

Chapter 3

1544
Fodor, J.
2000

THE MANIFEST MODEL AND THE PYTHAGOREAN INTUITION

The last chapter traced the genesis of the riddle of consciousness via the history of science. The riddle arises when we attempt to construct a scientific model of consciousness that includes our experience of the sensuous qualities of things. We encounter an apparent impossibility: the ontology of science is restricted to geometry, but geometry cannot model the obviously non-geometric sensuous qualia of our everyday experience. Thus the riddle: that in the very process itself of finding a way to explain the things around us, we have made ourselves impossible to understand. Whereas our intellectual forbears, including the founders of modern science itself, were willing to accept that impossibility, we are not. They were willing to postulate a realm outside the reach of science, a supernatural realm to house those phenomena of consciousness that cannot find a place in the natural world. Again, we are not. The evidence for materialism seems very convincing. And in any case, how could reality be split into two separate realms of mind and matter? The world must be one.

But no matter how strenuously we insist that there must be only one, undivided world, we still cannot see how it might be so. Even someone as devoted to the science of the mind as the cognitive scientist, Fodor, must sometimes admit this persistent failure:

Nor do we know, even to a first glimmer, how a brain (or anything else that is physical) could manage to be a locus of conscious experience. This last is, surely, among the ultimate metaphysical mysteries; don't bet on anybody ever solving it. (1998, p. 83)

It seems the world just will not go back together again. The only substance we have any understanding of, the good old physical substance studied by the physical sciences, simply cannot produce, hence cannot explain, the sensuous aroma, taste, and feel of that first sip of morning coffee. *Ex geometria solum geometria* — but never qualia. The dualism of the founders of science is not easily shaken off by its modern practitioners.

The business of this chapter is threefold. First, the *Complementarity Hypothesis* will be introduced and defended. Complementarity was briefly introduced in Chapter 1, and the Complementarity Hypothesis develops it further. This prepares the ground for the second task of this chapter, characterization of the *manifest model*, which is the complement of the *scientific model* described in the last chapter. The manifest model is the primordial view of the universe that predates science both for our species as a

whole and for each individual who comes to entertain the scientific model. This leads to a reconsideration of the origins of science, setting the stage for the third endeavor, beginning the job of deconstructing the riddle of consciousness. Defusing the apparent impossibility of including the sensuous qualities in the geometrical ontology of science will require distinguishing between the *method* of scientific modeling, and the *terrains* modeled. The crucial insight is that even though science must model things geometrically, this does not entail that it can only model geometric things, geometric properties, or geometric aspects of reality. As we saw in the last chapter, there are two intertwined ideas that go into the intuition that consciousness is anomalous from a scientific point of view: first, the *Pythagorean Intuition* that the geometric models of science penetrate to the essential, intrinsic properties of physical phenomena; and second, the *Galilean Intuition* that geometric models exclude the essential, intrinsic properties of sensuous phenomena. These two notions are the warp and the woof of the riddle of consciousness. Remove either one, and it unravels. In this chapter I will argue against the first, the Pythagorean Intuition. In the next chapter, I will argue against the second, the Galilean Intuition. Given that it is sufficient to remove either the warp or the woof alone, why will I be attempting to remove both? Better safe than sorry. More importantly, neither is true, and understanding why helps us see how the science of consciousness should go.

3.1 THE COMPLEMENTARITY HYPOTHESIS

Complementarity has three main, mutually supporting elements. First, there is the thesis that our current ways of thinking evidence the sometimes troubled, often puzzling, relationship between the scientific and the manifest models. This relationship becomes especially problematic when it comes to the project of understanding consciousness scientifically, that is, when it comes to the inclusion of the manifest model within the scope of the scientific model. The description of this relationship began with the characterization of the scientific model in the last chapter, and will continue with the introduction of the manifest model in this one. The second element is the *Complementarity Hypothesis*, to which we will now turn. In subsequent chapters we will consider the third element, the methodological aspects of complementarity.

To a first approximation, the Complementarity Hypothesis (Comp-Hy) is very much like the old fashioned idea that there is a correspondence between the states of our consciousness and the states of our brain. Not in the quaint and antique sense in which there was supposed to be a correspondence between mental processes and material processes, where

“mental” and “material” are understood as specifying ontological types, even types of substances. Comp-Hy is nothing like the occasionalism of Malebranche nor the parallelism of Leibniz. The correspondence of Comp-Hy focuses on method rather than metaphysics, on empirical observation rather than theory. It is a claim about how the *empirical* evidence must look. How should it look? This has already been defined for an ideal case by the three principles from Chapter 1, the Identity Principle (same brain, same consciousness), Difference Principle (same brain change, same conscious effect), and Relevant Identity Principle (certain brain properties will be more relevant than others in the application of the previous two principles). These claims also have empirical content. It is quite conceivable — however unlikely we judge it to be — that people in identical brain states will report quite different states of consciousness. So, Comp-Hy is not meant as a stipulation, or definition, or principle (though it implies those just stated), but as a hypothesis, indeed an *empirical* hypothesis, about the way the world will look according to scientific investigation. In fact, it is already partially confirmed by the evidence to date, which shows a correlation between brain states and states of consciousness. There is no doubt, for instance, that the brain of someone looking at something is active in a characteristic way, as is the brain of someone sleeping or dreaming.

But it is only about the evidence, not about its theoretical accommodation scientifically or metaphysically. It is meant to be a blandly ecumenical claim that philosophers from identity theorists to mysticians will grant, and which any scientist who thinks consciousness is a natural phenomenon must believe. Comp-Hy, in other words, is meant to spell out the empirical content of materialism, or physicalism, or whatever we call the view that says consciousness is not a supernatural phenomenon, that it does not require dualism for its explanation, that it somehow (however mysteriously) depends on the brain. The relationship between observable brain states and reportable states of consciousness cannot turn out just any old way if materialism is right (Foss 1987). Materialism would be a nonstarter if the brain had turned out to be nothing but a lump of bone marrow, for instance. The empirical evidence that the brain has the right sort of structure, and connections with the body, to be a candidate for the locus of mind and consciousness, is crucial to materialism.

For example, if you hear a sound, then there is a specific neural process going on in your head without which you would not have heard that sound. By *specific* neural process, I mean to rule out specious correspondences such as that between your hearing of the sound and such irrelevant neural processes as neural metabolism in general or neural spiking in general. The idea is that there is a set of neural processes essential to hearing — in fact, we know in a rough sort of way what they are, or at least

where they are: the ear, auditory nerve, auditory cortex, and so on. We identify which of these are essential to hearing by empirically testing for the proper correspondence. For instance, you are sitting in the crowd at a track and field competition, chatting with friends. You hear the starter's pistol and so turn to watch the race. In order for that specific sound to have been heard, a specific neural process had to occur. Conversely, given the neural process, the sound had to have been heard. This process of hearing and turning requires a set of neural processes in your body. Among them we discover those of your cochlea in which neural fibers are caused to alter their firing rate due to mechanical vibrations. What if these cochlear neurons had not fired? Is their activity the thing that corresponds with the sound in the sense just given? Would you still have heard the sound without them? That depends on what else changes given the absence of the cochlear neural activity. Of course, the entire set of neural events in the auditory nerve and cortex *caused* by the cochlear events would not have occurred in the absence of their usual cause, and so without the cochlear event you would not have heard the starter's pistol. But would you have heard the sound in the absence of the cochlear event if other things had remained the same, in particular if the auditory cortex had nevertheless been activated in just the same way? The evidence to date indicates that the answer is: yes. Direct stimulation of the auditory cortex is known to result in the hearing of sounds even in the absence of cochlear activity. So the evidence suggests that had the neural activity of your auditory cortex remained the same, you would still have heard the sound. Furthermore, nothing would be heard without it. So it is the activity of the auditory cortex that corresponds with the sound you hear, not the cochlear activity.

Obviously the correspondence mooted above is an empirical question, one which is independent of the various theoretical or metaphysical accounts of consciousness within the materialist fold. Comp-Hy is, for instance, neutral between reductive and non-reductive materialism. The former claims that a given *sort* of sound is nothing other than (is numerically identical with) a given *sort* of neural event. So reductive materialists expect that qualitatively identical sounds will always correspond to qualitatively identical activity in your auditory cortex. If we could somehow make the specific type of neural activity caused by the starter's pistol occur over and over again, you would repeatedly hear the same type of sound, that of a starter's pistol going off. Non-reductive materialists (functionalists, anomalous monists, etc.) do not expect that a given type of consciousness will generally be identical to a given type of neural event. So they allow for the evidence to turn out differently: perhaps different neural processes will give rise to the same type of sound, and perhaps different sounds will be heard each time the same sort of neural process occurs. Reductive and

non-reductive materialism have different empirical content. Comp-Hy does not pretend to adjudicate between these two positions, nor, indeed, between any other forms of materialism. Instead it leaves these determinations to adjudication by the empirical evidence — supposing, of course, that there always is an empirical difference between the various forms of materialism on offer. Even once we have identified the process in your auditory cortex (or wherever) that occurs if and only if you hear the sound of the starter's pistol, we may still not know whether to say the sound simply *is* the neural process, or is an emergent property of the process which nevertheless cannot be reduced to it, or supervenes on the process in some sense, or is merely the appearance of neural activity given our current, relatively benighted, folk-psychological ways of conceiving of them. Since Comp-Hy is silent about empirical questions that the evidence is so far inadequate to settle, it is completely mute about those questions that are independent of the evidence altogether.

You will also say that you turned to watch the race *because* you heard the sound. The scientific model will show that your turning to look at the race was *caused* by the effects of the starter's pistol on your auditory cortex. Again a correspondence. Interactive dualism implies a quite different pattern of evidence. According to the interactionist, the brain acts merely as a metaphysical two-way communication link between the body and your theater of consciousness. From there, your consciousness directs the activities of your body as it sees fit. Some conscious mental activities, therefore, occur entirely within the mental realm, and do not correspond to anything in the brain at all. In particular, your conscious volition to turn and see the race would not correspond to any process in your brain. Looking at the brain we should see a causal gap between the auditory neural processes caused by the firing of the starter's pistol and the motor neural processes that cause your body to turn towards the race. According to dualism, the processes of your auditory cortex alone are not sufficient to make you turn to see the race — assuming that turning to see races is not a reflex act, but a voluntary action.

It is interesting to note that given the present incompleteness of the physiological evidence, the dualist hypothesis has not been decisively ruled out. We cannot at this point trace the causal processes of the brain with sufficient accuracy to rule out any effect on the brain from outside, nor, therefore, to show that no such effect was necessary for your action to have occurred. Comp-Hy has untested, hence unproven, empirical content, which, as Hume showed us, is the norm for empirical claims in general. Comp-Hy requires that the processes of consciousness do not go unanswered by the processes of the nervous system. It is not mere faith, but faith backed by the evidence to date concerning sensory processes and sensory experience. So it is that Comp-Hy must be reckoned materialistic, if only in its modest way:

not as a doctrine about the metaphysical relationship between consciousness and brain, but rather as a claim about what scientific investigation will discover. Comp-Hy helps define the sense in which the materialism of 4M is metaphysically modest and methodological.

Comp-Hy observes that scientific models of the nervous system involved in its daily business of seeing, hearing, smelling, and the rest map onto our visual, auditory, olfactory, and other experiences — hence consciousness — in a straightforward way. It also predicts that further investigation will only add detail to this correspondence. From the scientific point of view, the sensory systems are in the business of providing information for the organism. The activity of your auditory cortex informed you, however opaquely, of the firing of the starter's pistol, and so you turned to watch the race. You will report that you heard a sound in such a case, and if our metaphysical scruples do not get in the way, we will agree that you did hear a sound. Thus the processes of consciousness and the processes of the brain form complementary sequences, complementary in the sense that each can be inferred from the other.

This licenses us to speak in a *metaphysically neutral* way about the whole process from either point of view. Because there is a correspondence between the moon and the tides, the sailor can look at the full moon and “see” the tide is high (should the shore be too far away to see the height of the water against the shore), or (should the clouds hide the moon) look at the high tide and “see” that the full moon is on the rise. The sailor can do this, moreover, without any understanding *whatever* about the causal connection between the moon and the tide (the tides could cause the moon to rise, as far as that goes), or indeed whether they are causally connected at all (they could merely be parallel sequences mandated by God or cosmic coincidence). Likewise, we can in this metaphysically neutral way “see” you hearing something by looking at your auditory cortex. In the ideal limit, should we ever completely model the brain, we would be able to tell all about your conscious experience on the basis of the model alone. We could “see” the contents of your consciousness, see what was showing in the theater of your consciousness. As we saw in Chapter 1, we can “see” whether the colors showing in your theater are the same as those showing in someone else's. The significance of this metaphysically bland observation is not generally appreciated, precisely because it is so bland — *metaphysically*. But it is crucial from a scientific point of view, precisely because science is essentially a practical affair, one that turns on method rather than metaphysics.

And if we modestly consent to look again at our metaphysical riddles in the resulting scientific illumination, we may discover their mystery somewhat diminished. Indeed, as we observed in Chapter 1, we can not only “see,” but even see (without the qualifying “scare quotes”) what you see. This

assumes, of course, that we are not blind, that we are gifted with the neural systems required for vision. Of course, there is a limit to the range of experience any of us can have, a limit that corresponds to the limits of our nervous system. But this is the *only* limit in principle to our seeing for ourselves what it is like inside someone else's Cartesian theater — whether or not we have an explanation of how the brilliant show of consciousness arises from the dull gray matter of the brain.

3.2 THE MANIFEST MODEL

By “manifest model” I mean our ordinary perceptual consciousness of the world around us via the senses. By calling our perceptual consciousness of the world the manifest model, I mean to imply that perception models the world in the same sense as scientific models do. Scientific models are structures that we create, and that carry information about the world by virtue of the specific details of that structure itself. Perception creates structures for us, neural structures that likewise carry information about the world by virtue of their structural details. The dot on the map indicating the city of Winnipeg carries information about the location of that city by virtue of its specific position on the map. Move the dot one way or the other, and the map will give different information, indeed misinformation, about the location of the city. Likewise, a specific structure of neural activation will occur in our visual cortex when we see something, or in our auditory cortex when we hear something, or whatever. Change that structure, and the information it contains will also be changed. The structure contains information *in virtue of* its specific structural details, and so constitutes a model. Of course, it is essential that this information be in a form that can be used by us, whether the model is an artifact of the scientific process as it has evolved culturally, or the natural product of the nervous system as it has evolved biologically. The footprint of the burglar in the soft mud outside the window through which he broke into the house carries information about his brand of footwear, his foot size, his weight, and so on. It carries this information, moreover, by virtue of its structural details. However, it is not a model in the relevant sense. It was not made in order to carry information. The detective's rubber cast of the footprint, by contrast, is a model, an artifact made to carry information. At a minimum, this requires that the information be extractable by its user. The information maps contain can be extracted when needed. They would not be maps unless this were so. Likewise, the information contained by the nervous system must be, and is, in a form suitable for extraction and use.

Observation of actual instances shows that perception requires that the organism be physically affected by the thing perceived in such a way that it is changed internally in a manner that meets the conditions specified above. When a mosquito perceives its prey, the nervous system of the mosquito is affected by its prey (via chemical and other intermediaries acting on the sensory mechanisms of the mosquito) in such a way that it goes into a different state, one with a particular salience for the mosquito in that it provides the mosquito with the information necessary (though perhaps not sufficient) for it to land on its prey, rather than somewhere else. In the case of more advanced creatures, processes in the things perceived cause changes in the nervous system that contain extractable information about those processes. Thus they model them. In the case of human beings, the state of the visual cortex, for instance, continuously changes as an effect of changes in the things seen, thereby allowing processes of change themselves to be perceived. Thus these changes in the visual cortex constitute the visual mode of the manifest model.

Scientific models often aim at prediction, and are of special interest insofar as they achieve this goal. Ptolemy's model of the heavens was not just a record of the past positions of the planets, but a prediction of future positions. The manifest model, by contrast, is not predictive. It tracks changes rather than anticipating them. This is not to say that the various forms of pattern completion achieved by sensory processes will never have any predictive content. For instance, the processes whereby the eyes track an object moving in a smooth trajectory may to some extent anticipate its motion. The claim that perceptual consciousness models processes by tracking them is meant as general description, rather than a definition, of the manifest model. With only trivial exceptions, perception itself merely reacts to, and thereby follows, events. The future may be predicted, but it is not usually perceived.

There is, of course, some plasticity in what is perceived, and some have made this plasticity an essential component of their philosophy of mind. Paul Churchland (1979), for instance, has claimed that we may someday be able to directly perceive, without any inference, the chemical composition of wine rather than its taste, and experience spiking frequencies in our brain rather than pain. No doubt perception has some plasticity, but it does not have unlimited plasticity. We cannot, with our unaided senses, see viruses or the gravitational field of the earth. This much, I trust, is too obvious from a straightforwardly scientific point of view to require further elaboration. I would further venture (contrary to Churchland) that we cannot, without training and inference, taste the chemical composition of our wine or perceive the activities of our brain. But inference is not perception. We must distinguish between perceiving, and perceiving *as* or perceiving *that*, as van

Fraassen proposes (1980). It is one thing to taste the wine, another to taste *that* it is wine. Persons with no experience of wine can nevertheless taste it. They do not realize that what they are tasting is wine, much less have any idea what its chemical composition might be. Those things must be learned, and that learning applied, even if its application is swift and unconscious. By the manifest model, I intend to refer to the model of the world provided by sense prior to, or independently of, learning. The manifest model is the information provided us by our *unaided* and *untutored* senses, the product of our biological evolution, rather than our cultural development.

I take it that the manifest model includes such disparate things as grass (and its colors), apples (and their tastes), garlic (and its smell), thunder (and other sounds), rainbows (and other non-substantial things), parades (and other processes), the sun (and other stars), people (and other mysteries). This is not to say that perceivers must see these things *as* grass, colors, or whatever. In addition, the manifest is not restricted to things that are outside the perceiver's body. Via proprioception one is aware of the relative position of arms and legs, whether one is walking or sitting, and so on. Bodily damage is also manifestly modeled, typically by a sensation of pain that carries information about the bodily location of the damage: a pain in one's finger caused by a thorn, a toothache, etc. There are also pleasures that have specific locations in the body. Needs of the body may be experienced as thirst, hunger, lust, and so on, while satisfactions of these appetites are manifest as pleasures. General states of the organism may be experienced as moods (sleepiness, boredom, restlessness) or emotions (anger, fear, hatred). In each of the cases listed, the neural state and corresponding state of consciousness arise without any aid from observational instruments or necessity for schooling.

Of course, in the educated human being, the information provided by education, whether the informal education of everyday experience or formal instruction by others, merges imperceptibly with the information provided by the senses. It is not always easy to tell where perception leaves off and education begins in ordinary acts of perceptual discrimination. The experienced outfielder simply "sees" where the ball is going and runs to intercept it, without any conscious inference. However, the fact remains that she could never have acquired this know-how in the first place without sensory experience. And though it may in many cases be practically impossible to draw a precise line between what the senses provide on their own, and what is learned by means of them, it does not follow that there is no distinction to be made. The fact that this distinction is sometimes successfully drawn, and quite precisely, is evidence that it is real. As Matthen (1988) points out, Bela Julesz (1960) and Edwin Land (1977) provide examples of how the distinction can be empirically determined. Julesz showed

experimentally that distance from the observer is, in effect, mapped by the manifest visual model, without the benefit of learned associations. He did this by demonstrating that slight lateral displacements embedded in randomly generated patterns of dots, viewed separately by each eye, will be resolved by the human visual system, and seen as different surfaces at different distances from the observer. Since there was nothing in the random patterns that could be distinguished or identified on the basis of past experience, the perception of depth by stereoscopic matching of patterns was shown to be independent of learning. Land demonstrated similar untutored abilities in color perception by experiments employing a similar logic. So these scientists have provided evidence that depth perception and color perception are included in the manifest image.

Depth perception raises another interesting issue. We know, due to the pioneering work of Hubel and Wiesel (1970), that the ability to perceive depth does require visual experience, if only in the minimal sense of light input to the eyes, during critical early stages after birth. Unless the eyes see some light during this early period, structures in the visual cortex critical to depth perception, namely ocular dominance columns, will not form. We are not born with the ability to perceive depth, but acquire it after birth. We cannot, therefore, simply identify manifest modeling with that done *innately* by perceptual systems, at least not if we want to include depth perception. And I think we do want to include it. The important fact, I suggest, is that the formation of ocular dominance columns is the norm, because the environment always contains some light. Though this particular development of the nervous system happens after birth, it is nevertheless a normal function of the visual system as shaped by evolution. It is automatically acquired under any of the wide variety of conditions normal for human babies, without the need for any special circumstances or activities on the part of those around it. A comparison can be drawn with feelings of sexual desire, and the pleasurable sensations of coitus. These, too, are hardly innate, since they arise long after birth. They do not, however, require any training, any special experiential input, in order to arise. They may be reckoned, therefore, as part of one's manifest modeling repertoire.

It is not, in any case, imperative from a logical point of view that the domain of the manifest be sharply demarcated. As van Fraassen points out (1980), a distinction, such as that between portable and non-portable televisions, may be perfectly sound even though there is no sharp boundary between the two. A pocket-sized television is clearly portable, while a very large screen model clearly is not. The distinction is not invalidated by the fact that there are console televisions that could be lugged around by the sufficiently robust, and small televisions that are designed to be carried but could not be carried by children. Likewise, the concept of the unaided and

untutored senses is sound, even if there is no obvious joint in normal perception where it can easily be cut away by the scalpel of analysis from the effects of such aids as eye glasses or such tutoring as provides the color vocabulary.

More importantly, except for the somewhat dicey case of newborn children, it is simply false that the manifest model is sharply demarcated from our other models of the world. Obviously we adults have many models of the world, such as commonsense models and religious models, and so on, though we have focused our attention here on the two models most germane to the problem of consciousness, the manifest and the scientific. For the sake of convenience, let us use the term *cultural model* as a tag for all of the models constructed with the aid of our naturally inherited manifest model. The scientific model is obviously a cultural model. Like any cultural model, it ultimately relies upon the natural cognitive endowment given us by evolution, including the manifest model. Before the scientific model existed, the manifest model was there, not only in our collective historical development, but also in the individual life of each of us. And between the manifest and scientific models are the innumerable cultural models past and present that have mediated our interchange with the world, and provided the bridge to the scientific model itself. So the fact that there is no gap demarcating the manifest and the scientific is no more surprising than the fact that there is no gap between the roots of a tree and its branches.

3.2.1 The Self-Centeredness of the Manifest

One crucially important fact about the manifest model is that it is *self-centered*. The scientific model, by contrast, is essentially social, shared, and public. It is made accessible to people in general by a judicious choice of *impersonal* coordinates. The scientific map of a flower or crystal, for instance, must include some specification of the scale of the map in terms of a widely shared standard of length, such as the meter. The manifest model is, by contrast, intensely personal: its coordinates are self-centered. To take an obvious, literal, example, in the manifest image the spatial positions of perceived objects are represented *relative to the perceiver*. It is as if there were a set of Cartesian coordinates in which the perceiver occupies the origin, with the x-axis extending in front, the y axis above and below, and the z axis from left to right. One sees things in terms of their distance from oneself, and their spatial relationship to oneself: in front, behind, above, below, left, or right. If you turn around, the entire visual array is re-mapped to accommodate the shift in these self-centered coordinates: what is in front is now behind, what was to the left is now to the right, and so on.

Consider our manifest mappings of thermal properties. Some things feel hot or warm to us, some cold or cool. However, these sensations do not map *temperature*. This is shown by the fact that if a piece of dry wood and a piece of iron are both at a temperature of 100° Celsius, the piece of iron will feel very hot indeed, while the wood will only feel very warm. Likewise, a piece of iron at 0° Celsius will feel very cold, while the wood will feel only cool. One reason for this is the difference in *heat conductivity* between the two substances, which is high for iron but low for wood. Thus a lot of heat flows from the hot iron to one's skin, making it feel very hot, and a lot of heat flows from one's skin into the cold iron, making it feel very cold. Wood has much lower heat conductivity, so heat flows are much smaller, making it seem less warm or cold as the case may be. Another factor enhances this difference between the two materials: their difference in *heat content*. Given the higher density of iron compared to wood along with its higher specific heat, a hot piece of iron contains more heat than a hot piece of wood. So not only does heat flow better between one's skin and the iron relative to wood, but there is more of it to flow from the volume adjacent to one's skin.

This sort of effect leads one to think that our sensations of hot and cold map *heat flow* to or from the skin, rather than temperature. Certainly this is closer to the truth. It allows us to explain, for instance, the phenomenon noted by Locke in declaring heat a secondary quality of objects (and later used by Berkeley to argue that heat is a mere idea): if one hand is warmed and the other cooled, then a single object at an intermediate temperature will feel cool to the warm hand and warm to the cool hand. This is due to the fact that heat flowing into the hand is experienced as warmth, while heat flowing out is experienced as coolness (assuming that the hand was neither terribly warm nor terribly cool to begin with). However, the heat flow model will not do for other reasons (see Akins 1996).¹ For one thing, there is the fact that various parts of the body have differing sensitivities to temperature. What feels pleasantly warm to one's hand may feel unpleasantly hot to one's foot. Generally speaking, the head and torso are more sensitive to heat and cold, as people become painfully aware when they wade into cold water: their progress is slowed when the cold water reaches their torso. There is roughly twice the density of thermal sensors on the torso as on the legs. The reason for this, presumably, is that it is far more essential to maintain normal body temperature in the torso and head than in the extremities. If you are exposed to dangerous cold, for instance, your body will respond by letting your arms and legs cool in order to conserve heat for the operation of essential organs such as heart and brain. So the feelings of warmth or cold are especially strong and salient when it comes to the torso and head.

This sort of phenomenon indicates that our sensations of heat and cold are keyed not to temperature, nor to heat flow, but to the potential for

harm or help for the perceiver. *Because* it is more important to maintain the temperature of the torso, the snowball in one's hand does not feel as cold as the snow stuffed inside one's shirt. In other words, the manifest thermal properties of such things as wood, iron, wading pools, and snow, that are mapped by our sensations of heat and cold may be properly defined only by reference to our interests and well-being. The manifest model is not of temperature, conductivity, specific heat, heat flow, or any other such properties of objects as are captured in the scientific model, but rather the salience, significance, or meaning of whatever these properties or their combinations have for one's own continued life and health. Our thermal sensations cleverly transform the complex thermodynamic properties (temperature, heat content, heat conductivity) of objects into a single dimension, the hot-cold dimension of the manifest model, that represents them solely and entirely in terms of our own interests.

That this is so is made obvious by the fact that a snowball feels not merely cold, but *uncomfortably* cold in one's hand, and *painfully* so down one's back. The hot-cold dimension is merged with various degrees of pain and pleasure: the warmth of the bath water simply *is* a sensation of pleasure; the sensation of extreme heat simply *is* a sensation of pain. This ensures that the mapping is self-centered in a straightforwardly selfish way. The coordinates of the manifest model merge the hot-cold dimension with degrees of pleasure and pain indicating in the most salient way possible the help or harm that the thermodynamic properties of things have for oneself. Thus the manifest model portrays thermodynamic properties only insofar as they are relevant to one's own interests and well-being. Representation could hardly be more greedily self-centered than this, and it is achieved by the creation of self-centered coordinates through the long process of evolution. The result is that the manifest model of sensations of hot and cold do not comprise models of the thermodynamic properties of things, but rather paint onto those things their thermodynamic relevance to oneself.

The concept of self-centered coordinates nicely explains a range of thermal experience.² First of all, we may note that the zero point of the manifest thermal scale is set to whatever concatenation of temperature, conductivity, and heat content that makes a thing innocuous to ourselves. The sensation of tepidity or indifference maps an air temperature of around 23° Celsius (unless the humidity is high, making it feel warm, or the humidity is low, making it feel cool). That is because air at that temperature poses neither hazard nor promise for us: it is cool enough to permit us to dissipate the heat generated by our metabolic processes, but warm enough to allow us to maintain the body temperature ideal for these processes. Water at the same temperature will, however, feel quite cool, since it will indeed cool us below ideal metabolic temperatures. Warmth and cold are experienced as properties

of bodies, whether or not these bodies touch our own. Either a campfire or the sun will feel pleasantly warm under the right circumstances, but painfully hot under others. But, though heat and cold are experienced as properties of external bodies, they are nevertheless felt as having a location in our own body. The warmth of the sun or campfire is felt precisely on those exposed parts of the body that face them. This tying of manifest heat and cold to bodily location has no exceptions. It is a further indication of the self-centeredness in the manifest mapping of this property. Just as a map is drawn on coordinates for the external compass directions of north, south, etc., the manifest map of thermodynamic properties is drawn on self-centered coordinates of good for me here, bad for me there, irrelevant for me anywhere (the way in which these coordinates may be combined will be taken up in 7.2). Once we see this, many things fall into place. It is only with very great difficulty that we can hold on to the hot mug that burns our hands: the sensation of great heat, which is indistinguishable from pain, is also seamlessly merged with the urge to let go of the mug.³ Very cold objects with high specific heat and conductivity, those which give us frostbite when touched, are felt as hot, rather than cold, stimulating the dropping response and thereby avoiding damage. Forms of heat that are extremely unusual in the environment in which we have evolved will not appear in the manifest mapping at all. X-rays, for instance, burn our tissues, but insensibly: they are so rare in nature as not to have been included in the manifest mapping. Thus the manifest model is self-centered in another sense: it is parochial, limited to those things that have happened to find their way into its locale and range of interests.

3.3 SCIENCE AS THE ESCAPE FROM SELF-CENTEREDNESS

To sum up, the manifest model is a process inside the perceiver that tracks self-centered aspects of the body and the external environment. It not only models the world, but remodels it in terms of its interest to oneself. It is not so much a re-presentation of the world, but rather an imposition upon it of our own selfish concerns. There was a time, long before we had developed science, when the manifest model was our only access to body and world. How, starting from this completely internal, self-centered model, could we ever have achieved a scientific model of what is external and independent of ourselves? Clearly the process was not an easy one, as is indicated by the fact that the scientific model is a recent innovation of the last few millennia, a tiny fraction of our time here on earth. Indeed, the achievement is not yet complete, and the riddle of consciousness locates the stage at which we have arrived in the process of developing the scientific model: the point where we

now contemplate a complete scientific model of ourselves, including the manifest image.

Let me preview the analysis of the riddle of consciousness that I will offer, though it is impossible to make the picture very clear at this stage. It is valuable to see our destination, if only vaguely on the distant horizon. The details will come into focus as we advance. The riddle of consciousness arises because we cannot completely transcend, progress beyond, or otherwise abandon, the manifest model. It is the onboard modeling system provided us by nature, and whatever else we may have achieved via the scientific model, it cannot enable us to completely escape this primordial system. It is *what we are*, and unless we become a quite different species of animal from the one we are, we are stuck with the manifest model as it is. Science permits us, at least to some degree, to escape the self-centered model of the world around us. But it does not permit us to escape our own nature, and our self-centered experience of the manifest world. This creates the illusion that *we* are somehow special, that we are unique among the things of this world, whereas we are special only to ourselves because we must be at the center of our own experience. We have a special epistemological significance to ourselves, rather than a special metaphysical status in the world at large. In the days of Galileo, Descartes, and Newton, the intuition that we are special led to the conviction that we essentially transcend the natural world itself, that our minds, which make us what we are, are not of this world. We no longer believe this, at least not officially. So now the sense that *we* are special is expressed in the mystic's conviction that we partially transcend not nature itself, but natural science: there can be no science of consciousness. Both are illusions born of the fact that the manifest model must hold a privileged position in our cognitive life. But to see this, we must first go back to the time before we devised the scientific model.

So how did we come as far as we have? How did we escape the self-centeredness of the manifest and achieve the scientific model? The answer is implicit in the previous chapter, but it is important to draw it out explicitly in light of our investigation of the manifest model. There are two interdependent themes: the replacement of self-centered coordinates by externally-centered ones and the restriction of scientific modeling to the spatial subset of manifest properties. Let us consider them in order.

The replacement of self-centered coordinates by externally centered ones would have arisen naturally in a communicative, highly social species such as our own. To consider an example, the self-centeredness of the manifest mapping results in the direction "left" for one observer corresponding to the direction "right" for someone facing in the opposite direction. We can overcome this relativity by using less selfish reference points: towards the sea, towards the mountains, and so on. These new

mappings of directions are not absolute, though they are obviously less relative. For everyone in a given locale, a mountain may provide a reference point that is free of the vagaries of self-centered "left" and "right." From a more distant point of view, the use of the mountain as a coordinate is seen to suffer from its own sort of self-centeredness. What direction the mountain marks depends on the location of the observer. More widely traveled individuals will prefer to use the cardinal points of the compass as defining direction: north, south, east, west. But from a point of view some millions of miles above the earth, it is plain that the cardinal points themselves are infected with relativity due both to the sphericity of the planet (what is east for you is nearly straight up for someone situated 90 degrees of longitude east of you) and to its motion (the direction due-east-for-you traces a path moving through the heavens as the Earth turns on its axis). That self-centeredness cannot be completely expunged, at least from the terms used for direction, seems to follow from the very logic of relative terms. To say that left, right, up, west, north, are relative, is to say that they are indeterminate until some point of reference is assigned. Something can be left or right only *relative to* some person or thing, up or down *relative to* some location on the surface of the Earth (as those at the antipodes are only too well aware), or west-by-northwest *relative to* compass coordinates.

Now, an absolute standard would be one that is completely independent of such relativity. Such standards would be, of course, the sort ideal for science. But ideal or not, it is impossible that direction be specified absolutely. Direction must be defined relative to something in the empirical world. Direction (like velocity, time, position, mass, etc.) is relative. And while it would seem that space, gravity, length, distance, time, etc., may be imagined absolute (as they were, indeed, imagined by Newton and Newtonians), their actual *use* in the world to represent and map things demands the introduction of measures defined relative to things in the world chosen as standards: the meter (relative to the established standard), the light year (relative to the time it takes our planet to circle its star), the erg (relative to the mass of standard measures of water), etc. Thus the idea that science has the complete *elimination* of self-centered relativism as its goal makes no clear sense. Nevertheless, we can make sense of the movement *away* from the complete self-centeredness of the manifest mapping towards a less provincial view. The movement away from the self-centered coordinates towards the compass coordinates permits vastly larger numbers of people to share the modeling of directions. Left, right, in front, and behind work for me, now, but must be translated if I need to communicate with you. North, south, west, east serve as a common system for all. The *ideal* would be a system that would serve for all intelligent beings whatever: a model of the world that is not merely true-for me, or even merely true-for us, but rather true-for every

possible thinker. Whether this ultimately is a cogent idea, it does not seem to offer much by way of *practical* direction, and science is a practical affair.

Fortunately, the practical advice is clear: replace the self-centered coordinates provided by biology and shaped by natural selection with externally anchored coordinates. The manifest model of the earth as a plane lying below the vaulting heavens was superseded in Aristotle's time by a spherical planet in which up and down became relativized to position — and so the apparent problem of what held the earth up disappeared. The manifest motion of the sun, moon, and stars around a stationary earth was replaced in Galileo's time by a scientific model in which we are hurtling through space aboard our spinning earth — and so the apparent problem of what turned the starry spheres of the heavens vanished. More recently, the flat, absolute space, time, and velocity of Newton has been remodeled by Einstein as a curved space-time sporting only relative positions, times, and velocities — and so the apparent problem of determining the Earth's velocity through the luminiferous ether evaporated. In each case we move farther away from the way things appear to us, farther away from the manifest model, farther away from our selves as the privileged point of view. Towards what? Ideally towards a model that, by not being tied in any way at all to our own point of view might accommodate any point of view, and in this universally intersubjective sense achieve the only practical measure of objective truth. Whether or not this ideal can be approached, whether, indeed, it even makes any sense at all to try to move towards it, there is nevertheless a clear *methodological* sense in systematically moving away from the self-centeredness of the manifest model. History teaches that this de-relativization of our maps and models is one essential thread of scientific progress (compare Nagel 1986, Foss 1993).

3.3.1 The Pythagorean Intuition

A second thread is the exclusive use of the spatial subset of the manifest properties. In the last chapter I argued that science models the world geometrically and defines its ontology geometrically. The question we need to consider here is not whether the spatial properties played this decisive role in the rise of science, but rather, *why*? Obviously the spatial properties are among those included in the manifest model. When the ancients discovered geometry, they did not thereby discover space or invent our form of perceiving it. The perception of space and spatial properties had already been achieved long before in the manifest model. The ancients developed geometry by building upon the spatial intuitions which that model provided. The reason that spatial properties could play such a powerful role in the rise

of science, not only helping us escape from self-centeredness but forming the permanent core of the scientific model, was described in the last chapter in terms of the *clarity* and *practicality* of geometry, which together give life to the Pythagorean Intuition. All of the great physicists had this intuition. Newton, in particular, even went so far as to claim that space (along with time) was an essential aspect of God (and hence of all creation):

He endures forever, and is everywhere present; and by existing always and everywhere, *he constitutes duration and space*. (Newton 1687, p. 545, my emphasis)

Duration and space are God's eternal and omnipresent Self! We modern thinkers would not hazard such an ambitious proposition. Nevertheless, it is only because we implicitly share the Pythagorean Intuition that our spatial perceptions allow us to directly see into the essential nature of physical things that we feel the sharp contrast between our understanding of them and our understanding of consciousness. We feel the Galilean Intuition that our comprehension does not penetrate beyond the surface of consciousness only because we feel we can, by contrast, look right into the heart of things by means of the scientific model. I cannot hope in a few pages to dislodge the powerful Pythagorean Intuition that is necessary to this contrast, and hence to the riddle of consciousness. Still, there is evidence lying ready to hand that is sufficient to budge it, if only slightly, from its sacred position in our thinking. This part of my argument is not absolutely necessary, inasmuch as the Galilean Intuition can be addressed separately, and will be, in the next chapter. But both intuitions, as I see it, lead us away from the true, pragmatic heart of science and into metaphysical speculation, speculation that gives rise to the riddle of consciousness. I would be remiss if I did not at least address them both.

We might begin by observing that our spatial intuitions, Pythagorean or otherwise, are not shared by all creatures capable of some form of perception. Many creatures are limited to olfactory awareness of odor and/or flavor as sensory modes, and hence have no awareness of spatial properties. So spatial awareness is not necessary for existence as such. But is it a necessary condition for science? Or was it merely a historical accident that we did not develop science on the basis of other sensible properties, such as the sensuous qualities? Is it inconceivable that science might be based on color, odor, or flavor, rather than length, breadth, and height? If canines had gotten ahead of us in the evolutionary race, could they have developed a sophisticated knowledge of the world based upon their exquisite sense of smell? We are troubled by the fact that odor, (and the rest of the sensuous qualities) cannot be reduced to geometry — why are we not troubled by the fact that geometry cannot be reduced to odor? One way to answer all of these questions at once is to rein in speculation, and refer them to the available

evidence. Geometry has proven itself *in practice*, the sensuous qualities have not. Spatial properties have proven to be the only ones adequate to the role of general-purpose modeling tool of physical science. This is, I think, ultimately the right approach to take. But this answer fails to address the Pythagorean Intuition. Instead, it just reformulates the questions just posed in a new form: *why* has geometry become *the* modeling tool of science? Is it because Pythagoras was right in thinking that reality itself is somehow deeply mathematical by nature?

Large parts of our brains are dedicated to visual processing, and the development of our rather large brains is in large part due to the disproportional growth, or hypertrophy, of the visual cortex among our ancestors. It was, likewise, the hypertrophy of the spatial subset of the manifest properties that engendered the rise of science. We cannot help but notice that geometry itself, the *sine qua non* of science, bears a special relationship to our faculty of sight. We simply *see* the truth of the more basic truths of geometry: that two straight lines meet in at most one point, that two solid spheres can touch at only one point, that the angles opposite the equal sides of an isosceles triangle are also equal, and so on. Indeed, the more recondite truths of geometry are likewise learned by means of "demonstrations" whereby one learns to see their truth by construction of simpler, more easily seen truths. In the *Gauss Contest 2000*,⁴ a national mathematics competition for gifted schoolchildren in Canada, the most difficult questions, the ones at the end of the question sheet after the many computational questions, algebra problems, and so on, are those that can only be solved by means of spatial imagination. They involve such things as finding the number of squares that can be formed by connecting quadruples of dots in a pattern of dots, finding the ratio of tiles of different colors in tiling patterns, discovering axes of symmetry, resolving complex geometrical figures into simpler components in order to calculate their area, and so on. In every case geometrical "insight" is required. Kant claimed that geometry simply *is* "the form of our sensible intuition." Be that as it may, there is obviously a close relationship between geometry and vision. On the other hand, there is more to geometry than is given in vision alone. This is implied in the Kantian dictum, which refers to the form of our *sensible* intuition, rather than the form of our *visible* intuition. A key piece of evidence concerning the special place of spatial properties in both the manifest and scientific models is the fact that they are not limited to any single sensory modality.

Some two millennia before Galileo, Descartes, and Newton had fastened on the spatial properties as the essence of the physical, Aristotle had identified them as *common sensibles*. The list of properties was the same: shape, size, number, motion, and rest. The need for a *common sense* that

would serve the purpose of linking the various sensory modalities was obvious to The Philosopher. If a dog hears a noise in the bushes, she will look at the source of the noise, and, if it seems interesting, move towards it. Aristotle reasoned that the dog must coordinate its vision, as well as bodily movement, with its hearing. He reasoned that the communication of the various sensory modalities was accomplished by means of their common spatial properties. Both hearing and vision have a spatial aspect, as do all of the other senses. In proprioception we are aware of the spatial arrangement of our own body. We feel ourselves moving, or sitting still, or turning on the spot, by means of our vestibular system. The pain of a thorn pricking the skin is felt in a particular spatial location. Whereas colors can only be seen, and heat or cold only felt, shape, size, number, motion, rest, can be seen *and* felt. Common sense was the faculty Aristotle postulated to explain the integration of the senses in the perceiving animal.

His concept of common sense was closely related to the original concept of consciousness. The word, "consciousness," comes from the Latin *conscientia*, which means to know together. If a group of us each knows that Caesar is about to be assassinated, and knows that each of the others knows this, this complex of communal knowledge comprises *conscientia*. The perceptual knowledge of manifest consciousness is a closely analogous thing. Each of the individual sensory modalities models the specific aspects of the world to which it has access, and common sense inter-coordinates (2.2.6) the resulting model. On Aristotle's approach, it is the spatial properties that make possible perceptual consciousness as such, as opposed to a mere collection of idiosyncratic sensory processes. Common sense transcends the self-centered view of each of the sensory modalities, and so achieves for a single individual what the Big Model does for the community of individual scientists. Common sense provides the *little model* for the individual animal. In both the Big Model and the little model, spatial properties are the common sensibles, and geometry provides the coordinate space of the model. Given a set of distinct sensory systems detecting quite disparate properties, common sensibles provide a common informational coin to permit information exchange and so integrate the systems in the set. The dog hears the sound, and this permits her eyes to look in the *direction* of the sound. Sound has a direction, just as sight has a direction, and direction is spatial. And it works the other way, too. The dog can listen in the direction of something she sees (this something that is most obvious in dogs with erect, movable ears). The various proprioceptive modalities report to body maps in the cortex, thus *locating* their referents in the body. The burning of the hot food is in your mouth, so manifest heat has a location in space. Likewise, the prickling of the thorn is in your big toe, which is visibly located down there on the ground as seen from the vantage point of your eyes. Thus the spatial properties thus

clearly qualify as common sensibles. This then puts vision in the driver's seat, since it is the sensory modality that is most thoroughly spatial.

Perhaps the crucial importance of spatial properties can be understood in this way: when it came to the business of integrating the sensory modalities, the task was best achieved by the modality with the greatest precision, accuracy, and scope. The information provided by hearing, smell, touch, taste, and the rest, could be embedded in the information provided by vision, but not conversely. This would explain why spatial properties became the common sensibles: every other property could be embedded in it. It would also explain why we are interested in the issue of whether or not the sensuous qualities can be reduced to the spatial ones, but not conversely: the sensuous qualities are ill-adapted to serve as a common sense. The most obvious handicap of the sensuous qualities is their narrow scope. Immediately disqualified are bodily sensibles like hunger, anger, pain, or lust, which simply do not apply to external things. The sensuous qualities of external things, colors, flavors, odors, temperature, and the rest, each have a restricted range of application. Things that may be seen may nevertheless have no flavor or odor (sun, stars, rainbows, distant things), while flavors and odors may persist on their own without the objects to give rise to them, colorlessly, invisibly. Thus, by a process of elimination, we are left with the visual sense as the best option for a common sense.

Perhaps the sheer size of the neural processes (the large optic nerves, the large visual cortex) dedicated to the most transparently spatial modality, sight, relative to those for the other senses, gives some indication of the evolutionary importance of spatial properties compared to others. And we must not overlook the fact that there is a brilliant logic of space, namely geometry, that finds no equal among the other sensory properties. Whether or not we are persuaded by any quasi-Kantian theory along the lines that geometry is but the form of our visual perception, there is nevertheless very good reason to think that the visual faculty is intimately connected to geometry. It is plausible that our development of geometry and intuitive feel for it depend on the logic of our visual cortex. But note that if this, or something like it, is the correct explanation of the special role of geometry in our modeling of the world around us, hence in our understanding of the world, it also explains our Pythagorean Intuition in a way quite opposite to the way Pythagoras himself intended. Pythagoras, like all of the other Pythagoreans before or since, believed that the clarity and practicality of geometry was due to the fact that nature itself was deeply mathematical. Nature was supposed to somehow be constructed out of geometric entities, or written in the language of geometry, which was sometimes taken to be the logic of God's own sensorium. But if the preceding explanation is correct, these Pythagorean ideas are nothing more or less than just one more case of

our manifest sense painting our own self-centered coordinates onto the world, rather than the other way around. Geometry is explained not as the logic imprinted upon our own models by the nature of reality itself, but the converse. We have imprinted the geometry of our visual system onto the world, just as we have painted it with sensations of heat and cold. Since our geometrical models have proven successful, we have adopted the view that the world is itself geometric — that is the Pythagorean Intuition. But before we begin to take this intuition seriously, we should note that the inference on which it is based has the same logic as concluding that the world is flat because we have constructed some very successful flat maps of it.

3.3.2 Metaphysical Modesty and the Pythagorean Intuition

Let me make it perfectly plain that I am not interested in deciding whether the neo-Kantian account just adumbrated is true, or whether the Pythagorean account is true, or which is better. I am not interested in determining the ultimate nature of reality, or anything nearly so ambitious. However, I do think it is essential to realize that the Pythagorean Intuition is an essential part of the riddle of consciousness, particularly as concerns the riddle of qualia. Because science is so profoundly saturated with the Pythagorean Intuition, we immediately find the idea of a scientific explanation of consciousness paradoxical. But there is plenty of reason to think that the Pythagorean Intuition is wrong, and therefore that the riddle of consciousness is based on a misunderstanding. And whereas demonstrating the truth of the countervailing neo-Kantian intuition is completely inconsistent with the metaphysical modesty of 4M, showing the unsoundness of the Pythagorean Intuition is not. Indeed, it is completely in the spirit of 4M to do so, and entirely germane to the business at hand.

So, in brief, the *individual* advantages of spatial properties over and above the other manifest properties are:

- a) they are common sensibles;
- b) they have a powerful logic, namely geometry; and,
- d) they have proven practicality.

These individual advantages, moreover, have a *synergistic* effect, making them far more powerful together than alone. Individually they represent necessary conditions for the rise of science, but collectively they amount to something like sufficiency. Of course, history might have conspired to keep us from any true science in any of a number of ways, perhaps by keeping us in small, solitary bands. But the spatial properties of the manifest model nevertheless provide the cognitive basis for the rise of science. Spatial properties provided the coordinates unifying the manifest model itself, and

thereby the basis for all scientific modeling and ultimately the generation of the Big Model. In brief, the combined effect of the three factors above is manifested in the power of the physical sciences. This power in turn confirms the Pythagorean Intuition on which it is based.

We can begin to challenge the Pythagorean Intuition by noting three things that *cannot* be included on the list above as advantages for manifest spatial properties:

- 1) they are essential properties of sensible things;
- 2) they are the only essential properties of sensible things;
- 3) they are, or include, the intrinsic properties of sensible things.

Let us consider why not.

As for the first point, the three-dimensional space of the Pythagorean Intuition simply is not the space of the world according to the scientific model. Odd as it may seem, the scientific project that was founded on the conviction that the physical world is constructed in the flat three-dimensional space of the manifest image has resulted in the rejection of that foundational claim. According to general relativity theory, the three-dimensional, Euclidean space of Galileo, Descartes, and Newton is a mere abstraction from — indeed a mere appearance of — the real state of affairs. The space of this world, according to modern physics, is a curved four-dimensional space-time continuum. The space of the manifest model, the space that is uncritically assumed in the Pythagorean Intuition to be the space of the real world, turns out to be an illusion, on a par with the illusion that the surface of the earth is flat. Though when you sit in a boat in the middle of the sea on a calm day you can see the surface of the water arrayed around you in an apparently endless plane disappearing asymptotically at the horizon, this is an illusion caused by the enormous size of the sphere of the earth compared to your own height on the boat. Should you try to construct a very large rectangle on the surface of the earth (such as the boundaries of Saskatchewan or Colorado), you would be forced to admit it was an illusion: you take a perfectly straight line as your baseline, extend two perfectly straight lines at right angles from it, and, lo and behold, they converge rather than stay parallel. In a closely analogous way, we have discovered that flat, three-dimensional space is an illusion. So the argument is very straightforward: the Euclidean, three-dimensional spatial properties of manifest sensibles is an illusion, and these illusory qualities can hardly be the essence of physical objects.

This argument is so brief that it may not be taken seriously, but brevity is quite compatible with validity, soundness, and truth. Nor should the argument be rejected on the grounds that it is seemingly paradoxical, inasmuch as it is prefaced with the claim that science was founded on the view that the universe is constructed in three-dimensional space, and then concludes quite the opposite. For one thing, science has since come to reject

its original foundation for perfectly good reasons. For another, there would be nothing wrong from a logical point of view if the proposition, that physical space is Euclidean, implies that it is not, since this entails that it is not.⁵ Various allied scientific considerations also support the conclusion, such as the fact that the manifest model is a product of evolution, hence not apt to be an ideal representation of the essence of things, but only sufficient for the purposes of survival and reproduction. Even so, it is extremely difficult to let go of the Pythagorean Intuition simply because we have no other way of conceiving of things. We have no alternative but to think in terms of the space of the manifest model. Imagine creatures whose perceptual mechanisms were confined to a two-dimensional plane, like the surface of a piece of paper. Their view of ordinary things, such as an orange, would be severely truncated. They would only see the roughly circular cross section where their perceptual plane sliced through the orange. Supposing that the plane was curved, they would have no way of directly perceiving that curvature. Indeed, they would have difficulty imagining what such a curvature could amount to, since conceiving it would require conceiving a third dimension through which the two that they perceive might bend. Of course, the creatures themselves could not be two-dimensional, since they must be made of atoms like the rest of us, and exist in the same (four-dimensional) space as the rest of us. It is just that they are perceptually unaware of their third dimension.

But we are just like these creatures, according to our standard scientific model of the universe: we are four-dimensional creatures in a four-dimensional world perceptually aware only of an apparently flat, three-dimensional slice of ourselves and our world. The book you hold is not really a three-dimensional object, but rather a four-dimensional object which your manifest model does not correctly capture. What you see is a projection of the book onto the three dimensions you can perceive, just as the two-dimensional creatures above see only the projection of the orange onto the space of their perceptual mechanisms. The manifest model is simply inadequate for fully revealing the reality it confronts. And in the most recent attempts to provide a unified theory of all of the forces and particles of physics, eleven-dimensional membranes are employed. Insofar as these scientific models succeed, to that extent the manifest model is deficient, and the Pythagorean Intuition false.

Still, it may be argued, though the space of reality is of a higher dimension than that of the common sense of the manifest model, reality is nevertheless essentially spatial in this higher-dimensional way. However many dimensions it contains, real space is, essentially, an extension of the space we perceive. Well, maybe, though this is obviously a very ambitious metaphysical thesis. What is more certain is that our *model* of reality, employing however many dimensions, is spatial — but this just brings us

back exactly where we started from. Yes, we model things geometrically. Yes, we have developed more sophisticated geometries to overcome the deficiencies in modeling reality via the geometry of manifest sense provided us by our biological nature. But, no, it does not follow that geometry must therefore be an *essential* property of things. Ptolemy might equally well have argued that perfect circular motion was an essential property of the heavens, since, given suitable sophistication via such devices as epicycles, deferents, and equants, it had been used successfully to model the movements of the planets. But clearly perfect circular motion is not the essence of the heavens, and the argument is invalid. By logical parity, then, we cannot conclude that geometry is the essence of reality. So we have no justification for the Pythagorean Intuition.

The second point is also clear from a scientific point of view: spatial properties are not the *only* properties essential to sensible things. This could be taken as a straightforward implication of the first point: given that spatial properties are not essential properties of things, they can hardly be the only essential properties. Despite the formal adequacy of this inference, it may not suffice to budge the Pythagorean Intuition. So it is well worth reflecting for a moment on the situation. On one hand science models whatever it confronts geometrically, but nevertheless, there is obviously more to real things than geometry alone. Consider, for instance, any two objects lying on your desk. There is a small gravitational force between the two objects that makes them attract each other with a force inversely proportional to the distance between them. But there is nothing in pure geometry that makes this the case. It is, rather, a fact that has been modeled geometrically, and only geometrically. How else could Newton's inverse square law be conceived? But no amount of geometry all by itself could reveal what it is, in the objects, that makes them exert this attractive force upon each other. Call this *The Riddle of Gravitation*. This is, surely, among the ultimate metaphysical mysteries; don't bet on anybody ever solving it. Of course, it is possible to *model* the gravitational force in terms of the bending of space-time by the masses of the objects. However, there is nothing in pure geometry that *requires* mass to have this effect on space-time, and no amount of pure geometry will ever explain why masses have this effect. The physical model simply represents the effect, in much the same way that a map of Hudson's Bay simply represents its shape. Just as the map of Hudson's Bay does not capture all of its properties, so too the model of these objects does not capture all of their properties. Clearly there is something about masses that makes them attract each other, or, equivalently, makes them bend space-time. Whatever it is, it is not a matter of the logic of space alone, not a matter of pure geometry. From a geometric point of view, masses could just as easily repel each other, or attract each other by an inverse cube law, or have no effect on each other.⁶

The third point follows close on the implicative heels of the first two: if the essence of things is neither captured nor exhausted by their manifest spatial properties, neither are their intrinsic properties. Presumably there is some intrinsic property, or properties, of masses themselves that makes them attract each other. If that is so, there is nothing in their geometry alone that tells us what it is. Indeed, whatever *intrinsic* property of masses there might be that makes them attract each other, or bend space-time, it can only be captured in a model via its *extrinsic* effects — and this is precisely what can be captured geometrically. Gravitational attraction can be understood and modeled only in terms of the way it tends to make bodies move, that is, via its extrinsic spatial effects. These clearly can be captured in a geometric model. Just as clearly, nothing more than this can be modeled geometrically. Therefore, the intrinsic gravitational properties of mass escape scientific modeling. Likewise, the intrinsic properties of the other three fundamental forces escape, as do, therefore, all of the remaining physical properties that can be defined in terms of them. Thus, scientific modeling does not capture intrinsic properties of things at all.

3.4 THE COUNTER-INTUITIVENESS OF THE RIDDLE

These three negative results entail that we must reject the false contrast between geometric and sensuous qualities expressed in the Pythagorean and Galilean Intuitions. We habitually think that ordinary physical science portrays the essential properties of physical things, that all of their essential properties are (or can be) so portrayed, and that we thereby see right into their intrinsic natures. By contrast, consciousness appears a dark mystery. We need to clearly realize — to see *intuitively* — that the true state of affairs is nothing like what this contrast assumes. Our intuitions need to be adjusted to accommodate the fact that our geometry is merely our best tool for modeling the universe, rather than the axis of the universe itself. We need to gain an intuitive feel for the fact that the intrinsic essence of physical things remains dark, mysterious, and unexplained despite the triumphs of science. Once this realization is achieved, we are equipped to recognize that the mysteries usually thought to be facing a science of consciousness are not in any way special. Our intuitions tell us, for instance, that the sensuous qualities are non-geometric. But so, too, are the fundamental forces and particles of physics — but that is no bar to understanding them scientifically. Science does not require that the things it models *be* geometric, but only that they can be modeled *geometrically*. And there is no reason to think, no immovable intuition indicating, that consciousness, including its many-colored qualia,

cannot be just as successfully modeled by science as any other natural phenomenon.

After many centuries of thinking of the physical world in geometric terms — and with such stunning success — we have become imbued with the spirit of Pythagoras. We may not explicitly endorse the Pythagorean Intuition that the world is, at its core, geometric. We may not even believe in essences, or be willing to speculate at all about what the essential properties of the physical world might be. Nevertheless, our thinking about the riddle of consciousness is colored by tacit Pythagoreanism. If you do not think this is so, then try to imagine what the problem of consciousness comes to when the Pythagorean Intuition is explicitly *denied*. Imagine for a moment that the world is not deeply geometric, that the “pure” science of geometry is but a model of that narrow part of the world open to our visual sense. Geometry is a powerful tool when it comes to modeling the universe around us — at least, the most powerful tool we have. But it does not reveal the essence of the physical. Now, what is the problem supposed to be as far as the science of consciousness goes? The problem is that there is some essence of conscious experience, particularly *sensuous* experience, which escapes geometry altogether. Well, okay. But what is supposed to follow? That the science of consciousness is impossible?

Well, that is just the Galilean Intuition. Sorry, and with apologies to Galileo, but it is hard to really feel any intuitive sympathy for this conclusion once the Pythagorean Intuition is denied. If science has been so successful in its modeling of the world around us despite the fact that the world is not essentially geometric, then the non-geometric properties of consciousness do not appear to be any barrier in principle to the science of consciousness. And if there is no barrier in principle, we will need some argument based on the difficulties faced *in practice* by the science of consciousness if we want to argue against its possibility. (I will argue in chapters 5 through 7 that from a practical, methodological point of view, the science of consciousness looks pretty healthy.)

The big payoff is this: if the science of consciousness is possible, then the riddle of consciousness is solved. The riddle of consciousness requires that there be *no* science of consciousness. If science can model consciousness, then there is no particular problem with consciousness as such. It is just a scientific problem, not a mystery. Of course, the scientific model of consciousness has to have explanatory power. It has to truly enable us to better understand consciousness, or else the riddle of consciousness endures. In the next chapter, I will argue that the scientific model of consciousness would have full explanatory vigor, that it could, indeed, explain the sensuous qualities themselves, even in their internal Lockean aspect as pure denizens of consciousness, that is, as qualia. Consciousness then is nothing more than a

blank area on the scientific map, like the part for Greenland on the old map on my wall. It is not a state-of-the-art map as far as Greenland is concerned, so it just portrays the bulk of the big island as a white spot, its geographical features uncharted, hidden as they were under the glaciers at the time the map was printed. If the science of consciousness is possible, the riddle of consciousness is like the conundrum of Greenland geography when my map was drawn — who knows what mysteries may lay in store for us down there under the ice?

NOTES TO CHAPTER 3

[1] Kathy Akins (1996) presents some of the scientific findings I will employ concerning our sensory processing of thermal information. I am also influenced by her telling characterization of this sensory process as “narcissistic.” But whereas she portrays *individual sensory systems* as responding only to those things in the world that have interest to *them* (the sensory systems themselves — her view a cousin, apparently, of Dennett’s 1991, pp. 237–42, “pandemonium” model of consciousness), it seems generally more accurate and plausible to me that separate perceptual systems each provide ranges of information relevant to the single organism or self to which they belong. Surely natural selection would favor such a centralized arrangement over anarchy of the senses. Akins also sees the evidence of sensory narcissism as indicating the necessity of distinguishing these systems “sensory motor” functions from their “ontological” (that is, representational) functions. Thus, she reads the evidence as placing the goal of a naturalistic account of intentionality even further out of reach. I will argue (Ch. 7), to the contrary, that the self-centeredness of manifest perception suggests a quite promising approach to a scientific account of intentionality.

[2] Indeed, I do mean to say that our current scientific model of our manifest modeling of thermal properties does explain thermal qualia, that is, *conscious* phenomena — if only in part. But this is not the place to get into this issue which will be the focus of Chapters 5 and 7.

[3] I will argue in Chapter 5 that under the influence of drugs such as Demerol, the apparently seamless blend of the sensation of heat-and-pain, along with the urge to release the object causing them, can come apart. This fact, however, does nothing to undermine the usual unity of these things in the manifest model.

[4] The Gauss Contest is composed and administered by The Centre for Education in Mathematics and Computation, University of Waterloo, Waterloo, Canada.

[5] From “If P, then not-P,” we can validly conclude “not-P.” After all, the first sentence is equivalent to “Either not-P or not-P,” which obviously entails “not-P.”

[6] It is an interesting fact — and a very relevant one as far as the riddle of consciousness is concerned — that explanations of qualia, intentionality, or consciousness are met with the response that they do not make qualia, intentionality, or consciousness *necessary*. Chalmers (1996) has the virtue of being quite explicit on this point. But this is to set an entirely inappropriate standard (as shall be argued in the next chapter).