

Drowning the Subject

The Stone Pharisee

Copyright ©2014 by The Stone Pharisee

Distribute freely.

ISBN ... Nope. This is a work in progress.

Of experience we can not speak. Measurement will
be negotiated.

Contents

1	Identification	3
2	Creation	19
3	Complementarity	37
4	Consensus	47

Prologue

In Alejandro Jodorowsky's film *The Holy Mountain*, ten pilgrims assemble to set out on a spiritual quest. Before they leave, they attempt a head count, but each of them can only see 9 pilgrims, so each says that someone is missing. A bowl of water is then brought out, and each looks into the bowl, sees his or her own reflection, and says that the tenth has now been found. "But he drowned" they add. The purpose of this small book is to be that bowl.

Version

This is a snapshot of a work in progress. But as it's been in progress for 10 years or so, it might as well be available as far as it has got. This version (Version 1) gets as far as consensus. I do not promise the end point will ever be reached.

Chapter 1

Identification

An anamorph is a form seen from a specific viewing position only. There are entertaining street artists who will create enormous chalk drawings that are properly viewed from one point only. When the correct point is chosen, a vast three-dimensional landscape reveals itself, with yawning crevices and cracks. Viewed from anywhere else, the picture distorts, cracks, and splits into unallied pieces. Most anamorphs work by exploiting the geometry of perspective, such that the optimal configuration can only be seen from one place. The best view of an anamorph is from a very specific viewing point. In exploring the world with our eyes, we move around a lot, and so a photograph, which can be taken from precisely one point, may better capture the optical gestalt than a live viewing. Figure 1 shows a modern anamorph by the Swiss artist, Felice Varini. Viewed from the designated viewing spot, or vantage point, a simple ellipse appears. Viewed from anywhere else, the form breaks apart.

An exceptional anamorph is found in the Neolithic burial chamber of Newgrange, in County Meath, Ireland. The burial chamber lies in the middle of a large mound, and is accessed only through a narrow tunnel. On one day in the year, the

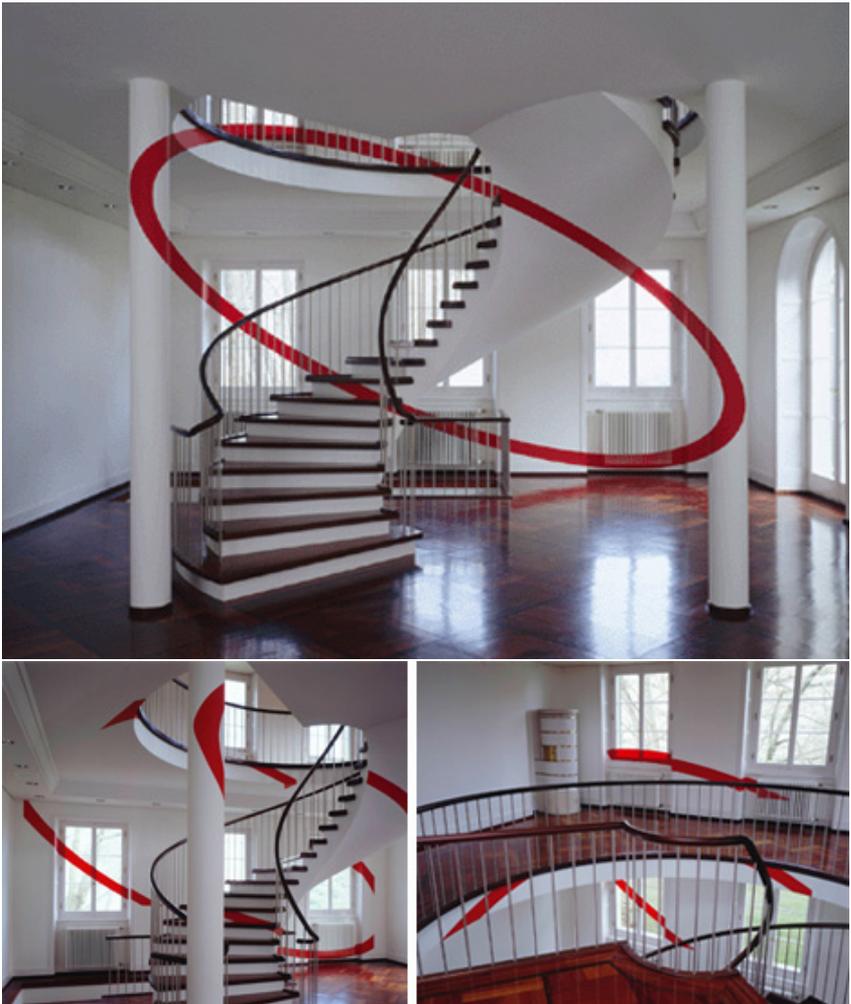


Figure 1.1: Example of anamorphic art by Felice Varini.

winter solstice, the light from the rising sun enters the passage of the tomb and reaches the central chamber. A remarkable feat of engineering, accomplished 2000 years before the pyramids of Egypt were built. Of course the chamber is small, and there are far more people who would like to see the sight than can be accommodated. The waiting list is many years long and subject to a lottery. This is, of course, not representational art. Unlike a merely spatial anamorph, this venerable instance requires the viewer to be in a specific place at a specific time—a 4-D anamorph, if you will!

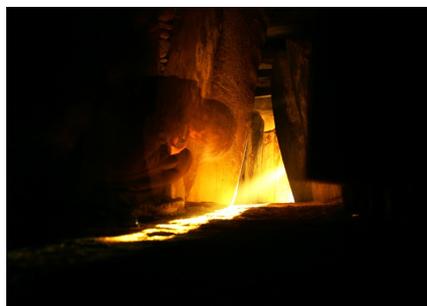


Figure 1.2: Newgrange Tomb. Left: from the outside. Right: sun enters the burial chamber on the Winter solstice.

The artists behind these works manage to do something that is rather difficult. They point you to your own location in space (and time). The work can not properly be viewed unless the subject has a specific here and now. The subject is thus drawn into the work, unlike, say the Venus de Milo, who will continue to stand gracefully when the tourist wearily moves on to the gift shop. For me, the magic of an anamorph is that it transcends any dichotomy between the subjective (taken as the here and now of an individual) and the objective (they exist “out there” in the world). They also blur the distinction between the individual (the viewing subject is incorporated in the work) and the collective (the work is, after all, made by

another).

The anamorph is a useful reminder that we always see the world, hear the world, experience the world, from a specific point of view. Consider the world as experienced by you right now—where “now” can be taken to be an arbitrarily small interval of time. You are reading this, so it is safe to assume that there is a text, viewed from a specific position. There will be some sounds, that extend out into the world only a little way, as far as walls, windows, and the passage of vibration through the medium of air allow. The sound of a passing train is not a freely existing entity; it is a sound as heard by you, at a specific distance, filtered through the surfaces that surround you. Each person who hears that same train at the same time will hear a specific version of that sound—an infinity of little anamorphic versions. There may be more in your little personal bubble of experience: The feel of the chair under you, the rumblings of an empty stomach, or perhaps the aftertaste of a recent cup of coffee. As attention roams over the here-and-now, other elements of experience become manifest. Your experiential domain has the character of an anamorph, arising in the relation between subject and world.

It is to the notion of a domain of experience that we now turn. In the language we use to describe our selves and our lives, there lies an assumption or presupposition that is somewhat embarrassing to admit to, but which we seem to be unable to do without. The assumption I speak of is the notion that there are experiences that are proper to one person, and unavailable to any other. Together these experiences necessarily constitute a domain. What shall we call this domain? What name is appropriate for the set of experiences that are proper to a single person at a single point in time?

The sensory qualities of immediate experience are perhaps the simplest facets of which we might uncontroversially speak. Pain is the clearest example. We all know that the expression

“this hurts me more than it hurts you” is not to be taken literally. There is *some* important sense in which the stinging felt after a slap belongs to the slapped, and not to the slapper. Feelings, memories, longings, and daydreams surely all belong here somewhere, but talking about them, and conveying their character to another, is far more difficult. If it were straightforward, we should have had no need of such writers as Proust, Joyce and Woolf. It was a psychologist and philosopher, William James, who gave us the term “stream of consciousness”, but it was his brother, Henry, and the novelists who came after him who put flesh on those descriptive bones in the modernist novels of the early 20th Century. In the narrative of the modernist novels, the reader is encouraged to be “inside the head” of the characters, as we see the world through their eyes, and follow the associative leaps of their thoughts. Events are imbued with significance that depends upon the past of the character; tell-tale details are interpreted with respect to a specific history.

In what follows, I will sketch a way of thinking about the “here and now” that attends every subject. This is perilous territory, and the reader is probably already weary of being confronted with unconvincing accounts of minds, consciousness, and cognition, for there are many such to choose from. One often hears that consciousness is the last and greatest of the scientific mysteries that so-far refuses to be shoe-horned into an increasingly sophisticated physical account of the cosmos. Or one hears that such concerns are misplaced, as there is nothing irreducibly different about consciousness and the mental. Marvin Minsky, one of the founders of the field of artificial intelligence, famously dismissed such concerns with the brash statement that “minds are what brains do”.

I do not have the one story to settle such disputes, nor do I believe a single account exists that will settle the matter. My purpose in this book is somewhat different. It is to

illustrate how one can think, talk, and reason about matters that have seemed to lie beyond the explanatory competence of the naturalistic enterprise, or science. And there is plenty of reason to want to do so. Bedazzled, as we all are, by the fantastic successes of the naturalistic attitude that has given us 21st Century science and technology, it is easy to overlook the important fact that there are very many, uncountably many, areas of human activity and living where science has been of absolutely no use whatsoever. The inbuilt inequities in every society, the persistence of wars of aggression, the clash of religions and ethnic groups, and the manifest difficulties that arise in planet husbandry all testify to the strict boundedness of our utilitarian science, and all suggest that we could readily do with more consensus and a better way of integrating the use of reason in the conduct of human affairs.

Some of the most impressive science we have is the basis for predicting the movements of inanimate objects. We can hurl a lump of metal the size of a small car at Mars, and hit it, have the object land gracefully, and have it call home. That is no mean feat. But at the same time as that is happening, our economies are crashing and burning, millions of people are displaced from their homelands, driven out by conflict, a new war rages out of control in the Middle East, and the relationship between our activities and the climate of our fragile planet lurks ominously in the background, threatening to intrude with extreme prejudice. We do not seem to have the same kind of precise control over human affairs. Of course one could properly object that control in human affairs is a threatening thing, smacking of dictatorship and the imposition of one world view on another. But utter impotence is not a good alternative.

When we discuss human affairs, we can not omit the notions of ethics and value. Unlike the sciences that deal with physical materials, claims and counterclaims in the realm of

human affairs must traffic in the currency of significance and meaning. The humanities, rather than sciences, have traditionally been home to much of this discussion. David Hume long ago drove a wedge between description of what *is*, and prescription of what *ought* to be, and by the latter half of the 19th Century, the scientific orthodoxy seemed to adopt a thorough-going materialism that persists obstinately today, even as our ideas of what matter and the physical world are have changed beyond all recognition. “Being” cannot be regarded as a simple matter, as if it needed no framing, and the notion that we can describe that which *is*, without troubling ourselves about how that description is inextricably convoluted with our own constitution and even beliefs, will no longer stand [5].

We need to harness the practice of science and apply it more effectively to human concerns. But human concerns necessarily encompass the realm of the subjective, and notions of value, significance, and meaning. These seem foreign within the domain of science as currently practiced. And so science itself must grow. Whether it is still called science as it grows is of no serious concern. The borders between science, religion and philosophy are notional and will forever defy precise delineation. There has been an alarming rise recently in a form of scientistic fundamentalism that has sought to draw firm boundaries between science and religion, insisting that they stand in opposition to one another. This shameful short-sightedness threatens to make science useless, incapable of addressing questions that concern **us**, whatever that pronoun may refer to. Daniel Dennett is a wonderful philosopher and Richard Dawkins is a marvelous biologist. Both have the knack of spinning a good yarn while conveying matters of genuine substance to a general audience. Yet both have regrettably fallen prey to this kind of intellectual provincialism, insisting that science, as they conceive it, is distinct from the frippery

and superstitions of religion, as they choose to characterise it [3, 2].

The kind of strident atheism both authors preach arises in a specific cultural context, where there are genuine problems in the interactions among belief sets, specifically those of Christian fundamentalists, predominantly in the United States of America, and scientists and educators, who wish to instill a fondness for learning and “truth” in subsequent generations. But the “religion” they have set in their sights is not Religion. For most of the world, religion includes the embedding of overt beliefs within the habits, rituals, and practice of daily life. The term points to the conduct of everyday life with a notion of purpose and assurance. It must include the background against which specific beliefs are held, and not just the overt, often counter-intuitive, statement of such beliefs. Dawkins and Dennett both have religious beliefs in this broader sense, as all humans do, but they do not appear to recognize them as such.

But we will not make progress here if we pick fights. We will instead try to chart a course towards a wider view of what science is, or might be. And at the center of this inquiry will be the issue of just what we mean when we employ those little words “I”, “we”, “you”, and “they”. Of all the elements in our vocabulary, the personal pronouns have the richest set of associations, connotations, and implications. Some possible referents of the term “I” will be of specific interest. When one of us uses the term, it might refer to the “mind”, conceived of as the activity of the brain, or it might be based instead on the bubble of present experience, where pain is felt by a first person subject, or to the corporeal body, or to the subject of an autobiography, a legal person, a social role, and more besides.

Which of these is the scientific account? Any of them? None of them? A diehard materialist will opine that the person is co-extensive with the body, and that any further at-

tributes are simply “what the body does”. Francis Crick, a brilliant biologist, went further and privileged the brain over the rest of the body when he insisted that “ ‘You’, your joys and your sorrows, your memories and your ambitions, your sense of personal identity and free will, are in fact no more than the behaviour of a vast assembly of nerve cells and their associated molecules” [1]. The identification of the person with the corporeal brain seems to persuade many people, especially when we consider the effect of removal of parts of the body. Remove my leg, and I am still me. Remove my brain, and I’m not so confident that I want to be associated with what is left.

Anyone who has had the bleak experience of seeing the person of a loved one slip away through Alzheimers, or through the struggle with “mental illness”, will surely agree that we cannot identify the person with the body alone. When we see somebody in a persistent vegetative state after brain trauma, we know that the person, in the fullest sense of the word, is not present. Perhaps this kind of awareness is part of the reason that many people seem comfortable with associating themselves with their brains, rather than with their bodies. For it is brain degeneration, brain trauma, brain pathology that seems to lead to this absence, leaving a barely animated corpse where there was once a whole person.

The philosopher’s notion of a brain in a vat makes the common commitment to the identification of the person with the brain perfectly explicit [9, 12]. In a far distant future, when we know more about such things, we can imagine a brain in a vat, with its neural connections hooked not into the world, but into a simulated world, by a clever computer system that provides just the right kinds of electrical signals at the sensory inputs, and that ‘reads’ the outputs of the system to generate future inputs. A rather horrifying scenario, made tangible by the awakening scene in the first Matrix movie.

Familiar stuff all this, no doubt. But I wish to suggest

that we do ourselves a disservice in allowing such accounts of the person to hide under the shield of science, with all the respectability that they thereby acquire. There is no scientific story of the person, and the identification of the person with the brain does not correspond to any fact whatsoever. Worse, by allowing such stories to masquerade as scientific stories, we shut ourselves off from very many other accounts of the person that could usefully help us out with our lives, with our dealings with each other, and our dealings with the fragile biosphere that supports all human activity. Science is harmed in many ways: we seek answers in brains for questions that can never be answered by neuroscience. We ignore perfectly real processes and structures that lie between us, rather than within individuals.

We will dance around some alternatives to a purely materialist view in this small work. In so doing, there will be no attempt to prove the materialist view ‘wrong’. The purpose of the exercise is rather to seek to become aware of the necessary limitations of such a view, and how we might hold complementary views that can serve us too. The materialist view will not suffice, and nor will any alternative. Pluralism is required for beings with many purposes, as we undoubtedly are.

The starting point was the domain of experience, as pointed to by anamorphic art. It is to experience that we will return, again and again. It is experience that has been shamefully neglected by the scientific endeavor to date. Experience, in the sense of that which is here, now, for a subject.

We can point to experience, through devices such as anamorphic art. Can we not discuss it directly? The philosopher Edmund Husserl believed passionately that we could. He strove to develop techniques for the direct interrogation of experience, and the means to report on it in a way that others could then compare to their own “here and now”. He tried to insist that one could bracket experience, hold it like precious spring

water in the hands, and examine it. I don't know to what extent he was aware of it, but Husserl was doing what thousands of others, in diverse traditions, have done before him, and since. And the answer is clear. No. You can not read out experience in a fashion that allows it to be shared or transmitted. And there is a very good reason why. We will build a story to explain this throughout what follows. It must be done slowly. Husserl, in his vain attempt, was fully subscribed to the notion that there is, for one person, a domain of private experience, a bubble of being. Appealing to the obvious first-person character of experience, it is absolutely clear that there must be some such. But when the claim that such a domain must exist is trotted out into the sunshine, it becomes just a little suspect. For if we could identify, delineate, and describe such a domain, would we not have opened the door to solipsism? Solipsism is tantamount to the belief that an individual has no direct knowledge or experience of anything but his or her own mind, and that therefore there is a logical possibility that one's own bubble of being is the only one that exists. Solipsism is, understandably, a hugely unpopular metaphysical position. To be fair, solipsism is usually treated as an extreme position that serves to bring home the limits to which experience is or can be shared. But if we admit of a private bubble of experience, available to one person and not to any others, it remains as a logical possibility, and can not be simply brushed aside on the grounds of implausibility.

Perhaps when I say "experience", you, the reader, think "mind". The M-word. I am not a fan of this word. I do not find it useful when I attempt to think clearly. Perhaps its every day uses serve to justify its existence, but I sometimes wish it would just go away. Much of what I will attempt in this book is to suggest a way out of the mind-trap. This is the belief that we are trapped within our own individual mental realms. We seem to be solipsistic by nature, believing ourselves

to be individual subjects who reign over this private singular mental domain within which feelings, thoughts, perceptions, ideas, imaginings, and the like arise and are witnessed and experienced by one person alone.

Minds and the mental are conventionally distinguished from matter and the physical. There is no shortage of literature on this distinction, which infuriates and delights philosophers in equal measure. Most of the effort is spent in one direction, trying to shoehorn the difficult and intangible components of Mind into a physical world view, presumed to be somehow less problematic. Material objects, such as tables, chairs and teapots, seem to admit of simpler treatment than ideas, dreams, feelings and convictions. I don't wish to make pronouncements, where vacuous pronouncements are ten-a-penny. But I suspect that any notion of Mind makes its complement, the physical, just as problematic, and the whole set of problems that arise then can never find resolution because, in setting up this dualism, reason has broken.

Why is the physical considered so much less problematic than the mental? We can readily discern two rather different senses of the term physical, each of which seems to be grounded in the need for certainty. The first appears simple and obvious. It is the belief that the physically real is tangible and indubitable. You thump the bar, or, like Doctor Johnson, kick the stone to verify its materiality. Doubting Thomas prodding a curious finger into Jesus's wounds was looking for a similar watermark of the real. If it feels real, then it is real. We might label this "bar-room physicality", to remind us that it is rooted in the familiar world of human experience.

A peculiar variant of this is a naive physicalism that believes that the chair is constituted of atoms, and heck, if atoms ain't real, what is? Naive physicalism tries to leverage both bar-room physicalism, with its reassuring commitment to common sense and tangibility (the indubitable reality of the chair

before our very eyes!), with a faith-based reliance on some version of physical science (atoms). We might call this usually unexamined faith “Simplistic Materialism”.

Which leads to the second sense of physical, which also comes in two varieties. Together they express the belief that the physical is whatever is revealed to us through the science of physics. The two variants we need to consider are those of Newtonian physics, dominated by the seductive ideal of the mechanism whose operation is governed by deterministic laws, and modern theoretical physics, whose models are entirely removed from any relation to the familiar world of tables and chairs.

Lets examine these, one by one. Start with Doubting Thomas or the man thumping the bar. Both of these examples are entirely experiential in nature. Sensorimotor engagement with the environment is required in order to thump the bar or feel the wound. Sensorimotor engagement means movement of the body, in conjunction with the world, generating thereby sensory information that reveals the world. We squeeze a sponge and declare it to be soft, a fact we can find out only by squeezing the object in question. I will tell a story in the next chapter that suggests how the causal basis for immediate experience and the origin of the phenomenal can be found in the sensorimotor coupling between a biological organism and its environment. In attempting to assume bar-room physics as a basis for reality, there is no meaningful distinction that can be made between the physical and the mental. Physicality is demonstrated through lived experience.

Simplistic materialism tries to have its phenomenal bar-room cake and eat it theoretically too. It assumes that matter, the stuff of atoms and molecules, provides a solid foundation for ontology. It tries to reassure its adherents that the solid phenomenal reality of tables and chairs is grounded in a science that reveals the world as it truly is. The simplistic materialist

has probably never felt the need to hang her hat on any specific physical theory, or if she did, heck, it could always relocate to another more likely theory if circumstances warranted, because that's what science is all about, right? This amounts to an abdication of responsibility, because the nature of the physical theory that is implicitly endorsed here really matters. Newtonian physics describes a mid-sized world, that chugs along at a constant rate of time, within a three dimensional space-like volume, and events are indexed by the action of massive bodies moved by energetic forces. It is the best physics in the world to understand the properties of massive bodies observed in a world that shares the spatiotemporal characteristics, including scale, of our familiar world of everyday experience.

There are good reasons that the Newtonian vision is so effective in accounting for the familiar world that bounces, crumbles, grows, and moves before our very eyes. I will have more to say on how this observation relates to the role of the senses, considered as perception–action interfaces, in a little while. This is, for many people, the real world, the world out there.

The relation of modern physics to the familiar world of experience could not be more different. First, data that inform the abstract mathematical models all comes from measurements that are made at completely different spatial and time scales to those of our human world. We peer back in space-time to a singularity at the cosmological level, billions of years ago. Or we go small, and uncover a weird, non-deterministic world of fluctuations at the sub-atomic level. When we stop to consider what kind of world is revealed by this, most recent, kind of physics, we find that any attempt to pin down a physical fact brings the observer back into the picture. There is deep mystery here, but it is not straightforward to make a direct link from that picture to the familiar, also-deserving-of-being-called-physical, world of cricket bats and song birds.

As this galloping summary illustrates, any version of physical science that may be called to court will bring with it elements of both the subjective and the objective. The subjective appears in the experiential basis of bar-room physicality, and in the role of the observer in determining the state of any physical system. One can not, therefore, appeal to physics to ground reality, and then seek the origin of the mental within such a worldview. The problem is not shoehorning the mental into the physical. Dualism, or the fracturing of the cosmos into appears once one admits of a physical domain, independent of experience.

Physicalism, materialism, realism, these are all terms used to describe ways of understanding that take the existence of a mindless, observer-independent, world as a simple fact, backgrounding all subsequent explanations and science one might care to build. And all of these fail to account for experience. The conventional move at this point in the discussion is to consider an idealist alternative that starts with the indubitable reality of lived experience, and therewith deny the apparent reality of tables and teapots. Let's not do that either. For rather obvious reasons, nobody is really happy with granting such primacy to Mind over Matter. We will here outline a middle way, that allows both materialism and idealism, and many other views to co-exist, without granting primacy of one over the others without bound. We are not the first. The closest relations to the story to be told here that I am aware of arise in the way that mind and matter are considered within the Middle Way school of Buddhism, and within the related view espoused by the late neuroscientist and biologist, Francisco Varela [13]. There are many other possible points of reference, but it is probably more efficacious to simply say what is to be said without dense referencing. Do we need to attach an -ism label to such an endeavor? Possibly. But any such label serves to make enemies and alienate, to stifle discussion and

to entrench old positions, so I would rather not. One might call it “bringing forth-ism”, as suggested by Proulx [8], or one might note that it lies close to an emerging philosophical stance known as enactivism [10]. But let us remember that any account developed here does not trump all others. It seeks rather to expand our rational discourse so that it can accommodate mutually contradictory positions, without thereby succumbing to the dictates of blind reason that insists that one account alone must prevail, before which all others must yield.

And so the question to be raised, and not answered, by this first chapter is this: with what do we identify? When we use words such as “I” and “we”, what assumptions are thereby made? Do we wish to identify our persons with our brains? If so, do we wish to do so always, or for some purposes only? When do we wish rather to be identified with our bodies? How should we develop an account of our own being? If we identify with our brains for some purposes, our bodies for others, why stop there? We are also family members, lovers, citizens, professionals, team members, and, of course, humans. Each of these might demand somewhat different consensus based accounts of our selves, none complete, but none negligible either. As humans, we also stand in some relation to other forms of life, and ultimately to the entire biosphere. Is there a case to be made that that, too, is *us*?

As we understand the world we live in better, we must necessarily also come to know and understand ourselves in new, richer ways. The journey of discovery is always about ourselves as much as it is about what we observe. I will suggest herein that we will ultimately come to understand ourselves, not by looking inwards, but by negotiation and consensus based on observation. As we see more and more, we shall ask not only “what do we *see*”, but “what are *we* that *we* should see such”?

Chapter 2

Creation

We are concerned in this little work with dancing around some questions that arise if we take seriously our everyday understanding of ourselves, our lives, our experiences, and our behavior. We saw before that a big issue we will be concerned with is the idea (belief? presumption? tacit acceptance?) that there exists for each of us a domain of private subjective experience, a phenomenal bubble of here and now, unique and isolated. As we will keep returning to this background assumption of our everyday language, we may as well give the notion a snappy name. I will call it the P-world – the personal, private, phenomenal bubble of experience centered in or about your head, that is proper to you. Naming such a notion is not the same thing as claiming it exists. If such a thing exists, we have all sorts of mind-body problems that no amount of fancy physics or philosophy can get around. If such a thing does not exist, we have all sorts of explanatory failure as we talk in our everyday language of pain, perception, feeling, value, meaning, memory, attention, and the whole psychological grab-bag of concepts that all depend upon the existence of some such entity.

The P-world I have just sketched is not co-extensive with

all senses of the word “mind”; nor is it quite the same as many senses of the word “consciousness”. Both of these words are mischief makers, each having so many senses that discussion couched in terms of either seems inevitably to lead back to the same old unresolved metaphysical and epistemological tangles. We met the word “mind” in the previous chapter, and we left it behind for now, as there we found no coherent way to distinguish between the mental and the physical. “Consciousness” is equally problematic. Reading much of the contemporary literature on the relation between brains and nervous system activity on the one hand and experience and behavior on the other, one could be forgiven for deducing that brains are distinguished from all other physical objects by virtue of exuding a special kind of substance called “consciousness”. Put like that, the notion is clearly risible. Yet something like this exists as a second embarrassing background assumption in much of our talk of our selves. A less ridiculous version notes that without brains, there does not seem to be any kind of consciousness, while those of us with human brains certainly have something we might call consciousness (where the word now means something like the P-world I have sketched above). And so the causal origin of experience is deduced to lie in the brain. This way of laying out the relation between brains and experience brings into focus a research agenda that inquires into the degree to which animals may be said to be conscious. It leads to futile attempts to identify those elements of human experience that are constitutive of consciousness, such as the ability to think about oneself, or the capacity for reason. Such attempts to justify the singular nature of human experience comes perilously close to theological argument for the existence of souls, which is not compatible with the naturalistic enterprise.

We proceed by querying, as ever, immediate experience. In experience we encounter a phenomenal world that is available

to a subject. The anamorph with which Chapter 1 opened illustrates the intertwining of the work of art and the subject. Something similar may help us in discussing the relation between subject and P-world. This relation can be better understood if we develop an alternative view to the rather messy business of brains secreting consciousness. We begin, not with nervous systems, but with life [11].

Life, life,life,life,life,life,life.

If we look back in time, considering the universe as it is conventionally understood, we can peer right back towards, but not as far as, the singularity at the start, often mistakenly called the big bang. There was no bang, of course, and the point just before the rapid expansion of matter that led to the formation of all that we see is better thought of as a blind spot. There, our theories are silent. There, we have nothing to say. Give us that, and we can follow the rest. We can track the expansion, the clumping of matter into stars, planets, galaxies, etc. But the singularity remains stubbornly outside the scope of our explanation. Our understanding of cosmology does not deliver us a picture of a universe with a starting point, but a universe that, once in motion, unfolds lawfully. The singularity is a blind spot in principle, not a detail we can fill in later¹.

We have another blind spot at the origin of life. Perhaps this is a fundamental blind spot—a necessary limit to the effectiveness of our theorizing, or perhaps it is a contingent ignorance and we might some day have an account of the evolution of the universe that includes the transition from the inanimate to life itself. Perhaps. For now, it is in many respects like the singularity at the origin of time. We know life came about, but there seems to be nothing prior to life that can predict, or

¹The presence of other singularities at the centers of black holes ought to relativize our understanding of the arrow of time. This will be expanded upon as this work grows.

motivate, its appearance.

But with the appearance of life on the stage, we have what appears to be the necessary preconditions for a concerned perspective and for significance. These two elements—a concerned perspective and significance—arrive together on the cosmic stage with the appearance of the first cell, and once they are there, we can begin to uncover a history of biology, with evolutionary change, multicellularity, the appearance of nervous systems, all the way up to humans. With humans, we also get language, and that will be essential in what is told here, but we can get there later. For now, it is the first cell that we wish to look at.

What arrives with the first cell that was not there before? The identity of a cell does not depend upon its material constituents, in that one can replace every atom and still have the same cell. Its identity is a dynamic identity that lies in the self-contained network of processes that persists from the moment of its formation until its extinction. But the weather also represents a bounded set of processes, and the Great Red Spot on Jupiter can properly claim to have a dynamic identity that lies in the relation among constituents, and not in the material identity of the constituents. What is it about the set of processes in a cell that makes them so very different?

The answer has impressed and confounded philosophers since at least Aristotle. The cell acts in its own interest. It exhibits what Kant called “natural purposes” [16]. The biologists Humberto Maturana and Francisco Varela have characterized the situation well in describing the cell as a cybernetic machine that acts to preserve its own continued existence [6]. If that sounds a little abstract, consider the humble thermostat. It is a cybernetic machine that acts to regulate a specific quantity, temperature, so that it does not fall outside a specific set range. The difference between the thermostat and the cell is that the regulation of the cell serves the cell’s own purposes,

or continued existence. It needs no externally imposed goal.

This seems to be the defining property of life—acting in its own interests. I say “seems to be”, because it is hard, if not impossible, to provide a definition of what it is to act in one’s own interest that does not incorporate the prejudices, opinions and biases of the observer in that statement. We can note this, for now, and do little more than ask again, not “what is this odd thing?”, but “what are we that we should see such a thing?”.

Machines express purposes, and in the normal run of things, the purposes they express are those of the designer: a vacuum cleaner is suited to vacuum cleaning because somebody made it that way with intent. Can we imagine a machine that is entirely machine-like, and yet that serves its own purposes, and not those of an external intentional agent? Indeed we can. We can concoct a picture of such a machine, and what we get is a simplified description of a single cell. The story is re-told again and again within the recent literature that connects the twin concepts of Mind and Life [11, 10]. It is the story of the humble single bacterium, swimming up a glucose gradient by a process known as chemotaxis, and we shall re-tell a brief version of it in a moment. But the story to be told is not really a biochemical, or biomechanical account of bacterial movement, even though the central figure is, nominally, a single *E. coli* cell. Some authors, in the re-telling, include more biology than is strictly necessary for the story. This story of the cell is a myth; it is not a scientific account, and we shall treat it here as the myth that it is.

Myths are muscular stories. They provide footholds for thought. With myths, we can frame new material, and make sense of it. Myths provide the thinking moves that make up the dance of comprehension. This myth is about no real cell, but its ingredients can help to frame many issues we must now face. Above all, the myth may help us to see what it is that we

do, when we insist on, or deny, the existence of the P-world.

The Myth of the Cell

The cell of the myth is a machine that serves its own purposes. It has a membrane, containing all the physiological soup necessary for metabolism. We simplify, as we must, and say that the cell's metabolic dependence on its environment lies in the availability of common sugar, or glucose. We imagine an aqueous environment that contains a single, distant, source of glucose, from which a glucose gradient emanates. It would be in the cell's interest to move upstream along this gradient. We give the cell a single simple sensor on the membrane that can detect the ambient sugar concentration. This is its entire sensory apparatus. It has a means of movement too, analogous to a real *E. coli*, comprising a set of whip-like flagellae that can move in two distinct modes. In one, the tails whip around without any coordination, generating random tumbling movement. In the other mode, they work together as a single rotor, propelling the cell along a straight line. The cell switches between the two modes probabilistically.

The last little piece of machinery we need here links the sensor and the flagellae. When two subsequent readings of the sensor reveal a decrease in sugar concentration, the probability of switching from the directed mode of locomotion to the tumbling mode is increased. In this way, whenever the cell is moving in the wrong direction, it will tend to change direction by tumbling a bit. Conversely, if it is going towards the source, the likelihood of transition from directed to tumbling mode is reduced. In this way, the long-term behavior of the cell is to approach the source of the sugar. It swims up the gradient in this manner, known as chemotaxis.

That's about all there is in the myth, if we strip it down.

A relatively simple piece of machinery that results in the continued existence of the cell, as it acts in “its own interests”. Now let’s use the myth to do some thinking.

Let us first contrast two perspectives. The first perspective is that of the detached scientific observer, who sees the cell in a petri dish. We can think of this as representing something like an objective, third-party account of the activity. This view can lean upon such notions as glucose, petri-dishes, cells, etc. From this vantage point, we look at the cell as outsiders, and we infer the purposes that must be invoked to make sense of the observed activity. Although this is hardly a God’s eye view, for our purposes here, it might as well be.

The second perspective is that of the cell itself. We do not have to insist that the cell manufacture sentience, or be conscious, in order to realize that the interaction with the glucose is of significance for the continued well-being of the cell itself. The interaction is, in this strict and limited sense, meaningful. Glucose, from the observer’s vantage point, is just another molecule, neither good nor bad. But with respect to the organizational unity that is the cell, it is valenced. It is a good thing.

The cell is, by hypothesis, ignorant of and insensitive to many aspects of its surroundings. If there are chemicals that do not affect its metabolism or its locomotion, they might as well not be there, just as we are insensitive to all kinds of radiation and energy that we are immersed in. The perspective of the cell thus gives us a way of partitioning the totality of the cosmos into those limited set of influences that can interact with the cell, and everything else.

Now we can go an imaginative step further. We can describe some aspects of the world that arises for the cell in its activity. Let’s be quite clear about what we are doing here though. We are not imagining that we are the cell. Thomas Nagle, the philosopher, famously opined that there simply was

no answer to the question “What is it like to be a bat” on the grounds that the experience of the bat arises from its sensorimotor engagement with the world that is mediated by organs with no human counterpart [7]. There is nothing in our experience that resembles the sonar sensory apparatus of a bat, and so there is no sensible way to imagine what it is “like” to be a bat. But before Nagle, the Estonian philosopher-scientist Jacob von Uexküll had asked a closely related question: could one say anything of substance about the world as it appears to a variety of creatures [15]. Von Uexküll considered, among others, the paramecium, the limpet, snail, tick, dog, fly, fish, and various classes of humans, including the scientist. He recognized that the same logic could be employed to make substantive statements about the form of the world that must arise. Thus, animals possessed of the semi-circular canals of the inner ear have a spatial understanding that critically depends upon that vestibular organ. Without it, the character of space as experienced would be very different. With it, the space as experienced is clearly organized into three dimensions.

In the spirit of Von Uexküll, then, we might say something about the world that arises for the cell of the myth, because we have taken the liberty of delimiting and defining the entire means by which the cell is coupled to its environment through sensorimotor activity. Its sensorimotor apparatus—perceptual system, if you like a bit of hyperbole—is capable of making a single distinction. It recognizes a gradient, and in so doing it makes a single distinction.

We are here sketching a structural account of the P-world that arises as a result of the sensorimotor coupling between an organism and its environment. When doing so, it is a good idea to express the situation *from the point of view of perception*, in which case one talks about how one might describe the world, and *from the point of view of action*, where one describes instead the opportunities for action. These are complementary

aspects of one and the same set of relations. Von Uexkill called these the *Merkwelt* (roughly: the world of signs and signification) and the *Wirkwelt* (roughly: the space of effective action), respectively. In the present case, the perceptual point of view serves to pick out one direction in an otherwise homogeneous field. Remarkably similar to the creation of light, or the partitioning of land and sea in the Christian creation myth, this sensorimotor apparatus partitions an otherwise entirely undifferentiated cosmos into a directed set—there is that way and its opposite.

So we have the beginnings of an account of the way in which sensorimotor coupling between an organism and its world allows us to consider the arising of something from an undifferentiated nothing. Interestingly, this is well known in many traditions. Beyond the Christian creation myth, it is found within the Buddhist tradition as “dependent arising” (Sanskrit: *pratītyasamutpāda*). One often finds the same fundamental insight arises in separate traditions, but its elaboration in story form differs, resulting in highly diverse religious narratives that nevertheless stem from the same source. For all humans at all times have had to deal with experience and the world, and the conundrums that arise. Each has used the means at their disposal to give voice to similar insights and ideas.

So much for the world that appears to arise for the cell, as if it were perceived. The flip side of that coin is the world of effective action. Within some schools of thought, this is referred to as the set of *affordances* that arise in the relation between an organism and its environment [4]. The minimal sensorimotor engagement we have provided the cell with shapes its capacity for action. Instead of being equally likely to move in all directions, it is now more likely to swim in one, preferred, direction. No matter what direction the cell is moving in, in the next instant some directions are more likely to occur than

others.

The notion of a perceived world and the notion of shaped behavior may appear unrelated, but they arise from the same simple sensorimotor coupling we provided the cell with. They are in fact two descriptions of the same thing. As with the anamorphic works of art we started out with, we can draw attention to a subjective facet or an objective facet as we wish, but these are different ways of characterizing something that arises in the relation between the cell and the world. The relation is fundamental. Its description may then be cashed out as something subjective (perceived world) or as something objective (observed behavior), depending on our purposes.

Beyond the Cell

The Myth of the Cell allows us to frame a discussion of subjective experience that circumvents the uncomfortable notion of an enclosed consciousness peering out at a pre-given world of which it, itself, is no part. It develops the hypothesis that it is life itself that provides the necessary conditions for the arising of a world of significance for an individual [11]. It thus places at life's door much of that which is traditionally ascribed to nervous systems and brains. In teasing out the myth, we begin to establish a vocabulary for addressing concerns that have hitherto sat uncomfortably within natural science. In particular, in the recognition of two perspectives—that of the observer and that of the cell—we have the beginnings of a non-mysterian account of significance or meaning from a concerned perspective. We use the terms *significance* and *meaning* here in a very restricted, and somewhat technical sense. If we, as observers, observe a system, S , we ask not only about our observations of an event E that can have an influence upon the state of S , but we ask also about their import *from the*

perspective of S. That is, we ask what is the relation between the acknowledged organizational unity of *S* and the observed perturbation.

But we do not get a free ride. This account requires an *a priori* acknowledgement of the self-serving organizational processes that constitute the phenomenon of life. There is thus an inbuilt teleological presupposition, without which the account does not get off the ground. Life is the condition within which significance arises, and without life, there would simply be no discussion to be had. The Myth of the Cell starts with a recognition of the purposes of the cell—metabolism serves to ensure the continuation of the organizational entity that is the cell itself. This starting move is non-scientific. It has more in common with myth or religion than science, yet without it, there is no science of life itself, and no science of meaning and perspective. One way to describe this pre-scientific positioning is to note that the cell is regarded as an *agent*, and the attribution of agency is one of the most fundamental ways we have of understanding causality and the world around us. Inanimate stuff happens, but animate goings-on are done by agents, with purposes and goals. The Myth of the Cell thus self-consciously injects agency into a basic account of the animate.

But it is a long way from a cartoon cell to the human condition. We need to consider how the appearance of multicellular bodies belonging to a single individual changes the account, and then we need to consider just what it is that nervous systems bring to the table, for they certainly matter, even if they do not bring secret powers of consciousness-secretion. In what follows, we elaborate on the myth of the cell, and we consider only a single relation—that of a single system, the organism, and its immediate surround. In this simple exegesis then, we may ask how the internal structure of the organism, and specifically the presence of a nervous system, affects the ability of it to respond to perturbations that affect the ongoing

maintenance of its own organization.

But in this elaboration of the myth to multi-cellular beings, to organisms with nervous systems, and ultimately to social forms of organization, we are staying deliberately within the myth. We are assuming the existence of system-intrinsic purposes. When we adopt that stance, we can treat the system as if it were distinct from its environment. We can speak of a response of the system to “external” events, and we can indulge, self-consciously, in a little psychologizing.

When we regard a multi-cellular being, whether it be a sea sponge or a human being, we see identity and meaning arising at multiple, inter-related levels. A cell in the liver is a full cell, with its own world, its own environment, and the goings on of the cell are partly to be understood with reference to its cellular identity. But the cell is enmeshed in a superordinate form of organization, the liver, and this provides a separate perspective, and can help us to understand more of the cell’s activity. The organ-level domain in turn is embedded within the entire multi-cellular entity, which constitutes a further, super-ordinate domain. In order to understand the observed activity of the cell, no single one of these layers can claim precedence. Each may be required in order to appreciate the significance of an observation, for significance is always *with respect to one or other domain*. We can zoom out all the way to the entirety of the biosphere, and we see layers upon layers of domains of relative autonomy. For each such domain, we may talk of significance *for* that autonomous organization.

On this view, there is nothing terribly special about multicellular organization, as no life form exists that is not suspended in a web of mutual relations with other life forms and inanimate processes. No cell is as independent as the cell of the myth. And no one of these layers corresponds to either the P-world, or to the person. If we identify with our bodies, we are insisting that meaning arises with respect to that do-

main. If we identify with our brains, we likewise fix the layer of reference with respect to which meaning arises. But, like the hoards of Satan in *Paradise Lost*, we are legion. As we progress in this little work, we will uncover a philosophy that denies that meaning can be indexed with respect to any single such domain.

What, then, is to be made of the arrival of nervous systems on the stage. If we refuse to grant them magical powers of consciousness-creation, do they change anything fundamentally? They do, but without magic.

We can view nervous systems as adding three innovations to the organisms that are so endowed. Each of these innovations concerns changes and refinements to the way in which the organism is embedded, through sensorimotor coupling, in its environment. That sounds terrifically technical and dull. But stop for a moment and take stock of your own subjective world at the present instant. It is defined, or brought into being, by your sensorimotor coupling to the environment. This, in a very strong and important sense, is what we mean by the term 'now'. The division of time into a past, a present, and a future, rests on the assumption that we can pick out a present. And consider a pain you felt yesterday, a pain you feel today, and a pain you may feel tomorrow. Which of those three is present in experience? It is that that arises from your sensorimotor embedding in your environment. So when we ask about the role of nervous systems, we are asking *inter alia* about how this biological innovation affects our view of the experiential world that arises for a subject. The gradient that is sensed by the bacterium dictates the entirety of the P-world, or the bubble of the "now", for the organism.

The first, and simplest, innovation is to allow the link between movement and the attendant sensory flux to be distributed in space. In consideration of this innovation, we do not have to concern ourselves with the subjective experience

of the organism, but rather with its manifest physical form. A toxin impinging at one point on the exterior of an organism can elicit a motor response at a distant location with a rapidity that is possible only to those endowed with the electrical signaling that nervous systems bring about. This innovation has enormous consequences for the opportunities available in the evolution of organisms. Bodies can now admit of a great degree of specialization. Receptors can be far removed from effectors. Concentrated groups of receptors can specialize in exploiting one or other energetic gradient. If we compare the visual system of a brittle star with that of a mammal, we see both have photo receptors, and both exploit the patterning of light and dark to structure their locomotion. But in the brittle star, the photosensitive cells are distributed over the whole body, and there is no single organ that can be called an eye. In the mammal, the cells are restricted to an area called the retina that itself is located in a biomechanical structure that can be moved specifically to facilitate the effectiveness of visual information in interacting with the world. Ears are a similarly specialized concentration of highly specialized receptors within a special-purpose device, the ear. Touch receptors, on the other hand, remain distributed throughout the body. So the de-linking of sensory and motor functions opens up vast spaces of complex body organization that are not possible without the rapid mediation of a nervous system to yoke movement and sensation. And it allows the development of distinct sensory modalities such as vision and audition.

The second innovation that nervous systems bring is that distinctions and discriminations that are made at distinct sites on the body can be combined non-linearly, thereby massively increasing the complexity of the discriminative power of an organism. Because energetic and chemical influences on the surface of the organism are translated into a common chemico-electrical vocabulary, different kinds of distinction can be made

simultaneously. Many, most even, of the distinctions we draw in the world around us are based, not on a single sensory modality, but on multi-modal distinctions that occur together. A fire, as Gibson pointed out, is not encountered as separate experiences of warmth, crackling, glowing, and a smell, rather it is a fire, and what it is to be a fire is distributed over these modalities [4]. There is a long history within psychology of treating the world as inferred constructively from distinct information sources, and much effort has been spent trying to understand how these disparate information streams are combined to generate a world, or a representation of a world. But this is nonsense. The division of complex multi-modal events into separate modalities is a *post hoc* analytic dismemberment of that which is, originally, whole.

The third novelty that nervous systems bring is the most far reaching. As well as distributing the relations between movement and sensory discrimination in space, that same set of relations becomes smeared in time. This happens because nervous systems are recurrently connected. That is, as we follow a path from the periphery towards the center, we almost immediately encounter neurons whose only functional connections are with other neurons, and these connection paths contain loops. This happens even with very simple animals who possess small ganglia, rather than brains. Thus any relation that obtains between sensory flux at one location, and muscle innervation at another, can involve an arbitrary time lag. The world that arises for an organism through its sensorimotor coupling with its environment now acquires a degree of temporal depth. The specious present within which we live, that persistent sense of the lived moment, arises precisely because the sensorimotor embedding that anchors us and creates a lived world has a great deal of temporal depth.

Figure 2.1 illustrates these fundamental changes to the sensorimotor relation that arise with the introduction of a nervous

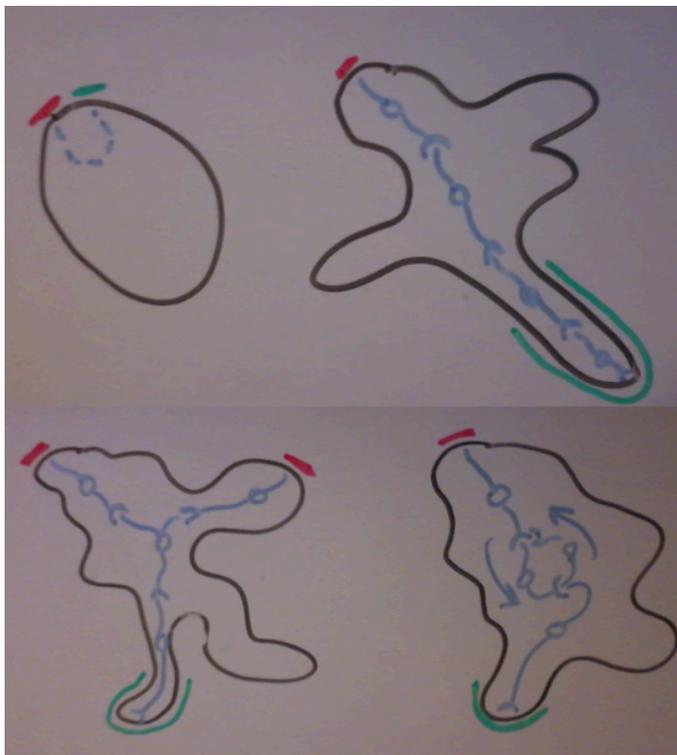


Figure 2.1: How nervous systems alter the sensorimotor relation.

system. At top left, we see a simple organism, not possessed of a nervous system. The sensorimotor relation is indexed by red (sensory discrimination) and green (muscle innervation), and the sluggishness of the processes of chemical diffusion ensures that a discrimination made at one spot has only local effects. At top right, we see the advantage of having a signaling system. A discrimination made at one spot can have a distal effect almost immediately, at least compared to the time constants of chemical diffusion. Elaborate body architectures now become possible. At bottom left, we include the possibility of non-linearly combining discriminations made at two distinct locations. The complexity of the distinctions the organism can now make is vastly increased. These two innovations between them give rise to the possibility of having distinct sensory modalities, and multimodal perception of complex events. Finally, the recurrent loop between the red and green shown at bottom right means that the temporal depth of the sensorimotor relation is now greatly extended, and the possibility of a subjective present moment in a complex world has now arisen.

Now

What have we learned from the above? Some frighteningly complex verbosity dances around a single concept that evades description. The only words I can find that point towards the reality of lived experience, without thereby bringing solipsistic prisons into being, are short words. Now. I. We. For what is the present, what is “now”, if not that which arises for a subject? What possible sense can we give this word? “Now” has no counterpart in physics. There is no physical “now”.

The reality of lived experience (to dress the pristine concept up in over-fancy ruffles), is relational. It is neither subject

nor object, but both arise thereby, as complementary poles. The myth of the cell provides us with a tangible plot line and argument structure for acknowledging this in words. But we are not the cell of the myth. The word “we” does not refer to a single entity, and nor does the word “I”. Neither is independent of the “now” which is thankfully impersonal. And so “you” are not the subject for whom a now arises, though breaking this habit of thought, and this mis-identification, is not a simple matter.

Within the Vedantic traditions of India, there is a useful thinking strategy, which has become something of a slogan. As a slogan, it is facile, but as a framework for thinking, it has its uses. The phrase “Tat Tvam Asi” has many translations, but roughly, it means, “you are that”. The import of the thought is to push from the subject, back towards the world, pointing out that the phenomenal world that arises in experience has no independent existence. To make sense, it requires that the subject give up completely, and accept that which is present in the immediacy of experience. We constantly identify with the subject, peering out at the world. But we are none such, and the world is not well characterized thus. It is in the co-arising of subject and object that we find anything that is to be found. But we will not find a self therein.

And so we have traversed the first two steps. **Identification** is the recognition that any story we tell ourselves comes laden with presumptions about what it is that we are, or that “I” am, or that “now” is. **Creation** provides us with a framework within which we can begin to talk of the nature of experience as relational, existing, as much as a relation can exist, in the co-dependent arising of subject and object. We turn now to **Complementarity**, and the task of learning how to speak of that which is relational, without unwarranted reification.

Chapter 3

Complementarity

We like certainty and truth. Things should be either so, or not so. Most cats, in most boxes, are either alive or dead, and I either exist, or I do not exist, for most purposes. It was ever so, and it will remain so. Most argumentation is settled by coming down on one side or the other, or, of course, by imposing a solution from one side on the other.

Discussion, argumentation, and laying out one's stall using language can be seen as a means of trying to establish that which is uncontroversially so, of overcoming controversy. But of course things don't always work out quite like that. We still fight wars, we perceive injustice and inequality in all societies, and we fail to agree on scores of issues of vital importance. Importantly, we also cherish some of our disagreements. Soccer would be a miserable exercise if the world simply decided that one team in particular merited support over all others.

In the heat of any argument, it is easy to forget that the conduct of that discussion is possible only within some local context, with some unstated but essential framing assumptions. Absent that framework, we are just gesticulating in sound at each other, without any hope of settling our affairs. To take a simple example, I might insist that because a par-

ticular ball went over a particular goal line in a particular direction, that a goal had been scored. The assertion would typically not need to make explicit the framing assumption that goals can only be scored within games of football. If the ball goes over the line when nobody is playing a game, no goal has transpired.

Any and every statement of fact is of a similar nature. Every assertion exists with some framing assumptions. This is simply the nature of linguistic communication; the rules of the game. We can sometimes point to the rules of the game (“something up is not down”, “three is greater than two”), but in making statements about the world, as opposed to statements about the words we use, there are necessarily framing considerations. They may be narrow, in which case they are easier to see. The framing assumption of the game of football is a narrow frame, and rather obvious. They may be broad and much harder to see. This is when things typically go wrong, and misunderstandings arise.

Which brings us to the vexatious problem of statements about “reality”. To many, “reality” seems to need no framing. It is that which *is*, in an uncontroversial sense of “is”. Under most circumstances, statements about reality are statements about an objective world, presumed to exist and have characteristics irrespective of whether anybody pays attention to it or not. Statements about reality are often co-extensive with claims about the “physical”. But we have seen problems with assuming that we can identify the “physical” in any simple sense, and we have seen that the distinction between the physical (presumed observable) and the mental (status unknown) is not something we can simply assume. The subjectivity of the observer seems to lie on one side of a divide, and the objectivity of the world on the other. This dualism between the observer and the observed, between the subject and the object, demands acknowledgment. It underlies any set of assertions

of fact about the world.

But it is not an ontological distinction. It is not the case that there are two different realms here that somehow interact. It is a distinction drawn in our description of things. We observe a person trembling, sweating, with wide eyes and restless movements. We say the person is afraid. When we observe the trembling and sweating, are we observing the fear, or does the fear lie beyond the appearance? If we take the subject–object distinction to be an ontological one, then our observations are not of the fear. We must infer the fear, which in turn must remain stubbornly unobserved. But there is another course available to us. This is to regard the physical and mental terms we bring to bear in our description as complementary, rather than distinct.

Complementarity obtains when we can describe things in one way, X , as fully as possible, and that still leaves aspects of the situation under description untouched. We can also describe it in another way, Y , and that may capture that which was absent the first time around, but it in turn must be silent about some characteristics that were patent and explicit under the first description. Physicists have become familiar with something analogous in the well known conundrum of trying to describe the position and momentum of an elementary particle. Neither position nor momentum is poorly defined. They are both perfectly fine things to ask about a particle. But if we attend to one, the other becomes unavailable. No single description of the particle is complete, and the description itself, or the stance of the observer, has become entangled with the phenomenon being described.

We previously noted that any attempt to provide a notion of the physical that was independent of the mental was doomed to failure, at least in as much as that physical account was connected to the everyday world we experience, filled with solid bodies that move and processes of growth and decay.

Physicality in any familiar sense is necessarily tangible, visible, audible, and hence intimately experiential.

When we consider the goings on of animate beings, there is no *a priori* ontological split between a mental and a physical domain. In discussing this or that phenomenon that arises through the actions of living beings, for whom the world is imbued with significance and value, it is inevitable that there will be aspects of the phenomenon that are observable, or public, and there will be aspects that must forever remain objects of conjecture. But this benign duality is not a pre-existing fissure in the nature of things. The separation of the observable from the unobservable and conjectured will be done differently as our discourse moves from this domain to that. The way that separation is done will change with our purposes, and with the context of the observation and the description.

We often talk as if there were a single world, referred to by verbs of existence, and over which our discourse ranges. Yet if we consider for a moment the multiplicity of the very real topics that we talk about, without for a moment questioning their reality, we might wonder whether such a simplistic form of objectivity can be maintained: A lovers' quarrel, rainbows, the stench of corruption, hope, the shared joy of a dance, a war, greengrocery and stamp collecting, to enumerate just a random few from the infinite set of things we occupy ourselves with and happily talk about. These are not rocks, but they are real. They are neither physical nor mental, nor both, nor neither. They are complex, and they belong to a variety of domains. There are domains nested within domains, and all are subsumed, it seems to me, within the world that arises for humans, the human world we inhabit and treat as if it were independent of our very constitution [5].

But science has no business going there, I hear someone object. This is the soft and squishy domain of the humanities, of social science, or of everyday life, but not suited to

the scalpel of analytical science. This is a grievous mistake. *Once science deals with the goings on of the animate, science must, perforce, traffic in the values of the animate.* There is no way around that. This is as true of microbiology as it is of anthropology. With life, value leaks in.

The myth of the cell provides us with one way of developing a consensus-based account of value. Within the myth, glucose is a value by virtue of its role in the self-sustaining metabolism of the cell. Another way of saying that is that the glucose *means* something to the cell. Biological systems, by virtue of their natural purposes, are enmeshed in webs of value, significance, or meaning. The invocation of something as quaint-sounding as “natural purposes” here is not an appeal to some external order of things, in which values are prescribed by a deity or fate. It is the defining property of life that it acts in its own interests.

But there is no single level of biology, or anything else, at which meaning is fixed. The cell in the liver gives rise to a domain of value by virtue of its own, cellular, organization. But it is a component in the liver too, and value arises with respect to the liver. The liver is within a body, and value arises with respect to the body. The body is a social component in ever shifting domains that arise in interaction, and disappear again. Meaning inheres in life and is not tied to any one level.

Of course we keep our stories about the goings on of ants and bacteria simple. We do not imagine them to inhabit baroque theaters of reflection, caverns of introspection, nor do we suppose that there is a trellis of beliefs underlying every act. We reserve such complications for the case of *Homo Sapiens Sapiens*, the knowing ape. Here, we feel, we must appeal to beliefs, to hidden motives, passions, desires, grudges, and character traits, from the impish to the cruel, to account for what we see the knowing ape do. Other animals are granted an intrinsic “nature”, to which we appeal when our beloved

pussycat slaughters a young bird and brings it home. Humans, by many of their own accounts, regard themselves as outside this world of nature, an escape trick surely the envy of Harry Houdini himself.

The essence of complementarity is to see that any given narrative leans on some ontological division of nature into agents and the affected world. Generally, we think of the agents as ourselves, but what that means seems to vary from one sentence to the next. We conceptualise the nature of the singular agent in different ways; sometimes it may be the body, sometimes the brain, more often we lean on an understanding of our “self” as the unseen hoard of beliefs and hang-ups that together make up our personal psyche. Other stories require recognition of group minds—stories of national identity, ethnic history, cultural affiliation—these all require different sorts of agents, and different referents of the term “we”. As we evoke one or other source of goings on, we bring into being ever-changing “minds”, with their intentions, beliefs, and goals. The singular minds, attached to your name or mine, are useful convenient fictions, no more and no less real than the seething anger of a lynch mob, the inscrutability of China, or the collective amnesia of a traumatized population.

For it is a defining characteristic of our historical and geographic location that our public discourse insists upon a single split between mind and stuff, a single locus of agency in the rational solipsist, and the strict attribution of one mind per person. We construct the fabric of our society around this bull-headed conviction. An act is done, and there must be a do-er of that act. Otherwise we have no way of according culpability, charging tax, or doffing our hat. We applaud genius and despise cruelty, assuming without question that both lie within the breast/heart/mind of a single individual. We bathe in narratives of individual doers of individual deeds. This myth is not as benign as the myth of the cell, for we rarely notice

that it is, indeed, a myth. We do not have a range of alternatives at our disposal that would allow us to countenance such stories as interesting interpretations, useful fictions.

We know that this foundational belief wears thin in places. We are aware of those borderlands where we are faced with failure of the cherished fiction of the solipsistic agent, happily doing deeds concocted within a hidden mind. We reserve the notion of insanity for some of these threadbare patches. When we choose to label someone insane, we no longer grant them membership of the group of rational agents. We no longer try to ascertain underlying reasons, and we simply deal with their actions. In doing so, we strip them of their personhood, and culpability. Not guilty, by reason of insanity, is a kind of failure, but a relief too.

Or consider the humble pedophile. A despised and hated character whose actions shame us all. But we know that these sad and depraved actions are almost always done by an individual who was, themselves, similarly abused. We simply choose to demand culpability, because we have not developed other ways of dealing with the actions of one so constituted. We demand responsibility for actions, as if we acted rationally, and in accordance with rules.

A toddler does not act rationally. So we defer until a magic threshold the bestowal of agency. Once the individual reaches an arbitrarily designated age, often 18, they go from ward of society to fully-fledged, morally culpable, agent. Overnight!

It is not difficult to acknowledge these examples in which we see the limitations of our insistence that acts are done by individual agents possessed of minds, in which intentions are formed underpinned by reasons. Most of the comforting stories we tell ourselves are not threatened by these rips in the fabric of the narrative. But we are in sore need of new stories, new ways of understanding who we are and what our relation to our world is.

A science of the inanimate does not, indeed can not, address the notion of value, significance, or function. The invocation of any of these notions is necessarily predicated upon a background of one or other system for whom events can be good or bad. A science of the inanimate would never be able to describe the “function” of the heart as being a pump for the circulation of the blood, for any such notion requires the agreed background of the body “for” whom this circulation exists. Only by framing our observations by the integrity of the body can we then distinguish between the healthy pumping of blood through an intact body and contrast that with the pathological emission of blood from arterial bleeding. Even at this everyday level, biology traffics in the values of the animate by allowing the body to support the definition of function.

But the body will not do as the domain of reference if we wish to discuss pedophilia, delusion, culpability, responsibility, or economic value. Each of these challenges requires framing of observations with some agreed background system. The adoption of one or other background is done by fixing, for the purposes of the conversation, the domain for whom value arises.

Recall that when I defined the P-world as the set of all that is resolutely first person, I insisted that the articulation of the definition was not the same thing as claiming that the P-world either exists or does not exist. The term is introduced precisely to allow us a degree of freedom in the construction of our consensus-based narratives, so that we are not shoe-horned into one fixed dualist position or another. This missing degree of freedom can be spent in attributing agency to John, Germany, or the lynch mob. It can be spent by taking the big bang as a beginning, in which case you can construct narratives that are neatly stretched out in calendar time. It can be spent by assuming that your experiences belong to you alone.

But each of these ways of constructing a narrative is finite

in its potential to underpin any discussion. Each depends upon us agreeing on the backgrounding assumptions, and hence the values that can be invoked. It is here that we seem to need help in moving beyond our current practices. Every time we fix a discussion that invokes values in one way, we close the door to other realities with equal claim to our attention. All discussions are local in the following sense: if we discuss the function of the heart, the value system we lean upon is the human body. That value system is incommensurable with discussion of the extinction of the spotted owl, the fluctuations of the dollar, or the insanity of war, but that does not make it in any sense wrong. It simply makes it local. And as our utterances range over the multiplicity of topics we encounter in a brief few minutes, so our framing assumptions careen violently from one set of assumptions to another, without our awareness. When we invoke one kind of mind, or agent, with that we conjure up one kind of world, or patient. Subject-hood and object-hood emerge as negotiable poles, and it is to the business of negotiation, or the construction of **Consensus**, that we now turn.

Chapter 4

Consensus

“Truth is the invention of a liar” (Heinz von Foerster, [14])

In the preceding chapters we have travelled a path into somewhat treacherous ground. By denying any absolute ontological split between the mental and the physical, we find ourselves challenged. The everyday language we use to describe ourselves, our values, and our allegiances will not serve uncritically. We have already acknowledged that we can not speak directly of experience, for the very notion of the private solipsistic domain, the P-world, cannot stand alone. We now turn to the diplomatic task of re-establishing some common ground, through the development of procedures and practices of measurement, that will allow us to extend the consensus-based account that we value so much from our scientific endeavours. In so doing, let us invoke another little aphorism:

Of experience, we can not speak.
Measurement will be negotiated.

Measurement is central to scientific practice. But it is easy to overlook the role of common ground among measurers that

is necessary for measurement to serve in the construction of a common narrative. Here is a little story about an experience I had that illustrates the fundamental problem with the business of measurement. One night not too long ago I was a little restless, which is unusual as I normally sleep well. I don't quite know what to do if I can't sleep, and so I listened to the quiet tick-tock of my alarm clock. And I lay there, listening to the incessant tick-tock, wondering how I would ever get used to the sound. I realized that there was both tick and tock. An alternating sequence of sounds. But I could not be sure that they were genuinely different. If I managed to divert my attention briefly, and then listened in again, were the ticks perhaps now the tocks? I know well that if one plays a sequence of identical sounds, they are not heard as such, but they start to form groups, usually in pairs, just like tick and tock, imposing structure where there is none in the origin. Was that it? Was the distinction between the subtly alternating sounds a result of my interrogation of this nuisance? I determined to look into the question at a later point, and lay there, miserable. It took a cup of cocoa and two cigarettes before weariness allowed me to relinquish my compulsive desire to make sense of the world and accept the peace that is sleep.

By chance, I had a little time after breakfast, so I took the abomination downstairs, and placed it over the tiny microphone in my computer and recorded a 30 second sequence of sounds. Here is what the sound wave looked like:

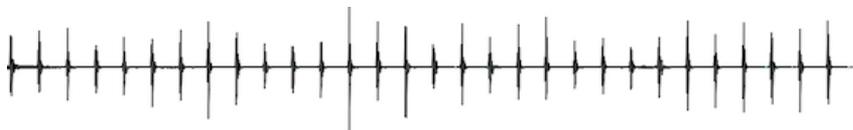


Figure 4.1: Tick and Tock, or is it just Tick?

And here, in excruciating detail, are a sample tick and tock:

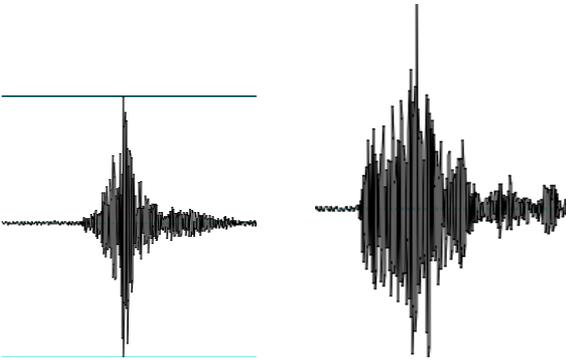


Figure 4.2: Left: Tick. Right: Tock

To be sure, there seems to be a difference between them. The one I have chosen to call 'tick' has a less abrupt onset than 'tock', and this pattern seemed to be relatively constant, with gentler-tick alternating with harsher-tock. But the microphone is, frankly, not hi-fidelity, and my recording method, which involved lying the clock down on top of my computer, is probably not what a studio engineer would do, if adding alarm clock sounds to a song. I was briefly satisfied I had not been hallucinating the alternation. But then I noticed that there was more visible in the picture of the sound than I had heard. There seemed to be a slower undulation, with ticks and tocks getting louder and softer. After a little signal processing jiggery pokery, I was able to draw it too. The pattern is shown in Figure 4.3.

How nice, I thought. I have not only discovered that tick is different from tock, but this seems to be showing me rather more about the detail of the escapement mechanism within the clock, if only I knew how to read it. I felt, briefly, like the operator of a sophisticated sonar device, using sound to explore unreachable landscapes, or mechanisms. Was I observing a slow build up of tension and release within the clock, perhaps a secondary mode of the main spring, or an oscillation of some

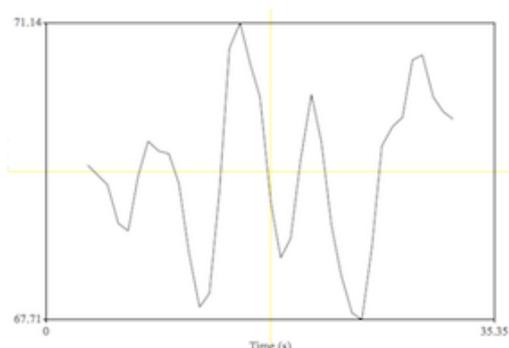


Figure 4.3: Contour of a longer sequence of ticks and rocks

part that was not perfectly fixed?

It was then I realized I could never know from my little breakfast recordings. This slow variation could just as well be a characteristic of the microphone and computer I was using to make the sound visible. I could be sure that the isolated sounds were of clock origin, because I had an independent source of information: I heard them, and I knew they would be there, one per second after second, incessantly. But I had no secondary source of information I could bring to bear upon this novel pattern.

And that, it occurred to me, is a fundamental problem of measurement; one from which there is no escape. Any single measurement has a dual nature; the thing to be measured is convolved with the characteristics of the device doing the measurement. Necessarily so. What I can observe (or, equivalently, what the act of measurement delivers) is a fusion of the two.

Take a simple measurement of length, made using a ruler. The measurement will be reliable up to a point, but if the ruler is warped, or the markings ever so slightly misaligned, the number I read off will contain characteristics of ruler-origin as well as those of the thing being subjected to measurement.

Rulers are simple measuring devices, with relatively little contamination of this sort, and, crucially, they do not change much from moment to moment, so a measurement made now will be about the same as a measurement made in an hour. The microphone on my computer is more complex. Do its characteristics change as the machine heats and cools? I presume so.

The microphone in this instance is made for specific purposes, and it is engineered to be relatively transparent. You speak into it, hit a button, and your words come back in recognizable format. If they sound a little muffled or tinny, so be it, it has nonetheless served the purpose for which it was designed. But in this case, it was being forced into service beyond its normal working range, and with the intent of extracting information it was never engineered to faithfully reproduce. And so I was left with a clock-microphone picture I could no longer call a signal, as the signal and channel had melded into a single time-varying quantity of which questions of origin were no longer appropriate.

As with my microphone, so too with the senses. We treat vision, audition and olfaction as if they delivered (to whom? where?) information about an exterior world; as if they were transparent channels through which we gain access to pre-existing properties of a world from which we are distinct. It takes a great deal of distance and disinterest to recognise that this is not at all an accurate picture. When the microphone is tasked with transmitting sounds, for example of the human voice, we quickly lose sight of the fact that the properties of the microphone have been engineered, and the notion of a voice was something we had before there were microphones. There is thus a sense in which the voice, or other sounds, that the microphone is engineered to be sensitive to are pre-existing structures. But the senses were not engineered. They arose without any preconception of the structure of the world, with-

out a direction or purpose. The processes of evolution shaped them, and the senses we have are wonderful at supporting our engagement in a world, but it is a human world, not a world from which we are distinct and removed.

To repeat the central point of this central chapter of a small work:

Of experience, we can not speak.
Measurement will be negotiated.

When we attempt to speak directly of experience, as the phenomenologists are wont to do formally, but as we all do informally, with certainty and insistence, we must be cautious. To speak of experience is to tacitly adopt the view that experience arises through the good action of the senses acting as measurement instruments that reveal properties of an external world. That is a poor picture that belies an ignorance on the part of the speaker; an ignorance of the vulnerability and groundlessness of the speaker's embedding in a world that co-arises with the self.

The senses are not engineered, and the channels of audition, vision, olfaction can not be treated as transparent channels that deliver information about a pre-existing world. The dynamic patterns they induce in a living being are convolutions of the form of the being with the form of the world, just as the oscillation I identified was a convolution of the properties of the measurement device and the sound of the alarm clock.

Recognition of this allows us to stand back from the picture we have all inherited of ghostly minds hidden inside (inside?) heads, peering out from the privileged bubble of the biosphere at an awe-inducing but dead universe. It allows us to live, where we are: Now.

But it presents us with a picture that demands negotiation.

Our condition of being can now be seen as lying at an indeterminate point along a continuum from solipsism, in which each of us encounters, in experience, a solitary discrete world available to a unique individual only, to one of complete exteriority in which we inhabit a common world (though we, as tangible forms, vanish in the process). And there is no privileged point along that continuum that corresponds to any fixed reality. Rather, each discussion we care to have is situated, knowingly or unknowingly, at a point along that continuum. At the solipsistic end, we are all utterly mad, isolated and private. At the public end of the continuum, we can speak freely, because our words manage to refer to things in a world, but we can no longer refer to ourselves.

We need to learn how to speak.

Bibliography

- [1] Francis Crick. *Astonishing Hypothesis: The Scientific Search for the Soul*. Simon and Schuster, 1995.
- [2] Richard Dawkins. *The God Delusion*. Random House, 2009.
- [3] Daniel Dennett. *Breaking the Spell: Religion as a Natural Phenomenon*. Penguin, 2006.
- [4] James J. Gibson. *The Ecological Approach to Visual Perception*. Houghton, Mifflin and Company, 1979.
- [5] Bruno Latour. *An Inquiry into Modes of Existence*. Harvard University Press, 2013.
- [6] H.R. Maturana and F.J. Varela. *Autopoiesis and cognition: The realization of the living*, volume 42. Springer, 1980.
- [7] Thomas Nagel. What is it like to be a bat? *The Philosophical Review*, pages 435–450, 1974.
- [8] J. Proulx. Some differences between maturana and varela's theory of cognition and constructivism. *Complexity: An International Journal of Complexity and Education*, 5(1), 2008.

-
- [9] Hilary Putnam. *Reason, Truth and History*. Cambridge University Press, 1981.
- [10] J. Stewart, O. Gapenne, and E.A. Di Paolo. *Enaction: Toward a New Paradigm for Cognitive Science*. MIT Press, 2011.
- [11] Evan Thompson. *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press, 2007.
- [12] Evan Thompson and Diego Cosmelli. Brain in a vat or body in a world?: Brainbound versus enactive views of experience. *Philosophical Topics*, 39(1):163–180, 2011.
- [13] Francisco J Varela, Eleanor Rosch, and Evan Thompson. *The Embodied Mind: Cognitive Science and Human Experience*. MIT press, 1992.
- [14] Heinz Von Foerster and Bernhard Pörksen. Wahrheit ist die Erfindung eines Lügners. *Gespräche für Skeptiker, Carl Auer Systeme-Verlag, Heidelberg*, 1998.
- [15] Jakob Von Uexküll. A stroll through the worlds of animals and men: A picture book of invisible worlds. *Semiotica*, 89(4):319–391, 1992.
- [16] Andreas Weber and Francisco J Varela. Life after Kant: Natural purposes and the autopoietic foundations of biological individuality. *Phenomenology and the Cognitive Sciences*, 1(2):97–125, 2002.