”.

Two thousand years ago, Plato created one of the most enduring possible-worlds tales, his famous “Allegory of the Cave” ([156], book VII). He asked us to imagine persons growing up, chained in a dark cave so that they could see only straight ahead. On the wall in front of them were the shadows of moving objects which themselves were out of sight. (A modern version would be a person in a movie theater who is restrained so as to be able to see only the images on the screen and nothing more.) In particular, the chained person is unable to move about, to touch any of the things he sees. He can talk to other persons who are similarly chained, but none of them can see or touch one another. Plato’s purpose in telling this particular possible-worlds tale (what he called an “allegory” or a “parable”) was to argue that ordinary perception stands to reality as shadows do to the objects which cast them. I want to make quite another use of his story.

There are, to be sure, profound problems with the story as originally told. Putting aside questions of the unethicalness of chaining innocent persons in such abominable conditions, it is very unlikely that persons who are raised from childhood prevented from moving about and examining physical things ever could learn a language. Persons raised in Plato’s cave would be worse off than feral children, i.e. children raised not by human beings but by animals. Such children do not acquire language ([77], 246-8). But we will ignore these complications and pretend that the chained prisoners can see, talk, and hear, and that they occasionally feel pain, anxiety, hunger, and the like. What would they make of such sensations? In particular, would they have any reason to think that visual sensations and auditory sensations, for example, were any more like one another than either was like the sensations of pain and hunger?⁴ In not being able to explore

4. Jean-Paul Sartre (1905-80) writes: “... if I hear voices, what proof is there that they come from heaven and not from hell, or from the subconscious, or a pathological condition? What proves that they are addressed to me? ... If a voice addresses me, it is always for me to decide that this is the angel’s voice” ([185], 19-20). Sartre’s pre-eminent concern is our ability, indeed our being condemned, to make choices. That overweening interest sometimes makes Sartre overlook some very obvious distinctions. Under certain circumstances it might be *rational* to attribute a voice to another person, or to an angel, or to the effects of a drug, etc. Perhaps for dramatic purpose in this passage Sartre has suppressed these distinctions; but surely they matter. How *do* we go about sorting out the origins of the voices we hear? Only by being able to interact physically with other persons and with material objects. When

the world, they would have no warrant whatever for associating their visual sensations with images on a distant wall or their auditory sensations with vibrations in the larynxes of other human beings. Would they have any reason, then, to believe that these visual and auditory sensations were of external objects, while pain and hunger were not?

Under the circumstances described, there would seem to be no grounds whatsoever to prompt the distinction between ‘internal’ and ‘external’: sensations would simply exist and that would be the end of it. Why should they attribute to other persons the ‘voices’ which answer their questions? These ‘voices’ just appear in their auditory spaces, just as hunger pangs occur from time to time in their gustatory spaces, and colors in their visual spaces.⁵ Without being able to move about the world, they would regard all sensations on an equal footing in that their sensations would flit in and out of consciousness. Colors would be different from sounds, both would differ from smells, and all would differ from pains. But other than the fact that colors, smells, sounds, and pains would all form distinct categories, there would be no grounds on which to regard any of them as ‘internal’ or any as ‘external’.

If a conclusion can be drawn from our use of Plato’s myth, it would be that there is nothing inherent or intrinsic in the having of a sensation which marks it as ‘external’ or ‘internal’, as being better explained by a relational theory of experience than by an adverbial theory. Why we treat colors and sounds as ‘objects’ and pains as ‘manners of sensation’ has to do, not with the phenomenology of these sensations, but with certain empirical facts we discover about this world. We discover empirically – experimentally and not introspectively – that other persons share our visual sensations but do not share our pain sensations.

A.J. Ayer put the point this way, in his own possible-worlds tale:

Suppose, for example, that people’s feelings were very much

the latter opportunities are denied to us, then – and only then – might our decision as to the cause of the voice come down to a matter of ‘choice’. For ordinary circumstances, Sartre has considerably overstated the role of choice. But for Plato’s cave, where the denizens are physically restrained and unable to probe their environment, something like Sartre’s (arbitrary) choice is their only recourse.

5. We will have more to say of this use of “space” in the next chapter.

more uniform than they actually are, so that whenever anyone felt bored, or happy or angry, or depressed, his neighbours nearly always felt the same. In that case, we might very well find use for saying that there was not a multiplicity of feelings, one to each person, but a single feeling, one and the same for all, which different people experienced in different ways. Certain people might fail to experience it at all, just as certain people fail to perceive physical objects which are in their neighbourhood. There might be illusions of feeling, corresponding to illusions of perception. But the feeling would still be there, just as the physical object is there whatever illusion someone may be having. To make the analogy with physical objects closer still, one might make it possible [by telling a possible-worlds tale] for feelings to exist when no one was actually feeling them. This might be said in cases where the normal conditions in which the feeling habitually occurred were present, but some special factor, such as the drugging or hypnotizing of the person in question, intervened. ...

The point of this fantasy is to show how *the distinction between what is public and what is private depends upon a contingent matter of fact*. We do not find it useful to publicize [i.e. regard as being public] feelings, or sensations, or thoughts, or images, because they vary so much from person to person: we do find it useful to publicize physical objects because of the extent to which the perceptions of different people agree. But it is not difficult to imagine that the two should be on a level, or even that the position should be reversed. ([18], 201-2; italics added)

For Ayer, what makes for the ‘internality’ of pains and the ‘externality’ of tables and chairs is something we discover, not by examining our own pain sensations, our own visual sensations, our own tactile sensations, etc., but by examining the world and the reports of other persons.

For Ayer, it is a contingent fact about this world, nothing inherent in the nature of pain sensation itself, that pains are not public objects like tables and chairs, and that pains do not exist unperceived. Ayer is quite prepared to allow (as I have argued in section 7.1) that pains could be public, shareable sensations. And Ayer goes one step further. He suggests that, were the world a certain way, we would want to acknowledge the existence of unperceived pains. If a pain can be

experienced by several persons (again, see section 7.1), then were one of those persons to be shielded from the pain, by medication, hypnosis, etc., the situation would be perfectly analogous to a person's being shielded from perceiving a physical object by, e.g., blindfolding him or erecting a wall across his field of view. Just as the obscured physical object would still be regarded as existing, although not perceived, so too would it be reasonable to regard the pain as still existing although it, too, was not being perceived.

In short, unfelt pains are no logical impossibility. That they do not exist is an empirical discovery we have made about this particular world. There is nothing inherent in the concept of pain, or in the sensation of pain, to preclude public, shareable pains, pains which can exist without being felt by anyone. The alleged differences which have prompted the creating of two side-by-side theories of experience, the relational and the adverbial theories, are differences which are not inherent in the nature of sensation itself but only in empirical facts we have discovered about other persons.

7.3 Case study: Pains in the absence of nervous systems

In the last one hundred years, we have learned enough about neurophysiology to be able to say confidently that our pains are, with virtual certainty, attributable to chemical and electrical activity (i.e. certain specific physical 'goings-on') in specific substructures of our central nervous systems. Even in cases of so-called psychosomatic pains, we have good reason to believe that often, if not always, such pains are, again, attributable to physical goings-on, although, unlike the cases of physical injuries and trauma, not originating in lesions, bruises, infections, etc. In short, we are much inclined – being knowledgeable of some of the data of modern medical research – to attribute the sensing of pains to physical activity in our bodies.

The possible-worlds tale of section 7.1 – which told of the sensing of an out-of-body pain (the strange affair of Michael Robins and others) – still capitalized upon, indeed invoked, the very account being reviewed here. Although the 'initiating cause' of the distant pain was not the firing of the sensitive ending of a peripheral nerve, the pain sensed did come about because of certain activity in the spinal cord and the brain. To that degree, that tale did not depart from contemporary scientific accounts.

Nevertheless, in spite of this familiar scientific background, I want to insist upon a *conceptual* distinction between pains or, better, 'the

felt aspect of pains' and their physical origins. This particular distinction, between pains (or the felt aspect of pains) and their specific physical origins in the central nervous system, is one of the most troublesome for some persons to grasp, and takes us right to the heart of one of the most difficult and controversial distinctions in philosophy, one which some writers insist upon and which others reject.

Suppose, as I am perfectly willing to concede as being highly probable, that all pains without exception are attributable to certain kinds of physical goings-on in a living creature, and suppose, further, that whenever there is an occasion (episode) of such physical goings-on there is a felt pain. At least two theories, as we have earlier seen, have been proposed to explain such correlations: either the pain literally *is* the physical goings-on themselves (this is part of the so-called identity theory of mind and brain) or the pain is *caused by* the physical goings-on (the so-called causal theory). But whichever of these theories we might want eventually to settle upon, I would still want to insist upon the *conceptual* difference between one's pains and physical goings-on in one's body.

Now this latter claim – that pains are conceptually distinct from the states of the central nervous system which 'account for' the presence of those pains – is bound to strike some persons as particularly strange. For how can I, on the one hand, allow that pains might literally *be* physical goings-on and yet, on the other hand, insist on the *conceptual* difference between pains and physical events or states? It all depends, of course, on what exactly is meant by 'a conceptual distinction'.

In the analysis I have been proposing, two concepts are distinct if, for each of them, it is *logically possible* for there to be a situation (thing, event, state, etc.) to which that concept applies and not the other one. According to this analysis, then, one can say that there is a conceptual distinction between having a pain and having the brain-state which 'gives rise to' (or even 'is') that very pain, if it is *logically possible* to apply the one concept to an organism and not the other.

But as so often happens in philosophy, we have answered one question only in turn to have prompted another. For now the question arises, "How shall we tell when attributing one concept to a thing while withholding another is logically possible?" We can of course *say* such a thing as "He is in pain, but nothing is happening to him physically to account for that pain", but we also know that it is possible to say things which (if taken literally) describe logically impossible situations, e.g. "He traveled across town in an empty taxi" or

“Her elder brother is an only child.” Being able to *say* something does not make what is said logically possible. The test for logical possibility must be something else.

Again, we invoke the technique of constructing a possible-worlds tale. To show that there is a conceptual difference between the two concepts, we might attempt not just to utter a single sentence or two, but to fill in details, to tell a more robust tale, in which it is appropriate to apply one of the two concepts to the situation and explicitly withhold the other. If this can be done *without contradiction*, then this may be taken as evidence that the two concepts are – after all – distinct.

Thus, in the case of pain, if we want to argue for the *conceptual* difference between pains and physical goings-on in a central nervous system, we might try to construct a possible-worlds tale in which creatures experienced pain but in which those same creatures did not have the physical goings-on that are usual when you and I feel pain. Here is such a tale.

Once upon a time there was a universe in which there was a planet, Htraenon, very like Earth in certain respects but very unlike it in others. Outwardly its creatures resembled those of Earth remarkably, but internally they were surprisingly different. The ‘human beings’ of that planet looked pretty much as human beings do on Earth: each had a head, two eyes, two ears, a nose, a mouth, two arms, a torso, two legs, ten fingers, and ten toes. By our own standards, we would consider all of them extremely good-looking. Their social life was much like ours as well: they laughed, told jokes, complimented one another, prayed, sang songs, elected officials, instituted moral codes, educated their children, exchanged goods and services for money, and occasionally indulged in metaphysical speculations. They were also subject to many of the same sorts of ailments and frailties that befall the likes of you and me. They had illnesses, they suffered occasional pains, they sought medical treatment, they grew old, and each of them eventually died. But in one respect they were very different from us.

Up until about the time they began to use microscopes and electronic devices in medicine, their history and science were unfolding in a parallel fashion to our own. But once they began to examine the internal functionings of their own bodies, using devices much more sensitive than their own eyes and ears,

they discovered a physiology which is remarkably different from yours and mine. For one thing, they had no central nervous system. It is not that they failed to find such a thing. It went beyond that: they really had no central nervous system at all. Now failing to find a central nervous system did not surprise them in the slightest. After all, they had had no reason to hypothesize, and even less reason to believe, that such a thing existed. They had no beliefs attributing their pains to specific goings-on in a central nervous system, no more so than they had beliefs at the beginning of their modern period of chemistry about glands in their bodies releasing hormones to regulate, for example, the level of glucose in their bloodstreams.

In time, as their science grew in sophistication and they continued to perform physiological, chemical, and microscopic examinations of their own bodies, they discovered within themselves a hormonal system, came to understand its regulatory functions, and even managed to synthesize all of these hormones in their pharmaceutical laboratories. But they never made comparable discoveries about a central nervous system, and this for the aforementioned reason that they did not have a central nervous system. Thus, it was inevitable that these people simply persisted with their timeworn theory that pains, thoughts, musings, etc. were somehow features or characteristics of themselves, but did not think it necessary, or warranted, to attribute them to any particular physical goings-on in their bodies.

When, in due course, several centuries later, scientists from a distant galaxy arrived on Htraenon for the first time and discovered that the Htraenonites had no central nervous systems, they were initially dumbfounded. But soon, the incredulity of the visiting scientists crumbled. After only a few weeks adjusting to such an unexpected jolt to their beliefs about the physiology of sensations, most of these latter scientists found themselves working handily with the idea that some creatures (themselves, for example) feel pain as a result of certain goings-on in their central nervous systems, but that other creatures (e.g. the Htraenonites) feel pain without, so far as could be ascertained, anything in particular happening within their bodies, besides the injuries themselves, which could be correlated with those sensations.

Although the world just described is merely a possible world, i.e. is not this, the actual world, it is meant to mirror the history of the actual world up to the point in our own history when we began to discover the existence of, and the secrets of the functioning of, the central nervous system. Until that point in our own history, we (actual) human beings did not associate pains with goings-on in a central nervous system. So ignorant were our forebears of the existence and operations of the central nervous system that we find in Aristotle, for example, the perfectly serious hypothesis that the brain was nothing more than an organ to ‘cool the blood’. Aristotle knew nothing, nor could he have, of the manner of connection of certain sensory organs, e.g. of touch and taste, with the brain. And thus he thought it demonstrable that the brain was *not* involved in sensation: “This brain ... has a character peculiar to itself, as might indeed be expected. That it has no continuity with the organs of sense is plain from simple inspection, and is more clearly shown by the fact, that, when it is touched, no sensation is produced. ([12], book II, 652b1-10) ... Nature has contrived the brain as a counterpoise to the region of the heart with its contained heat, and has given it to animals to moderate the latter ... The brain, then, tempers the heat and seething of the heart” (652b20-7).

Did Aristotle – knowing nothing of modern neurophysiology – have a concept of *pain* different from ours? Certainly we may suppose that he, like us, had experienced pains. He, doubtless, from time to time, had pricked his finger, cut his hand, stubbed his toe, suffered a toothache, and endured a headache. He was, we may be sure, familiar with pain in many of its forms. But he knew nothing of peripheral nerves, of Ad- and C-fibers, of electrical pathways in the spinal cord, of the firing of nerve cells, or even for that matter, of the existence of nerve cells. As a matter of fact, he did not even have the concept of *nerve*, of *cell*, of *electricity*, of *endorphin*, etc. Could he, then, have had *our* concept of pain? I want to suggest that he did, that he could have understood, as well as any of us, claims about persons being in pain, about certain pains being more intense than others, about certain medications’ ability to relieve pain, about most persons trying to avoid pain, etc. All he would lack is a twenty-first-century scientific explanation of the physiology of pain. But that information ought not, I suggest, to be regarded as part of the *concept* of pain itself. Scientists, as they pursue neurophysiology, are not refining the concept of *pain*, but are furthering our knowledge of pain, its causes, and its relief.

7.4 Case study: Must pains hurt?

Squares have four sides; the edible parts of pineapples are yellow. But whereas it is part of what we have called the “narrow” analysis (i.e. it is *analytic*) of the concept of *square* that squares have four sides, it is not part of the analysis of the concept of *pineapple* that the edible parts of pineapples are yellow. After all, cabbages may be white or purple; ripe apples may be red, yellow, or green; grapefruits may be white or pink; etc. In this age of hybrid fruits and vegetables, we have grown used to the appearance on grocers’ shelves of produce in an ever-increasing variety of colors. To date, all the pineapples on the market have yellow flesh, but few of us would be unduly surprised to discover one day a product identical to present-day pineapples but which differed in hue, being orange or pink perhaps. Given how we have handled analogous cases in the past, we probably would not in the slightest be tempted to argue that these orange-colored fruits were not pineapples; we would simply allow that pineapples could come in more than one color.

Insofar as it is analytic of the concept *square* that all squares have four sides, it is impossible to tell a possible-worlds tale, without contradiction, in which there are squares having other than four sides, e.g. five sides. But insofar as being yellow is *not* analytic of the concept *pineapple*, it is perfectly possible – as I have just done in previewing a possible future state of this very universe – to describe without internal contradiction a situation in which there are non-yellow (orange perhaps) pineapples.

What about pain? Is *being hurtful* or *causing hurt* analytic of the concept of *pain*? Would a sensation even be describable as pain if it did not hurt? Is a non-hurtful pain logically impossible in the way, for example, a five-sided square would be? Or, is a non-hurtful pain rather more like an orange-colored pineapple – unusual to say the least, perhaps never even part of one’s own experience, but logically possible nonetheless? Once again we take recourse to telling a possible-worlds tale, this time in an attempt to describe a pain which does not hurt.

Mr. J.R. had very advanced cancer of the neck and jaw. When pain became unbearable despite huge doses of narcotics, a frontal lobotomy was discussed with his family and finally performed in an effort to make his last few months comfortable. Under local anesthesia, small drill holes were made in the skull

over the frontal lobes, and *the fibers connecting the frontal cortex with the thalamus were severed on both sides by an instrument lowered into the brain*. For several days after the operation he was sleepy and confused, did not know where he was, and had trouble recognizing his family. This cleared, however, and he seemed cheerful and alert. *He did not complain of pain and stopped asking for pain medication, but when asked if he still had pain he said, "Of course I do, it's the cancer, right here", pointing to his diseased jaw*. Examination revealed his perception of temperature and pin [pricking] to be acute – he actually “jumped” [in response] to the pin prick, and complained of being “tortured”, something he had never said before.

Most [similarly treated, i.e. lobotomized] patients said, unemotionally, that pressure on the tumor was still painful but they were obviously not disturbed by the pain. There was a chasm between the affective emotional aspect of pain [i.e. its hurting] and the pure sensation of pain.

The sensory and emotional aspects of pain can [also] be dissociated [from one another] by certain drugs. For example, high doses of antianxiety agents such as certain tranquilizers do not seriously impair discrimination between stimuli ranging from painless to extremely painful. However, even though subjects may report certain stimuli as being excruciatingly painful they do not seem to care. The drug appears to leave the sensory aspects of pain intact, while almost completely suppressing the emotional aspects.

Is such a tale really possible? Is it really possible – as this tale alleges – to feel a pain *and* yet not have it hurt? It turns out that it *is* possible, indeed that it can be proven to be possible. And it can be proven in the strongest fashion in which we can prove anything possible: by showing that the alleged possibility is *actual*.

The possible-worlds tale just told – of patients who could feel pain but who experienced no hurt – is not the product of a philosopher’s imagination as were the tales of the three preceding sections, but is compiled from actual clinical studies.⁶ This assembled story is, of

6. The first of the three paragraphs indented just above is a direct quotation from *Basic Human Neurophysiology*, by Lindsley and Holmes ([123], 117); the second, from *The Nervous System*, by Peter Nathan ([142], 105); and the

course, still a possible-worlds tale. Its only difference from the usual possible-worlds tales of philosophical texts is that this one happens to be *true* (i.e. is true not just in some other possible worlds, but is true in this, the actual world, as well).

Someone might object: “Pains which do not hurt are a logical impossibility. There cannot be, in this world or any other, pains which do not hurt. The patients who described themselves as feeling pains which did not hurt had *misdescribed* the situation. They certainly had some sensation; but just insofar as it did not hurt, then it was not pain. Perhaps we currently lack a term for such sensations in our language, the phenomenon being so rare. Nonetheless, such sensations cannot be described as pains. Whether we have a name for them or not, they must be regarded as something other than pains.”

The possible-worlds tale just told, then, will not be convincing to everyone. As we can see, it is possible for someone to reject the apparent conclusion of such a tale by arguing that the persons in the tale are misdescribing their sensations. How, now, can we possibly resolve this latter debate? The possible-worlds tale may at first seem convincing to some persons (it is to me, for one), but others can – if they are inclined – find grounds to reject it.

At this point we must be very careful not to think that there is some one ‘right’ or ‘wrong’ answer to the question. There would be only if there were some *independent* way, other than our agreeing to use a concept in a certain way, to ascertain when a concept is used correctly or incorrectly. But there is no such way. Our concepts are our own inventions. We do not discover them. If concepts were not of our making, but the sorts of things we could examine to see what is ‘really’ entailed by them and what not, then we could – theoretically – discover, for example, that we have had the concept of *square* wrong all these many years. We had thought that all squares are four-sided, but having now examined the concept of square we see that we had made a mistake: squares may have either five or eleven sides. It is clear that such a notion would be nonsense. There is no such independent concept of *square* that, if we are careful and attentive, some of us will get right, but if we are careless, all of us might get wrong. It is impossible for *everyone* to have a mistaken notion of *square*.

The question boils down to this. Virtually any pain any of us has

third, from *The Neurosciences and Behaviour*, by Atrons and Curthoys ([14], 93).

ever felt has hurt. Is the hurting to be identified with the pain – in the sense that nothing logically could be a pain if it did not hurt – or might – given the pressures of certain kinds of peculiar circumstances – we want to allow that not all pains hurt?

The cancer sufferers who underwent the frontal lobotomies outlined above described their sensations, at the sites where they formerly had pain, as still being pain. The difference, they alleged, was not in their ceasing to feel pain but in that pain's no longer hurting.

The concept of *pain*, like so many other concepts – e.g. of *person*, of *fairness*, of *duty*, of *consciousness* – is in a state of flux. If we are to judge by the on-the-spot linguistic behavior of medical patients and medical practitioners who are involved with the amelioration of pain, all the indicators are that our language is evolving toward making a distinction between sensing a pain and experiencing hurt. The two concepts are coming apart. If, someday, a safe medication is developed which – unlike the opiates, which suppress the sensation of pain altogether – acts like the chemical equivalent of a frontal lobotomy or like a massive dose of certain tranquilizers in that it (here I must use current terminology) suppresses not the pain, but only the accompanying feeling of hurt, we will be hastened toward making a sharp distinction between the two concepts. It may even happen that persons in the future will wonder how it was that we ever confused the two, so used will they have become to taking a tablet when they have a headache to relieve, not the pain, but the hurt.

What, then, shall we finally conclude? Does the concept of *pain* include the concept of *hurting*? I think the answer must be something like this. So frequently are pains accompanied by hurting that we invariably tend to conflate the two concepts. Even so, they can (in my opinion), reasonably, be regarded as logically, or conceptually, distinct. And given certain as yet unrealized developments in medical research, we would in time virtually be forced to use these concepts separately. We have enough empirical data now in hand to suggest that our language just might evolve in that direction. We cannot foretell with any certainty that it will, however. From a philosophical point of view, all we can do is to prepare ourselves for that eventuality.⁷ We must be sufficiently flexible in using our concepts so that

7. "When the hypothetical upheaval is sufficiently radical we have to go through the agonizingly innovative process of rebuilding part of our conceptual scheme from the ground up. Genuine conceptual innovation is necessary

we are prepared for the kind of intellectual reorientation that might be required.

Summary of sections 7.1-7.4: For the purposes of illustrating one way in which philosophers probe concepts, I have subjected the concept of *pain* to analysis by invoking it in a number of possible-worlds tales in which some quite extraordinary situations have been described. I have tried to show four things: that although pains typically occur within our skins, that although all pains are felt, that although pains typically are the product of nervous systems, and that although pains typically hurt, none of these features is logically entailed by the concept of *pain*. It is possible, I have tried to show, that pains could (logically speaking) occur outside our skins, and even for that matter might be public in the way in which tables, sounds, and aromas are public, more particularly, might exist unperceived; that pains could (logically speaking) occur in creatures which lacked central nervous systems; and that, finally, pains could (logically speaking) occur without a sensation of accompanying hurt. My expectation is that of the several things being claimed, the last – that pains could occur without there being a sensation of hurt – will be the most difficult for many readers to accept. And yet, of the several, it is the one for which there is the best empirical evidence.

The *point* of pursuing such exercises is threefold. First and foremost, some persons find such conceptual explorations fascinating. For such persons, that alone merits the pursuit. The second justification lies in the sharpening of our conceptual tools. Only in understanding what is and what is not entailed by our concepts can we aspire to use them well. If we are going to think, then our thinking cannot be any better than the tools we use. If the tools are dull, the product will be rough and crude. If the tools are well made, and their interrelations understood, then at least we have a chance of making something beautiful and useful with them. We can no more do philosophy well without critically examining our concepts than we could do genetics without a microscope or physics without mathematics.

The third point of the exercise is to unfetter our imaginations. In

and there is no way of predicting its outcome. To the question ‘What would you say if ...?’ we would in such cases have to reply: ‘We just wouldn’t know what to say . . . We’ll just have to cross that bridge when we get there’” (Rescher [169], 114).

arguing that certain concepts are not essential to the concept of pain, I have been attempting to *enlarge* the possible application of the concept. By peeling away inessentials, I have tried, not to narrow our concept of *pain*, but to stretch it. Although pains typically, perhaps always, occur within the bodies of creatures having nervous systems, and although pains usually hurt, might we not want to consider the possibility of pains outside of bodies, in things quite unlike us in structure, and unaccompanied by hurt? Only in imagining possibilities beyond the commonplace have we any hope of understanding this world. For understanding this universe (or any other one) consists in large measure in seeing the difference between what *might be*, what *is*, and what *must be*.

Postscript (Added, second edition)

This chapter, in which I have argued for (among other things) the conceptual possibility of pains existing outside of one's body, indeed even being shared by two or more persons, has proven exceedingly troublesome for some readers.

One reviewer of *Beyond Experience* singled out this chapter for special comment:

... a (rare) place where Swartz's discussion loses credibility is his alleged possible world (see pp. 118 ff.) where we would call something a pain that was literally outside of our body. While it is fun and often enlightening to envisage possible worlds, surely the experience of pain, the having a pain, is always in one's brain. In the actual world, pains are private because they are brain states. If Swartz is trying to show only that they might not have been brain states, he would be right. But so long as pains are experiences and not objects of perception, they cannot sensibly be said to be shared (unless by Siamese twins sharing a brain) or public. Pains could not literally be out of the body and at the same time experiences. If we imagine a world where pains are analogous to colours, of course, this is not necessarily true. —Jack Ornstein, *Canadian Philosophical Review*, vol. XII (2)-(5) (April 1992-Oct. 1992), p. 355.

By way of defense, I would reply that if pains can be likened (there are of course differences) to tactile sensations, then there is not only a

conceptual possibility of having pain sensations outside of one's body, there has been, more recently, actual laboratory evidence bearing on the question. I reproduce here an article (appearing seven years after this chapter was first published and six years after Ornstein's critique) reporting an experiment that demonstrates, not merely the possibility, but the *reality* of out-of-body tactile sensations. (Sometimes, it turns out, philosophical theorizing presages scientific discovery.)

This Here Hand Is My Hand, I Think

Participants in a recent psychological study will probably never look at mannequins – or their own bodies – in quite the same way again. Before the study, they knew their arms belonged to them and synthetic ones didn't, simply because seeing is believing. Now they're not so sure.

Researchers at Carnegie Mellon University in Pittsburgh asked subjects to keep their eyes on a rubber arm that was sitting on a table in front of them. With the subject's left arm hidden from view by a screen, the researchers simultaneously stroked both the rubber hand and the subject's hand with a paintbrush. Even though they knew their own hand was being stroked behind the screen, nearly all the subjects experienced the same bizarre sensation: they felt the fake hand was actually their own.

According to Matthew Botvinick, the Ph.D. psychology student who coauthored the study with advisor Jonathan Cohen, awareness of self seems to depend on intricate conversations between the brain and a range of sensory inputs that it constantly receives. If those conversations become garbled by contradictory messages, the brain is even willing to stretch the bounds of where the body ends and the outside world begins in order to draw a coherent picture.

"It's like ventriloquism," says Botvinick, who was so spooked by the illusion when he tested it on himself that he let out a yelp and threw the fake hand across the room. "In the experiment, when something touches the fake hand, you feel it, so the rubber hand appears to be an object with which you sense. And when there is an object of that kind, it's usually part of you. That seems to be one basis of self-identification."

To confirm that the subjects were experiencing a true shift in

their perception of themselves, researchers asked them to run their right index finger along the underside of the table until it was directly underneath their left one. Those who had experienced the rubber-hand illusion invariably missed their real finger altogether and pointed more closely to the fake hand.

“When you look at your hand, it doesn’t feel as if your brain might be going through all kinds of complicated computations to arrive at the conclusion that this thing is yours,” says Botvinick. “You just know it’s your hand.” —Jennifer Van Ezra, in the column “Nexus” in *Equinox*, no. 99 (July 1998), p. 14.