

KNOWLEDGE

Nagel, T. 1986. *The View from Nowhere*. Oxford: Oxford University Press.

1. *Skepticism*

The objective self is responsible both for the expansion of our understanding and for doubts about it that cannot be finally laid to rest. The extension of power and the growth of insecurity go hand in hand, once we place ourselves inside the world and try to develop a view that accommodates this recognition fully.

The most familiar scene of conflict is the pursuit of objective knowledge, whose aim is naturally described in terms that, taken literally, are unintelligible: we must get outside of ourselves, and view the world from nowhere within it. Since it is impossible to leave one's own point of view behind entirely without ceasing to exist, the metaphor of getting outside ourselves must have another meaning. We are to rely less and less on certain individual aspects of our point of view, and more and more on something else, less individual, which is also part of us. But if initial appearances are not in themselves reliable guides to reality, why should the products of detached reflection be different? Why aren't they either equally doubtful or else valid only as higher-order impressions? This is an old problem. The same ideas that make the pursuit of objectivity seem necessary for knowledge make both objectivity and knowledge seem, on reflection, unattainable.

Objectivity and skepticism are closely related: both develop from the

idea that there is a real world in which we are contained, and that appearances result from our interaction with the rest of it. We cannot accept those appearances uncritically, but must try to understand what our own constitution contributes to them. To do this we try to develop an idea of the world with ourselves in it, an account of both ourselves and the world that includes an explanation of why it initially appears to us as it does. But this idea, since it is we who develop it, is likewise the product of interaction between us and the world, though the interaction is more complicated and more self-conscious than the original one. If the initial appearances cannot be relied upon because they depend on our constitution in ways that we do not fully understand, this more complex idea should be open to the same doubts, for whatever we use to understand certain interactions between ourselves and the world is not itself the object of that understanding. However often we may try to step outside of ourselves, something will have to stay behind the lens, something in us will determine the resulting picture, and this will give grounds for doubt that we are really getting any closer to reality.

The idea of objectivity thus seems to undermine itself. The aim is to form a conception of reality which includes ourselves and our view of things among its objects, but it seems that whatever forms the conception will not be included by it. It seems to follow that the most objective view we can achieve will have to rest on an unexamined subjective base, and that since we can never abandon our own point of view, but can only alter it, the idea that we are coming closer to the reality outside it with each successive step has no foundation.

All theories of knowledge are responses to this problem. They may be divided into three types: *skeptical*, *reductive*, and *heroic*.

Skeptical theories take the contents of our ordinary or scientific beliefs about the world to go beyond their grounds in ways that make it impossible to defend them against doubt. There are ways we might be wrong that we can't rule out. Once we notice this unclosable gap we cannot, except with conscious irrationality, maintain our confidence in those beliefs.

Reductive theories grow out of skeptical arguments. Assuming that we do know certain things, and acknowledging that we could not know them if the gap between content and grounds were as great as the skeptic thinks it is, the reductionist reinterprets the content of our beliefs about the world so that they claim less. He may interpret them as claims about possible experience or the possible ultimate convergence of experience among rational beings, or as efforts to reduce tension and surprise or to increase order in the system of mental states of the knower, or he may even take some of them, in a Kantian vein, to describe the limits of all

possible experience: an inside view of the bars of our mental cage. In any case on a reductive view our beliefs are not about the world as it is in itself—if indeed that means anything. They are about the world as it appears to us. Naturally not all reductive theories succeed in escaping skepticism, for it is difficult to construct a reductive analysis of claims about the world which has any plausibility at all, without leaving gaps between grounds and content—even if both are within the realm of experience.

Heroic theories acknowledge the great gap between the grounds of our beliefs about the world and the contents of those beliefs under a realist interpretation, and they try to leap across the gap without narrowing it. The chasm below is littered with epistemological corpses. Examples of heroic theories are Plato's theory of Forms together with the theory of recollection, and Descartes' defense of the general reliability of human knowledge through an a priori proof of the existence of a nondeceiving God.¹

I believe, first of all, that the truth must lie with one or both of the two realist positions—skepticism and heroism. My terminology reflects a realistic tendency: from the standpoint of a reductionist, heroic epistemology would be better described as quixotic. But I believe that skeptical problems arise not from a misunderstanding of the meaning of standard knowledge claims, but from their actual content and the attempt to transcend ourselves that is involved in the formation of beliefs about the world. The ambitions of knowledge and some of its achievements are heroic, but a pervasive skepticism or at least provisionality of commitment is suitable in light of our evident limitations.

Though a great deal of effort has been expended on them recently, definitions of knowledge cannot help us here. The central problem of epistemology is the first-person problem of what to believe and how to justify one's beliefs—not the impersonal problem of whether, given my beliefs together with some assumptions about their relation to what is actually the case, I can be said to have knowledge. Answering the question of what knowledge is will not help me decide what to believe. We must decide what our relation to the world actually is and how it can be changed.

Since we can't literally escape ourselves, any improvement in our beliefs has to result from some kind of self-transformation. And the thing we can do which comes closest to getting outside of ourselves is to form a detached idea of the world that includes us, and includes our

1. A fourth reaction is to turn one's back on the abyss and announce that one is now on the other side. This was done by G. E. Moore.

possession of that conception as part of what it enables us to understand about ourselves. We are then outside ourselves in the sense that we appear inside a conception of the world that we ourselves possess, but that is not tied to our particular point of view. The pursuit of this goal is the essential task of the objective self. I shall argue that it makes sense only in terms of an epistemology that is significantly rationalist.

The question is how limited beings like ourselves can alter their conception of the world so that it is no longer just the view from where they are but in a sense a view from nowhere, which includes and comprehends the fact that the world contains beings which possess it, explains why the world appears to them as it does prior to the formation of that conception, and explains how they can arrive at the conception itself. This idea of objective knowledge has something in common with the program of Descartes, for he attempted to form a conception of the world in which he was contained, which would account for the validity of that conception and for his capacity to arrive at it. But his method was supposed to depend only on propositions and steps that were absolutely certain, and the method of self-transcendence as I have described it does not necessarily have this feature. In fact, such a conception of the world need not be developed by proofs at all, though it must rely heavily on a *priori* conjecture.²

In discussing the nature of the process and its pitfalls, I want both to defend the possibility of objective ascent and to understand its limits. We should keep in mind how incredible it is that such a thing is possible at all. We are encouraged these days to think of ourselves as contingent organisms arbitrarily thrown up by evolution. There is no reason in advance to expect a finite creature like that to be able to do more than accumulate information at the perceptual and conceptual level it occupies by nature. But apparently that is not how things are. Not only can we form the pure idea of a world that contains us and of which our impressions are a part, but we can give this idea a content which takes us very far from our original impressions.

The pure idea of realism—the idea that there is a world in which we are contained—implies nothing specific about the relation between the appearances and reality, except that we and our inner lives are part of reality. The recognition that this is so creates pressure on the imagination to recast our picture of the world so that it is no longer the view from here. The two possible forms this can take, skepticism and objective

2. The idea is much closer to what Bernard Williams calls the absolute conception of reality, which is a more general description of Descartes' idea of knowledge. See Williams (7).

knowledge, are products of one capacity: the capacity to fill out the pure idea of realism with more or less definite conceptions of the world in which we are placed. The two are intimately bound together. The search for objective knowledge, because of its commitment to a realistic picture, is inescapably subject to skepticism and cannot refute it but must proceed under its shadow. Skepticism, in turn, is a problem only because of the realist claims of objectivity.

Skeptical possibilities are those according to which the world is completely different from how it appears to us, and there is no way to detect this. The most familiar from the literature are those in which error is the product of deliberate deception by an evil demon working on the mind, or by a scientist stimulating our brain *in vitro* to produce hallucinations. Another is the possibility that we are dreaming. In the latter two examples the world is not totally different from what we think, for it contains brains and perhaps persons who sleep, dream, and hallucinate. But this is not essential: we can conceive of the possibility that the world is different from how we believe it to be in ways that we cannot imagine, that our thoughts and impressions are produced in ways that we cannot conceive, and that there is no way of moving from where we are to beliefs about the world that are substantially correct. This is the most abstract form of skeptical possibility, and it remains an option on a realist view no matter what other hypotheses we may construct and embrace.

2. *Antiskepticism*

Not everyone would concede either this skepticism or the realism on which it depends. Recently there has been a revival of arguments against the possibility of skepticism, reminiscent of the ordinary language arguments of the fifties which claimed that the meanings of statements about the world are revealed by the circumstances in which they are typically used, so that it couldn't be the case that most of what we ordinarily take to be true about the world is in fact false.

In their current versions these arguments are put in terms of reference rather than meaning.³ What we refer to by the terms in our statements about the external world, for example—what we are really talking about—is said to be whatever *actually* bears the appropriate relation to the generally accepted use of those terms in our language. (This relation is left undefined, but it is supposed to be exemplified in the ordinary world by the relation between my use of the word 'tree' and actual trees, if there are such things.)

3. See for example Putnam (2), ch. 1.

The argument against the possibility of skepticism is a *reductio*. Suppose that I am a brain in a vat being stimulated by a mischievous scientist to think I have seen trees, though I never have. Then my word "tree" refers not to what we now call trees but to whatever the scientist usually uses to produce the stimulus which causes me to think, "There's a tree." So when I think that, I am usually thinking something true. I cannot use the word "tree" to form the thought that the scientist would express by saying I have never seen a tree, or the words "material object" to form the thought that perhaps I have never seen a material object, or the word "vat" to form the thought that perhaps I am a brain in a vat. If I were a brain in a vat, then my word "vat" would not refer to vats, and my thought, "Perhaps I am a brain in a vat," would not be true. The original skeptical supposition is shown to be impossible by the fact that if it were true, it would be false. The conditions of reference permit us to think that there are no trees, or that we are brains in a vat, only if this is not true.

This argument is no better than its predecessors. First, I can use a term which fails to refer, provided I have a conception of the conditions under which it would refer—as when I say there are no ghosts. To show that I couldn't think there were no trees if there were none, it would have to be shown that this thought could not be accounted for in more basic terms which would be available to me even if all my impressions of trees had been artificially produced. (Such an analysis need not describe my *conscious* thoughts about trees.) The same goes for "physical object". The skeptic may not be able to produce on request an account of these terms which is independent of the existence of their referents, but he is not refuted unless reason has been given to believe such an account impossible. This has not been attempted and seems on the face of it a hopeless enterprise.

A skeptic does not hold that all his terms fail to refer; he assumes, like the rest of us, that those that do not refer can be explained at some level in terms of those that do. When he says, "Perhaps I have never seen a physical object," he doesn't mean (holding up his hand), "Perhaps *this*, whatever it is, doesn't exist!" He means, "Perhaps I have never seen anything with the spatiotemporal and mind-independent characteristics necessary to be a physical object—nothing of the kind that I take physical objects to be." It has to be shown that he couldn't have *that* thought if it were true. Clearly we will be pushed back to the conditions for the possession of very general concepts. Nothing here is obvious, but it seems clear at least that a few undeveloped assumptions about reference will not enable one to prove that a brain in a vat or a disembodied spirit couldn't have the concept of mind-independence, for example. The main issue simply hasn't been addressed.

Second, although the argument doesn't work it wouldn't refute skepticism if it did. If I accept the argument, I must conclude that a brain in a vat can't think truly that it is a brain in a vat, even though others can think this about it. What follows? Only that I can't express my skepticism by saying, "Perhaps I'm a brain in a vat." Instead I must say, "Perhaps I can't even *think* the truth about what I am, because I lack the necessary concepts and my circumstances make it impossible for me to acquire them!" If this doesn't qualify as skepticism, I don't know what does.

The possibility of skepticism is built into our ordinary thoughts, in virtue of the realism that they automatically assume and their pretensions to go beyond experience. Some of what we believe must be true in order for us to be able to think at all, but this does not mean we couldn't be wrong about vast tracts of it. Thought and language have to latch onto the world, but they don't have to latch onto it directly at every point, and a being in one of the skeptic's nightmare situations should be able to latch onto enough of it to meet the conditions for formulating his questions.⁴

Critics of skepticism bring against it various theories of how the language works—theories of verifiability, causal theories of reference, principles of charity. I believe the argument goes in the opposite direction.⁵ Such theories are refuted by the evident possibility and intelligibility of skepticism, which reveals that by "tree" I don't mean just anything that is causally responsible for my impressions of trees, or anything that looks and feels like a tree, or even anything of the sort that I and others have traditionally called trees. Since those things could conceivably not be trees, any theory that says they have to be is wrong.

The traditional skeptical possibilities that we can imagine stand for limitless possibilities that we can't imagine. In recognizing them we recognize that our ideas of the world, however sophisticated, are the products of one piece of the world interacting with part of the rest of it in ways that we do not understand very well. So anything we come to believe must remain suspended in a great cavern of skeptical darkness.

4. There is perhaps one form of radical skepticism which could be ruled out as unthinkable, by an argument analogous to the *cogito*: skepticism about whether I am the kind of being who can have thoughts *at all*. If there were possible beings whose nature and relation to the world was such that nothing they did could constitute thinking, whatever went on inside them, then I could not wonder whether I was such a being, because if I were, I wouldn't be thinking, and even to consider the possibility that I may not be thinking is to think. But most forms of skepticism are not this extreme.

5. This is a theme of Clarke's and Stroud's work on skepticism. See Stroud, pp. 205–6. Stroud's book is a highly illuminating discussion of skepticism and the inadequacy of most responses to it. He is nevertheless slightly more optimistic than I am about the possibility of finding something wrong with skepticism—and with the desire for an objective or external understanding of our position in the world that leads to it.

Once the door is open, it can't be shut again. We can only try to make our conception of our place in the world more complete—essentially developing the objective standpoint. The limit to which such development must tend is presumably unreachable: a conception that closes over itself completely, by describing a world that contains a being that has precisely that conception, and explaining how the being was able to reach that conception from its starting point within the world. Even if we did arrive at such a self-transcendent idea, that wouldn't guarantee its correctness. It would recommend itself as a possibility, but the skeptical possibilities would also remain open. The best we can do is to construct a picture that might be correct. Skepticism is really a way of recognizing our situation, though it will not prevent us from continuing to pursue something like knowledge, for our natural realism makes it impossible for us to be content with a purely subjective view.

3. Self-transcendence

To provide an alternative to the imaginable and unimaginable skeptical possibilities, a self-transcendent conception should ideally explain the following four things: (1) what the world is like; (2) what we are like; (3) why the world appears to beings like us in certain respects as it is and in certain respects as it isn't; (4) how beings like us can arrive at such a conception. In practice, the last condition is rarely met. We tend to use our rational capacities to construct theories, without at the same time constructing epistemological accounts of how those capacities work. Nevertheless, this is an important part of objectivity. What we want is to reach a position as independent as possible of who we are and where we started, but a position that can also explain how we got there.

In a sense, these conditions could also be satisfied by a conception of the world and our place in it that was developed by other beings, different from us; but in that case the fourth element would not involve self-referential understanding, as it does in the understanding of ourselves. The closest we can come to an external understanding of our relation to the world is to develop the self-referential analogue of an external understanding. This leaves us in no worse position than an external observer, for any being who viewed us from outside would have to face the problem of self-understanding in its own case, to be reasonably secure in its pretensions to understand us or anything else. The aim of objectivity would be to reach a conception of the world, including oneself, which involved one's own point of view not essentially, but only instrumentally, so to speak: so that the form of our understanding would be specific to ourselves, but its content would not be.

The vast majority of additions to what we know do not require any advance in objectivity: they merely add further information at a level that already exists. When someone discovers a previously undetected planet, or the chemical composition of a hormone, or the cause of a disease, or the early influences on a historical figure, he is essentially filling in a framework of understanding that is already given. Even something as fruitful as the discovery of the structure of DNA is in this category. It merely extended the methods of chemistry into genetics. Discoveries like this may be difficult to make, but they do not involve fundamental alterations in the idea of our epistemic relation to the world. They add knowledge without objective advance.

An advance in objectivity requires that already existing forms of understanding should themselves become the object of a new form of understanding, which also takes in the objects of the original forms. This is true of any objective step, even if it does not reach the more ambitious goal of explaining itself. All advances in objectivity subsume our former understanding under a new account of our mental relation to the world.

Consider for example the distinction between primary and secondary qualities, the precondition for the development of modern physics and chemistry. This is a particularly clear example of how we can place ourselves in a new world picture. We realize that our perceptions of external objects depend both on their properties and on ours, and that to explain both their effects on us and their interactions with each other we need to attribute to them fewer types of properties than they may initially appear to have.

Colin McGinn has argued convincingly that this is in the first instance an a priori philosophical discovery, not an empirical scientific one. Things have colors, tastes, and smells in virtue of the way they appear to us: to be red simply *is* to be the sort of thing that looks or would look red to normal human observers in the perceptual circumstances that normally obtain in the actual world. To be square, on the other hand, is an independent property which can be used to explain many things about an object, including how it looks and feels. (McGinn, ch. 7)

Once this is recognized and we consider how the various perceptible properties of objects are to be explained, it becomes clear that the best account of the appearance of colors will not involve the ascription to things of intrinsic color properties that play an ineliminable role in the explanation of the appearances: the way in which the appearances vary with both physical and psychological conditions makes this very implausible. Objective shape and size, on the other hand, enter naturally into an account of variable appearance of shape and size. So much is evident

even if we have only a very rough idea of how as perceivers we are acted upon by the external world—an idea having to do primarily with the type of peripheral impact involved. It is then a short step to the conjecture that the appearances of secondary qualities are caused by other primary qualities of objects, which we can then try to discover.

The pressure to make an objective advance comes, here as elsewhere, from the incapacity of the earlier view of the world to include and explain itself—that is, to explain why things appear to us as they do. This makes us seek a new conception that can explain both the former appearances and the new impression that it itself is true. The hypothesis that objects have intrinsic colors in addition to their primary qualities would conspicuously fail this test, for it provides a poorer explanation of why they appear to have colors, and why those appearances change under internal and external circumstances, than the hypothesis that the primary qualities of objects and their effects on us are responsible for all the appearances.

Consider another example. Not all objective advances have been so widely internalized as this, and some, like general relativity and quantum mechanics, are advances beyond already advanced theories that are not generally accessible. But one huge step beyond common appearance was taken by Einstein with the special theory of relativity. He replaced the familiar idea of unqualified temporal and spatial relations between events, things, and processes by a relativistic conception according to which events are not without qualification simultaneous or successive, objects are not without qualification equal or unequal in size, but only with respect to a frame of reference. What formerly seemed to be an objective conception of absolute space and time was revealed to be a mere appearance, from the perspective of one frame of reference, of a world whose objective description from no frame of reference is not given in a four-dimensional coordinate system of independent spatial and temporal dimensions at all. Instead, events are objectively located in relativistic space-time, whose division into separate spatial and temporal dimensions depends on one's point of view. In this case it was reflection on electrodynamic phenomena rather than ordinary perception that revealed that the appearances had to be transcended. There was also, as with the primary-secondary quality distinction, an important philosophical element in the discovery that absolute simultaneity of spatially separated events was not a well-defined notion, in our ordinary system of concepts.

These examples illustrate the human capacity to escape the limits of the original human situation, not merely by traveling around and seeing

the world from different perspectives, but by ascending to new levels from which we can understand and criticize the general forms of previous perspectives. The step to a new perspective is the product of epistemological insight in each case.

Of course it is also the product in some cases of new observations that can't be accommodated in the old picture. But the satisfactoriness of a new external perspective depends on whether it can place the internal perspective within the world in a way that enables one to occupy both of them simultaneously, with a sense that the external perspective gives access to an objective reality that one's subjective impressions are impressions of. Experience is not the sole foundation of our knowledge of the world, but a place must be found for it as part of the world, however different that world may be from the way it is depicted in experience.

Only objectivity can give meaning to the idea of intellectual progress. We can see this by considering any well-established objective advance, like the examples discussed already, and asking whether it could be reversed. Could a theory which ascribed intrinsic colors, tastes, smells, feels, and sounds to things account for the appearance that these are to be explained as the effects on our senses of primary qualities? Could a theory of absolute space and time explain the appearance that we occupy relativistic space-time? In both cases the answer is no. An objective advance may be superseded by a further objective advance, which reduces it in turn to an appearance. But it is not on the same level as its predecessors, and may well have been essential as a step on the route to its successors.

Still, the fact that objective reality is our goal does not guarantee that our pursuit of it succeeds in being anything more than an exploration and reorganization of the insides of our own minds. On a realist view this always remains a possibility, at least in the abstract, even if one isn't thinking of a specific way in which one might be deceived. A less radical point is that whatever we may have achieved we are only at a passing stage of intellectual development, and much of what we now believe will be overthrown by later discoveries and later theories.

A certain expectation of further advance and occasional retreat is rational: there have been enough cases in which what was once thought a maximally objective conception of reality has been included as appearance in a still more objective conception so that we would be foolish not to expect it to go on. Indeed we should want it to go on, for we are evidently just at the beginning of our trip outward, and what has so far been achieved in the way of self-understanding is minimal.

4. Evolutionary Epistemology

Because self-understanding is at the heart of objectivity, the enterprise faces serious obstacles. The pursuit of objective knowledge requires a much more developed conception of the mind in the world than we now possess: a conception which will explain the possibility of objectivity. It requires that we come to understand the operations of our minds from a point of view that is not just our own. This would not be the kind of self-understanding that Kant aimed for, that is, an understanding from within of the forms and limits of all our possible experience and thought (though that would be amazing enough, and there is no reason to suppose that it could be arrived at a priori). What is needed is something even stronger: an explanation of the possibility of objective knowledge of the real world which is itself an instance of objective knowledge of that world and our relation to it. Can there be creatures capable of this sort of self-transcendence? We at least seem to have taken some steps in this direction, though it is not clear how far we can go. But how is even this much possible? In fact, the objective capacity is a complete mystery. While it obviously exists and we can use it, there is no credible explanation of it in terms of anything more basic, and so long as we don't understand it, its results will remain under a cloud.

Some may be tempted to offer or at least to imagine an evolutionary explanation, as is customary these days for everything under the sun. Evolutionary hand waving is an example of the tendency to take a theory which has been successful in one domain and apply it to anything else you can't understand—not even to apply it, but vaguely to imagine such an application. It is also an example of the pervasive and reductive naturalism of our culture. 'Survival value' is now invoked to account for everything from ethics to language. I realize that it is dangerous to enter into discussion of a topic on which one is not an expert, but since these speculations can't be ignored, and since even when they come from professional biologists they are in the nature of obiter dicta, let me try to say something about them.

The Darwinian theory of natural selection, assuming the truth of its historical claims about how organisms develop, is a very partial explanation of why we are as we are. It explains the selection among those organic possibilities that have been generated, but it does not explain the possibilities themselves. It is a diachronic theory which tries to account for the particular path evolution will take through a set of possibilities under given conditions. It may explain why creatures with vision or reason will survive, but it does not explain how vision or reasoning are possible.

These require not diachronic but timeless explanations. The range of biological options over which natural selection can operate is extraordinarily rich but also severely constrained. Even if randomness is a factor in determining which mutation will appear when (and the extent of the randomness is apparently in dispute), the range of genetic possibilities is not itself a random occurrence but a necessary consequence of the natural order. The possibility of minds capable of forming progressively more objective conceptions of reality is not something the theory of natural selection can attempt to explain, since it doesn't explain possibilities at all, but only selection among them.⁶

But even if we take as given the unexplained possibility of objective minds, natural selection doesn't offer a very plausible explanation of their actual existence. In themselves, the advanced intellectual capacities of human beings, unlike many of their anatomical, physiological, perceptual, and more basic cognitive features, are extremely poor candidates for evolutionary explanation, and would in fact be rendered highly suspect by such an explanation. I am not suggesting, as Kant once did (Kant (2), pp. 395–6), that reason has negative survival value and could from that point of view be replaced by instinct. But the capacity to form cosmological and subatomic theories takes us so far from the circumstances in which our ability to think would have had to pass its evolutionary tests that there would be no reason whatever, stemming from the theory of evolution, to rely on it in extension to those subjects. In fact if, per impossibile, we came to believe that our capacity for objective theory were the product of natural selection, that would warrant serious skepticism about its results beyond a very limited and familiar range. An evolutionary explanation of our theorizing faculty would provide absolutely no confirmation of its capacity to get at the truth. Something else must be going on if the process is really taking us toward a truer and more detached understanding of the world.

There is a standard reply to skepticism about evolutionary explanation of the intellect, namely that Darwinian theory doesn't require every feature of an organism to be separately selected for its adaptive value. Some features may be the side effects of others, singly or in combination, that have been so selected, and if they are not harmful they will survive. In the case of the intellect, a common speculation is that rapid enlargement of the human brain occurred through natural selection after the development of an erect posture and the capacity to use tools made brain size

6. Stephen Jay Gould reports that Francis Crick once said to him, "The trouble with you evolutionary biologists is that you are always asking 'why' before you understand 'how' " (Gould (2), p.10).

an advantage. This permitted the acquisition of language and the capacity to reason, which in turn conferred survival value on still larger brains. Then, like an adaptable computer, this complex brain turned out to be able to do all kinds of things it wasn't specifically "selected" to do: study astronomy, compose poetry and music, invent the internal combustion engine and the long-playing record, and prove Gödel's theorem. The great rapidity of civilized cultural evolution requires that the brains which took part in it have been developed to full capacity from its beginning.

Since this is pure speculation, not much can be said about its consistency with the empirical evidence. We know nothing about how the brain performs the functions that permitted our hunter-gatherer ancestors to survive, nor do we know anything about how it performs the functions that have permitted the development and understanding of the mathematics and physics of the past few centuries. So we have no basis for evaluating the suggestion that the properties which were necessary to fit the brain for the first of these purposes turned out to be sufficient for the second as well, and for all the cultural developments that have led to it.

Spinoza gives this description of the process of intellectual evolution:

As men at first made use of the instruments supplied by nature to accomplish very easy pieces of workmanship, laboriously and imperfectly, and then, when these were finished, wrought other things more difficult with less labour and greater perfection; and so gradually mounted from the simplest operations to the making of tools, and from the making of tools to the making of more complex tools, and fresh feats of workmanship, till they arrived at making, with small expenditure of labour, the vast number of complicated mechanisms which they now possess. So, in like manner, the intellect, by its native strength, makes for itself intellectual instruments, whereby it acquires strength for performing other intellectual operations, and from these operations gets again fresh instruments, or the power of pushing its investigations further, and thus gradually proceeds till it reaches the summit of wisdom. (Spinoza (1), p. 12)

The question is whether not only the physical but the mental capacity needed to make a stone axe automatically brings with it the capacity to take each of the steps that have led from there to the construction of the hydrogen bomb, or whether an enormous excess mental capacity, not explainable by natural selection, was responsible for the generation and spread of the sequence of intellectual instruments that has emerged over the last thirty thousand years. This question is unforgettably posed by

the stunning transformation of bone into spaceship in Stanley Kubrick's *2001*.

I see absolutely no reason to believe that the truth lies with the first alternative. The only reason so many people do believe it is that advanced intellectual capacities clearly exist, and this is the only available candidate for a Darwinian explanation of their existence. So it all rests on the assumption that every noteworthy characteristic of human beings, or of any other organism, must have a Darwinian explanation. But what is the reason to believe that? Even if natural selection explains all adaptive evolution, there may be developments in the history of species that are not specifically adaptive and can't be explained in terms of natural selection.⁷ Why not take the development of the human intellect as a probable counterexample to the law that natural selection explains everything, instead of forcing it under the law with improbable speculations unsupported by evidence? We have here one of those powerful reductionist dogmas which seem to be part of the intellectual atmosphere we breathe.

What, I will be asked, is my alternative? Creationism? The answer is that I don't have one, and I don't need one in order to reject all existing proposals as improbable. One should not assume that the truth about this matter has already been conceived of—or hold onto a view just because no one can come up with a better alternative. Belief isn't like action. One doesn't have to believe anything, and to believe nothing is not to believe something.

I don't know what an explanation might be like either of the possibility of objective theorizing or of the actual biological development of creatures capable of it. My sense is that it is antecedently so improbable that the only possible explanation must be that it is in some way necessary. It is not the kind of thing that could be either a brute fact or an accident, any more than the identity of inertial and gravitational mass could be; the universe must have fundamental properties that inevitably give rise through physical and biological evolution to complex organisms capable of generating theories about themselves and it. This is not itself an explanation; it merely expresses a view about one condition which an acceptable explanation should meet: it should show why this had to happen, given the relatively short time since the Big Bang, and not merely that it could have happened—as is attempted by Darwinian proposals. (I think an explanation of the original development of organic life should meet

7. See Gould (1) for details.

the same condition.) There is no reason to expect that we shall ever come up with such an explanation, but we are at such a primitive stage of biological understanding that there is no point in making any predictions.⁸

5. Rationalism

One image of self-reconstruction that has appealed to philosophers is Neurath's: that we are like sailors trying to rebuild our ship plank by plank on the high seas. This can be interpreted in more than one way. We might think of ourselves as simply rearranging and perhaps reshaping the planks, making small alterations one at a time, and using the materials we find ready to hand.⁹ Such an image may fit the mundane case where knowledge is accumulated gradually and piecemeal, at a given objective level. But if we wish to depict the great objective advances on which real progress depends, we need a different image. Though we may incorporate parts of the original ship in the new one we are about to create, we call up out of ourselves most of the materials from which we will construct it. The place which we occupy for this purpose may be one we could not have reached except on the old ship, but it is really in a new world, and in some sense, I believe, what we find in it is already there. Each of us is a microcosm, and in detaching progressively from our point of view and forming a succession of higher views of ourselves

8. It might be argued that the observation that the universe contains intelligent beings does not have to be explained in terms of fundamental principles which show it to be inevitable, because it has a much simpler explanation: if there were no such beings, there would be no observers and hence no observations. No general inferences can therefore be drawn from their existence. I am not persuaded by this argument. The fact that an observation can be predicted on this sort of ground does not mean that it needn't be explained by other, more fundamental principles as well.

It may be worth mentioning an analogy, the application of the anthropic principle in cosmology. The anthropic principle states that "what we can expect to observe must be restricted by the conditions necessary for our presence as observers" (Carter, p. 291). A special case of this is the strong anthropic principle: "the Universe (and hence the fundamental parameters on which it depends) must be such as to admit the creation of observers within it at some stage" (p. 294). About this Carter says that "even an entirely rigorous prediction based on the *strong* principle will not be completely satisfying from the physicist's point of view since the possibility will remain of finding a deeper underlying theory explaining the relationships that have been predicted" (p. 295). In other words, predictability does not always eliminate the need for explanation.

9. As Neurath puts it, we are "never able to dismantle it in dry-dock and to reconstruct it there out of the best materials" (Neurath, p. 201).

in the world, we are occupying a territory that already exists: taking possession of a latent objective realm, so to speak.

I said earlier that the position to which I am drawn is a form of rationalism. This does not mean that we have innate knowledge of the truth about the world, but it does mean that we have the capacity, not based on experience, to generate hypotheses about what in general the world might possibly be like, and to reject those possibilities that we see could not include ourselves and our experiences. Just as important, we must be able to reject hypotheses which appear initially to be possibilities but are not. The conditions of objectivity that I have been defending lead to the conclusion that the basis of most real knowledge must be a priori and drawn from within ourselves. The role played by particular experience and by the action of the world on us through our individual perspectives can be only selective—though this is a very important factor, which makes the acquisition of such knowledge as we may have importantly subject to luck: the luck of the observations and data to which we are exposed and the age in which we live. Also important, for possession of the a priori component, are the possibilities and questions that are suggested to us and that we might not have formulated for ourselves—like the boy in Plato's *Meno*.

If the possibilities, or at least some of them, are available a priori to any mind of sufficient complexity, and if the general properties of reality are fairly uniform throughout, then the pursuit of objective knowledge can be expected to lead to gradual convergence from different starting points. But this limit of convergence is not the definition of truth, as Peirce suggests: it is a consequence of the relation between reality and the mind, which in turn must be explained in terms of the kind of part of reality the mind is. Obviously the capacities of different minds, and of different species of mind, differ. But in our case the capacities go far beyond the merely adaptive. A reasonably intelligent human being is capable of grasping, even if it cannot generate on its own, an extraordinary and rich range of conceptual possibilities, as we know from what has been learned already. There is no reason to think our mental capacities mirror reality completely, but I assume we all carry potentially in our heads the possibilities that will be revealed by scientific and other developments over the next few thousand years at least: we just aren't going to be around for the trip—perhaps it should be called the awakening.

This conception of knowledge is in the rationalist tradition, though without the claim that reason provides an indubitable foundation for belief. Even empirical knowledge, or empirical belief, must rest on an a

priori base, and if large conclusions are derived from limited empirical evidence a large burden must be carried by direct a priori formulation and selection of hypotheses if knowledge is to be possible at all.¹⁰

This accounts for the extremely high ratio of rational to empirical grounds for great theoretical advances like Newton's theory of gravitation or the special and general theories of relativity: even though the empirical predictions of those theories are enormous, they were arrived at on the basis of relatively limited observational data, from which they could not be deduced. And I would maintain that even induction, that staple of empiricism, makes sense only with a rationalist basis. Observed regularities provide reason to believe that they will be repeated only to the extent that they provide evidence of hidden *necessary* connections, which hold timelessly. It is not a matter of assuming that the contingent future will be like the contingent past.

The capacity to imagine new forms of hidden order, and to understand new conceptions created by others, seems to be innate. Just as matter can be arranged to embody a conscious, thinking organism, so some of these organisms can rearrange themselves to embody more and more thorough and objective mental representations of the world that contains them, and this possibility too must exist in advance. Although the procedures of thought by which we progress are not self-guaranteeing, they make sense only if we have a natural capacity for achieving harmony with the world far beyond the range of our particular experiences and surroundings. When we use our minds to think about reality, we are not, I assume, performing an impossible leap from inside ourselves to the world outside. We are developing a relation to the world that is implicit in our mental and physical makeup, and we can do this only if there are facts we do not know which account for the possibility. Our position is problematic so long as we have not even a candidate for such an account.

Descartes tried to provide one, together with grounds for certainty that it was true, by proving the existence of the right sort of God. While he was not successful, the problem remains. To go on unambivalently

10. Both Chomsky and Popper have in very different ways rejected empiricist theories of knowledge and emphasized the incomprehensibility, at present, of our capacities to understand and think about the world. Chomsky in particular has argued that our innate capacity to learn languages is contrary to the empiricist conception of how the mind works. This is one aspect of his general attack on reductionism with respect to the mind. I believe that the scientific gaps between data and conclusions are of much greater importance to the theory of knowledge than the gap between the fragmentary linguistic data of early childhood and the grammar of the language that is learned from it, remarkable as that is. Somehow we call up whole worlds out of our heads, not just languages whose form has presumably evolved in part to suit our ability to learn them.

holding our beliefs once this has been recognized requires that we believe that something—we know not what—is true that plays the role in our relation to the world that Descartes thought was played by God. (Perhaps it would be more accurate to say that Descartes' God is a personification of the fit between ourselves and the world for which we have no explanation but which is necessary for thought to yield knowledge.)

I have no idea what unheard-of property of the natural order this might be. But without something fairly remarkable, human knowledge is unintelligible. My view is rationalist and antiempiricist, not because I believe a firm foundation for our beliefs can be discovered *a priori*, but because I believe that unless we suppose that they have a basis in something global (rather than just human) of which we are not aware, they make no sense—and they do make sense. A serious rationalist epistemology would have to complete this picture—but our beliefs may rest on such a basis even if we cannot discover it. There is no reason to assume that even if we are so organized as to be capable of partly understanding the world, we can also gain access to these facts about ourselves in a way that will fill the blanks in our understanding.¹¹

A theory of reality with pretensions to completeness would have to include a theory of the mind. But this too would be a hypothesis generated by the mind, and would not be self-guaranteeing. The point is made by Stroud with reference to Quine's proposal of naturalized epistemology, which is essentially an empiricist psychological theory of the formation of empirical theories (Stroud, ch. 6). It applies equally to a possible rationalist theory of the mind's capacity for *a priori* theorizing. But of course we have neither of these theories: we don't even have a hypothesis about our capacity to transcend the phenomena. The idea of a full conception of reality that explains our ability to arrive at it is just a dream.

Nevertheless, it's what we aim toward: a gradual liberation of the dormant objective self, trapped initially behind an individual perspective of

11. It may be that those areas of knowledge that are entirely *a priori* permit greater access to their sources in us than do other types of knowledge—that we can develop a better understanding of how our thoughts can lead us to the truths of arithmetic than of how they can lead us to the truths of chemistry. It is possible to make discoveries about something *a priori* if our representation of the thing has so intimate a relation with the thing itself that the properties to be discovered are already buried in the representation. Thus we can think about mathematics because we are able to operate with a system of symbols whose formal properties make it capable of representing the numbers and all their relations. This system can itself be mathematically investigated. To that extent mathematics gives us a partial answer to the question of how the world that it describes can contain beings who will be able to arrive at some of its truths.

human experience. The hope is to develop a detached perspective that can coexist with and comprehend the individual one.

6. Double Vision

To summarize, what we can hope to accomplish along these lines is bound to be limited in several ways. First, we are finite beings, and even if each of us possesses a large dormant capacity for objective self-transcendence, our knowledge of the world will always be fragmentary, however much we extend it. Second, since the objective self, though it can escape the human perspective, is still as short-lived as we are, we must assume that its best efforts will soon be superseded. Third, the understanding of the world of which we are intrinsically capable—leaving aside limitations of time and technology—is also likely to be limited. As I shall argue in the next chapter, reality probably extends beyond what we can conceive of. Finally, the development of richer and more powerful objective hypotheses does nothing to rule out the known and unknown skeptical possibilities which are the other aspect of any realist view.

None of this will deter us from the effort to make objective progress so far as our minds, our culture, and our epoch may permit. But there are other dangers in the pursuit of that goal, dangers not of failure but of ambition. These dangers are of three kinds: excessive impersonality, false objectification, and insoluble conflict between subjective and objective conceptions of the same thing.

The first comes from taking too literally the image of the true self trapped in the individual human perspective. This is a compelling image, and many have succumbed to its attractions. If the real me views the world from nowhere, and includes the empirical perspective and particular concerns of TN as merely one of myriad sentient flickers in the world so viewed, then it may seem that I should take as little interest in TN's life and perspective as possible, and perhaps even try to insulate myself from it. But the discovery and awakening of the objective self with its universal character doesn't imply that one is not also a creature with an empirical perspective and individual life. Objective advance produces a split in the self, and as it gradually widens, the problems of integration between the two standpoints become severe, particularly in regard to ethics and personal life. One must arrange somehow to see the world both from nowhere and from here, and to live accordingly.

The second danger, that of false objectification, is one I have already discussed in connection with the philosophy of mind—though it arises

also in other areas. The success of a particular form of objectivity in expanding our grasp of some aspects of reality may tempt us to apply the same methods in areas where they will not work, either because those areas require a new kind of objectivity or because they are in some respect irreducibly subjective. The failure to recognize these limits produces various kinds of objective obstinacy—most notably reductive analyses of one type of thing in terms that are taken from the objective understanding of another. But as I have said, reality is not just objective reality, and objectivity itself is not one thing. The kinds of objective concepts and theories that we have developed so far, mostly to understand the physical world, can be expected to yield only a fragment of the objective understanding that is possible. And the detachment that objectivity requires is bound to leave something behind.

The third problem, that of insoluble subjective-objective conflict, arises when we succeed in constructing an objective conception of something and then don't know what to do with it because it can't be harmoniously combined with the subjective conception we still have of the same thing. Sometimes an internal conception can't acknowledge its own subjectivity and survive, nor can it simply disappear.

Often an objective advance will involve the recognition that some aspects of our previous understanding belong to the realm of appearance. Instead of conceiving the world as full of colored objects, we conceive it as full of objects with primary qualities that affect human vision in certain subjectively understandable ways. The distinction between appearance and objective reality becomes the object of a new, mixed understanding that combines subjective and objective elements and that is based on recognition of the limits of objectivity. Here there is no conflict.¹²

But it may happen that the object of understanding cannot be so cleanly divided. It may happen that something appears to require subjective and objective conceptions that cover the same territory, and that cannot be combined into a single complex but consistent view. This is particularly likely with respect to our understanding of ourselves, and it is at the source of some of the most difficult problems of philosophy, including the problems of personal identity, free will, and the meaning of life. It is also present in the theory of knowledge, where it takes the form of an inability to hold in one's mind simultaneously and in a consistent form the possibility of skepticism and the ordinary beliefs that life is full of.

12. This is McGinn's point; the scientific image doesn't on reflection conflict with the manifest image over secondary qualities.

What should be the relation between the beliefs we form about the world, with their aspirations to objectivity, and the admission that the world might be completely different from the way we think it is, in unimaginable ways? I believe we have no satisfactory way of combining these outlooks. The objective standpoint here produces a split in the self which will not go away, and we either alternate between views or develop a form of double vision.

Double vision is the fate of creatures with a glimpse of the view *sub specie aeternitatis*. When we view ourselves from outside, a naturalistic picture of how we work seems unavoidable. It is clear that our beliefs arise from certain dispositions and experiences which, so far as we know, don't guarantee their truth and are compatible with radical error. The trouble is that we can't fully take on the skepticism that this entails, because we can't cure our appetite for belief, and we can't take on this attitude toward our own beliefs while we're having them. Beliefs are about how things probably are, not just about how they might possibly be, and there is no way of bracketing our ordinary beliefs about the world so that they dovetail neatly with the possibility of skepticism. The thought "I'm a professor at New York University, unless of course I'm a brain in a vat" is not one that can represent my general integrated state of mind.¹³

The problems of free will and personal identity yield similarly unharmonious conclusions. In some respects what we do and what happens to us fits very naturally into an objective picture of the world, on a footing with what other objects and organisms do. Our actions seem to be events with causes and conditions many of which are not our actions. We seem to persist and change through time much as other complex organisms do. But when we take these objective ideas seriously, they appear to threaten and undermine certain fundamental self-conceptions that we find it very difficult to give up.

Earlier I said it was impossible fully to internalize a conception of one's own personal identity that depended on the organic continuity of one's brain. Ordinarily, an objective view of something with a subjective aspect does not require us simply to give up the subjective view, for it

13. There is a further problem. In the course of arriving at a skeptical conclusion, we pass through thoughts to which we do not simultaneously take up a skeptical stance—thoughts about the relation of the brain to experience, for example. These appear in the skeptic's reasoning in unqualified form. In order to draw skeptical conclusions from the objective standpoint, we have to engage in the kind of direct thought about the world that skepticism undermines. This is like the Cartesian circle in reverse: Descartes tried to prove the existence of God by the use of reasoning on which we can rely only if God exists; the skeptic reaches skepticism through thoughts that skepticism makes unthinkable.

can be reduced to the status of an appearance, and can then coexist with the objective view. But in these cases that option seems not to be available. We cannot come to regard our ideas of our own agency or of the purity of our self-identity through time as mere appearances or impressions. That would be equivalent to giving them up. Though our intuitive convictions about these things emerge very much from our own point of view, they have pretensions to describe not just how we appear to ourselves but how we are, in some as yet unspecified sense which appears to conflict with the objective picture of what we are. This problem arises even if the objective picture does not claim to take in everything—for what it willingly omits is only subjective appearance, and that is not good enough. The claims of both the objective and the subjective self seem to be too strong to allow them to live together in harmony.

This problem will reappear in later chapters, but let me mention one further example: Wittgenstein's unacknowledged skepticism about deduction. I believe his view is rightly regarded by Kripke as a form of skepticism because the external account it gives of what is really going on when we apply a formula or a concept to indefinitely many cases—what the apparently infinite reach of meaning really rests on—is not an account we can take on internally. For example we can't think of the correct application of 'plus 2' as being determined by nothing more than the fact that a certain application is natural to those who share our language and form of life, or by anything else of the sort. In employing the concept we must think of it as determining a unique function over infinitely many cases, beyond all our applications and those of our community and independent of them, or else it would not be the concept it is. *Even if Wittgenstein is right*, we can't think of our thoughts this way while we have them. And even in the philosophical act of thinking naturalistically about how language and logic work, we can't take the Wittgensteinian stance toward *those* thoughts, but must think them straight.

I think a view deserves to be called skeptical if it offers an account of ordinary thoughts which cannot be incorporated into those thoughts without destroying them. One may be a skeptic about *x* no matter how sincerely one protests that one is not denying the existence of *x* but merely explaining what *x* really amounts to.¹⁴

14. See Kripke (2), p. 65.

THOUGHT AND REALITY

1. Realism

I have at various points expressed commitment to a form of realism, and must now say more about it. In simple terms it is the view that the world is independent of our minds, but the problem is to explain this claim in a nontrivial way which cannot be easily admitted by everyone, and thereby to show how it conflicts with a form of idealism that is held by many contemporary philosophers.

Realism makes skepticism intelligible. In the last chapter we discussed skepticism with regard to knowledge. Here I want to introduce another form of skepticism—not about what we know but about how far our thoughts can reach. I shall defend a form of realism according to which our grasp on the world is limited not only in respect of what we can know but also in respect of what we can conceive. In a very strong sense, the world extends beyond the reach of our minds.¹

The idealism to which this is opposed holds that what there is is what we can think about or conceive of, or what we or our descendants could come to be able to think about or conceive of—and that this is necessarily true because the idea of something that we could not think about or conceive of makes no sense. The “we” is important. I leave aside

1. This thesis is also defended by Fodor (pp. 120–6).

views, also called idealist, that hold reality to be correlative with mind in a much wider sense—including infinite minds, if there are such things. Perhaps, given any type of world, there could be a mind capable of conceiving it adequately; I have no idea of the limits on possible minds. The realism I am defending says the world may be inconceivable to *our* minds, and the idealism I am opposing says it could not be.

There are more radical forms of idealism than this, such as the view that to exist is to be perceived, or that what exists must be an object of possible experience for us, or that what exists or is the case must be an object of possible knowledge for us, or must be verifiable by us, or that it must be something about which we could have evidence. These views have all found adherents, but I believe that they depend finally on the more general form of idealism I have described, together with different specific views about the conditions of human thought. The common element is a broadly epistemological test of reality—which has never lost its popularity despite the supposed death of logical positivism.

I want to oppose to this general position a different view of the relation of our thoughts to reality, particularly those thoughts that attempt to represent the world objectively. In pursuing objectivity we alter our relation to the world, increasing the correctness of certain of our representations of it by compensating for the peculiarities of our point of view. But the world is in a strong sense independent of our possible representations, and may well extend beyond them. This has implications both for what objectivity achieves when it is successful and for the possible limits of what it can achieve. Its aim and sole rationale is to increase our grasp of reality, but this makes no sense unless the idea of reality is not merely the idea of what can be grasped by those methods. In other words, I want to resist the natural tendency to identify the idea of the world as it really is with the idea of what can be revealed, at the limit, by an indefinite increase in objectivity of standpoint.

It has already been argued that in various respects the pursuit of objectivity can be carried to excess, that it can lead away from the truth if carried out in the wrong way or with respect to the wrong subject matter. That is one way in which objectivity does not correspond to reality: it is not always the best mode of understanding. But human objectivity may fail to exhaust reality for another reason: there may be aspects of reality beyond its reach because they are altogether beyond our capacity to form conceptions of the world. What there is and what we, in virtue of our nature, can think about are different things, and the latter may be smaller than the former. Certainly *we* are smaller, so this should not be surprising. Human objectivity may be able to grasp only part of the world, but when it is successful it should provide us with an understand-

ing of aspects of reality whose existence is completely independent of our capacity to think about them—as independent as the existence of things we can't conceive.

The idea that the contents of the universe are limited by our capacity for thought is easily recognized as a philosophical view, which at first sight seems crazily self-important given what small and contingent pieces of the universe we are. It is a view that no one would hold except for philosophical reasons that seem to rule out the natural picture.

That picture is that the universe and most of what goes on in it are completely independent of our thoughts, but that since our ancestors appeared on Earth we have gradually developed the capacity to think about, know about, and represent more and more aspects of reality. There are some things that we cannot now conceive but may yet come to understand; and there are probably still others that we lack the capacity to conceive not merely because we are at too early a stage of historical development, but because of the kind of beings we are. About some of what we cannot conceive we are able to speak vaguely—this may include the mental lives of alien creatures, or what went on before the Big Bang—but about some of it we may be unable to say anything at all, except that there might be such things. The only sense in which we can conceive of them is under that description—that is, as things of which we can form no conception—or under the all-encompassing thought “Everything,” or the Parmenidean thought “What is.”

I am not claiming that much of what we find *positively inconceivable*—what we can see to be *impossible*, like round squares—may nevertheless be possible. Though there may be cases where our strong convictions of positive inconceivability cannot be relied on as evidence of impossibility, I assume they are rare. I am concerned here rather with the admission of possibilities and actualities which are *negatively* inconceivable to us in the sense that we have and can have no conception of them. (This is different from seeing positively that no such conception could be coherent, because, for example, it would involve a contradiction.)

Not everything about the universe must lie in the path of our possible cognitive development or that of our descendants—even if beings like us should exist forever. It is a philosophical problem how we are able to think about those aspects of reality that we can think about. There is also the question of whether we can think about those things ‘as they are in themselves’ or only ‘as they appear to us’. But what there is, or what is the case, does not coincide necessarily with what is a possible object of thought for us. Even if through some miracle we are capable in principle of conceiving of everything there is, that is not what makes it real.

2. Idealism

The philosophical argument against this natural view is simple. It parallels one of Berkeley's arguments for the view that for unthinking things, to exist is to be perceived. Berkeley claimed that this became evident if we tried to form the idea of an unperceived object. It turns out to be impossible, he said, because as soon as we try to think, for example, of an unperceived tree, we find that all we can do is to call up a perceptual image of a tree, and that is not unperceived.²

It would be generally recognized now that this argument involves the mistake of confusing perceptual imagination as the vehicle of thought with a perceptual experience as part of the object of thought. Even if I employ a visual image to think about the tree, that does not mean I am thinking about a visual impression of the tree, any more than if I draw a tree, I am drawing a drawing of a tree (cf. Williams (1)).

A similar mistake would be to argue that we cannot form the thought of something that no one is actually thinking about or the conception of something that no one is conceiving of. Clearly we can think and talk about the possible state of affairs in which no one is thinking or talking about Bishop Berkeley. The fact that we must talk about Berkeley to talk about the situation in which he is not being talked about doesn't make that situation either inexpressible or impossible.

But the form of idealism with which I am concerned isn't based on this mistake: it is not the view that what there is must be actually conceived or even currently conceivable. Rather it is the position that what there is must be possibly conceivable by us, or possibly something for which we could have evidence. An argument for this general form of idealism must show that the notion of what *cannot* be thought about by us or those like us makes no sense.

The argument is this. If we try to make sense of the notion of what we could never conceive, we must use general ideas like that of something existing, or some circumstance obtaining, or something being the case, or something being true. We must suppose that there are aspects of reality to which these concepts that we *do* possess apply, but to which no other concepts that we *could* possess apply. To conceive simply that such things may exist is not to conceive of them adequately; and the realist would maintain that everything else about them might be inconceivable to us. The idealist reply is that our completely general ideas of what

2. Berkeley, secs. 22–3. This is not Berkeley's only argument for idealism, but he says he is "content to put the whole upon this issue."

exists, or is the case, or is true, cannot reach any further than our more specific ideas of kinds of things that can exist, or be the case, or be true. We do not, in other words, possess a completely general concept of reality that reaches beyond any possible filling in of its content that we could in principle understand.

Or to put the same point in terms of language, as Davidson does, we do not possess a general concept of truth that goes beyond the truth of all possible sentences in any language that we could understand, or that could be translated into a language that we or others like us could understand. Our general idea of what is the case does not transcend the sum of what we could truly assert to be the case. Here is Davidson rejecting the idea of a conceptual scheme that meets the conditions for applying to the world but is different from our own: "The criterion of a conceptual scheme different from our own now becomes: largely true but not translatable. The question whether this is a useful criterion is just the question how well we understand the notion of truth, as applied to a language, independent of the notion of translation. The answer is, I think, that we do not understand it independently at all."³

So the argument that parallels Berkeley's is that if we try to form the notion of something we could never conceive, or think about, or talk about, we find ourselves having to use ideas which imply that we could in principle think about it after all (even if we cannot do so now): because even the most general ideas of truth or existence we have carry that implication. We cannot use language to reach beyond the possible range of its specific application. If we attempt to do so, we are either misusing the language or using it to refer to what is conceivable after all.

This argument is not guilty of Berkeley's error. It does not ascribe to the object of thought something which is only an aspect of the vehicle of thought. It does not claim that to exist is to be thought about, or to have been or to be going to be thought about. Nevertheless, it works in a similar way, for it claims that certain attempts to form significant thoughts fail because they run up against boundaries set by the conditions of the possibility of thought. In Berkeley's argument the hypothesis of existence without the mind is said to conflict with the conditions of thought, and the same is said here about the hypothesis of inconceivability. Realists are deluded, in other words, if they think they have the idea of a reality beyond the reach of any possible human thought except

3. Davidson (3), p. 194; note that some realists might wish to accept the "largely true" condition for a different conceptual scheme, provided it could be shown not to conflict with the "not translatable" condition.

that one. If we examine carefully what they take to be that idea, we will discover that it is either the idea of something more fully within our reach, or no idea at all.

To answer this objection it is necessary to dispute the view of thought on which it depends—as with Berkeley's argument. But first let me try to make clear how paradoxical the conclusion of the idealist argument is. An examination of what is wrong with the conclusion may shed some light on what is wrong with the argument. I must say at the outset that I do not have an alternative theory of thought to offer in place of those that support idealism. My argument will be essentially negative. I believe that the statement of a realist position can be rejected as unintelligible only on grounds which would also require the abandonment of other, much less controversial claims. My position is that realism makes as much sense as many other unverifiable statements, even though all of them, and all thought, may present fundamental philosophical mysteries to which there is at present no solution.

It certainly seems that I can believe that reality extends beyond the reach of possible human thought, since this would be closely analogous to something which is not only possibly but actually the case. There are plenty of ordinary human beings who constitutionally lack the capacity to conceive of some of the things that others know about. People blind or deaf from birth cannot understand colors or sounds. People with a permanent mental age of nine cannot come to understand Maxwell's equations or the general theory of relativity or Gödel's theorem. These are all humans, but we could equally well imagine a species for whom these characteristics were normal, able to think and know about the world in certain respects, but not in all. Such people could have a language, and might be similar enough to us so that their language was translatable into part of ours.

If there could be people like that coexisting with us, there could also be such people if we did not exist—that is, if there were no one capable of conceiving of these things that they cannot understand. Then their position in the world would be analogous to the one which I have claimed we are probably in.

We can elaborate the analogy by imagining first that there are higher beings, related to us as we are related to the nine-year-olds, and capable of understanding aspects of the world that are beyond our comprehension. Then they would be able to say of us, as we can say of the others, that there are certain things about the world that we cannot even conceive. And now we need only imagine that the world is just the same, except that these higher beings do not exist. Then what they could say

if they did exist remains true. So it appears that the existence of unreachable aspects of reality is independent of their conceivability by any actual mind.

Does this analogy work? Or is there some asymmetry between our situation and that of the hypothetical nine-year-olds?

An objection might be that in thinking about them I have all along been conceiving of the world—even the world from which we are absent—in terms of what we actually know about it. The features they can't conceive are fully specifiable in our language. It might be held illegitimate to try to explain simply by analogy with this the idea of our own situation in a world that in some respects we can't think about. We cannot talk sense simply by saying that these people are related to the laws of general relativity as we are to some other features of the world—unless independent significance has been given in this statement to the general expression “features of the world.” How can an analogy give sense to something that has no sense on its own, apart from the analogy?

In response, let me extend the analogy to make room for this general concept, by an addition to the story about the congenital nine-year-olds. Suppose that, in the world in which we do not exist, one of them, call him Realist junior, develops philosophical leanings (why not?) and wonders whether there may be things about the world that he and others like him are incapable of ever finding out about or understanding. Is this impossible? That is, if he were to utter these words (supposing in other respects his language were like part of ours), would it be a mistake to take them as expressing a hypothesis which would in fact be true in that situation? Would he simply be talking nonsense without realizing it? Would he be incapable of thinking in general terms what we know to be true about his situation (what we have stipulated to be true)? Here the analogy goes in the other direction. If we would be talking nonsense by engaging in such speculation, so would he.

The question is whether we can attribute to Realist junior a general concept of reality which applies, though he can never know it, to the laws of general relativity and all the other features of the universe that humans may be capable of comprehending. Can he have, on the basis of the examples of reality with which he is acquainted, a general concept which applies beyond everything with which he and his like could conceivably become acquainted? If he can, then we ourselves can have the same concept, which will apply to features of the universe that we are incapable of comprehending.

Suppose Realist junior expands on the idea by speculating that there might be other beings, with capacities that the nine-year-olds lack, who could understand aspects of reality that are inaccessible to them, though

intellectual distance would make it impossible for the higher beings to communicate their understanding to the lower. (I assume an intellectual distance great enough so that the lower beings couldn't form a vicarious conception of something by relying to some extent on the judgments of the higher beings. Much of our ordinary conception of the world is vicarious in this sense—depending on the greater expertise or intelligence of other humans. But merely believing there might be something others could understand is not yet a conception of it, not even a vicarious one.)

Admittedly the idea of a higher form of understanding depends for its significance on the idea of something to be understood, and it is in dispute whether they have the latter idea. But it seems very artificial to deny that someone in this position could believe something we know to be not only significant but true: that there are concepts usable by other types of minds, which apply to the world and can be used to formulate truths about it, but which cannot be translated into his language or any language that he can understand. Wouldn't a nine-year-old Davidson who arose among them be wrong?

In fact, wouldn't Davidson himself have to say that this lesser Davidson was wrong in denying the intelligibility of realism? How would Davidson's principle of charity be applied to the dispute between Davidson junior and Realist junior? Wouldn't it imply that Realist junior was right, because what he said could be translated into something we would assert, whereas what Davidson junior said couldn't be translated into something we would assert? I'm not sure. The problem is that Davidson's notion of translation seems to be asymmetrical. I might be able to translate a sentence of someone else's language into a sentence of my language, even though he cannot translate my sentence into his. According to Davidson (so far as I can see) I could say that Realist junior was right and Davidson junior wrong, but Realist junior would be wrong to agree with me—as Davidson junior would no doubt point out to him. This doesn't make the doctrine any less paradoxical. And if these consequences are unacceptable with regard to the nine-year-olds, they are unacceptable with regard to us.

This issue recalls the one discussed in chapter 5, section 2, about the intelligibility of skepticism. In both cases the question is how far we can go in forming the idea of a world with which our minds cannot make contact. The general point I wish to make against restrictive theories of what is thinkable is this. Every concept that we have contains potentially the idea of its own complement—the idea of what the concept doesn't apply to. Unless it has been shown positively that there cannot be such things—that the idea involves some kind of contradiction (like the idea

of things that are not self-identical)—we are entitled to assume that it makes sense even if we can say nothing more about the members of the class, and have never met one.

To be the value of a variable in our universal or existential quantifications it is not necessary to be the referent of a specific name or description in our language, because we already have the general concept of *everything*, which includes both the things we can name or describe and those we can't. Against that background we can speak of the complement of any concept whatever, unless it has been shown to be positively inconceivable. We can speak of 'all the things we can't describe', 'all the things we can't imagine', 'all the things humans can't conceive of', and finally, 'all the things humans are constitutionally incapable of ever conceiving'. The universal quantifier does not have a built-in limitation to what can be designated in some other way. It could even be used to form the idea, 'all the things *no* finite mind could ever form a conception of'. Naturally the possibility of forming these ideas does not guarantee that anything corresponds to them. But in the nature of the case it is unlikely that we could ever have reason to believe that nothing does.

Creatures who recognize their limited nature and their containment in the world must recognize both that reality may extend beyond our conceptual reach and that there may be concepts that we could not understand. The condition is met by a general concept of reality under which one's actual conception, as well as all possible extensions of that conception, falls as an instance. This concept seems to me adequately explained through the idea of a hierarchical set of conceptions, extending from those much more limited than one's own but contained in it to those larger than one's own but containing it—of which some are reachable by discoveries one might make but others, larger still, are not. (The hierarchy could also include parallel conceptions, not intersecting our own but joined with it only in a larger one.) We could ascribe this concept to the philosophical nine-year-olds in our example, and I maintain that it is the same as our general concept of what there is. It seems to me so clear that we have this concept that any theory of thought or reference which implies that we can't have it is for that reason deeply suspect.

The pursuit of a more objective view, by which we place ourselves in the world and try to understand our relation to it, is the primary method of extending and filling out our particular conception of this reality. But the general concept implies that there is no guarantee that the whole of what there is coincides with what we or beings like us could arrive at if we carried the pursuit of objectivity to the limit—to the convergence of views that would come at that mythical point of stupefaction, "the end of inquiry."

There are limits of objectivity as a form of understanding that follow from the fact that it leaves the subjective behind. These are inner limits. There are also outer limits of objectivity that fall at different points for different types of beings, and that depend not on the nature of objectivity but on how far it can be pursued by a given individual. Objectivity is only a way of extending one's grasp of the world, and besides leaving certain aspects of reality behind, it may fail to reach others, even if more powerful forms of objectivity could encompass them.

3. Kant and Strawson

As with the topic of knowledge, so with the topic of thought, I believe there is a middle ground between skepticism and reductionism. In the case of knowledge, skepticism arises when we reflect that our beliefs inevitably claim to go beyond their grounds. The effort to avoid skepticism by eliminating this gap can lead to reductionism, a reinterpretation of the content of our beliefs in terms of their grounds.

With thought, the problem is the relation between our conceptions and what is possible. Thought purports to represent facts and possibilities beyond itself, and skepticism is the view that our thoughts themselves give us no way of telling whether they correspond enough to the nature of actual and possible reality to be able to make contact with it at all—even to the extent of permitting false beliefs about it. To escape from such skepticism the reductionist reinterprets the domain of possibilities as the domain of what is or could become conceivable to us—thus guaranteeing (finally) that our thoughts can make contact with it.

To explain the intermediate position let me try to locate it in relation to two opposed views: Kant's skepticism and Strawson's reductionism. (Needless to say, this is not a term Strawson would use to describe his own view.) They are related to one another in that Strawson's view is offered as an empiricist criticism of Kant. Since Strawson's reductionism is quite generous, admitting a great deal into the universe of possibilities, it is important to explain what I think it leaves out.

Kant's position is that we can conceive of things only as they appear to us and never as they are in themselves: how things are in themselves remains forever and entirely out of the reach of our thought. "Doubtless, indeed, there are intelligible entities corresponding to the sensible entities; there may also be intelligible entities to which our sensible faculty of intuition has no relation whatsoever; but our concepts of understanding, being mere forms of thought for our sensible intuition, could not in the least apply to them" (Kant (1), pp. B308–9).

Strawson wishes to remove the Kantian opposition between thinking

of things as they appear to us and thinking of them as they are in themselves by declaring the latter idea (in its Kantian version) nonsensical. He believes there is an appearance-reality distinction *within* what Kant regards as the world of appearance, but it is basically the distinction between how things appear to us at any particular time or from a particular vantage point, and how they would come to appear as the result of an improved view or further investigation. Application of the distinction depends, he says, on identity of reference plus a corrected view: the world that now appears to me in one way might come, as a result of procedures of corrective revision, to appear to me or others like me in another way—a way that could be seen as a correction of the first. No idea of reality is left standing by this account which could have application to anything outside the range of possible human conception, evidence, or discovery.

Strawson does not claim that the real is coextensive with what we can *actually* form a conception of. Here is what he says:

In rejecting the senseless dogma that our conceptual scheme corresponds at no point with Reality, we must not embrace the restrictive dogma that Reality is completely comprehended by that scheme as it actually is. We admit readily that there are facts we do not know. We must admit also that there may be *kinds* of fact of which we have, at present, no more conception than our human predecessors had of some kinds of fact admitted in our conceptual schemes but not in theirs. We learn not only how to answer old questions, but how to ask new ones. The idea of the aspects of Reality which would be described in the answers to the questions we do not yet know how to ask is one which, like the idea of the realm of the noumenal, though not in the same drastic style, limits the claim of actual human knowledge and experience to be "co-extensive with the real".

This seems to be the necessary, and not very advanced, limit of sympathy with the metaphysics of transcendental idealism.⁴

This is, as I have said, a broad interpretation of the range of possibilities. But the extension of reality beyond what we can now conceive is here still thought of as what *we* might come to if we developed new experiences or ways of thinking, and therefore conceptions of new types of individual, property, and relation that would have application on the basis of possible experience. The implied reference to ourselves and our world remains, even if the series of discoveries is thought of as incom-

4. Strawson (3), p.42. See also Strawson (4), where he explicitly endorses the position that scientific realism is acceptable so long as we recognize its relativity to a particular intellectual standpoint.

pletable. That is the basis for Strawson's claim that we cannot be said to know things only as they appear or might appear to us, because the contrasting idea of what never could appear to us is meaningless.

I want to agree with Strawson in denying that we know things *only* as they appear to us, but agree with Kant in holding that how things are in themselves transcends all possible appearances or human conceptions. Our knowledge of the phenomenal world is partial knowledge of the world as it is in itself; but the entire world can't be identified with the world as it appears to us because it probably includes things of which we cannot and never could conceive, no matter how far the human understanding is expanded, as Strawson suggests, in directions we cannot now imagine. The difficulty is to state this in a way that defies idealist reinterpretation; realists always find it hard to say anything with which idealists cannot arrange to agree by giving it their own meaning.

Let me first explain my disagreement with Kant.⁵ I hold the familiar view that secondary qualities describe the world as it appears to us but primary qualities do not. To be red simply is to be something which would appear red to us in normal conditions—it is a property whose definition is essentially relative. But to be square is not simply to be such as to appear square, even though what is square does appear square. Here the appearance of squareness is significantly explained in terms of the effect on us of squareness in objects, which is not in turn analyzed in terms of the appearance of squareness. The red appearance of red things, on the other hand, cannot be noncircularly explained in terms of their redness, because the latter is analyzed in terms of the former. To explain why things appear red we have to go outside the circle of color qualities.

The Kantian view that primary qualities, too, describe the world only as it appears to us depends on taking the entire system of scientific explanations of observable phenomena as itself an appearance, whose ultimate explanation cannot without circularity refer to primary qualities since they on the contrary have to be explained in terms of it. Primary qualities are nothing more on this view than an aspect of our world picture, and if that picture has an explanation, it must be in terms of the effect on us of something outside it, which will be for that reason unimaginable to us—the noumenal world.

The view that this must be how things are results from a refusal to distinguish between two ways in which the human point of view enters into our thoughts—as form and as content. The content of a thought

5. I am indebted here to Colin McGinn, though he does not explicitly consider the possibility of things we could never conceive of. See McGinn, chs. 6 and 7.

may be quite independent of its particular form—independent, for example, of the particular language in which it is expressed. All of our thoughts must have a form which makes them accessible from a human perspective. But that doesn't mean they are all about our point of view or the world's relation to it. What they are about depends not on their subjective form but on what has to be referred to in any explanation of what makes them true.

The content of some thoughts transcends every form they can take in the human mind. If primary qualities were conceivable only from a human point of view, then the ultimate explanation of why the world appears to us in that way could not refer to primary qualities of things in themselves. But if primary qualities can equally well be grasped from a point of view that has nothing subjectively in common with ours, then the description of the world in terms of them is not relative to our point of view; they are not merely aspects of the phenomenal world but can on the contrary be used, by us or by others, to explain the appearance of that world.

The question is whether every possible explanation of our awareness of primary qualities, either perceptual or theoretical, must refer to the primary qualities of things outside us, or whether they disappear from the final explanation by being entirely included, as appearances, in what is to be explained. Unless the latter is true, the analogy with secondary qualities is meaningless. The mere fact that any thoughts we have about primary qualities must be formulated in language and images that we understand does not settle the matter, nor does the fact that any explanation we accept will be ours. It has to be claimed that at the limit, whether or not we are able to reach it, primary qualities *drop out* of the explanation of their appearance.

But there is no reason to believe that reference to primary qualities will disappear from the explanation of the appearance of primary qualities, no matter how complicated we make the "appearance" to be explained. We can't explain the fact that things look spatially extended except in terms of their being extended. And we can't explain the fact that *that* explanation seems true except again in terms of things being extended, their extension affecting us perceptually in certain ways, and the existence of that relation affecting the results of our investigation into the causes of our perceptual impressions of extension. And so on. If each explanation of the appearance of spatial extension in terms of extension in the world is counted as a new, higher-level appearance of extension, then that too has to be explained in terms of extension in the world. However far up we escalate in the series of "appearances" of

extension, the extension of things in themselves will keep one step ahead and recur in the explanation of those appearances.⁶

The only thing that can refute this view is a better alternative. There might be one—at some level the explanation of the fact that so far all our theories of the physical world involve spatial extension might conceivably be explained in terms of something entirely different, something which we might or might not be able to grasp. But Kant's nonexplanation in terms of the inconceivable noumenal world is not that better alternative. It is just a placeholder for something beyond our comprehension, and there is no reason to accept it unless the available realist position, ascribing extension to things in themselves, is ruled out as impossible.

But what reason could there be to hold that, granting the intelligibility of the notion of things in themselves, they couldn't be spatially extended? There is no good reason—only a bad Berkeleyan reason: the move from subjectivity of form to subjectivity of content. In other words, it would have to be claimed that because *we have* the conception of primary qualities, detect them through observation, and use them in explanations, they are essentially relative to our point of view, though in a more complex way than secondary qualities are: relative not just to our perceptual point of view, but to our entire cognitive point of view. I believe there is no defense of this position which does not beg the question.

But it is necessary to mention one Kantian argument for doubting that things in themselves have primary qualities, even though it does not show this to be impossible. The argument is this. Suppose Kant is right to claim that the primary qualities are essential features of our world picture, so that we cannot conceive of a world without them. This doesn't imply that they can't also be features of the world as it is in itself. But it does mean that we could not understand any explanation of the appearance of primary qualities which did not involve the ascription of primary qualities to things in the world. And if any such alternative would be inconceivable to us, the fact that every explanation we come up with involves primary qualities is not very good evidence that no better explanation of some other kind exists.

That is true. On the other hand, whatever may be the limits of our world picture (whether or not it has to include primary qualities), there

6. By contrast, the colors of things drop out of explanations very early. We can explain something's looking red to *me* in terms of its being red, but we can't nonvacuously explain the fact that red things generally look red to human beings in terms of their being red.

is no guarantee that we will be able to find credible explanations of the appearances within those limits at all. To the extent that we do, there is some reason to think that the picture does describe things in themselves, as far as it goes. At any rate it may. Moreover, there is remarkable flexibility in our conception of the kinds of primary qualities there are—far beyond anything Kant would have imagined to be conceivable—and it has enabled us to formulate theories of the physical world further and further removed from immediate experience. If we agree with Kant that the idea of the world as it is in itself makes sense, then there is no reason to deny that we know anything about it.

Let me now turn to the other aspect of the view I am trying to defend—the claim that there are probably things about the world that we (humans) cannot conceive. This follows naturally from what has already been said. If our conception of primary qualities is a partial conception of things as they are in themselves, its existence is a side-effect of the existence of those things: it results from their effects on each other and on us, together with our mental activity. What we are able to understand about the world depends on the relation between us and the rest of it. It appears to us naturally in certain ways, and with the help of reason and controlled observation we can form hypotheses about the objective reality underlying those appearances. But how much of what there is we can in principle reach by these methods is contingent on our mental makeup and the hypotheses it enables us to understand. Our capacity for understanding what there is may be only partial, because where it exists it depends not only on how things are but on our constitution, and the former is independent of the latter. In this global picture we are contained in the world and able to conceive some of it objectively, but much of it may remain constitutionally beyond our reach. This too, of course, is *our* conception, but that doesn't mean it is only a higher-order description of the world as it appears to us. To insist otherwise is to assume that if any conception has a possessor, it must be about the possessor's point of view—a slide from subjective form to subjective content. If there are other intelligent beings whose point of view is incommensurate with ours, there is no reason why they should not also see us as contained in the world in this way.

My disagreement with Strawson, then, is with the way he interprets the idea that we are embedded in a world larger than we can conceive. What lies beyond our current understanding is not adequately captured in the idea of answers to questions we do not yet know how to ask. It may include things that we or creatures like us could never formulate questions about.

Whether such things are possible is absolutely central, for if they are,

they set a standard of reality independent of the mind, which more familiar things may also meet. We may then say that the reality of the features of things in themselves that we have discovered is just as independent of our capacity to discover them as is the reality of whatever may lie outside our conceptual reach, actual or possible. In that case what we know or think about when we think about the structure of matter or the physical nature of light or sound is something whose appearance to us in any form is incidental, and whose existence is not merely that of a character in our best theories.

4. Wittgenstein

Let me turn finally to one of the most important sources of contemporary idealism: Wittgenstein.⁷ His later views on the conditions of meaning seem to imply that nothing can make sense which purports to reach beyond the outer bounds of human experience and life, for it is only within a community of actual or possible users of the language that there can exist that possibility of agreement in its application which is a condition of the existence of rules, and of the distinction between getting it right and getting it wrong. This appears to rule out not only languages which could be understood by only one person, but also the use of language—even the general language of existence and states of affairs—to talk about what we cannot in principle make any judgments about.

Wittgenstein doesn't propose sociological truth-conditions for propositions about rules and meaning any more than he proposes behavioral truth-conditions for propositions about the mind. The criterial conditions he describes are assertability-conditions—truth-conditions having been abandoned as a tool of analysis in the *Philosophical Investigations*.⁸ Still, if Wittgenstein is right, it makes sense to say that someone is or is not using a concept correctly only against the background of the possibility of agreement and identifiable disagreement in judgments employing the concept. What implications does this have for my claim that the concepts of existence and reality can be significantly applied in the thought that there may be things that we could never conceive of?

As a desperate measure one might argue that Wittgenstein's claim about rules, if true, does not have such restrictive consequences as it is often thought to have. The possibility of agreement in judgments is a very broad condition that can be met in many ways. It need not imply,

7. See Williams (5).

8. See Dummett (1), and Kripke (2), pp. 73 ff.

and was not thought by Wittgenstein to imply, that we could understand only what we could verify or confirm. That would be true only if the judgments in which there had to be agreement were specific judgments of fact. But that need not be so. In regard to the possibility that there are things about the world that we couldn't conceive and that what there is very likely extends beyond the range of possible human thought, it might be suggested that Wittgensteinian conditions of publicity and possible agreement could be met in another way. Agreement about what we do not know and cannot conceive, and about what is possible, are as important as agreement about what we do know and what is true. That too is agreement in judgments.

The defense of my position would then be this: if it is a natural consequence of things that human beings are generally willing to admit, that what there is and what is true may extend beyond what they or their human successors could ever discover or conceive of or describe in some extension of human language, then that casts doubt on philosophical arguments purporting to show that such words transgress the rules of our language. Such arguments must appeal to truth-conditions for meaning that are stronger than Wittgensteinian assertability-conditions. It might even be suggested that Wittgenstein's occasional evocation of possible forms of life incomprehensibly different from our own was a gesture in the direction of admitting the reality of what we cannot conceive.

What convinces me of the hopelessness of this attempt to reconcile my realism with Wittgenstein's picture of language is that the interpretation supports a similar argument against some, though by no means all of the conclusions Wittgenstein himself draws concerning linguistic illegitimacy. I claimed in the last chapter that certain kinds of skepticism, for example, cannot be ruled out as violations of the language because they are part of the data about how people are naturally inclined to use the language, the judgments—of ignorance, of possibility, of doubt—which they naturally agree to in certain circumstances. If these are errors, they are not linguistic errors, and the disposition to make them must be taken as important evidence of how the language works: what factual judgments and claims to knowledge mean—specifically, how ambitious they are to reach beyond their grounds.

The position I am trying to defend is not really compatible with Wittgenstein's picture of the relation of language to the world, even if some of what he says about rules can be interpreted in a way which seems to allow it. His view of how thought is possible clearly implies that any thoughts we can have of a mind-independent reality must remain within the boundaries set by our human form of life, and that we can't appeal

to a completely general idea of what there is to defend the existence of kinds of facts which are in principle beyond the possibility of human confirmation or agreement. We fall into nonsense, he thinks, if we try to take language too far from these conditions. We can't think of our world as part of a larger universe that also contains things revealed only to forms of life inaccessible to us; we can't apply the concepts of belief or truth to a point of view totally unreachable from our own.

The radical problem to which Wittgenstein's position is a response—how finite beings can grasp concepts and rules with infinite applications—has been forcefully expounded by Kripke, and I have no idea how to deal with it. If Wittgenstein is right, then my claim to have formed a significant thought about what is entirely beyond the reach of our minds clearly won't stand up. The assertion that there are such transcendent "facts" can have no foothold in human life. But though I have no alternative, I find it completely impossible to believe Wittgenstein's view—psychologically impossible. For example, it is a clear sign of something wrong that Wittgenstein is led to doubt that there must be an answer to the question whether the sequence 7777 occurs in the decimal expansion of π (Wittgenstein (2), sec. 516). If he is right, *nothing* in my mind determines the infinite application of any of my concepts. We simply do apply them, unhesitatingly, in certain ways, and correct others who do not.

It seems to me that to accept this as the final story is to acknowledge that all thought is an illusion. If our thoughts do not have infinite reach in a much stronger sense than this, then even the most mundane of them is not what it pretends to be. It is as if a natural Platonism makes the attempt to view the world in any other way look phony. In sum, the Wittgensteinian attack on transcendent thoughts depends on a position so radical that it also undermines the weaker transcendent pretensions of even the least philosophical of thoughts. I can't imagine what it would be like to believe it, as opposed to subscribing to it verbally.

I don't know how thought is possible. But assuming it is, I continue to believe that the general concept of reality, reaching beyond anything we could do to fill it out, involves no mistake, linguistic or otherwise. Like the unrecoverable details of the past, which cannot be given verificationist or even evidential significance, or the truth value of undecidable propositions in mathematics, it is part of our common system of concepts. Whether or not we will ever be able to conceive, let alone know, what things were like before the Big Bang, for example, the idea of what was the case then stands without the support of such elaboration, or even the remote possibility of it. The general idea of what we can never conceive is just another form of this kind of extension.

Realism is most compelling when we are forced to recognize the existence of something which we cannot describe or know fully, because it lies beyond the reach of language, proof, evidence, or empirical understanding. *Something* must be true with respect to the 7s in the expansion of π , even if we can't establish it; there must be something it is like to be a bat, even if we can never conceive it adequately. But once it is accepted in these cases, realism becomes a possible position with regard to what we *can* understand as well.

The idea of objectivity always points beyond mere intersubjective agreement, even though such agreement, criticism, and justification are essential methods of reaching an objective view. The language that we can have because of our agreement in responses enables us to reach beyond the responses to talk about the world itself. As almost everyone would concede, it enables us to say, truly or falsely, that rain was falling on Gibraltar exactly fifty thousand years ago, even if there is no way of reaching agreement in the application of the term to such a case. My position is simply a radical extension of this straight antiverificationism, and I do not see on what grounds Wittgenstein can draw the line between legitimate and illegitimate extensions beyond the range of actual agreement in judgments. Language reaches beyond itself, whether in the concept of rain or in the concept of what there is, though what it reaches can only be designated by using language or some other form of representation.

Any conception of the world must include some acknowledgment of its own incompleteness: at a minimum it will admit the existence of things or events we don't know about now. The issue is only how far beyond our actual conception of the world we should admit that the world may extend. I claim that it may contain not only what we don't know and can't yet conceive, but also what we never could conceive—and that this acknowledgment of the likelihood of its own limits should be built into our conception of reality. This amounts to a strong form of antihumanism: the world is not our world, even potentially. It may be partly or largely incomprehensible to us not just because we lack the time or technical capacity to acquire a full understanding of it, but because of our nature.

This may strike some as an extreme example of objective overreach. For it gives priority over the view from within our lives and our language to a view so far outside that we cannot even take it up directly, but must form a secondary conception of it, as something closed off to us, though perhaps accessible to other beings who could observe us from outside. But in fact the reference to an external view is inessential: it is just a device to make vivid the idea of the world in which we are contained. That world is not dependent on our view of it, or any other view: the

direction of dependence is the reverse. Even though we must use language to talk about the world and our relation to it, and even though certain conditions of agreement enable us to have this language, the areas of such possible agreement are a limited part of the world.

There may be some portions of discourse where rock bottom is the language game and the shared responses on which it depends, with no real reference in the world outside of our responses and no objectivity except what comes from agreement. There is a real issue about this with respect to ethics and aesthetics, and room for disagreement about where to draw the line between realism and mere intersubjectivity without external reference. But there seems to me still much of language and thought that must be interpreted in a strong sense "realistically."

As Wittgenstein points out (Wittgenstein (2), secs. 241–2), there is a parallel between conditions of meaning and conditions of measurement. We cannot measure temperature with a thermometer unless there is a certain constancy in the results of such measurement. But that doesn't mean that temperature is nothing but a phenomenon of agreement among thermometer readings. It would exist even if there were no thermometers, and we can explain the actual agreement among thermometers by the uniform effect on them of temperature. In giving this explanation we use the concept of temperature, and a condition of our having the concept of temperature we have is that we can measure it. But that doesn't make the explanation circular, any more than a lecture on the operation of the larynx is circular. To use something that one is trying to explain is not to explain it in terms of itself.

In the employment of language we are ourselves a bit like measuring instruments, able to respond consistently to certain aspects of the world, and therefore able to talk about them. We can then explain these responses without circularity in terms of the things that produce them, using the concepts that the responses enable us to have. My contention is that just as our concepts allow us to recognize that these things would exist even if we lacked the responses that enable us to talk about them, so they allow us to recognize that there are other things we haven't been able to connect with yet, and that there may be still others (I should say there almost certainly are) with which creatures like us could never make such a connection, because we couldn't develop the necessary responses or the necessary concepts. I do not see how this supposition violates the conditions of significant thought.

In fact, to deny it shows a lack of humility, even if the denial is defended by reference to a theory of the concepts of existence and truth. Idealism—the view that what exists in the widest sense must be identified with what is thinkable by us in the widest sense—is an attempt to cut the universe down to size.

BIRTH, DEATH, AND THE MEANING OF LIFE

1. Life

One summer more than ten years ago, when I taught at Princeton, a large spider appeared in the urinal of the men's room in 1879 Hall, a building that houses the Philosophy Department. When the urinal wasn't in use, he would perch on the metal drain at its base, and when it was, he would try to scramble out of the way, sometimes managing to climb an inch or two up the porcelain wall at a point that wasn't too wet. But sometimes he was caught, tumbled and drenched by the flushing torrent. He didn't seem to like it, and always got out of the way if he could. But it was a floor-length urinal with a sunken base and a smooth overhanging lip: he was below floor level and couldn't get out.

Somehow he survived, presumably feeding on tiny insects attracted to the site, and was still there when the fall term began. The urinal must have been used more than a hundred times a day, and always it was the same desperate scramble to get out of the way. His life seemed miserable and exhausting.

Gradually our encounters began to oppress me. Of course it might be his natural habitat, but because he was trapped by the smooth porcelain overhang, there was no way for him to get out even if he wanted to, and no way to tell whether he wanted to. None of the other regulars did anything to alter the situation, but as the months wore on and fall turned to winter I arrived with much uncertainty and hesitation at the decision

to liberate him. I reflected that if he didn't like it on the outside, or didn't find enough to eat, he could easily go back. So one day toward the end of the term I took a paper towel from the wall dispenser and extended it to him. His legs grasped the end of the towel and I lifted him out and deposited him on the tile floor.

He just sat there, not moving a muscle. I nudged him slightly with the towel, but nothing happened. I pushed him an inch or two along the tiles, right next to the urinal, but he still didn't respond. He seemed to be paralyzed. I felt uneasy but thought that if he didn't want to stay on the tiles when he came to, a few steps would put him back. Meanwhile he was close to the wall and not in danger of being trodden on. I left, but when I came back two hours later he hadn't moved.

The next day I found him in the same place, his legs shriveled in that way characteristic of dead spiders. His corpse stayed there for a week, until they finally swept the floor.

It illustrates the hazards of combining perspectives that are radically distinct. Those hazards take many forms; in this final chapter I shall describe some that arise in connection with our attitude toward our own lives.

The pursuit of objectivity with respect to value runs the risk of leaving value behind altogether. We may reach a standpoint so removed from the perspective of human life that all we can do is to observe: nothing seems to have value of the kind it appears to have from inside, and all we can see is human desires, human striving—human *valuing*, as an activity or condition. In chapter 8 I alluded to the fact that if we continue along the path that leads from personal inclination to objective values and ethics, we may fall into nihilism. The problem is to know where and how to stop, and it shows itself in some of the more personally disturbing questions of philosophy.

The uneasy relation between inner and outer perspectives, neither of which we can escape, makes it hard to maintain a coherent attitude toward the fact that we exist at all, toward our deaths, and toward the meaning or point of our lives, because a detached view of our own existence, once achieved, is not easily made part of the standpoint from which life is lived. From far enough outside my birth seems accidental, my life pointless, and my death insignificant, but from inside my never having been born seems nearly unimaginable, my life monstrously important, and my death catastrophic. Though the two viewpoints clearly belong to one person—these problems wouldn't arise if they didn't—they function independently enough so that each can come as something of a surprise to the other, like an identity that has been temporarily forgotten.

The subjective view is at the core of everyday life, and the objective develops initially as a form of extended understanding; much of what it reveals can be used instrumentally in the pursuit of subjective aims. But taken far enough, it will undermine those aims: to see myself objectively as a small, contingent, and exceedingly temporary organic bubble in the universal soup produces an attitude approaching indifference. My attitude toward TN's life from that perspective is the same as toward any other creature's life. My attitude toward myself is quite different, and the two collide. The same person who is subjectively committed to a personal life in all its rich detail finds himself in another aspect simultaneously detached; this detachment undermines his commitment without destroying it—leaving him divided. And the objective self, noticing that it is personally identical with the object of its detachment, comes to feel trapped in this particular life—detached but unable to disengage, and dragged along by a subjective seriousness it can't even attempt to get rid of.

Some of the attitudes that lead to these conflicts may be mistaken; others may be modifiable by a process of reconciliation in which detachment is reduced and engagement is modified, so that the clash is lessened; but some of the conflicts are impossible to eliminate. No doubt many who have experienced the discomfort of objective detachment from themselves simply forget about it and live inside the world as if there were no external view. Some may dismiss these existential worries as bogus or artificial. I find both these responses unacceptable, because the objective standpoint, even at its limit, is too essential a part of us to be suppressed without dishonesty and because the effort to find a form of individual life that acknowledges and includes it can be fruitful, even if complete integration inevitably eludes us.

The wish to live so far as possible in full recognition that one's position in the universe is not central has an element of the religious impulse about it, or at least an acknowledgment of the question to which religion purports to supply an answer. A religious solution gives us a borrowed centrality through the concern of a supreme being. Perhaps the religious question without a religious answer amounts to antihumanism, since we cannot compensate for the lack of cosmic meaning with a meaning derived from our own perspective. We have to accomplish what reconciliation we can inside our own heads, and the possibilities are limited. The problem is not merely intellectual. The external standpoint and the contemplation of death lead to loss of equilibrium in life. Most of us have felt suddenly giddy at the thought of the extreme unlikelihood of our own birth or the thought of the world sailing on after we are dead. Some of us feel a constant undertow of absurdity in the projects and

ambitions that give our lives their forward drive. These jarring displacements of the external view are inseparable from the full development of consciousness.

Let me start with the problem of birth—the problem of what attitude to take toward the fact that we exist at all. This is less familiar than the problems of death and the meaning of life, but its form is similar. From an objective standpoint two things, neither of them easy to assimilate, strike me about my own birth: its extreme contingency and its unimportance. (The same can be said of anyone's birth; and if it is someone I care about, many of the same difficulties will arise as a result.) Let me discuss these two points in turn, contingency first.

Subjectively, we begin by taking our existence for granted: it is a given of the most basic kind. When in childhood each of us first learns of the contingency of his existence, even the simple fact that it depends on his parents having met, the result is a lessening of his unreflectively secure footing in the world. We are here by luck, not by right or by necessity.

Rudimentary biology reveals how extreme the situation is. My existence depends on the birth of a particular organism that could have developed only from a particular sperm and egg, which in turn could have been produced only by the particular organisms that produced them, and so forth. In view of the typical sperm count, there was very little chance of my being born given the situation that obtained an hour before I was conceived, let alone a million years before, unless everything that happens in the world is determined with absolute rigidity—which appears not to be the case. The natural delusion of my own inevitability collides with the objective fact that *who* exists and has existed is radically contingent, my own existence in particular being one of the most inessential things in the world. Almost every possible person has not been born and never will be, and it is sheer accident that I am one of the few who actually made it.

The subjective effects of this information are complex: it is not just a sobering thought. There is amazement and relief, plus the giddiness or retrospective alarm that comes from learning about a very close call only after the risk is past. There can also be a taste of survivor's syndrome, guilt toward all those others who will never be born. Yet at the same time, the sense of subjective inevitability doesn't completely disappear; these indisputable objective facts about myself provoke emotional incredulity. I can imagine having died at the age of five, but it is not easy to grasp in full consciousness the fact that the history of the universe might have run its course without my ever putting in an appearance *at all*. When I do get my mind properly around this idea, it produces a sinking feeling which reveals that a powerful but unnoticed support has been

removed from my world.

My own existence looms large at the center of my prereflective world picture, since this life is the source and avenue for my understanding of everything else. It is unnerving to be led through it to the discovery that it is totally inessential—one of the least “basic” things in the world. A world without me at any point in its history seems like a world with a crucial piece missing, a world that has suddenly lost its moorings. If you concentrate hard on the thought that you might never have been born—the distinct possibility of your eternal and complete absence from this world—I believe you too will find that this perfectly clear and straightforward truth produces a positively uncanny sensation.

There is a solipsistic tendency at work here, which attaches curiously enough to the objective rather than to the subjective viewpoint—like the solipsism of Wittgenstein’s *Tractatus*. The objective standpoint, which considers everything *sub specie aeternitatis*, easily slips into regarding itself as the condition of the world or the frame of all existence, rather than as an aspect of a particular individual within the world.

If I forget who I am, I can imagine TN’s never having existed without a qualm; but there is a risk of doing this only halfway, by imagining *my world* without TN in it—the objective self being held in reserve, so to speak. To fully imagine the world without me, I have to get rid of the objective self as well, and this begins to feel like getting rid of the world itself rather than of something in it. It is as if there were a natural illusion that the world is not completely detachable from my conception of it.

The appeal of this illusion is mysterious. I know perfectly well that the world is not essentially my world: it is not a necessary truth that it can be or could ever have been referred to or thought about by me. Just as the room I am now in, which I can refer to as this room, might never have borne the relation to me that makes that reference possible—because I might never have been in it—so it is with this world and everything about it. The fact that I must exist in it to think that it might not have contained me does not make the possibility any less real. Even if some of the respects in which the world can be objectively described are essential to it, I as a particular possessor of the objective view am not, and I can think of the world in abstraction from my connection with it.

But even if we firmly reject the solipsistic assumption as a mistake, a pale version of it remains. Although the world is not essentially my world, the objective recognition of my contingency has to coexist in my head with a total world picture whose subject is inescapably me. The person whose contingency I recognize is the epicenter not just of the world as it looks from here, but of my entire world view. To suppose that he should never have existed is to suppose that my world should never have existed.

This makes it seem as if the existence of my world depends on the existence of something in it. But of course that is not the case. The real me is not merely part of *my* world. The person who I am is a contingent bit of a world that is not just mine. So my world depends for its existence on me, I depend for my existence on TN, and TN depends on *the world* and is inessential to it. This is another of the discomforts of being someone in particular: my world depends for its existence on his birth, even though he also appears in it as a character. It is eerie to see oneself and one's entire world in this way as a natural product.

The second thing that emerges from an objective view of my birth is its unimportance. While importance and unimportance will be more thoroughly discussed in the following section, let me say a bit about it here.

Setting aside broader questions about the value of human existence as such, when we look at the world from a general vantage point it seems not to matter who exists. My own existence or that of any other particular person is entirely gratuitous. There may be some reason for my continued existence now that I am here, but there is absolutely no reason why I should have come into existence in the first place: if I hadn't, the world would have been none the worse; I certainly wouldn't have been *missed*! There may be a few people like Mozart and Einstein whose non-existence would have been a real loss, but for the most part there is no reason why anyone in particular should exist. We might go further and say there is no reason why human beings and their form of life should ever have existed: if they hadn't, it would not have been necessary to invent them—in any case there might have been other types of beings instead. But the problem I am talking about here doesn't depend on that wider claim. Let us suppose that the world would be a less interesting and valuable place if there were no people in it. The narrower claim is that objectively it doesn't matter which *particular* persons have come into existence.

This collides with the view we naturally take from within the world. Subjectively we feel that we and those we love belong here—that nothing could undermine our right of admission to the universe. Whatever others may think, the last thing we expect is that *we* may come to see the world in a way that drains our birth of value. But as objectivity increases, detachment sets in and the existence from which all of our concerns and motives and justifications begin becomes a matter of indifference. Anyone else would have done as well as you and I. And because nearly everything that matters to each of us depends on actually existing lives, what matters is firmly rooted in something whose nonexistence wouldn't have mattered in the least.

The obvious reply is that something can matter even if it didn't matter

in advance whether it would exist and wouldn't have mattered if it hadn't: however gratuitous its original appearance, once it exists it brings its own value with it and its survival and well-being become important.

There is something in this. When we look at the actual world, even the particular trees in it appear to have a value that is not canceled by their intersubstitutability or their gratuitousness. But it isn't enough to reconcile the two standpoints. From an external view of the universe, which abstracts from our own position in it, it still wouldn't have mattered if we had never existed—and that is not something we can simply accept from the standpoint of real life. It forces on us a kind of double vision and loss of confidence which is developed more fully in doubts about the meaning of life. It is easy to have these thoughts about someone else, and they are not disturbing when we abstract from who in particular we are. But when we bring them back home they cannot be assimilated. Neither as objective self nor as TN can I comfortably think of myself as utterly inessential.

There is a recognizable human desire to find our existence significant no matter how cosmic a view we take, and a consequent discomfort with the partial disengagement that objectivity induces. But perhaps these are mistakes, both intellectual and emotional. Perhaps we are just too demanding, and confer too much authority on the objective standpoint by permitting its lack of independent concern for our existence to support a judgment of insignificance. Perhaps the detachment of the objective self is empty, because justifications come to an end within life and if it makes no sense to seek them outside it we cannot be disappointed by the failure to find them there. Perhaps, as Williams claims, the view *sub specie aeternitatis* is a very poor view of human life, and we should start and end in the middle of things. Or perhaps there is something to be said on both sides.

2. *Meaning*

In seeing ourselves from outside we find it difficult to take our lives seriously. This loss of conviction, and the attempt to regain it, is the problem of the meaning of life.

I should say at the outset that some people are more susceptible to this problem than others, and even those who are susceptible to it vary over time in the degree to which it grips them. Clearly there are temperamental and circumstantial factors at work. Still, it is a genuine problem which we cannot ignore. The capacity for transcendence brings with it a liability to alienation, and the wish to escape this condition and to

find a larger meaning can lead to even greater absurdity. Yet we can't abandon the external standpoint because it is our own. The aim of reaching some kind of harmony with the universe is part of the aim of living in harmony with ourselves.

To the subjective view, the conditions that determine whether life makes sense are simply given, as part of the package. They are determined by the possibilities of good and evil, happiness and unhappiness, achievement and failure, love and isolation that come with being human, and more specifically with being the particular person you are in the particular social and historical setting in which you find yourself. From inside no justification can coherently be sought for trying to live a good and meaningful life by those standards; and if it were needed, it couldn't be found.

Serious problems about the meaning of a life can arise entirely within it, and these should be distinguished from the completely general philosophical problem of the meaning of life, which arises from the threat of objective detachment. A life may be absurd, and felt to be absurd, because it is permeated by trivia or dominated by a neurotic obsession or by the constant need to react to external threats, pressures, or controls. A life in which human possibilities for autonomy and development are largely unrealized and untested will seem deficient in meaning; someone faced with such a life may lack the significant will to live as a purely internal matter, not because of any objective detachment. But all these forms of meaninglessness are compatible with the possibility of meaning, had things gone differently.

The philosophical problem is not the same, for it threatens human life even at its subjective best with objective meaninglessness, and with absurdity if it cannot stop taking itself seriously. This problem is the emotional counterpart of that sense of arbitrariness which the objective self feels at being someone in particular.

Each of us finds himself with a life to lead. While we have a certain amount of control over it, the basic conditions of success and failure, our basic motives and needs, and the social circumstances that define our possibilities are simply given. Shortly after birth we have to start running just to keep from falling down, and there is only limited choice as to what will matter to us. We worry about a bad haircut or a bad review, we try to improve our income, our character, and our sensitivity to other people's feelings, we raise children, watch Johnny Carson, argue about Alfred Hitchcock or Chairman Mao, worry about getting promoted, getting pregnant, or becoming impotent—in short we lead highly specific lives within the parameters of our place, time, species, and culture. What could be more natural?

Yet there is a point of view from which none of it seems to matter.

When you look at your struggles as if from a great height, in abstraction from the engagement you have with this life because it is yours—perhaps even in abstraction from your identification with the human race—you may feel a certain sympathy for the poor beggar, a pale pleasure in his triumphs and a mild concern for his disappointments. And of course given that this person exists, there is little he can do but keep going till he dies, and try to accomplish something by the standards internal to his form of life. But it wouldn't matter all that much if he failed, and it would matter perhaps even less if he didn't exist at all. The clash of standpoints is not absolute, but the disparity is very great.

This kind of detachment is certainly possible for us, but the question is whether *it* matters. What am I doing out there, pretending to be a visitor from outer space—looking at my life from a great height in abstraction from the fact that it is mine, or that I am human and a member of this culture? How can the unimportance of my life from that point of view have any importance for *me*? Perhaps the problem is a purely philosophical artefact, and not real.

I'll return to this objection later, but first let me pose another: Even if the problem can't be dismissed as unreal, it may have a simple solution. Is it so certain that the attitudes really conflict as they appear to? Since the two judgments arise from different perspectives, why isn't their content appropriately relativized to those perspectives, rendering the conflict illusory? If that were true, it would be no more problematic that the course of my life should matter from inside but not from outside than that a large mouse should be a small animal, or that something should look round from one direction and oval from another.

I do not believe this solution is available, logical as it may seem. The trouble is that the two attitudes have to coexist in a single person who is actually leading the life toward which he is simultaneously engaged and detached. This person does not occupy a third standpoint from which he can make two relativized judgments about his life. If all he had were two relativized judgments, they would leave him with no attitude toward his life at all—only information about the appropriate attitude from two points of view, neither of which was his. But in fact he occupies both of the conflicting points of view and his attitudes derive from them both.

The real problem is with the external point of view, which cannot remain a mere spectator once the self has expanded to accommodate it. It has to join in with the rest and lead this life from which it is disengaged. As a result the person becomes in significant part detached from what he is doing. The objective self is dragged along by the unavoidable engagement of the whole person in the living of a life whose form it recognizes as arbitrary. It generates a demand for justification which it

at the same time guarantees to be unsatisfiable, because the only available justification depends on the view from inside.

Some philosophers have held that the soul is trapped in the body. To Plato this meant not just that the soul is housed in the body, but that the needs of the body invade and compromise the soul, threatening it with domination by its basest portion, the appetites. I am talking about something different but analogous, the inevitable engagement of the objective self in a particular contingent life, whose consequence is not depravity but absurdity. It isn't so much the animal aspects of life which generate the absurd, for no judgment of importance need attach to an instinctive effort to survive. The problem comes especially with those more developed human projects that pretend to significance and without which life would not be human. It is heightened by the equally human involvement we all have with our *own* lives and ambitions.

This invites a quick solution to the problem. Perhaps we can avoid the absurd if we devote ourselves to providing only for the basic needs of everyone. There is a great deal of misery in the world, and many of us could easily spend our lives trying to eradicate it—wiping out starvation, disease, and torture.

Such aims do indeed seem to give life a meaning that is hard to question. But while they are certainly worthy and perhaps imperative goals, they cannot eliminate the problem. Granted, one advantage of living in a world as bad as this one is that it offers the opportunity for many activities whose importance can't be questioned. But how could the main point of human life be the elimination of evil? Misery, deprivation, and injustice prevent people from pursuing the positive goods which life is assumed to make possible. If all such goods were pointless and the only thing that really mattered was the elimination of misery, that really *would* be absurd. The same could be said of the idea that helping others is the only thing that really gives meaning to life. If no one's life has any meaning in itself, how can it acquire meaning through devotion to the meaningless lives of others?

No, even if the problem of meaning can be postponed until all misery is wiped out, it will not go away permanently. In any case, most of us face it in our own lives. Some of us have bigger egos than others, and a dominant obsession with personal standing or success is recognized as absurd even without the benefit of philosophy. But everyone who is not either mystically transformed or hopelessly lacking in self-esteem regards his life and his projects as important, and not just to him.

From outside we do indeed tend to see most of our pursuits as important only relatively. Watching the human drama is a bit like watching a Little League baseball game: the excitement of the participants is per-

fectly understandable but one can't really enter into it. At the same time, since one is one of the participants, one is caught up in the game directly, in a way that cannot include an admission of relativity. When you are considering a career, marriage, children, or even whether to go on a diet, review a book, or buy a car, the external standpoint is excluded and you face the matter directly, from the internal standpoint of ordinary life. The detached external view just has to come along, and accommodate itself to the unqualified concerns that it can't internalize. At the same time there is a temptation to resist this detachment by inflating the sense of objective importance inside one's life, either by overvaluing one's own significance or by attributing a wider significance to one's pursuits.

It is the same phenomenon we have discussed in other connections—in epistemology, for example. The internal view resists the reduction to a subjective interpretation of its contents which the external view tries to force on it. But this puts the objective standpoint in conflict with itself. Finding my life objectively insignificant, I am nevertheless unable to extricate myself from an unqualified commitment to it—to my aspirations and ambitions, my wishes for fulfillment, recognition, understanding, and so forth. The sense of the absurd is the result of this juxtaposition.

This is not an artificial problem created by a philosophical misstep, any more than epistemological skepticism is. Just as we can't evade skepticism by denying the pretensions of our beliefs about the world and interpreting them as entirely relative to a subjective or personal point of view, so we can't evade the impact of objective detachment by denying the objective pretensions of our dominant aims in life. This would simply falsify the situation. The problem of the meaning of life is in fact a form of skepticism at the level of motivation. We can no more abandon our unqualified commitments at will than we can abandon our beliefs about the world in response to skeptical arguments, however persuasive we may find them, as Hume famously observed. Nor, I believe, can we avoid either problem by refusing to take that step outside ourselves which calls the ordinary view into question.

Several routes might be attempted out of this impasse. I don't believe there is a way out, though there are adjustments we can make to live with the conflict. But it is worth considering what would be required to eliminate it entirely. I'll discuss two proposals which try to meet the problem head-on and one which tries to dissolve it.

The first solution to consider is the most Draconian: to deny the claims of the subjective view, withdraw from the specifics of individual human life as much as possible, minimize the area of one's local contact with the

world and concentrate on the universal. Contemplation, meditation, withdrawal from the demands of the body and of society, abandonment of exclusive personal ties and worldly ambition—all this gives the objective standpoint less to be disengaged from, less to regard as vain. I gather this response is recommended by certain traditions, though I don't know enough to be sure that it isn't a caricature: the loss of self in the individual sense is thought to be required by the revelations of an impersonal view, which takes precedence over the view from here. And apparently it is possible for some individuals to achieve this withering away of the ego, so that personal life continues only as a vehicle for the transcendent self, not as an end in itself.¹

I cannot speak from experience, but this seems to me a high price to pay for spiritual harmony. The amputation of so much of oneself to secure the unequivocal affirmation of the rest seems a waste of consciousness. I would rather lead an absurd life engaged in the particular than a seamless transcendental life immersed in the universal. Perhaps those who have tried both would laugh inscrutably at this preference. It reflects the belief that the absurdity of human life is not such a bad thing. There are limits to what we should be prepared to do to escape it—apart from the point that some of these cures may be more absurd than the disease.

The second solution is the opposite of the first—a denial of the objective unimportance of our lives, which will justify full engagement from the objective standpoint. While this response to detachment has some merit, the truth in it is not enough to resolve the conflict.

As I argued in chapter 8, an impersonal perspective doesn't necessarily lead to nihilism. It may fail to discover *independent* reasons to care about what subjectively concerns us, but much that is of value and significance in the world can be understood directly only from within the perspective of a particular form of life, and this can be recognized from an external standpoint. The fact that the point of something can't be understood from the objective standpoint alone doesn't mean it must be regarded objectively as pointless, any more than the fact that the value of music is not directly comprehensible to someone deaf from birth means he has to judge it worthless. His knowledge of its value must depend on others. And the objective standpoint can recognize the authority of particular points of view with regard to worth as it can with regard to essentially perspectival facts. This includes recognizing the worth of what is of value only to a particular creature—who may be one-

1. Or perhaps alternatively, each particular element, whatever its character, is seen as a manifestation of the universal.

self. We might say that absolute value is revealed to the objective view through the evidence available to particular perspectives, including one's own. So even if there is no externally appreciable reason for the existence of any particular form of life, including my own, at least some of the values, positive and negative, that are defined by reference to it can be externally acknowledged. Playing in a Little League baseball game, making pancakes, or applying a coat of nail polish are perfectly good things to do. Their value is not necessarily canceled by the fact that they lack external justification.

This is not enough, however, to harmonize the two standpoints, because it doesn't warrant a particular objective interest in the individual life that happens to be mine, or even in the general form of human life of which it is an instance. These things have been handed to me and they demand my full attention. But to the external view, many different actual and possible subjective values must be acknowledged. Those arising within my life may evoke sympathy, but that is not the same as true objective engagement. My life is one of countless many, in a civilization that is also not unique, and my natural devotion to it is quite out of proportion to the importance I can reasonably accord it from outside.

From there I can accord it no more importance than it merits in a global view which includes all possible forms of life and their value on an equal footing. It is true that my life is the one among all these that I am in the best position to devote attention to, and it could be argued that the traditional principle of division of labor warrants my concentrating on it in the usual way as the best method of contributing to the cosmic pool. But while there is something in this, it should not be exaggerated. The argument would not really justify us in engaging fully with our personal aims from an objective standpoint, and such engagement as it warranted would be on sufferance from an objective concern for the whole of which we were a part. This is at best a method of partial reconciliation between inner and outer views: we can try to avoid assigning ourselves a personal importance grotesquely out of line with our objective value, but we can't realistically hope to close the gap completely. So while the acknowledgment of objective worth inside human life may make the conflict of standpoints less extreme, it doesn't eliminate it.

The third candidate solution I want to discuss can be thought of as an argument that the problem is unreal. The objection is that to identify with the objective self and find its detachment disturbing is to forget who you are. There is something deranged in looking at one's existence from so far outside that one can ask why it matters. If we were actually detached spirits about to be thrown into the world by embodiment in a

particular creature whose form of life had so far been only externally observed by us, it would be different: we might well feel a threat of impending captivity. But it isn't like that. We are first of all and essentially individual human beings. Our objectivity is simply a development of our humanity and doesn't allow us to break free of it. It must serve our humanity and to the extent that it does not we can forget about it.

The point here is to force withdrawal of the external demand which gives rise to the problem. This is a natural and in some ways appealing response, but as a conclusive argument it will not work. Objectivity is not content to remain a servant of the individual perspective and its values. It has a life of its own and an aspiration for transcendence that will not be quieted in response to the call to reassume our true identity. This shows itself not only in the permanent disaffection from individual life that is the sense of the absurd, but in the demands for objective justification which we sometimes *can* meet, as in the development of ethics. The external standpoint plays an important positive role in human motivation as well as a negative one, and the two cannot be separated. Both depend on the independence of the external view and the pressure it puts us under to bring it into our lives. The sense of the absurd is just a perception of the limits of this effort, reached when we ascend higher on the transcendental ladder than our merely human individuality can follow, even with the help of considerable readjustment. The objective self is a vital part of us, and to ignore its quasi-independent operation is to be cut off from oneself as much as if one were to abandon one's subjective individuality. There is no escape from alienation or conflict of one kind or another.

In sum, I believe there is no credible way of eliminating the inner conflict. Nonetheless, we have a motive for reducing it, and it is possible to promote a degree of harmony between the two standpoints without taking drastic measures. The attitude toward one's own life is inevitably dominated by the fact that it, unlike all those others, is one's own. But the domination should not be so complete that the objective standpoint treats the values defined by that life as ultimate. While objective reason does naturally fall into the service of the subjective passions, it can retain its recognition, however sympathetic, that they are the passions of a particular individual and whatever importance they may have derives from that. So objectivity itself is split into spectator and participant. It devotes itself to the interests and the ambitions, including the competitive ambitions, of one person while at the same time recognizing that he is no more important than anyone else and that the human form of life is not the embodiment of all value.

One of the devices by which these two attitudes are combined is moral-

ity, which seeks a way to live as an individual that affirms the equal worth of other individuals and is therefore externally acceptable. Morality is a form of objective reengagement. It permits the objective assertion of subjective values to the extent that this is compatible with the corresponding claims of others. It can take various forms, some of which I have discussed. All of them involve, to one degree or another, occupying a position far enough outside your own life to reduce the importance of the difference between yourself and other people, yet not so far outside that all human values vanish in a nihilistic blackout.

But there is more to integration than that. The most general effect of the objective stance ought to be a form of humility: the recognition that you are no more important than you are, and that the fact that something is of importance to you, or that it would be good or bad if you did or suffered something, is a fact of purely local significance. Such humility may seem incompatible with full immersion in one's life and in the pursuit of those enjoyments and goods that it makes possible. It may sound like a form of deadening self-consciousness, or self-denigration, or asceticism; but I don't think it has to be.

It does not create self-consciousness but simply gives it content. Our capacity for taking an external view of ourselves poses the problem; we cannot get rid of it, and we must find some attitude or other that reckons with it. Humility falls between nihilistic detachment and blind self-importance. It doesn't require reflection on the cosmic arbitrariness of the sense of taste every time you eat a hamburger. But we can try to avoid the familiar excesses of envy, vanity, conceit, competitiveness, and pride—including pride in our culture, in our nation, and in the achievements of humanity as a species. The human race has a strong disposition to adore itself, in spite of its record. But it is possible to live a complete life of the kind one has been given without overvaluing it hopelessly. We can even resist the tendency to overvalue the historical present, both positively and negatively; what is going on in the world right now is not for that reason especially important. The present is where we are, and we cannot see it only in timeless perspective. But we can forget about it now and then, even if it won't forget about us.

Finally, there is an attitude which cuts through the opposition between transcendent universality and parochial self-absorption, and that is the attitude of nonegocentric respect for the particular.² It is conspicuous as an element in aesthetic response, but it can be directed to all kinds of things, including aspects of one's own life. One can simply look hard at a ketchup bottle, and the question of significance from different stand-

2. I am grateful to Jacob Adler for making me see this.

points will disappear. Particular things can have a noncompetitive completeness which is transparent to all aspects of the self. This also helps explain why the experience of great beauty tends to unify the self: the object engages us immediately and totally in a way that makes distinctions among points of view irrelevant.

It is hard to know whether one could sustain such an attitude consistently toward the elements of everyday life. It would require an immediacy of feeling and attention to what is present that doesn't blend well with the complex, forward-looking pursuits of a civilized creature. Perhaps it would require a radical change in what one did, and that would raise the question whether the simplification was worth it.

Apart from this, the possibilities for most of us are limited. Some people are genuinely unworldly, but if it doesn't come naturally, the attempt to achieve this condition is likely to be an exercise in dishonesty and self-distortion. Most of us care a great deal about forms of individual success that we can see from an impersonal standpoint to be much less significant than we cannot help taking them to be from inside our lives. Our constitutional self-absorption together with our capacity to recognize its excessiveness make us irreducibly absurd even if we achieve a measure of subjective-objective integration by bringing the two standpoints closer together. The gap is too wide to be closed entirely, for anyone who is fully human.

So the absurd is part of human life. I do not think this can be basically regretted, because it is a consequence of our existence as particular creatures with a capacity for objectivity. Some philosophers, such as Plato, have been unhappy that the higher self was trapped in a particular human life, and others, such as Nietzsche, have denigrated the role of the objective standpoint; but I believe the significant diminution of either of them in force or importance would lessen us and is not a reasonable aim. Repression can operate effectively and damagingly not just against the instincts but against the objective intelligence. These civil wars of the self result in an impoverished life. It is better to be simultaneously engaged and detached, and therefore absurd, for this is the opposite of self-denial and the result of full awareness.

3. Death

The desire to go on living, which is one of our strongest, is essentially first-personal: it is not the desire that a particular, publicly identifiable human being survive, though its fulfillment of course requires the survival of someone like that, and therefore it collides with objective indif-

ference about the survival of anyone in particular. Your relation to your own death is unique, and here if anywhere the subjective standpoint holds a dominant position. By the same token, the internal standpoint will be vicariously dominant in your attitude toward the deaths of those to whom you are so close that you see the world through their eyes.

Some people believe in an afterlife. I do not; what I say will be based on the assumption that death is nothing, and final. I believe there is little to be said for it: it is a great curse, and if we truly face it nothing can make it palatable except the knowledge that by dying we can prevent an even greater evil. Otherwise, given the simple choice between living for another week and dying in five minutes I would always choose to live for another week; and by a version of mathematical induction I conclude that I would be glad to live forever.

Perhaps I shall eventually tire of life, but at the moment I can't imagine it, nor can I understand those many distinguished and otherwise reasonable persons who sincerely assert that they don't regard their own mortality as a misfortune.³

I can't take the kind of metaphysical consolation offered by Parfit (who observes that his view has parallels with Buddhism). By breaking down the metaphysical boundaries between himself and other persons, and loosening the metaphysical bonds that connect him now with his future self, he claims to have become less depressed about his own death, among other things. His death will be the termination of a certain connected sequence of activities and experiences, but not the annihilation of a unique underlying self. "Instead of saying, 'I shall be dead', I should say, 'There will be no future experiences that will be related, in certain ways, to these present experiences'. Because it reminds me what this fact involves, this redescription makes this fact less depressing" (Parfit (2), p. 281). As I've said in chapter 3, I can't accept the metaphysical revision, but I'm not sure that if I did, I'd find the conclusion less depressing. I actually find Parfit's picture of *survival* depressing—but that of course is by comparison with what *I* take survival to be. By comparison with Parfitian survival, Parfitian death may not seem so depressing; but that may owe as much to the deficiencies of the former as to the advantages of the latter. (See his remarks on p. 280.)

I am not going to concentrate here on explaining why death is a bad thing. Life can be wonderful, but even if it isn't, death is usually much worse. If it cuts off the possibility of more future goods than future evils for the victim, it is a loss no matter how long he has lived when it hap-

3. For example Williams (3), ch. 6, "The Makropulos Case; Reflections on the Tedium of Immortality." Can it be that he is more easily bored than I?

pens. And in truth, as Richard Wollheim says, death is a misfortune even when life is no longer worth living (p. 267). But here I want to say something about what it means to look forward to our own deaths, and how, if at all, we can bring the inner and outer views together. The best I can hope to offer is a phenomenological account. I hope it is not simply idiosyncratic. Like the contingency of our birth, the inevitability of our death is easy to grasp objectively, but hard to grasp from within. Everyone dies; I am someone, so I will die. But it isn't just that TN will be killed in a plane crash or a holdup, have a stroke or a heart attack or lung cancer, the clothes going to the Salvation Army, the books to the library, some bits of the body to the organ bank and the rest to the crematorium. In addition to these mundane objective transitions, my world will come to an end, as yours will when you die. That's what's hard to get hold of: the internal fact that one day this consciousness will black out for good and subjective time will simply stop. My death as an event in the world is easy to think about; the end of my world is not.

One of the difficulties is that the appropriate form of a subjective attitude toward my own future is expectation, but in this case there is nothing to expect. How can I expect nothing *as such*? It seems that the best I can do is to expect its complement, a finite but indeterminate amount of something—or a determinate amount, if I am under definite sentence of death. Now a good deal could be said about the consequences of the finiteness of my future, but that is relatively banal and something most of us automatically allow for, particularly after reaching the age of forty. I am concerned with the adequate recognition of my eventual annihilation itself. There will be a last day, a last hour, a last minute of consciousness, and that will be it. Off the edge.

To grasp this it isn't enough to think only of a particular stream of consciousness coming to an end. The external view of death is psychological as well as physical: it includes the idea that the person who you are will have no more thoughts, experiences, memories, intentions, desires, etc. That inner life will be finished. But the recognition that the life of a particular person in the world will come to an end is not what I am talking about. To grasp your own death from within you must try to look *forward* to it—to see it as a *prospect*.

Is this possible? It might seem that there could be no form of thought about your own death except either an external view, in which the world is pictured as continuing after your life stops, or an internal view that sees only this side of death—that includes only the finitude of your expected future consciousness. But this is not true. There is also something that can be called the expectation of nothingness, and though the mind tends to veer away from it, it is an unmistakable experience, always

startling, often frightening, and very different from the familiar recognition that your life will go on for only a limited time—that you probably have less than thirty years and certainly less than a hundred. The positive prospect of the end of subjective time, though it is logically inseparable from these limits, is something distinct.

What is the specific object of this feeling? In part, it is the idea of the objective world and objective time continuing without me in it. We are so accustomed to the parallel progress of subjective and objective time that there is some shock in the realization that the world will go calmly on without me after I disappear. It is the ultimate form of abandonment.

But the special feeling I am talking about does not depend only on this, for it would be there even if solipsism were true—even if my death brought with it the end of the only world there was! Or even if, reversing the direction of dependence, my death was going to occur *in consequence* of the end of the world. (Suppose I came to believe a crackpot scientific theory that as the result of a spectacular and inexorable rise in matter/antimatter collisions, the universe was going to annihilate itself completely six months from now.) It is the prospect of nothingness itself, not the prospect that the world will go on after I cease to exist, that has to be understood.

It hardly needs saying that we are accustomed to our own existence. Each of us has been around for as long as he can remember; it seems the only natural condition of things, and to look forward to its end feels like the denial of something which is more than a mere possibility. It is true that various of my possibilities—things I might do or experience—will remain unrealized as a result of my death. But more fundamental is the fact that they will then cease even to be possibilities—when I as a subject of possibilities as well as of actualities cease to exist. That is why the expectation of complete unconsciousness is so different from the expectation of death. Unconsciousness includes the continued possibility of experience, and therefore doesn't obliterate the here and now as death does.

The internal awareness of my own existence carries with it a particularly strong sense of its own future, and of its possible continuation beyond any future that may actually be reached. It is stronger than the sense of future possibility attaching to the existence of any particular thing in the world objectively conceived—perhaps of a strength surpassed only by the sense of possible continuation we have about the world itself.

The explanation may be this. In our objective conception of the world, particular things can come to an end because the possibility of their non-existence is allowed for. The possibility of both the existence and the

nonexistence of a particular object, artefact, organism, or person is given by actualities which underlie either possibility and coexist with both of them. Thus the existence of certain elements and the truth of the laws of chemistry underlie the possibility of synthesizing a particular chemical compound, or of decomposing it. Such possibilities rest on actualities.

But some possibilities seem themselves to be basic features of the world and not to depend on more deep-seated actualities. For example the number of possible permutations of m things taken n at a time, or the number of possible Euclidean regular solids, are possibilities whose existence is not contingent on anything.

Now the various possibilities some of which make up my life, and many of which I will never realize, are contingent on my existence. My existence is the actuality on which all these possibilities depend. (They also depend on the existence of things outside me which I can encounter, but let me leave that aside for the moment.) The problem is that when I think of myself from the inside, there seems to be nothing still more basic which reveals the actuality of my existence as in turn the realization of a possibility of existence which is correlative with a possibility of non-existence based on the same foundation. In other words the possibilities which define the subjective conditions of my life seem not to be explainable in turn, within a subjective view, as the contingent realization of deeper possibilities. Nothing is subjectively related to them as the existence of the elements is related to the possibility of a compound.

To explain them we have to go outside of the subjective view for an objective account of why TN exists and has the characteristics that determine his subjective possibilities. They rest on an external actuality.

But this gives rise to an illusion when we think about our lives from inside. We can't really make these external conditions part of the subjective view; in fact we have no idea how they generate our subjective possibilities on any view. It is as if the possible contents of my experience, as opposed to the actualities, themselves constituted a universe, a domain within which things can occur but which is not itself contingent on anything. The thought of the annihilation of this universe of possibilities cannot then be thought of as the realization of yet another possibility already given by an underlying subjective actuality. The subjective view does not allow for its own annihilation, for it does not conceive of its existence as the realization of a possibility. This is the element of truth in the common falsehood that it is impossible to conceive of one's own death.

All this is fairly obvious, but I think it explains something. The sense each person has of himself from inside is partly insulated from the exter-

nal view of the person who he is, and it projects itself into the future autonomously, so to speak. My existence seems in this light to be a universe of possibilities that stands by itself, and therefore stands in need of nothing else in order to continue. It comes as a rude shock, then, when this partly buried self-conception collides with the plain fact that TN will die and I with him. This is a very strong form of nothingness, the disappearance of an inner world that had not been thought of as a contingent manifestation at all and whose absence is therefore not the realization of a possibility already contained in the conception of it. It turns out that I am not the sort of thing I was unconsciously tempted to think I was: a set of ungrounded possibilities as opposed to a set of possibilities grounded in a contingent actuality. The subjective view projects into the future its sense of unconditional possibilities, and the world denies them. It isn't just that they won't be actualized—they will *vanish*.

This is not just a realization about the future. The submerged illusion it destroys is implicit in the subjective view of the present. In a way it's as if I were dead already, or had never really existed. I am told that the fear of flying typically has as its object not just the possibility of crashing, but flying itself: hurtling along in a smallish vehicle miles above the surface of the earth. It's a bit like that, only in this case it's something you can repeatedly forget and rediscover: all along you have been thinking you were safely on the ground and suddenly you look down and notice that you're standing on a narrow girder a thousand feet above the pavement.

I have said nothing so far about the most perplexing feature of our attitude toward death: the asymmetry between our attitudes to past and future nonexistence. We do not regard the period before we were born in the same way that we regard the prospect of death. Yet most of the things that can be said about the latter are equally true of the former. Lucretius thought this showed it was a mistake to regard death as an evil. But I believe it is an example of a more general future-past asymmetry which is inseparable from the subjective view.

Parfit has explored the asymmetry in connection with other values such as pleasure and pain. The fact that a pain (of ours) is in prospect rather than in the past has a very great effect on our attitude toward it, and this effect cannot be regarded as irrational (Parfit (3), sec. 64).

While I have no explanation of the asymmetry, I believe it must be admitted as an independent factor in the subjective attitude toward our own death. In other words it can't be accounted for in terms of some other difference between past and future nonexistence, any more than the asymmetry in the case of pain can be accounted for in terms of some

other difference between past and future pains, which makes the latter worse than the former.

It is a fact perhaps too deep for explanation that the cutting off of future possibilities, both their nonactualization and their obliteration even as possibilities, evokes in us a very different reaction from any parallel nonrealization or nonexistence of possibilities in the past. As things are, we couldn't have come into existence earlier than we did, but even if we could, we wouldn't think of prenatal nonexistence as the same kind of deprivation as death. And even though our nonexistence two hundred years ago faces us with the fact that our subjective existence is the realization of a possibility grounded in objective facts about the world, this does not affect us as the prospect of our annihilation does. The sense of subjective possibility does not project itself into the past with the same imaginative reality with which it faces the future. Death is the negation of something the possibility of whose negation seems not to exist in advance.

The incongruity between this and the objective view of death is clear; much of what was said about birth and the meaning of life applies here and need not be repeated. My death, like any other, is an event in the objective order, and when I think of it that way, detachment seems natural: the vanishing of this individual from the world is no more remarkable or important than his highly accidental appearance in it. That applies both to the full inner life of the individual and to the objective standpoint itself. Granted that the death of something that exists seems worse than its not coming into existence, it still seems not a matter of great seriousness, considered as part of the general cosmic flux.

Another reason to regard death without too much concern is that everyone's mortality is part of the general cycle of biological renewal which is an inseparable part of organic life. Particular deaths may be horrible or premature, but human death itself is a given which, like the fact that hawks eat mice, it makes no sense to deplore. This is no more consoling to someone about to die than it would be to a mouse about to be eaten by a hawk, but it is another obstacle to closing the subjective-objective gap.

One could try to close it in the opposite direction by arguing that the impersonal standpoint should take its view of each death from the attitude of the one whose death it is. If for each person his own death is awful, then every death should be regarded objectively as awful. Detached indifference would then be a form of blindness to what is clear from an internal perspective—and not the only example of such blindness.

There is something right about this; certainly it would be a good thing if some people took death more seriously than they do.⁴ But if we try to do justice to the fact that death is the ultimate loss for everybody, it isn't clear what the objective standpoint is to do with the thought of this perpetual cataract of catastrophe in which the world comes to an end hundreds of thousands of times a day. We cannot regard all those deaths with the interest with which their subjects regard them: sheer emotional overload prevents it, as anyone who has tried to summon a feeling adequate to an enormous massacre knows. The objective standpoint simply cannot accommodate at its full subjective value the fact that everyone, oneself included, inevitably dies. There really is no way to eliminate the radical clash of standpoints in relation to death.

None of this means one can't subordinate one's life to other things—sometimes it would be indecent not to. People are willing to die for what is external to themselves: values, causes, other people. Anyone incapable of caring enough about something outside himself to sacrifice his life for it is seriously limited. Moreover such external concerns, while they may require the loss of life, often have the effect of diminishing that loss and can even be cultivated for the purpose. The more you care about people and things outside your own life, the less by comparison will be the loss from death, and you can, to a degree, reduce the evil of death by externalizing your interests as it approaches: concentrating on the welfare of those who will survive you and on the success of projects or causes that you care about independently of whether you will be around to see what happens. We see this kind of disinvestment in mortal, individual life all the time—and more ambiguously in the personal desire for posthumous fame, influence, or recognition.

But the effect of these measures should not be exaggerated. There is no way to achieve a fully integrated attitude no matter how much you expand your objective or posthumous interests. The objectively unremarkable death of this creature will terminate both its stream of conscious experience and the particular objective conception of reality in

4. The widespread willingness to rely on thermonuclear bombs as the ultimate weapon displays a cavalier attitude toward death that has always puzzled me. My impression is that whatever they may say, most of the defenders of these weapons are not suitably horrified at the possibility of a war in which hundreds of millions of people would be killed. This may be due to monumental lack of imagination, or perhaps to a peculiar attitude toward risk which leads to the discounting of probabilities of disaster substantially below 50 percent. Or it may be a mechanism of defensive irrationality that appears in circumstances of aggressive conflict. But I suspect that an important factor may be belief in an afterlife, and that the proportion of those who think that death is not the end is much higher among partisans of the bomb than among its opponents.

which its death is included. Of course from the objective standpoint the existence or nonexistence of any particular objective self, including this one, is unimportant. But that is a limited consolation. The objective standpoint may try to cultivate an indifference to its own annihilation, but there will be something false about it: the individual attachment to life will force its way back even at this level. Here, for once, the objective self is not in a position of safety. We may see more clearly, but we cannot rise above death by occupying a vantage point that death will destroy.

The objective standpoint can't really be domesticated. Not only does it threaten to leave us behind, but it gives us more than we can take on in real life. When we acknowledge our containment in the world, it becomes clear that we are incapable of living in the full light of that acknowledgment. Our problem has in this sense no solution, but to recognize that is to come as near as we can to living in light of the truth.