

**MEANINGFULNESS AND ITS LINKS WITH
EXPERIENCE, LANGUAGE AND EMOTION**

Sharron A. Kraus
Magdalen College,
Oxford



A thesis submitted for the degree of D.Phil. at the University
of Oxford
Trinity Term, 1997

ABSTRACT

Sharron A. Kraus
Magdalen College, Oxford

D.Phil. Thesis
Trinity Term, 1997

MEANINGFULNESS AND ITS LINKS WITH EXPERIENCE, LANGUAGE AND EMOTION

In this thesis I am concerned with the possibility of meaningfulness of life: I examine sceptical problems which threaten the possibility of meaningfulness and then outline the basis of an account of meaningfulness.

In part I I look at sceptical problems concerning the relationship between experience and reality. I argue that global scepticism about the external world can be ruled out. I then offer an inclusivist account of reality according to which the relationship between experience and reality is taken to be a relationship of part to whole and a corresponding 'containment' account of objectivity, enabling a number of problems concerning objectivity to be resolved and scepticism about the objectivity of meaningfulness and value to be refuted.

In part II I introduce language into the picture, considering the relationship between language and experience and also that between language and reality. In particular I look at sceptical problems concerning the relationship between our conceptual world and reality.

In part III I argue that there is a need for emotional experience as a basis for meaningfulness of life, that without some deep or emotional experiences, whether personal, religious, aesthetic, epiphanic, or just wonder at the nature of the world, someone's life will have no real grounding; whatever they do will be instrumental - there will be nothing done for its own sake. This leads into the discussion of sceptical problems about the reliability of the emotions and the subconscious.

In part IV I deal with questions specific to meaningfulness including the idea of absurdity, the relationship between life's seeming meaningful and its being meaningful, and the links between meaningfulness and happiness, goodness, understanding, unity and destiny. I argue that the meaningfulness of life is an objective matter and that the possibility of living a meaningful life requires the possibility of there being a right way of life.

ACKNOWLEDGEMENTS

The writing of this thesis has been supervised by three very different philosophers: Bill Child, Ralph Walker and Michael Inwood. Having such diverse feedback has been interesting and helpful.

I have enjoyed and greatly benefited from discussing the subject matter of this thesis with the following: Viva Bartkus, Alex Berman, Dominic Camus, Philomena Edwards, Ben Gilbert, Reid Hoffman, Ward Jones, Scott LaMorte, Richard Sutcliffe, Roger Teichmann and Helen Tookey. I only hope that the result is worthy of them.

As well as the people mentioned above, there are a number of people without whom this thesis might never have been written. I wish to thank the following for their unerring love and support: Ivor Davies, Chris Forester, Steven Galbraith, Jack Harris, Ruth Killick, Debbie Kraus, Jon Kraus, Dylan Owen, Paul Ratcliffe, Moni Sheehan and Rachael Solomon.

CONTENTS

Acknowledgements	i
INTRODUCTION	1
I: SCEPTICISM, EXPERIENCE AND OBJECTIVITY	9
1. Scepticism and the world around us	9
1.1 <i>Cartesian scepticism: Stroud's reconstruction of the argument</i>	10
1.2 <i>Virtual reality - a more threatening possibility?</i>	19
1.3 <i>Underlying reality and incompleteness</i>	28
2. Objectivity	34
2.1 <i>Nagel's view of objectivity as detachment</i>	36
2.2 <i>Subjective truths, subjectively acquired truths and truths with subjective content</i>	40
2.3 <i>Doing away with subjective truth and relativism</i>	43
2.4 <i>Is any detachment good?</i>	48
2.5 <i>Objectivity as containment</i>	53
2.6 <i>Objectivity as mind-independence</i>	60
II: REALITY, LANGUAGE AND EXPERIENCE	66
1. The relationship between language and experience	67
1.1 <i>The emergence of objectivity</i>	68
1.2 <i>From mere experience to experience of the world</i>	70
1.3 <i>Essence and naming</i>	78
1.4 <i>Different ways of seeing the same thing</i>	83
1.5 <i>Conceptual experience and the dawning of aspects</i>	87
2. The relationship between language and reality	94
2.1 <i>Scheme-independent or preconceptual reality</i>	96
2.2 <i>Carving nature at the joints</i>	103
2.3 <i>Creating systems and worlds</i>	109
III: INVOLVEMENT, OBJECTIVITY AND EMOTION	115
1. The importance of emotional involvement	116
1.1 <i>Feelings and objectivity</i>	118
1.2 <i>Bare goods and emotion</i>	127
1.3 <i>The possibility of error about the content of emotion</i>	130
2. Emotion and its relation to evaluative belief	134
2.1 <i>Are emotions rational?</i>	135
2.2 <i>An alternative account of emotion</i>	144
2.3 <i>Emotional knowledge in Aristotle's Nicomachean Ethics</i>	147
3. The causes of emotion and association in emotional life	150
3.1 <i>Induced emotions and scepticism about the evaluative element of emotion</i>	150
3.2 <i>The causes of emotion</i>	152
3.3 <i>Real and projected values</i>	157
3.4 <i>Association</i>	161
3.5 <i>Emotional seepage and transference</i>	165

IV: ELEMENTS OF MEANINGFULNESS	171
1. Meaningfulness, objectivity and optimism	172
1.1 <i>Meaningfulness and absurdity</i>	172
1.2 <i>The relationship between meaningfulness of life and life's seeming meaningful</i>	176
1.3 <i>Meaningfulness and optimism</i>	183
1.4 <i>Is optimism justified?</i>	187
2. The substance of meaningfulness	196
2.1 <i>Arbitrariness of values, concerns and projects</i>	197
2.2 <i>Meaningfulness and unity of life</i>	206
2.3 <i>Meaningfulness, destiny and the right way of life</i>	209
 BIBLIOGRAPHY	 216

INTRODUCTION

‘There is but one truly serious philosophical problem and that is suicide. Judging whether life is or is not worth living amounts to answering the fundamental question of philosophy ... the meaning of life is the most urgent of questions.’

-ALBERT CAMUS, *THE MYTH OF SISYPHUS*

When one asks ‘Is life meaningful?’ what is one asking? This can seem a strange question. It is relatively easy to ask about the meaning of a symbol in a particular system and find out *what* it means. To then ask ‘But is it meaningful?’ seems redundant: its meaning is given, therefore it has meaning; it is a meaningful symbol - within a system of course. “But does it have *intrinsic* meaning; does it have meaning independently of the system?” The answer seems to be no: there are no intrinsically meaningful *signs*.

“What about the system it is part of, does that have meaning?” Some would say that this is a bad question, that it doesn’t make sense. But they would be wrong, too hasty: it does make sense to ask this of some systems. For instance, we can ask whether a formal system can be used for something not contained within it, whether there is an application for the system that yields a semantics as well as a syntax or we can look at a logical system and ask if it accurately mirrors thought. The meaning of symbols is tied up with interpretation - we explain the meaning of a symbol by giving a definition in terms of other symbols.

‘What is the meaning of life?’ cannot be answered so simply. When we ask ‘Is life meaningful?’ we are not asking ‘can we find an *interpretation* for life?’. It is a different kind of question. How are we to approach this question?

There is a sense in which meaning exists in the world, there is natural meaning wherever there is causal regularity: those heavy clouds ‘mean’ rain, those

spots 'mean' measles, etc. There are natural reasons and functions too: eyes evolved because they enabled their possessors to see, colourful flowers evolved because they attracted pollinators, etc. Can the question of what makes life meaningful be answered with reference only to natural meaning, to features of our nature, our biological function, and the causal relations between us and the world? No, this natural meaning is not enough. There is a sense in which the lives of non-human animals are devoid of meaning. Non-human animals and the rest of the natural world, including Mother Nature, the blind watchmaker, lack self-conscious intentionality, and so lack the ability to see themselves as meaningful, to see anything as meaningful, and this seems to be a prerequisite for meaningfulness. The question I ask - 'what makes life meaningful?' - is not a question that can be asked by hippopotami, genes or gravity.

*

What is meaningfulness? What is the phenomenon I am interested in? It seems harder to specify what meaningfulness *is* than to say what it *is not*, to give examples of meaninglessness. I have already mentioned that the lives of animals lacking in self-consciousness seem to be incapable of meaningfulness: animals may live good or bad lives, they may thrive or suffer, but even the most contented animal life still seems to lack meaningfulness.

So meaningfulness is a characteristically human phenomenon. But it is not universal among humans. Even given the requisite self-consciousness, meaningfulness is not guaranteed. With self-awareness comes the possibility of meaningfulness, but also the possibility of meaninglessness, of life's seeming futile, repetitious, empty, etc. Camus uses the myth of Sisyphus as an example of meaninglessness: Sisyphus is condemned to roll a rock to the top of a mountain,

watch it fall back down and then return to roll it back up again, over and over, endlessly. Perhaps the first time he rolls the rock to the mountain top, he experiences a sense of purpose and achievement. But once the struggle has been shown to come to nothing and yet must still be carried on, meaninglessness sets in. For at least some of us, this is an accurate metaphor for the way life is.

Even if our lives needn't be like that of Sisyphus, even if it's possible for our labours to bear fruit, there are still threats. Supposing that one has clear aims and goals, then the meaningfulness of one's life is still not guaranteed. If one's aims are thwarted and one's projects fail miserably, then the meaning will be drained out of one's life. If one succeeds and reaches one's goals but then finds that the endpoint is an empty anticlimax - that the satisfaction one expected is just not there - then one's struggles come to seem meaningless. Or suppose that one lives according to certain principles or assumptions, one lives according to one's religious beliefs, for example. Then if one's principles or religious beliefs turn out to be mistaken then one comes to see one's life as based on a lie.

It might be that in these cases one can at least console oneself with the fact that one has lived the best way one could, but that somehow life got the better of one, that one had sought the truth but had been misled. In that instance, one's struggle could be seen as a noble one. But what if one comes to doubt whether any of the choices one must make could be made in any meaningful way? What if one comes to see one's decisions as arbitrary? If one is presented with many different possible lives, none of which seem inherently better than the others, then the question 'How should I live?' may seem pointless to one. If one encounters admirable people whose lives are based on radically different principles to one's own then one's choice of principles is called into question. In Herman Hesse's

Narziss and Goldmund, the love and mutual respect between Narziss the ascetic scholar and Goldmund the hedonistic artist is in tension with the sense of opposition in their ways of life. Narziss says to Goldmund ““Only now do I begin to perceive how many paths lead us to knowledge, that study is not our only way to it .. I see you by the opposite way, the way which leads through the senses, reach as deep a knowledge as that most thinkers achieve ...”¹. But, Hesse tells us, Narziss also comes to realise that ‘Goldmund had not only enriched his life; he had made him poorer too, poorer and weaker ... This world in which he had his home, this cloister, his learning and his office, the whole well-grounded structure of his thought - had it not been shaken to its base, his faith in it almost destroyed, by his life with Goldmund?’²

Can one even know what would make one’s life meaningful? One may sense that something is missing when one’s life doesn’t seem meaningful, but can one really say what it is that is missing? Will one know it when one sees it? Will one ever come across it, or is it a grail one chases endlessly and never succeeds in reaching? We yearn for things unspecified, our hearts pull us in certain directions, but are these yearnings ever really satisfied? If they are not, then meaningfulness eludes us. But if they are satisfied then surely it seems that we must say that life is doomed again to meaninglessness, since once they have been satisfied there will be nothing further to live for. This seems to show that meaningfulness is impossible.

In this thesis I will cover a number of sceptical and other threats to the meaningfulness of life, concentrating on four main areas. In part I I look at the

¹ *Narziss and Goldmund*, p. 280.

² *Ibid.*, p. 286.

threats that arise as a result of the relationship between experience and reality. In part II I move on to consider problems concerning the relationship between language or conceptualisation and reality. In part III the importance of emotion and the subconscious in matters of meaningfulness is explored and problems arising as a result of the apparently irrational nature of emotion are tackled. Finally in part IV the basis of an account of meaningfulness is given and some problems concerning the possibility of living a meaningful life are resolved.

Some of the questions I ask are necessarily located contextually, being typically existentialist or postmodern, assuming the rejection of absolutism. Once we come to see that absolutism, i.e. the view that there is only one right (God-given) way, one intrinsic structure to the world, is false, that God is dead, as Nietzsche put it, and come to realise that there are many valid views, we can be thrown too far in the opposite direction, to subjectivism and the idea that ‘anything goes’, that there are no rules, and thus that there can be no real meaning. This is an unstable stopping point - one is rendered dumb by the recognition of the limited nature of all one wishes to say, for everything one wants to say is only ‘true for one’. For those who don’t reject absolutism, who never question the norms they adhere to, the problems I am attempting to solve will not even arise. The problems only arise if one accepts the possibility of different acceptable ways. From the acceptance of valid difference there are two options: one either resigns oneself to nihilism or one aims to find something to grasp onto in order to climb out of despair. This thesis can be seen as an attempt at the latter.

*

Is asking questions about meaningfulness itself a meaningful activity? In short, is philosophy meaningful? This question lurks in the background of this study.

When my work is going well I feel enriched by it; when it is going badly I feel despair and wish I hadn't the desire to ask the questions I ask at all and could just leave them unanswered - for they seem unanswerable anyway! - and do something else instead. The question 'what makes life meaningful?' is important to me because I find myself unable to answer the question 'what should I do with my life; should I spend my time doing philosophy?' without answering it, without doing philosophy.

'Why do philosophy?' Why ask this question? Why not just get on with doing philosophy if one is going to do it? Philosophy is caricatured as being a constant asking *why*: where others do, philosophers ask 'why do this?'. The caricature is based on the fact that in doing philosophy one is not just trying to answer questions about reality but also about one's knowledge of reality, so it is essential that philosophical enquiry be self-reflexive, that one asks how one can be sure of the answers generated by this particular form of enquiry. If one did not question one's own philosophising then one would not be a very good philosopher - one would have a blind spot.

When one asks the question 'why do this?' about a certain activity, there are various possible answers. If there is some good intrinsic to the activity then that good can be an answer. If there are aims or ends of the activity then one can answer that one is doing that activity in order to achieve those ends. Or there may be good side effects or by-products of the activity that make it worthwhile. Consider music: if one lives the life of a musician then there are many things to be said in favour of one's chosen vocation. There is pleasure in the playing both for oneself and for one's audience: playing is an end in itself. There are also goals one can work towards and attain, goals which continue to unfold as one's abilities and

sensibilities grow. The benefits can be described in terms of creating something original, developing one's talents, bringing pleasure to people, etc. Because the important ends are not detachable from the process, there is value to be found in the process itself and not just in one final end towards which one is moving.

But philosophy seems different. In philosophy one asks specific questions and tries to answer them, and the answers themselves are not necessarily tied to the method of discovery. For example, if there is an answer to the question 'what makes one a good person?', then it wouldn't matter if one found the answer as a result of philosophising or if one intuitively knew it. Or would it? Is there something about the philosophical method that bears on the answers found using it? If one seeks understanding in order to live one's life well, then perhaps it doesn't matter if one's understanding is not fully articulated. Perhaps it's only if one seeks understanding for its own sake that one must philosophise. In this case having the answers is not enough, one must know that what one knows is valid, and why, and one must have justification.

It seems necessary to answer the question 'how is meaning possible?' as a preliminary to living: in order to get on with living a meaningful life one must first work out what it is for a life to be meaningful. This is the question I aim to answer through philosophy. But if it is through philosophy that one answers this question then philosophy seems to be a means to an end which, once reached, will enable one to live one's life (with or without philosophy). Once one has answered this fundamental question one will be able to see whether a life of philosophy is right for one, or whether one should then get on with something else instead. Once one has answered the fundamental philosophical question it doesn't seem likely that one will feel the need to carry on doing philosophy: once one has answered the

philosophical question that is most difficult and important to one then tackling smaller questions would seem trivial. If this is right then it seems that one would be unable to see philosophy as a lifelong activity, and would view it instead as a tool or a ladder, to be thrown away when one reaches one's goal.

If the answer one seeks is independent of the method of discovery, then the method of discovery - philosophy - has only instrumental value. And this means that philosophy is a doomed enterprise: if the end one aims at through philosophy is truth or knowledge or the answer to a fundamental question such as the question 'what makes life meaningful?', then philosophy cannot be meaningful unless one finds the answer, since to imagine an endless seeking without finding is to imagine a frustrated life. (In contrast, to imagine an endless creating is to imagine a fruitful life.) But if one can find the answer one seeks then philosophy succeeds in its purpose, but yet is still doomed. Like the dolls a child plays with and talks to in the absence of human playmates and then throws away when she grows up, philosophy, having shown us what constitutes a meaningful life, the way to live, will be discarded in favour of this way of life, whatever it is.

This paradox concerning the possibility of philosophy being a meaningful activity is a more specific version of the paradox noted earlier concerning the possibility of any striving being meaningful. I don't hope to be able to solve this paradox here, but do hope to prove to myself that it is only an *apparent* paradox by doing philosophy.

I

SCEPTICISM, EXPERIENCE AND OBJECTIVITY

‘I suppose, therefore, that whatever things I see are illusions; I believe that none of the things my lying memory represents to have happened really did so; I have no senses; body, shape, extension, motion, place are chimeras. What then is true? Perhaps only this one thing, that nothing is certain.’

-DESCARTES, *SECOND MEDITATION*

‘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 observe: nothing seems to have value of the kind it appears to have from inside.’

-THOMAS NAGEL, *THE VIEW FROM NOWHERE*

The idea that knowledge of the world is impossible is the sceptical problem posed by Descartes. The idea that even if knowledge of the objective world is possible, such knowledge would undermine the possibility of life’s being meaningful is posed by Nagel. In this chapter I address these two problems, arguing, first, against global scepticism and, secondly, against scepticism about the objectivity of meaningfulness and value. Both forms of scepticism arise as a result of a certain view of the relationship between the world and our experience of it, a view of experience as (mere) appearance and the world as underlying reality. I will argue against this view and present an alternative.

1. Scepticism and the world around us

If it were to turn out that all the things we believe about the world are false, that the things we attach importance to are not the way they seem, then the meaningfulness of life would be threatened. If one devotes one’s life to the pursuit of scientific knowledge as a physicist and then finds out that there are no physical objects, that the physical world is an illusion, then that would be a terrible blow; if

one devotes oneself to the good of humankind as an aid worker and then finds out that there are not really any human beings, that one is just a disembodied brain in a vat¹, then the purpose of one's life would disappear. If scepticism about the world around us cannot be ruled out, then the meaningfulness of our lives is threatened. I will argue that the threat can be met, that global scepticism is based on a fundamental mistake which can best be pinned down as an ambiguity in the use of the phrase 'the world around us'.²

1.1 Cartesian scepticism: Stroud's reconstruction of the argument

Barry Stroud, in *The Significance of Philosophical Scepticism*, discusses Descartes' sceptical problem, i.e. the problem of 'how we can have any knowledge of the world at all.'³ The problem arises in an attempt to evaluate the reliability of perceptual knowledge, which accounts for a significant part, if not all, of our knowledge of the world. Noting that in some non-standard cases we make mistakes on the basis of our perceptions is not enough to generate scepticism, so Descartes takes what seems to be an example of the most reliable perceptual knowledge - that he is sitting by the fire in his dressing gown, holding a piece of paper - and asks whether this can be doubted. Descartes raises a doubt about this perception on the basis of having had dreams in which he seemed to be sat in his

¹ The brain in a vat sceptical possibility is a modern version of Descartes' 'evil demon' possibility: one is to imagine the possibility that one's brain has been removed from one's body by mad scientists, kept in a vat of nutrients and stimulated in such a way as to make it seem as if one is still living in the everyday world, eating, walking around, talking, etc.

² In the following I use the phrases 'the world around us', 'reality' and 'objective reality' interchangeably, meaning those parts of the world that are mind-independent in the sense that we can be wrong about them. I hope that this is intuitive enough to not require justification. I will argue against the equating of this idea with the idea of the external world in contrast with the internal world of experience.

³ See *ibid.*, chapters I and II.

chair by the fire. From this starting point Descartes argues that we cannot know anything about the world around us.

Stroud sees Descartes' argument as relying on the obtaining of three conditions.⁴ The first condition is that the possibility that one might be dreaming be a threat to one's knowledge about the world. The second is that one must know that one is not dreaming in order to know things about the world. The third is that it is impossible to ever know that one is not dreaming. Do these conditions obtain? And if they do, do they add up to global scepticism, to the conclusion that knowledge about the world around us is impossible?

The first condition is easily seen to obtain: Stroud draws out the plausible requirement that one's experience be connected with a state of affairs in order for it to yield knowledge of that state of affairs. If one dreams that something is happening and it is in fact happening, then one does not have knowledge, since one's dream only matches reality coincidentally. Stroud says 'When you dream that something is going on in the world around you you do not thereby know that it is.'⁵

Whilst accepting the point Stroud is making we should note the unjustified assumption he is making about the nature of the world around one. For Stroud 'the world around one' is just taken as equivalent to 'the external world' and this implicitly denies the possibility of there being other worlds or aspects of the world to be taken into consideration. When one is dreaming, the world around one is not obviously the everyday world. Likewise when one is reading a novel, watching a film, playing a game of chess, etc. It seems a mistake to always characterise 'the

⁴ Ibid., p. 14.

⁵ Ibid., p. 14.

world around one' as the physical world, even when one is dreaming or immersed in a fictional world. This is a kind of chauvinism - only the physical world is real and everything else is second rate. To neglect the richness and variety of experience in this way is to be blind to many of the things that make up 'the world around us'. I shall return to this theme later on. For now what is important is the idea that experiential knowledge of a world requires one to be in experiential contact with that world. Stroud's point is valid - any true beliefs one acquires coincidentally do not count as knowledge.

The second condition is more contentious. Stroud suggests that the need to know that one is not dreaming in order to have perceptual knowledge follows from a requirement of our ordinary conception of knowledge - if one knows p one must know the falsity of all those things incompatible with one's knowing p . Stroud concludes that 'if that requirement is a 'fact' of our ordinary conception of knowledge we will have to accept the conclusion that no one knows anything about the world around us.'⁶ Is this requirement a fact about our ordinary conception of knowledge? This is not obvious: it seems quite normal for us to talk of knowing things without having an opinion either way about some of the consequences of our knowledge. The requirement seems too strong - it seems to require us to have acquired knowledge about knowing that p before acquiring knowledge that p , and if this is universal then we can never acquire knowledge at all. I will not pursue this point though, even if we grant Stroud this condition, there are problems with the third condition and also with the way the three conditions are thought to come together to yield global scepticism.

⁶ Ibid., p. 31.

The third condition to be satisfied if the Cartesian argument is to work is that it is impossible for one to know one is not dreaming. What reason do we have for accepting this? It seems highly implausible. Someone might say that they doubt that anyone could dream that they were sitting in their chair asking questions of the kind Descartes asks and be having an experience identical to the real thing. Not that it is a logical impossibility, but just that dreams are not like that - there are differences between dreams and reality and even if we can't specify exactly what they are, we can recognise the difference, at least when we are awake, and know that we are awake and not dreaming.

Stroud answers this line of thought as follows: 'However obvious and undeniable it might be that we often do know that we are not dreaming, I think this straightforward response to Descartes' challenge is a total failure.'⁷ He argues that there could be 'no certain indications by which we may clearly distinguish wakefulness from sleep'⁸. The structure of the argument is as follows: in order to differentiate between wakefulness and sleep, we require some test for wakefulness which can only be performed successfully if one is not dreaming. In order for such a test to be useful we must know that it is reliable, but how could we acquire *that* knowledge? Knowledge of the test is threatened by the possibility that we were dreaming when we acquired it. In other words, we would have had to have known we were not dreaming when we acquired knowledge of the test for wakefulness and this is precisely what we required the test for in the first place.

This is invalid for two reasons, firstly, it might be that we can distinguish between two kinds of thing without using any such test in the way some people are

⁷ Ibid., p. 19.

⁸ Ibid., p. 19.

able to sex chickens intuitively without appealing to any rule. Secondly, it is assumed that since it would be impossible to gain knowledge of such a test perceptually that therefore it would be impossible to know about it at all - i.e. it is assumed without any justification that perceptual knowledge is the only kind of knowledge we may appeal to. Stroud started off looking at perceptual knowledge, leaving aside the question of whether all knowledge is perceptual,⁹ and now needs to be assuming that the question has been answered and that it is settled that all knowledge is perceptual, or at least that knowledge about the nature of perception must be perceptual.

Suppose we did have such a test. Regardless of how we acquired it, would we then be able to know we were not dreaming? There does seem to be a reliable test practised in order to master lucid dreaming: to have lucid dreams, i.e. dreams in which one knows one is dreaming, one performs 'reality checks' - one asks oneself 'Am I dreaming?', as often as possible during the day with the aim of prompting the same question during one's dreams, and then being able to realise that one is dreaming. Asking the question 'Am I dreaming?' does seem to be a valid test for wakefulness, then.

The Cartesian will then shift ground, acknowledging the possibility of such tests, but suggesting that for all we know they merely differentiate between different kinds of dreams: we cannot be sure that we have *ever* actually been awake. It remains possible that the whole of our life is really just a bigger, more elaborate dream. Stroud says 'Our sensory experience gives us no basis for believing one thing about the world around us rather than its opposite ... The

⁹ See *ibid.*, p. 7 where he says 'It seems possible, then, to acknowledge their importance [i.e. the importance of the senses] and to assess the reliability of that source, quite independently of the

world around us, whatever it might be like, is in that way completely beyond our grasp.’¹⁰ Just as we are unable to rule out the possibility that all we experience is one big dream, we are unable to rule out the possibility that we are brains in vats, or living in a virtual reality world.

But take note of the use of the phrase ‘the world around us’. There is a tension between the description ‘the world around us’ and the idea of a world completely beyond our grasp: how can a world completely beyond our grasp be thought of as being around us? And how can the world around us be completely beyond our grasp? What is ‘the world around us’? The world around us can either be thought of as the experiential world or as the world underlying the experiential world or some combination of the two. Stroud is making a firm division between the experiential world and the underlying real world and equating ‘the world around us’ with the latter, so for Stroud, the world around us is the external world, contrasted with the internal world of experience. He talks about the kind of response to scepticism that suggests we forget about an unknowable world and make do with the world of experience, but says that this would be like a prisoner reconciling himself to life behind bars, clearly settling for second best. Furthermore, he argues, the second best we would be settling for would not leave the world we live in intact - ‘we would not be left with the familiar objects of our everyday experience - tables and chairs, trees and flowers, bread and wine. If Descartes is right we know nothing of such things. What we perceive and are in direct sensory contact with is never a physical object or state of affairs, but only a

difficult question of whether *all* our knowledge comes to us in that way.’ Now, it seems, we are not able to do so.

¹⁰ Ibid., p. 32.

representation - something that could be just the way it is even if there were no objects at all of the sort it represents.’¹¹

The idea is that objects like tables and chairs, trees and flowers, exist in the objective world independently of us (if they do exist and we’re not dreaming) and that our perceptions might be completely erroneous about their nature. Experience is thought of as a veil between us and the world, our experiences are clearly caused by something, but the question of what kind of thing is left completely open. According to this way of thinking, it could be the case that the way we perceive things is systematically in error, it might be the case that the colour spectrum we see is inverted, that the things we see as green are really red, etc. Or it could be the case that the way we perceive things is not systematically related to the way things are at all, that some of the time we see things as they are and some of the time we don’t, it might be that some of the things we perceive as round are really round and others are elephant-shaped! There is no reason to think that the things we group together as being similar on the basis of experience, e.g. the round-looking things, are really similar at all.

Natural things like trees and flowers if they do exist, would have existed regardless of whether human beings or other perceivers had existed. There is a difference between these things and artefacts though - artefacts would not have existed if we had not existed. And yet artefacts have just as objective an existence as natural things. It might make sense to suppose that if the objective world is not the way we think it is, then flowers might not really exist, that the things we see that look like flowers are not really flowers. But does it make sense to suppose that tables might not really exist? Surely what makes something a table is its

¹¹ Ibid., p. 37.

having been designed to function as a table and its being used as a table? We might be completely mistaken about the nature of flowers - their nature is something we must investigate, a scientific matter. But with tables there is no possibility of a complete gap between what they are and what we take them to be. This can be loosely characterised by saying that the essence of flowers is biological whereas the essence of tables is functional. Tables and other functional objects couldn't exist independently of humans (or other artefact creators and users) and the things we use as tables are, by our using them as tables, real tables. What makes something a table is its functioning as a table for us, whatever its underlying structure.¹²

The assumed veil of perception that prevents us from knowing the underlying reality does not take away *all* 'the familiar objects of our everyday experience', then. Not being able to know the underlying structure of everyday things does not lead to global scepticism. This might be put as follows: the world around us is not to be taken as equivalent to the world underlying our experiences, since some of the things that make up the world around us would still exist if the world underlying our experiences was different. But denying that the world around us is equivalent to the reality underlying our experiences is not to embrace idealism - we can and must deny that the world around us is equivalent to the world of experience as well - what we experience is just the tip of the iceberg. The world around us, the real world, must be seen to straddle and include both experience and what underlies experience.

¹² If this is not convincing, then imagine there to exist naturally structures physically identical to tables. Would these be tables? Surely not unless there were creatures of some sort using them as tables. A tree that has fallen across a river is not a bridge unless used as one.

The sceptic could reply that the reality of tables would be threatened if the reality of the physical world is threatened. If it is possible that there are no physical objects, then surely this rules out the possibility of there being any real tables. If one is dreaming, or is just a brain in a vat, then not only are there no tables but there are no people. If it is possible that one is dreaming or is a brain in a vat, then one should not conclude that tables exist but are much weirder than one thought, one should conclude that tables do *not* really exist. Is this right? I will argue that we should say that there are tables and people regardless of the nature of the underlying reality.

Suppose it turns out to be the case that we are non-physical beings, that the world around us is non-physical. If this were the case then it would turn out that our beliefs about physical objects qua physical are wrong - beliefs such as the belief that trees are made of matter would be false, for example.¹³ But would our beliefs about ourselves and our artefacts be wrong? No: claims like the claim that tables are often found in offices and that generosity is a virtue would still be true. Unless we are given a reason to think that the discovery that our artefacts are not physical would be the discovery that they are not real, such a possibility would leave a large part of our world intact. Equating 'real' with 'physical', as philosophers often do, requires justification. The mere fact that the world is physical (if it is) is not enough.

The possibility of finding out that the world is non-physical is not radical enough to generate global scepticism - we would be finding out that we were radically mistaken about the nature of the world around us, but would not be

¹³ It might even be more accurate to say that in this situation it would turn out that trees, water, and other natural kinds didn't really exist.

finding out that the world around us didn't exist. What would happen if one were to discover that one had been dreaming all one's life, or that one was a brain in a vat? Suppose one were to wake up in a strange place and realise that one's life till then had been a dream, that would be to find out that none of the things one had thought had happened to one had actually happened. Or would it? Not necessarily. If the world one woke up in was devoid of any other life, completely empty, or if it was so chaotic and incomprehensible that one couldn't make any real life for oneself, and the dream world one had lost had been rich and consistent, then one would long for the dream world one had lost.¹⁴

1.2 Virtual reality - a more threatening possibility?

The dream possibility, although based in reality, doesn't seem a threatening enough sceptical possibility - one's dreams and reality are not easily confused. The brain in the vat possibility, although set up to be undetectable and so more damaging than the dream possibility if it is convincing, does not have a firm enough basis in reality. There is a more plausible example the sceptic can use - the possibility of virtual reality. As virtual reality technology becomes more sophisticated, the difference between some real experiences and their virtual equivalents becomes negligible, the experience of flying a real plane and the experience of 'flying' a simulation is a good example. It might seem unlikely that a simulation could ever be as realistic as real food, real sex, real walks in the mountains, etc., but the possibility must be taken seriously. So the sceptic can ask: how can we be sure that we're not just living in a virtual world?

¹⁴ Cf. H.G. Wells's 'The Door in the Wall' in which the paradise behind the door that the narrator finds and then loses is what is most real to him, his life in the everyday world being a poor

The more convoluted the causal connection between oneself and a part of the world, the greater the possibility of error or deception. Whilst it currently seems farfetched to suppose that it's possible that the person one is sitting next to might not really be there, there is a real possibility that the person one thinks is on the other end of a telephone is not. If the line is bad and the conversation is short, it's quite easy to mistake one voice for another, to be deceived into thinking one is talking to one's best friend even. And it's possible for a machine to do the fooling, for there to be no one there at all. So the sceptical possibility is a real possibility in particular cases. Whatever we experience indirectly, over the telephone, on the radio, etc., might be a deception. And if technological developments come to make direct experience and simulation indistinguishable, then it would become true that whatever we were experiencing might be a deception. And then, if Stroud is right about the dream argument, that would lead to the possibility that the real world is in fact just another virtual world.

Many of the things we currently experience *are* simulations. Many of the things making up the world around us are not any more 'real' than the best virtual reality possibilities, we are just used to them and now accept them as part of life. Listening to a radio broadcast or a recording of an orchestral performance is experiencing a simulation. Talking to someone on the telephone is having a virtual conversation with them. One day it might become possible to have virtual sex with someone on the other side of the world. Or with a simulated person. It might become impossible to tell the difference between having sex with a real person and having virtual sex with a computer-generated virtual person.

Some simulations are unproblematic: talking on the telephone is not thought of as unreal, watching the news on television is a reliable source of information, etc. Similarly, long-distance virtual sex could become acceptable - it would be a way of relating to a real person. The difference between virtual reality of this kind, with the emphasis on 'reality' and the idea of plugging into a self-contained virtual scenario in which one only 'relates' to virtual 'people' is significant - we don't think 'having sex' with a piece of software is really having sex at all. There seems to be a significant difference between a virtual chain which links one to a part of reality and a virtual reality situation which is completely virtual, then.

What is it about the real world that makes it real, if it is real? There are a number of popular answers to this question, one being that to be real is to be physical and another that to be real is to be experienced. Also, there is the view according to which there is some objective fact of the matter about what is real and what isn't, even though it may not be possible for us to know the truth. I won't attempt to answer this question decisively but will draw attention to an important feature of reality.

Reality is what matters. This seems an undeniable truth, although the question of whether something is real because it matters or whether something matters because it's real seems hard to answer. Why does the truth matter? Is Truth an objectively real property? The answer must be 'yes, in some sense': truth is real - certain beliefs really are true and others really are false. The question 'what makes truth real?' might be answered in two ways, either by saying that truth is part of the fabric of the universe or by saying truth is a necessary concept

(for us). A third answer that ties the two together is plausible, saying that truth, like anything else fundamental to our lives, i.e. anything that matters, is part of the real world, since the real world contains us and our concepts are formed in response to it.

One of the things which differentiates between reality and virtual reality is that reality matters to us and virtual reality doesn't - virtual reality is thought of as a trivial game, a waste of time. And if a form of virtual reality becomes important to us, like telephone conversations and radio broadcasts, then it becomes accepted as part of reality. What currently differentiates reality from the computer generated virtual reality worlds we have is that reality has a depth, richness, complexity and consistency that the best virtual reality games or simulations lack. Computer games are less real than a good book or film in spite of being interactive.

One might think that the fundamental difference between the real world and virtual worlds is that the real world is objective and virtual worlds are not. But this cannot be true: although virtual and fictional worlds are dependent on us for their creation, as are tables and chairs, there are truths about characters and situations within virtual and fictional worlds which obtain regardless of whether we know them to obtain.¹⁵ In the most basic of virtual worlds there are objective truths about what happens to you if you do certain things, e.g. if you take the right hand path from the entrance you encounter a two-headed monster.

Suppose we were able to create a complete virtual world such that the people who entered it would be able to forget that they were in a virtual world.

¹⁵ For an excellent discussion of objectivity within fictional worlds see Ronald Dworkin's 'How Law is Like Literature'. Dworkin quotes John Fowles: 'Only one reason is shared by all of us

Suppose a person's bodily needs could be taken care of or removed, so that they wouldn't have to keep leaving the virtual world and coming back to the real world. If someone took the option of living in the virtual world, would this be a valid life? There are two conflicting intuitions about this - one, that such a life would be a waste, that they would have lost touch with reality and would be living a lie, a fantasy. Robert Nozick appeals to this intuition when he argues that a life plugged in to his 'experience machine' would not be a good life.¹⁶ The other, that if the virtual world has enough sophistication and depth, then it may well be a valid world, just as the world of the chessboard is a rich enough world for some people to live in.

Some might think that a chessplayer who spends all his days cooped up playing chess, not getting out and meeting non-chessplayers, not getting to see the world, etc., is wasting his life. But this is unwarranted - chess is surely not a waste of time. On the contrary, the single-minded, focused life of the dedicated chessplayer is admirable. If the chessplayer's world is a good world, then this is an amazing thing - that with the help of just a square board and thirty-two pieces we can create a whole world! If it takes so little to make a real world, then virtual worlds good enough to be real are certainly possible.

[novelists]: we wish to create worlds as real as, but other than, the world that is. ... We also know that a genuinely created world must be independent of its creator.' (p. 157).

¹⁶ In *Anarchy, State and Utopia*, pp. 42 ff. Nozick is not considering a virtual world of quite the kind I am considering - in his example you would be able to choose what to experience and then 'superduper neuropsychologists could stimulate your brain so that you would think and feel you were writing a great novel, or making a friend, or reading an interesting book'. What could possibly be an experience as of reading an interesting book which wasn't actually an experience of reading an interesting book is baffling - surely if the book has been written by someone at some point in time, and then the book has been translated from hard copy into a suitable form for your brain to be fed directly, then the unorthodox method of getting the story is still a real and valid experience. And if not, then what do we say about the experience of being read a story or listening to a radio broadcast of a story? Maybe one is supposed to be able to experience a 'raw feel' which is somehow 'just like' reading a good book.

If one had the choice between living in a barren world with no intelligent beings and living in an exotic world populated with interesting creatures, the latter would seem the best option. Even if the first is the future of the world we know - the physical world - and the second is a possible virtual world. Imagine that due to environmental damage and technological developments, our earthly world became a ghost world, no life remaining on it except in that everyone still alive had a physical body in this world, but was wired up to the new virtual world permanently.¹⁷

Someone might object that there is an important difference between this virtual world and the sceptical possibility that one is living in a virtual world or is a brain in a vat: there would still be other human beings for one to interact with in this world, and that makes the difference between a real life and a deception. In virtual worlds like those internal to popular computer games where one interacts with characters, shooting them willy-nilly, etc., one is not genuinely relating at all. Suppose, though, that one could survive the death of one's body and continue one's 'life' as a cyberbeing. Then oneself and one's 'friends' would be nothing more than characters within the virtual world. Would one have reason to mourn one's loved ones? Or wouldn't one just get used to them as fellow cyberbeings?

There is a gut feeling we have that staying on earth, or maybe moving to some other planet in the universe, is preferable to and is a more real existence than living in any virtual reality world. Whether that feeling can be allayed depends on how much our existence depends on our actual world. Whether it can or cannot, global scepticism is defeated: if we conclude that the best virtual reality world

¹⁷ To anyone who believes in an afterlife, this should seem an attractive option; this nearly non-physical reality must surely be closer to heaven than our earthly life.

would still lack something that the natural world offers us, then we must reject the sceptical possibility, that for all we know the natural world just is one big virtual world. If we conclude that a good virtual world could become real, then we must also reject the sceptical conclusion that the world might turn out not to be a real world.

My approach to global scepticism shares some features with Hilary Putnam's. I will clarify my position with reference to his. Putnam presents an argument against the brain in a vat possibility in *Reason, Truth and History*.¹⁸ The argument is a transcendental argument concerning the preconditions of reference and thought and is as follows: the things (at least some of) our words refer to must be accessible to us - there must be a causal or other empirical connection between a word such as 'tree' and what it refers to. If we were brains in a vat then we wouldn't be able to refer to real trees or real vats, since there wouldn't be any connection between us and trees and vats in the external world. So, if we were brains in a vat our words 'tree', 'vat' would have to refer to trees and vats in the vat-world. Thus if a brain in a vat were to say 'We are brains in a vat' it would have to mean 'We are brains in a vat in vat-world'. But this is false, since ex hypothesi in vat-world we are human beings in a normal world. Putnam's conclusion is 'if we are brains in a vat, then 'We are brains in a vat' is false. So it is (necessarily) false.'¹⁹

This argument cannot be valid, though, since if it were then we wouldn't be able to talk about unexperienced underlying reality at all. If it does make sense to speculate about how things beyond our experience are, then Putnam's argument is

¹⁸ See chapter 1.

¹⁹ Ibid., p. 15.

to be rejected. We can refer to underlying reality quite easily, e.g. by using the expression ‘whatever causes these experiences’.²⁰ Putnam presents an argument against the possibility of underlying reality being radically different to the way we imagine it because he doesn’t see that a threat to underlying reality needn’t be a sceptical threat to reality. The way he presents the sceptical possibility is telling:

‘Instead of having just one brain in a vat, we could imagine that all human beings ... are brains in a vat. ... let us suppose that the automatic machinery is programmed to give us *a collective hallucination*, rather than a number of separate unrelated hallucinations. Thus when I seem to myself to be talking to you, you seem to yourself to be hearing my words. ... when I produce my words, what happens is that the efferent impulses travel from my brain to the computer, which both causes me to ‘hear’ my own voice uttering those words and ‘feel’ my tongue moving, etc., and causes you to ‘hear’ my words, ‘see’ me speaking, etc.’²¹

Isn’t this the kind of thing that happens in reality anyway? We may not be brains in a vat, but we are collections of particles in a whirl of causation. If we replace Putnam’s ‘automatic machinery’ with ‘the laws of nature’ we have a strangely worded but nonetheless accurate description of human interaction. If we don’t need to use the word ‘hallucination’ and scare quotes in talking about our perceptions in the world as we know it, then why should we do so in the brains in a vat scenario? Putnam goes on to say ‘In this case, we are, in a sense, actually in communication. I am not mistaken about your real existence ... you do, after all, really hear my words when I speak to you, *even if the mechanism isn’t what we*

²⁰ Putnam would probably reply with an attack on the reference of the word ‘cause’, saying that by the word ‘cause’ we refer to causes in the world - so even in a vat-world it would be true to say that trees cause tree-experiences. I agree, but there are different causal explanations compatible with this one. As well as saying that tree-experiences are caused by trees we say that tree-experiences are caused by certain retinal stimulations followed by certain brain states. These two different senses of ‘cause’ make it seem perfectly reasonable to include talk of the relationship between what the scientists controlling the brains in a vat do to the brains and the brains’ experiences under the heading of causality.

²¹ Ibid., p. 6-7.

*suppose it to be.*²² What this should lead Putnam to see is that the brains in a vat possibility is not a global sceptical possibility.

To summarise then, the argument I have presented started off acknowledging the possibility that the world around us might be a virtual world (or a dream, or induced in us by scientists) and then considering whether that would mean that nothing was real. I argued that the discovery that the reality underlying the things in our world is not what we thought it was would not be the discovery that the world around us doesn't exist, since we shouldn't equate the world around us, or reality, with *underlying* reality. I stressed that this is not to equate reality with experience: underlying reality and experience are parts of the whole of reality.

My view of reality is of an inclusive whole: reality is neither contrasted nor equated with the world of experience, but is thought of as the totality of which our experience gives us access to parts.²³ This effectively dissolves the realism/idealism debate. Consider an example: suppose we see something that looks like an elephant. Suppose that the perception is actually caused by a dog and that the light travelling from the dog gets distorted in some strange way, like reflections in curved mirrors. Suppose this is what happens whenever we seem to see elephants. When we seem to be perceiving elephants, in this case, are we perceiving elephants are we perceiving dogs? An idealist would say that we are perceiving elephants whereas a realist would say that we are really perceiving dogs but our perceptions are misleading. The idealist leaves out an important part of the

²² Ibid., p. 6, my italics.

²³ The view I develop shares important features with F.H. Bradley's view of the relationship between appearance and reality. For Bradley, absolute reality is such that 'it contains everything phenomenal in a harmonious form' and appearance is only defective in its incompleteness. (See

story - the relationship between elephant-shapes and dogs, whereas the realist only sees the first part of the story as part of the nature of things. According to the inclusivist, we can't leave out either side of things if we are to fully understand reality.

1.3 Underlying reality and incompleteness

One could assert the existence of an unknown underlying reality for one of two reasons: the first would be due to a realisation of the possibility of gaps in our knowledge of the world, a realisation that there is more happening than our everyday experience of the world reveals, the second would be the drawing of a fundamental distinction between the way the world seems and the way the world is. Global scepticism relies on a fundamental appearance/reality distinction, i.e. not just on there being some part of reality that is actually unknown, but on there being a part of reality that is in principle unknowable.

Someone who misunderstood the sceptical point might think that once we filled in the gaps in our scientific understanding of the world, the sceptic would be defeated and would concede that we do have knowledge of the real world. This would be to understand scepticism as pointing to a particular flaw in our current state of knowledge. But the sceptic's claim will not disappear no matter how much knowledge of the world around us we acquire, for the realm we lack knowledge of is not within reach of us, no matter how extensive the breadth and depth of our knowledge, it is in the wrong zone. The sceptic may as well say "Find out all you like about this realm of appearances, you will not be one jot

Appearance and Reality, chapters XIII., XIV. and XVI.) I part company with Bradley on the question of whether reality is to be equated with experience.

closer to knowledge of the real world.” By definition, anything we can find out about is in the realm of appearance, so we may as well give up trying.

But just as it doesn't make sense to talk about knowledge if there is no room for error, if there is no room for finding out the truth behind appearances, if we are *always* in error, then there is no sense to be made of the idea. There must be a gap between what we think is correct and what is correct for objectivity, but if the gap is so wide as to prevent us from ever having access to what is correct, what is true, then the truth is not guiding our behaviour at all. The underlying reality that our knowledge is supposed to be answerable to has no more of a claim on us than some set of happenings on Mars. Our experiences and thoughts are no more about these hidden things, these aspects of 'underlying reality' than they are about happenings on Mars. They are about *these* things, the things we experience, whatever the whole truth about them may be.

Suppose there to be beings who inhabit or have access to the unknowable reality: what is their relationship to the world around us? They might see it as a deception, they might be able to see that what we see and what they see are two different things. What they see is that the real patterns governing the universe are nothing like the patterns we see. But the problem is, we are unable to see the patterns they see. Why is this? Either we are completely cut off from them - there is no causal or other linking connection between their world and ours - or anything they tried to communicate to us would fail because of our lack of ability to understand things at their level. If the former, it just seems wrong to say that they are in the world at all: they are in another world and what happens in that world is nothing to do with us - we can be interested in the existence of their world as

explorers are interested in the possibility of life elsewhere, but we should dismiss their claim to know better about our world than we do. If the latter, then either they view us in the way we view less intelligent animals - there is a world beyond us that we can't reach, but the crude, unsophisticated judgements we do make are not useless. Just as frogs can know enough about the world to catch flies, our crude judgements are in harmony with reality, it's just that there are huge gaps to be filled in. In this case, the real world is not completely out of our reach. Or if this is not the way they view us, then they view our experiential world as unintelligible. It is not just that we wouldn't be able to understand their world, they wouldn't be able to understand our world either.

There are two kinds of failure to understand, the first, failure of a less intelligent being to understand something a more intelligent being understands. Second, failure of interpretation on both sides. In the first case we can be comforted by the fact that we are on the right lines (even though the thought that we might have so much further to develop epistemically would be daunting). In the second case, we have no reason to think that the other beings are right about our world and we are wrong.

We must always allow for the possibility of being in error about the nature of the world. But the move from this possibility to the possibility of there being some things that we can never know is invalid. Even if there must always be some unknown, there is no reason to think that there must be an Unknown, i.e. an unknowable, which can never be known. We can be reassured that we are not completely cut off from the world, that the best characterisation of 'the world around us' will include some of the things we experience, as we experience them.

But even if we have shown that what we experience is real, it could be that reality as a whole is much vaster than our experience would lead us to think, that our experience of the world is a tiny fraction of the whole.

Even if not, even if the world is not that much more complex than we take it to be, it is impossible for us to ever know all of reality. The impossibility is logical not merely physical. It's not just that we are finite knowers and limited by the capacities of our brains, there is more to it than that. The problem is that knowing everything about reality would mean knowing everything about everything known about reality. The problem is one of inescapable incompleteness. Unless knowledge is self-justifying, then one lacks the knowledge that one's knowledge is justified in at least one case. For example, suppose I want to check that my eyes are not deceiving me about the colour of the bird in the tree in front of me. If I ask someone else to report what they see, and they tell me that they see a black bird in a tree and they see me looking at the bird, then, if their description matches mine, I will be able to conclude that my eyes are not deceiving me, that they are working as they should be. With the help of another person I have been able to construct a meta-view of my relationship with the world, a view with myself in it, I have been able to see myself in a mirror and thereby step outside of myself, or expand outwards. So perhaps I now know all the things about the earlier me, but I have changed in gaining more perspective, and so I don't know everything about myself, as I am now. It is like trying to catch your shadow, with every step towards complete knowledge, you create new unknowns.

It might be thought that reality is what is given 'from nowhere', that the truth about the world is given by God or physics. Why is this wrong? It is not that there cannot be such a description - it could be something like this: 'There are atoms, there are forces, such as gravity, there are solid objects composed of groups of atoms arranged in molecules, there are properties possessed by these objects such as mass and size, ...' But the important question is not whether such a description is possible but whether it would be the definitive truth about the world if it did exist. If God were to reveal the description to us would we have grounds to accept it as the whole truth about the world? Unless we define 'real' as 'what is accessible to God' there is no reason to think that we have anything other than one more viewpoint on the world. Similarly with science: there can be no reason for thinking that science tells us what is real other than a prior belief that this is what science does.

Suppose that God were to reveal himself as a mischievous demon, saying "I can't keep it up any longer, I have to break it to you all: Descartes and Berkeley were on the right lines - the physical world is not what it seems; well, some of it is - there are people, values, colours, etc. but all the things you think of as mind independent, all the physical properties of things, atoms, electrons and such - all the things physics tells you - these are just appearances I conjured up for you. From my vantage point, everything is really just like a film, images that flit in front of my eyes. It keeps me entertained for hours, for aeons even!" What should our reaction be? Should we then throw out physics? A modest physicist might argue that no matter that there are really no atoms, no nuclei, no electromagnetic forces, etc., still these fictions are useful. We shouldn't eliminate talk of these

things in favour of talk of non-physical things but can continue to rely on physics as a predictive tool since the predictions physics makes help us to get around in the world.

Suppose that then an evil demon appeared saying “Don’t listen to him, he’s making it all up. I am the one who created all this. There are physical properties - of course there are - but the thing you should know is this: everything is really the opposite of the way it seems. This of course means that the world is contradictory. Ha ha ha!” At this point we would either despair completely or start to question the authority of these demons. We might have been prepared to believe the first demon - his revelation might have been accommodated in some way without bringing our world crashing down around us. (It might in time have come to seem no more radical an adjustment than any other paradigm shift.) We can’t accept what the second demon says without going mad, though. And if we start to doubt the authority of the demons then we will start to wonder whether we have more cause to believe any of them than to believe our own eyes. The question will still be “How do we make sense of the world?” and even if the first demon is being honest with us there is a third option - we don’t have to choose between believing him and seeing our own views as radically deficient or trusting our own views and so disbelieving him. The third and most sensible option is to take his view as one among many and try to reconcile them all.²⁴

²⁴ As a child one may be struck by the thought that one could never know what one was really like, because if one asked other people, one would not know whether they were being truthful; if they said nice things they were probably being tactful, if they said nasty things they might be trying to be hurtful. Then one might wish one could listen in on other people talking about one without them knowing one was there, thinking that then one would find out the truth. But as one grows up and interacts with people with diverse outlooks and opinions, one realises that no amount of listening in would give one The Truth - not because people are never honest but because there is no one truth. Thus the general scepticism is dissolved and the only remaining curiosity is the desire to know what certain people think of one.

The point made above can be summarised in the following way: there is no greater authority on the nature of reality than us; there is no ultimate viewpoint from which all other viewpoints can be rendered false. This is not to say that we can know everything about reality, but just that what we do know must be a valid part of the whole.

2. Objectivity

There are two notions of objectivity often thought to be equivalent or intimately related: the idea of objective reality as mind-independent reality of which our beliefs may be in error, and the idea of objective reality as contrasted with subjective appearance. There are two sorts of mind-independence which are run together when these two notions are confused: the first is an epistemic idea - mind-independence is supplied by the existence of an epistemic gap between the way things are and the way things seem to us: something is mind-independent if it can exist or have certain properties independently of whether we know it to exist or have those properties. The second is a more demanding ontological mind-independence - something is mind-independent if it exists completely independently of minds, if its nature can be understood without reference to minds. If objectivity is thought of as the first kind of mind-independence then all kinds of things will be contenders for objectivity - physical objects, colours, numbers, tunes, meanings, experiences, memories, and so on. There seem to be objective facts of the matter to be found out about the nature of most things, including those things thought of as subjective in the sense of being experiential. Unless we believe that a person has infallible and complete knowledge of their mental states,

which is implausible even of current mental states, then experiences will have objectivity. On the other hand, if objectivity is thought of as the second kind of mind-independence then only the underlying physical properties of things will be objective. This would lead to a distinction between the primary qualities of things, which are definable purely in terms of their effects on physical objects, and their secondary qualities, which are only definable in terms of their effects on perceivers, and the idea that only primary qualities are objective. If we equate objectivity with this kind of mind-independence then objective reality is contrasted with subjective appearance and so experience is denigrated.

If these two ideas of objectivity are conflated then it can seem impossible to deny the subjective (appearance)/objective (reality) distinction without losing objectivity completely, without falling into idealism. Once the two ideas are separated, however, it becomes clear that we can accept one and reject the other. I eschew the idea of objective reality as contrasted with subjective appearance and will explore the remaining idea of objectivity as epistemic independence. This will be the kind of objectivity I accept throughout. I will take the account of objectivity given by Thomas Nagel in *The View From Nowhere* as a starting point because I see Nagel as being at a philosophical cross-roads: he aims to give an account of objectivity which doesn't invalidate the subjective, but he is still in the grip of the subjective (appearance)/objective (reality) distinction. If a satisfactory account of objectivity can be given without reference to this distinction, then we can offer Nagel what he wants - a realist view which embraces the subjective.

2.1 Nagel's view of objectivity as detachment

We ascribe objectivity to aspects of the world and also to viewpoints. If we know how to determine which aspects of the world are objective then we can define a viewpoint as objective if it gives us access to truths about (some of) those aspects of the world. Conversely if we can determine whether or not a viewpoint is objective then we can define objective aspects of the world as those aspects of the world available to an objective viewpoint. Nagel starts with viewpoints and talks of the objective and subjective facts available to objective and subjective viewpoints respectively. He defines objectivity of viewpoint without reference to objectivity of reality but aims to be answerable to our intuitions about objective reality. *The View From Nowhere* is an attempt to defend and critique objectivity of viewpoint: objectivity is useful and important since it enables us to transcend our particular viewpoints, but it must be kept in check lest it trample things of importance. Nagel says 'A great deal is essentially connected to a particular point of view, or type of view, and the attempt to give a complete account of the world in objective terms detached from these perspectives inevitably leads to false reductions or to outright denial that certain patently real phenomena exist at all.'²⁵

I will argue that this is wrong, that it is not objectivity which must be kept in check, that objectivity of viewpoint properly understood poses no threat to the things Nagel sees as important. I will attack Nagel's main characterisation of objectivity as detachment but will then show that there is a better characterisation of objectivity to be found in *The View From Nowhere* - objectivity as containment.

Nagel defines objectivity of viewpoint as follows: 'A view or form of thought is more objective than another if it relies less on the specifics of the

individual's makeup and position in the world.'²⁶. Nagel sees this as a good definition because it makes the difference between subjective and objective a matter of degree - we get a more objective view as we become more detached from our own particular idiosyncratic view. Another good thing about it is that it doesn't automatically link degree of objectivity with degree of reality or truth. A more objective viewpoint is not necessarily better than a less objective one, a more objective view doesn't automatically override a more subjective one. Nagel is keen to stress the importance of the subjective.

If objective reality is thought of as completely mind-independent and contrasted with subjective experience then the idea of objectivity of viewpoint as detachment seems intuitive - in order to get at underlying mind-independent reality we slough off the particularities of our subjective viewpoints. If we think of objective reality in this way, though, we are on our way to the false reductions Nagel warns against, and, I will argue, if we don't think of objective reality in this way then detachment doesn't seem a natural way of thinking of objectivity of viewpoint.

There is a problem with Nagel's characterisation: either degree of objectivity is related to degree of reality or not. If it is, then Nagel's attempt to hang on to the validity of the subjective is doomed. If it is not, then we have a problem trying to justify objectivity at all: if objectivity is detachment but is not related to reality and truth then why is objectivity to be valued? If objectivity is to be taken as a normative constraint on viewpoints then Nagel's characterisation of

²⁵ *The View From Nowhere*, p. 7.

²⁶ *Ibid.*, p. 5.

objectivity as detachment either fails or it succeeds at the expense of the subjective.

It would be strange to think that it is primarily viewpoints, not aspects of the world, which are objective or not, that the objectivity of things in the world is a function of viewpoint. This seems to get things the wrong way round - the reason we care about objectivity is that we link objectivity with the way things really are, with the truth. If a viewpoint is objective then surely this is because it is likely to yield the truth about the world? That is, the question 'What makes a view more objective?' is answered by saying 'A view is more objective if it is conducive to presenting the truth about objective reality.' And the question 'What is objective reality?' is not to be answered by saying 'Objective reality is what is given to us by the most objective view'. In other words, whatever objectivity of viewpoint is, detachment, for example, it ought to be truth-conducive. Objective reality is prior to objectivity of viewpoint: we value detachment or neutrality of viewpoint where we do because we think that attachment or bias blinds us to the truth, i.e. because we equate the subjective with the biased and don't see it as truth-conducive.

For Nagel, objective features of the world are those features of the world available from an objective point of view. He contrasts objective facts about one's life with the subjective facts available to one 'from the inside': '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.'²⁷ How are we to understand this? What are subjective facts? If we conclude, as he wants us to, that the subjective view as well as the objective view yields facts or truths, that whilst it is

objectively true that my life is pointless, it is also subjectively true that my life is important, then we are left with a form of relativism. If the objective truths of the matter do not override the subjective truths, then we have two sets of truths each relative to viewpoint. This makes the word 'objective' unimportant in the question 'Is x an objective feature of the world?' because x is to be thought of as real even if x is only a feature of a subjective view of the world. So either we accept that subjective views are just as valid as more objective ones, in which case, objectivity loses its value, or we value objective views over subjective ones insofar as they are more likely to yield truth about objective reality. Unless there is reason to think that a more objective viewpoint gives us a better, e.g. less biased, view of the world, there is no reason to think that how objective a viewpoint is, i.e. how detached it is, is epistemologically relevant.

Nagel wants objectivity to be normative, i.e. he wants it to be a standard by which to measure viewpoints against, but he doesn't want to have to apply it critically to the subjective, to the experiential, the ethical, etc. As a result he is forced to seal off subjective views from more objective views and to fall back on subjective truths, which are said to be just as important as objective truths. I will now attempt to make sense of the idea of subjective facts or truths and argue that it is an unhelpful notion and that unless we accept relativism about truth we should think of all truth as objective. This will not inevitably lead to the feared denigration of the subjective though.

²⁷ Ibid., p. 209.

2.2 Subjective truths, subjectively acquired truths and truths with subjective content

The distinction Nagel implicitly draws between subjective and objective facts is a distinction between the things presented to one when one takes a subjective view and the things presented to one when one takes an objective view respectively. From a subjective view - from the inside - things are presented as being one way, whereas from a more objective view - from the outside - things are presented as being another way. How are we to characterise the truth status of the presentations of subjective and objective views? Nagel disassociates himself from the tradition of equating the subjective with appearance and the objective with reality, that is, he rejects the idea that all our experiences are merely subjective and so incapable of presenting us with truths about reality. Experience, human standpoints, subjective views in general must be thought of as capable of presenting us with truths, according to Nagel. But not with truths with the same standing as those presented by objective views such as physics.

Should we think that where there is a difference in the way truths are ascertained, e.g. between truths ascertained subjectively and those ascertained objectively, there must be a corresponding difference in truth status? Must truths ascertained subjectively be seen as (only) subjectively true and those ascertained objectively be seen as objectively true? I suggest not: I suggest that we distinguish between the way truths are ascertained and their status.²⁸

Nagel fails to separate the way a belief is acquired from its truth status in this way and as a result confuses his main view according to which the conflict

²⁸ This parallels the distinction between knowledge and belief - this distinction doesn't make the truth status of true beliefs lesser than the truth status of knowledge.

between objective detachment and subjective involvement is real, i.e. the truths acquired subjectively and those acquired objectively are both competitors for truth, with a form of relativism according to which subjective truths are those which are true relative to a subjective viewpoint and objective truths are those which are true relative to an objective viewpoint and the two do not come into conflict.²⁹

We should differentiate between subjective truths, which are thought of as truths subjectively true, and truths ascertained subjectively. Another confusion is that between subjective truths and truths with subjective content. Truths with subjective content are to be contrasted with truths with more general content. Truths which apply to me alone, including truths about how things seem to me, need not be subjective in either of the previous senses, the truth that I love a particular person is an objective fact ascertainable by anyone who observes me for long enough, for example. If there are truths about what I should do, moral truths about me in particular, then they are of this kind, they are objective truths with subjective content. Diverse cultural beliefs are sometimes of this kind - it is objectively true that some of the things that are good things for us to do are not good things for those of a different culture to do. To describe things in this way is to do away with the need for relativism in these cases: by acknowledging the existence of truths about human beings which are not general truths about all human beings we avoid having to say that truths about what is good for someone are only true relative to their culture.

²⁹ Nagel does not say anything about the relationship between his view and a relativist view of the kind I am considering. The only thing he says about relativism in *The View From Nowhere* is that it cannot provide a solution to the subjective/objective collisions concerning the meaning of life. He says that relativising truths to viewpoint will not help to dissolve the problems since both viewpoints are available to one and the same person. (Cf. p. 216.)

When one gives a report of one's perceptions, one reports that a book is red, for example, and then is told that the book is not really red, one might revert to saying something like that it is subjectively true that the book is red, meaning that the book seems red to one. But to say that this is a situation in which something is subjectively true is wrong, it is to confuse subjective content with truth status. In this case it is an objective truth that the book seems red to one. I will call truths like this 'truths with subjective content'. There will be no general reason why truths with subjective content cannot be ascertained objectively. Then, truths which can be ascertained objectively I will call 'objectively ascertainable truths'. Truths which can only be ascertained subjectively will thus be those truths which are not objectively ascertainable.

A good example of truths only ascertainable subjectively are psychological truths. Daniel Dennett argues that there are intentional facts which can only be ascertained if we take the intentional stance. He says 'while belief is a perfectly objective phenomenon ... it can be discerned only from the point of view of one who adopts a certain predictive strategy [i.e. the intentional stance].'³⁰ This generalises to all psychological items - desires, thoughts, ideas, etc. Dennett imagines a super-intelligent Martian able to predict our behaviour without postulating intentionality, i.e. without talking about beliefs and desires. We are to suppose that human beings are as predictable physically to the Martian as simple physical objects like stones are to us. If such a being is intelligible, then there is a possible point of view from which we are not seen as intentional. Dennett asks whether this threatens the objectivity of our intentionality, answering that it

³⁰ *The Intentional Stance*, p. 15.

doesn't. Such a view would be missing 'something perfectly objective; the *patterns* in human behaviour'³¹.

Even truths about me (i.e. with subjective content) which can only be ascertained by me (i.e. subjectively) are true objectively if they are true. But then this suggests that the idea of subjective truth is redundant. Subjective ways of knowing and understanding, are important, but when we talk about whether something is objective, i.e. true independently of us, we are not asking whether it can be known objectively, but, regardless of how it is ascertained, whether it holds true independently of us.

The distinction between subjective and objective truths as a distinction between truths that are subjectively (or relatively) true and truths which are objectively true is to be rejected. This means that 'objective truth' and 'truth' are to be taken as equivalent. I will examine some candidates for subjective truths and show that they are not in fact best thought of as such but best thought of either as truths with subjective content or as truths only subjectively ascertainable (or as falsehoods). This will not provide a conclusive argument against the existence of subjective truth but will make a strong case against it.

2.3 Doing away with subjective truth and relativism

Subjectivism and relativism are often put forward in order to protect the subjective, experience, ethics, etc., from the tyranny of objectivism. Such protection is not good enough, though: subjective or relative truth is a poor substitute for real truth. We can better protect these things by seeing their claims as objectively true although only subjectively ascertainable. I suggest that in all

³¹ Ibid., p. 25.

cases where subjectivism or relativism is thought to be appropriate, things are better understood in terms of truths with subjective content or truths only ascertainable subjectively. Let us consider some plausible candidates for subjective or relative truth, then.

In order to understand a sentence one must understand the meaning of the terms. In order to understand an English sentence like 'Grass is green' one must understand English. This doesn't make the truth of the proposition 'grass is green' dependent on us, relative to our language. That grass is green is an objective fact about the world which would have held true if the English language had never existed. Nagel makes a distinction between form and content: although the form of a thought, an English sentence in this example, depends on us, the content expressed by the thought transcends this form and so is objective.³² In general, the fact that understanding a sentence relies on specific features of an individual or a practice doesn't imply that the truth of the sentence is reliant on the individual or practice in any way. So just the fact that in order to understand certain truths one needs to experience the world in the right way or be part of a particular culture or practice doesn't make those truths subjective.³³ The same seems to apply to those features of the world which can only be perceived by creatures with particular perceptual equipment - colours, for example.

We might be happy to say that the colour of an object is an objective fact even though colour can only be detected by certain perceivers, but what about features of the world which don't seem to be as independent of us as colours -

³² See *ibid.*, pp. 101-2.

³³ What is suggested by this discussion is that understanding a sentence is more than knowing its truth-value, that acquiring understanding is not to be thought of as obtaining a list of truths. The fact that certain truths can only be ascertained subjectively, because they can only be understood

smells, tastes or values, for example. It sometimes seems that there is no objective fact of the matter about these things, only subjective facts. If it seems to one that gruyere tastes better than stilton, then this must surely be a subjective truth, since it is false to many people. There might seem to be no reason for denying that truths about taste are subjective, but we are more reluctant to say the same about truths about values.

Whatever the subjective element of matters of taste, it is at least objectively true that things seem a certain way to a particular subject - gruyere tastes better than stilton to one, for example. We can think of all truths about matters of taste as truths with subjective content, since they are not strictly truths about things, cheeses, for example, but truths about the subject's relation to things, cheeses, for example. Truths about moral values are different, though - we don't think of moral truths as truths with subjective content. One's saying that it is morally better to eat vegetables than animals, for example, is not saying something personal, that it is better *for one* to eat vegetables than meat. But then neither can moral truths be thought of as subjective truths: in making a moral judgement one is aiming to express an objective truth. So there is a fork: either a truth thought to be a subjective truth is actually better expressed as an objective truth with subjective content, or it is not to be translated into a truth with subjective content but is actually an objective truth as it stands (if it is in fact true), not a subjective truth.

There is a confusion between truths only available to specific viewpoints, to certain individuals or groups, i.e. truths which can only be ascertained subjectively, and truths relative to an individual or culture. Regardless of how

subjectively, from within, shows that even if an objective list of the facts could be known, knowing this would not necessarily constitute any real understanding of the world.

specific one's viewpoint and attributes must be to find out the truth of a proposition, once one finds out its truth-value, one has found its objective truth value. If one person believes that God is everywhere, for example, how can it be true that God is everywhere 'for them'? If God is everywhere 'for them' then God is everywhere. It is not that God is everywhere from their perspective and there are other places available from different perspectives where God isn't, or if it is, then we'd better take that person to those places to show them the error of their belief. The fact that two people can be in the same place and one perceive something that the other cannot doesn't necessarily mean that the thing is not 'out there' to be seen.³⁴ Any viewpoint or conception of the world may yield truths about the world. The fact that this is so doesn't imply that the facts yielded by a particular viewpoint or conception can then be assessed from any viewpoint, by anyone. The fact that a truth is objective doesn't guarantee that it is universally available.³⁵ Someone who claimed that there was no real conflict between one person who believed in an omnipresent God, another who believed in a transcendent God and one who didn't believe in a God at all, that each of their claims could be relativised or that the contradiction between their assertions is only apparent, would have to argue their case - there certainly seems to be serious disagreement.

³⁴ An overly physicalist view of things might seem to suggest that something real must be perceivable by any perceiver, but in fact it doesn't - there are things we can perceive but which are too big for tiny creatures to perceive, for example. And even if it did, there are good reasons for rejecting such a view.

³⁵ This is not to say that anything anyone asserts is objectively true, just that anything *true* anyone asserts is objectively true. We are not forced to embrace as objectively true everything anyone asserts any more than a relativist is forced to embrace as true relative to something everything anyone asserts. Presumably not even the most open-minded relativist would say that all utterances are true relative to some view: sometimes people are just mistaken.

What shall we say about the conflicts between subjective and objective that Nagel presents? If from an objective viewpoint ‘my birth seems accidental, my life pointless, and my death insignificant’, but from my subjective point of view ‘my never having been born seems nearly unimaginable, my life monstrously important, and my death catastrophic’³⁶, what is the truth of the matter? The conflict can be resolved but not if we assume that one or other view must be completely right and the other completely wrong. It might be that the truth about my birth is given by the objective view - that there is no reason why *I* was born and why I was born when I was. Presumably my death will be similar, I will die when I die and it will not necessarily be when I’m ready. But the importance of my life, the value it has to me and to the world I live in, is to a certain extent independent of these facts; the meaning my life acquires it acquires regardless of the fact that its beginning and end do not themselves have meaning. So it looks like a solution is possible to this conflict by separating things, by acknowledging the insights of both subjective and objective views. This will be followed up in section 2.5 below.

Nagel presents us with examples of irreducibly subjective facts: if I give objective descriptions of a group of people, telling the story of their lives in the third person, and I am one of the people in the group, then a person reading the description will have no way of knowing which person is me. I will have left out one fact, the fact that one of them is me. Facts like the fact that I am *SK*, that the time now is 10.30 a.m., that it is raining here, seem to be subjective facts. Even though we can find objective truth conditions for them - ‘I am *SK*’ is true if and only if uttered by *SK*, for example - their subjective meaning doesn’t seem to be

³⁶ Ibid., p. 209.

translatable into objective terms. Nagel puts this clearly: ‘indexicals in general are untranslatable into objective terms, because they are used to refer to persons, things, places and times from a particular position within the world, without depending on the user’s objective knowledge of that position. It is elementary that one can’t translate a sentence whose truth depends on its context of utterance into one whose truth does not.’³⁷ This is why an objective view is necessarily incomplete.

Even if these truths are irreducibly subjective in some sense, it must be stressed that they are objectively true in that their truth is independent of whether we take them to be true or not. I can be wrong about what time it is now just as I can be wrong about what time something occurs. I can be wrong about who I am, although this is unlikely. So the irreducible subjectivity of these truths should not be thought of as subjectivity of truth status.

I hope to have made a good case for the possibility of dealing with all candidates for subjective truth in this way. If all truths yielded by subjective viewpoints can be thought of either as (objective) truths with subjective content or as (objective) truths ascertained subjectively then the worry that important subjective aspects of our lives are threatened by objectivity can be met.

2.4 Is any detachment good?

Nagel’s view of objectivity is motivated by the thought that greater objectivity, greater detachment, is good. There is then a problem between two apparently irreconcilable goods - the good of being neutral - objectivity - and the good of experiencing the world as a subject - subjectivity. In order for objectivity to be

³⁷ Ibid., p. 59.

important it must be true that at least some things are better apprehended objectively, i.e. through detachment.

On the detachment view of objectivity, we move from a particular viewpoint towards greater objectivity by transcending the particular, by abstraction. So the maximally objective view, Nagel's view from nowhere, is the intersection of all possible views. In abstracting from the particular we achieve generality but leave behind the concrete. Thus in one sense the objective is less real than the subjective.³⁸ Also it is obviously incomplete - the intersection of a collection of non-identical sets is smaller than at least one of the sets. If we see the objective in this way then we need no argument to the conclusion that there are features of the world which go unnoticed from the most objective view.

It might be thought that the most important things in the world are those things which are most objective in this sense, that the intersection we are left with is the set of core features of the world - ultimate reality. But this is an assumption requiring support. It might well be the case that some of the most important things in the world are available to all, but if so, this is not just *because* they are. Thinking that objectivity automatically grants status is appealing to the lowest common denominator. If we really believed this line of thought then we would give up philosophy and just spend all our time eating, sleeping and having sex! (Or, taking things further, we'd refuse to think of ourselves as persons at all and view ourselves as collections of particles, rejecting agency, ethics and all other human concerns.) It seems obvious that there are some things of the utmost importance and reality which are not accessible to all.

Physics is thought to give us the most objective view of reality on the detachment account.³⁹ This is because it strips away all viewpoint-relative properties and deals with matter qua matter. Since all things in the world are physical and interact physically, the causal descriptions offered by physics will apply to the interactions of all things.⁴⁰ What this should lead us to see is that physics is the highest common factor, not that ultimate reality is what is described by physics.

Scientific realism is the view that the best account of underlying reality is given by science, usually meaning physics. The scientific realist is attempting to answer the sceptic. He contrasts objective properties such as size, mass and shape with subjective properties such as colour, smell and goodness, and claims that objective properties are real and subjective properties are not 'part of the fabric of the universe'. But this equating of objective with real is assumed on the basis of a subjective (appearance)/objective (reality) distinction and the objection that science gives us an incomplete view of reality is answered by dismissing as unreal the things it leaves out (since they don't appear in it). This is obviously circular. The question of what the irreducibility of certain experiential qualities to physical properties shows is an interesting one: scientific realists take irreducibility as a sign of unreality, whereas those on the other side of the argument take it as a sign that physics gives an incomplete account of reality. If a scientific realist is going to dismiss non-physical properties as unreal, what is his justification for doing so?

³⁸ Cf. Nancy Cartwright's arguments that the fundamental explanatory laws of physics are less true than phenomenological laws: there is 'a trade-off between factual content and explanatory power'. See *How the Laws of Physics Lie*.

³⁹ This is not strictly true - it is folk physics, not Physics as practised by Physicists, which is thought to be accessible to all views.

⁴⁰ Physics will not necessarily best describe an interaction, but a physical description will apply.

And does physics really give us the most objective, the most detached view? What exactly is a viewpoint? Nagel talks about the subjective viewpoint, meaning the experiential world, he talks about the scientific/physical conception of the world, the ethical standpoint, the internal and external viewpoints it is possible to take towards one's life, etc. His aim is to reconcile conflicting viewpoints, to work out how to deal with conflicts which arise between different viewpoints. The problem is, he says, 'how to combine the perspective of a particular person inside the world with the objective view of that same world, the person and his viewpoint included.'⁴¹

A viewpoint needn't be occupied to be actual, we can imagine a view of the earth from outer space, for example, and when Nagel talks about the viewpoint of physics he is not supposing the existence of a being who occupies that viewpoint, a being who sees the world purely in terms of physics, observing only physical properties and relations of objects and being blind to secondary qualities, values, etc. So a viewpoint can be a construction of the imagination.

Are there not, then, possible views even more detached than physics? The move from the experiential world of primary and secondary qualities to the world of primary qualities, the world of physics, can be repeated, taking the process of stripping away a stage further. We can imagine the view of a weighing machine, or a speedometer; a view in which only one primary quality exists. A weighing machine cannot 'perceive' size, shape or motion, only weight. The weighing machine's view is more detached than the viewpoint of physics, having abstracted from some of the specifics physics hangs on to.

⁴¹ Ibid., p. 3.

But such a view surely cannot be considered to be more objective than physics. Why not? One response might be that the weighing machine's view detaches from *relevant* specifics, not just from the irrelevant ones: there is a difference between the move from the world of primary and secondary qualities to the world of physics and the move from the world of physics to the world of the weighing machine in that secondary qualities are irrelevant, the colour of an object is irrelevant to how it behaves physically. But in response to this it can be pointed out that from the viewpoint of the weighing machine the other primary qualities are irrelevant - the shape and size of an object is irrelevant to its weight.⁴²

At this point it seems clear that if objectivity is equated with detachment, it is not to be valued. And, if we don't equate objectivity with detachment, if we have another way of defining objectivity, then greater detachment won't always yield greater objectivity. The amount of detachment we require is dependent on what we are trying to find out, but objectivity doesn't seem to be like this. Nagel unconvincingly talks about the problem of being too objective, but can one really be too objective?⁴³ Nagel seems to be thinking of the choice between subjective and objective viewpoints as being like the choice between different scale maps for different purposes, saying things like if we are only interested in the physical interactions of objects then physics is what we need, if we are interested in what different people would do in different situations then we need a human viewpoint, like choosing different scale maps depending on whether one is walking round a town or driving across a country. But this analogy doesn't work: anything on a

⁴² Of course large objects will often be heavier than smaller ones, but then there are links of this kind between primary and secondary qualities, something's apparent colour and its distance from the observer are linked in this way.

⁴³ One can certainly be too *obsessed* with objectivity, with truth, etc., but this is a different point. One cannot be criticised for being too accurate, too open-minded, etc.

map of whatever scale is real: some maps include x 's and some don't - this doesn't threaten the objectivity or reality of x 's. All good maps are ways of representing objective reality and a map of the solar system is not more objective than a map of Oxford.

Nagel assumes that there is a series of views each one more objective than the previous. If objectivity is detachment then this is not true, there are countless different ways of creating a more detached view from one starting view. The viewpoint of the weighing machine and the viewpoint of a speedometer are both detachments from the physical viewpoint. More relevantly, we can abstract from our experience and model the world in many different ways, we can take up a view in which what is important about things is their colour and shade, in order to plan and decorate a room, for example, or we can take up a view in which time of day and gravity are irrelevant, if we are preparing to go out into space, etc. None of these models is more objective than the others.

It seems unintuitive to say that the weighing machine's viewpoint is objective at all, it would not contain re-identifiable particulars, for the same reasons that Peter Strawson's imagined purely auditory world would not.⁴⁴ If the definition of objectivity as detachment leads to the conclusion that the weighing machine's viewpoint is an objective advance on the viewpoint of physics then we should reject that definition.

2.5 Objectivity as containment

It seems clear that there is a significant difference between the viewpoint of physics and the viewpoint of the weighing machine. We do not feel a tension

analogous to the tension we feel between our subjective view and the neutral view of physics between the view of physics and the view of the weighing machine. The fact that from the viewpoint of the weighing machine most of the primary qualities are irrelevant does not threaten their objectivity. For there to be any conflict between two views, both views must be important to us. What makes a view important to us? There are many possible views or models we could adopt, what makes those we do adopt better than the ones we see as contrived and pointless, like the view of the weighing machine? Whatever the answer to this question, it draws our attention to the fact that objectivity of viewpoint, unlike detachment, cannot be thought of as neutral: objectivity is a good thing.

Nagel talks about the possibilities concerning the conflict between the subjective and objective views: we can either allow the objective view to override the subjective view, or we can refuse to accept the claims of the objective view and stay within the realm of the subjective. Neither of these options is satisfactory, since both views are essential: 'it seems impossible to give the authority to either [view] in deciding conflicts between them'.⁴⁵ This means that there is no possible resolution. Nagel assumes that either the objective view must give us the whole truth of the matter, the subjective view being completely erroneous, or vice versa, but this does not exhaust all the options. In most disagreements truth is not allocated so unevenly but is split between the conflicting views, both having some contribution to make. When this is so, then there is the possibility of a better way of seeing things, a third viewpoint which encompasses the conflicting views, from which a resolution is possible. Consider a perceptual disagreement: a coin appears

⁴⁴ See *Individuals: An Essay in Descriptive Metaphysics*, chapter 2.

⁴⁵ *Ibid.*, p. 8.

to be round from one viewpoint and elliptical from another. What is the truth of the matter? According to Nagel's way of putting things, we must either say that the whole truth of the matter is that the coin is round and the view which presents it as elliptical is a distortion, or that the whole truth of the matter is that the coin is elliptical and the round appearance is a distortion. But surely a better understanding of the nature of things would explain the relationship between roundness and ellipticity, why it is that round things look elliptical from certain angles. This would be to include the elliptical appearance as part of the whole truth about the nature of the coin, to relate appearances of roundness and ellipticity in a way that put them on a par, as different aspects of a phenomenon.⁴⁶ It seems natural to think of the emergent third view as a more objective view than either of the first two; it gets to the bottom of the apparent conflict and tells us how things really are and why there appeared to be a conflict.

Nagel takes his view of objectivity as detachment to imply a conception of knowledge acquisition or 'objective advance' such that 'an advance in objectivity requires that already existing forms of understanding should themselves become the object of a new form of understanding.'⁴⁷ He gives two examples to illustrate the idea of objective advance: the discovery of the distinction between primary and secondary qualities and the discovery of special relativity. In both of these epistemic advances, we develop a new theory which explains why things appear the way they do and also why things are not exactly as they appear. Nagel says 'The hypothesis that objects have intrinsic colors in addition to their primary qualities ... provides a poorer explanation of why they appear to have colors, and

⁴⁶ It is not true that the appearance of ellipticity of a round thing is just a matter of our deceptive experience. Ellipticity is more importantly related to roundness - round objects often cast elliptical

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.’⁴⁸

If we think of objectivity of viewpoint as the result of objective advancement in this way, we equate objectivity of view with containment: one view is seen to be more objective than another if it contains it. We don’t necessarily equate greater objectivity with greater viewpoint-independence or detachment. Nagel assumes that objectivity as detachment and objectivity as containment coincide and this is not always the case. According to the containment model, physics (or a more detached viewpoint like that of the weighing machine) is not the most objective view: the view of the philosopher of science who examines physics and its methods, etc., is more objective than the view of the physicist.

The main idea for Nagel - the idea of a view from nowhere - is that given by the detachment model. Greater objectivity is equated with greater detachment from specifics of viewpoint, less bias, a higher degree of neutrality. Even though Nagel rejects the association between ‘more objective’ and ‘more real’, this model does naturally lead into that equation. The contrast between detached objectivity and biased subjectivity has built into it the idea that one gets a better, less clouded view of the real world if one’s view is more objective. On the other hand, the containment model of objectivity doesn’t contrast objectivity with subjectivity or experience at all. Objectivity is not an endpoint but an ongoing opening up and questioning of one’s current position, a getting to the bottom of things. In one’s

shadows, for example.

⁴⁷ Ibid., p. 75.

experience one learns to transcend previous views by developing new views which are just as specific but which are wider, which contain the previous views and explain them. Objectivity as containment works against small-mindedness, not against subjectivity. When I ask a friend what she sees and find out that my eyes have deceived me, I find an explanation for my mistake in terms of my new experiential view, I realise that I do need glasses after all, for example. Nagel runs together containment and detachment assuming that they are equivalent and failing to notice that in the examples he gives of objective advances the resulting, i.e. more objective, view isn't more detached. Relativity theory is not more objective than Newtonian physics by being less viewpoint dependent, although it is in fact less viewpoint dependent, but by being more complete, by containing Newtonian physics.

Nagel seems to assume that a more objective view corrects the subjective view it contains, he doesn't consider examples in which a more objective view explains why the subjective view is correct. If it is the case in general that the objective corrects the subjective, then there is a tension created between this and the value he places on the subjective. It makes it look like one can only value the subjective by ignoring the objective advance that would invalidate it. It may just be that Nagel is just focusing on conflicts, but he does not consider the possibility that the conflicts he looks at may not be real conflicts, that his scientific view might be wrong in some ways.

The example I gave of an objective advance, of finding out that my vision is defective through asking someone else, could easily have gone differently - if my friend had reported seeing the same as me then that would be an example of an

⁴⁸ Ibid., p. 76.

objective advance that yielded no conflict at all between the prior, less objective view and the resultant, more objective view. Nagel says 'Every objective advance creates a new conception of the world that includes oneself, and one's former conception, within its scope; so it inevitably poses the problem of what to do with the older, more subjective view, and how to combine it with the new one. A succession of objective advances may take us to a new conception of reality that leaves behind the personal or merely human perspective further and further behind.'⁴⁹ A succession of objective advances *may* take us beyond the human perspective, but it needn't. And in fact once we stop thinking of objective advances as detachment this comes to seem unlikely. On the containment model of objectivity, an objective advance is an advance in self-understanding: coming to see that one makes certain mistakes or is limited in certain ways leads one to search for an explanation of why things aren't the way they seem, or for another source of information to supplement one's current knowledge. Better, more obvious examples of objective advances than those Nagel gives are to be found in psychology, Freud's discoveries concerning the extent of the effect of the unconscious on the conscious, for example.

The containment model of objectivity is more helpful than the detachment model. It yields a better understanding of the way conflicts lead to advances in knowledge and doesn't equate the ongoing quest for objectivity, the desire to get to the bottom of things, with disengagement and disillusion. We can continue to strive for objectivity without worrying that the quest will destroy or leave behind what is most important to us. In some cases containment is equivalent to detachment - a theory which contains another may do so by being less tied to a

⁴⁹ Ibid., p. 5-6.

specific set of conditions, for example. If we develop a theory which contains a previous one it will often do so by taking the specific axioms of the previous theory to be just part of a wider class of axioms, to apply to more situations than the previous set, i.e. to be less tied to specifics. But containment is perhaps better thought of as a widening of viewpoint rather than a casting off of specifics.

I hope to have shown that an acceptable view of objectivity of viewpoint is compatible with the valuing of the subjective, that subjectivity - experience - is not in opposition to objectivity. Through our subjectivity we discover objective truths about the world, and, further, through considering or imagining different subjectivities, abstracting from our own experience, and having open minds we develop more objective views of the world. Thus the problem Nagel set out to solve, i.e. the problem of conflict between viewpoints, can be solved without the subjective being threatened.

The view of objectivity as containment, unlike the view of objectivity as detachment, doesn't lead to 'false reductions or outright denials' - it does not suggest that we dismiss things we experience but that we embrace everything and try to find an explanation of how apparent conflicts can be resolved. 'The way things really are' is to be taken to be closer to 'the totality of way things can be' than to 'the way things are underneath'. The epistemological gap between the way things are and the way things seem, which is the main hallmark of objective reality, can be cashed out in terms of incompleteness of view rather than in terms of an opposition between our view which only presents us with appearances and unknowable underlying reality. So we can reject the appearance/reality distinction without rejecting objectivity.

2.6 Objectivity as mind-independence

So far the idea of a mind-independent world, the objective world, has lurked in the background. We have been considering the objectivity of viewpoints. It is time to look at the link between the two. A viewpoint is assessed as more objective than another insofar as it leads to a better understanding of the world, to discovery of the truth. Objective viewpoints are to be contrasted with biased or distorting viewpoints and also with viewpoints which are narrow or which fail to be self-critical. Objectivity of view when thought of on the containment model is something which requires constant work. A view must constantly be revised in the light of new slants on things. The quest for objectivity is thus not to be thought of as the quest with an end, it is not to be thought of as a journey to a land - the objective realm - which once completed would mean one could settle and not travel any further.

If the way things really are is not to be contrasted with the way things are in relation to us, which it isn't - relational properties such as gravity are just as real as non-relational properties such as mass - then the kind of mind-independence that is important is not a kind of distancing or sealing off of reality, an embalming of the real world, protecting it from our interference. This would be to fall back on the idea of reality in itself. We ask about the way things are independently of us because we want to know about their whole nature, about how they are in relation to other things too. And the things we want to know about are not just physical objects, so the scientific investigation conducted by physicists is not enough. We also want to know about thoughts, political systems, time, values, numbers, emotions, meaningfulness, etc. All of these things are such that there are unknown

aspects of them for us to find out, all of these things can have objectivity. Anything that is the subject of our scrutiny must have a certain amount of objectivity, for if not, we'd already know all there was to know about it and no longer be puzzled by it. If we seriously ask the question: 'Given that x 's exist, are they objective?' we are assuming the reality of x 's. If x 's exist and we are asking a question about their nature, then there is a gap between our knowledge of them and the way they are, hence we are assuming that they are objective.

If the quest for objectivity is a quest for greater and greater understanding, not for the final true description, then where does the idea of objective, mind-independent reality fit in? I take mind-independence to be an epistemic relationship; something is mind-independent in the relevant sense if our beliefs about it could be wrong (or our perceptions of it misleading). If something is objective then there is a possible gap between what we take to hold true of it and what does hold true of it. Why is being mind-independent in this way important? If there were some things which were not mind-independent in this way, i.e. things we knew perfectly so that there was no possible gap between the way they were and the way we took them to be, would their status as objective parts of reality be challenged? If so, why? If our knowledge of something is not complete to start with and we investigate it, ending up with complete and perfect knowledge of it, then does it cease to be objective?

If there were an omniscient God who knew what we would do before we did it, then would we be free? Surely if God knows that we will ϕ then we are not free to choose not to ϕ ? Similarly, if there are things in the world which can't surprise us, things we have complete knowledge of, then they are not independent

of us. In both of these situations, the worry is only apparent, though: once we point out that it is not the knowledge which constrains the world but the world which constrains the knowledge, the problem dissolves. We are free to do as we choose, and whatever we choose is what God will believe (albeit in advance). God is the one who is not free in this instance - God is not free to believe something false, since he is omniscient. (If he were free to believe anything and then his having chosen to believe something would make it true, then we would not be free.) Similarly with the things we know completely, it is their nature that brings about our beliefs, not vice versa. The relationship is perhaps best thought of in terms of the direction of explanation: something is objective if it is such that our beliefs about it are explained in terms of its nature and its nature is not (fully) explained in terms of our beliefs. This is to say, the things which are objective are the things which affect us and are not completely within our control.⁵⁰

Ralph Walker distinguishes three senses of ‘objectivity’: the weakest is the existence of an epistemological gap: ‘to say [certain things] are objective is to say that one can make mistakes about them - the truth about them is not just whatever one happens to believe, but has to conform to some standard.’⁵¹ To equate the possibility of making mistakes with the idea that it is one’s beliefs that conform to the way things are not vice versa is wrong; the possibility of making mistakes implies but is not implied by the fact that our beliefs answer to the way things are. But the important thing is the idea of beliefs conforming to a standard, not

⁵⁰ Maybe the reason we value the gap between what we know and the way things are is that in order for something to become real to us, to be worthy of our interest, it must be capable of surprising us; it must challenge us, stimulate us and be bigger than us. Or do we value things which surprise us because they are objective?

⁵¹ *The Coherence Theory of Truth*, p. 67.

themselves determining the nature of things. As noted above, all of the things we take seriously are objective in this sense.

On Walker's second, stronger sense of objectivity, something is objective if it exists independently of our perceptions of it. Most of the things we perceive are objective in this sense.⁵² If we are considering things that are not perceptual objects, such as thoughts and numbers, the question of whether they are objective in this sense must be refined. It is not good enough to just say that non-perceptual things are of course objective in this sense - they are completely independent of perception in general and so obviously independent of our particular perceptions. We should talk of experiential independence rather than perceptual independence, to include all forms of awareness.

On Walker's third and strongest form of objectivity, things are objective if they are public, i.e. if they 'are not just private to one individual but common to everyone'⁵³. I reject this characterisation of objectivity since it conflicts with the possibility of truths only ascertainable subjectively. Objectivity is not to be thought of as accessibility from any viewpoint: the fact that our perceptual apparatus allows us to see colours other animals are unable to see doesn't threaten the objectivity of colours - we discover that the world is coloured. This applies in

⁵² Whether things like beauty are objective is a question I will return to. Beauty is an essentially relational property even if it is not completely in the eye of the beholder. Even if beauty is not completely in the eye of the beholder, then the existence of beauty is not independent of perception of it. Consider attractiveness to the opposite sex, to be more specific. If sexual selection is a good way of understanding why peacocks have large tails and why men are strong and women slim and curvy, then the objective characteristics of a particular peacock's tail are brought about by the preferences of previous generations of peahens. At any point, it is an objective fact whether or not a feature is attractive to a perceiver or group of perceivers, (an objective fact with subjective or relative content) and this objective fact is causally relevant to the feature's success or demise. Even if beauty is objective, then, there would just be no point in being beautiful if there's no one to appreciate it.

⁵³ Ibid., p. 67.

general - general accessibility should not be thought of as necessary for objectivity.

According to the view of objective reality I am endorsing, i.e. objectivity as being that which is beyond our control, thoughts, numbers and emotions, and most of the things we experience in a non-perceptual way, will be seen to be objective. If a thought can be shared or unconscious, then it is independent of my being aware of it. Emotions, like pains, are generally thought to be such that one cannot have one without being aware of it, but this is not always so. One often has to struggle to understand one's feelings, to find out whether the way one feels is love or something else, for example. Just as one may believe things without being conscious of doing so, one may feel things without knowing. The unconscious is objective in that it is made up of things we don't know nearly enough about, things that affect the way we are and what we do. In order to gain self-understanding we must draw these things out into consciousness, to examine them, much as we examine physical objects we don't fully understand.

Does anything fail to be objective according to this view? I will mention two kinds of things, which will be discussed later: the basic axioms of a system and those indeterminate aspects of ourselves and the world that we make determinate by our choices. There are certain principles we may choose to live by, and as a result of our choosing them, our lives become a certain way. To a certain extent the way one answers the question 'what kind of person am I?' determines the kind of person one is or becomes. If one sees oneself as hopeless, then one will be hopeless. If one sees oneself as confident and capable, then one will be confident and capable. It is not that who one is is completely open, but there is a

certain amount of indeterminacy/possibility that is converted into an actuality as a result of one's self-image. The same kind of thing happens with one's outlook on the world, if one is happy, optimistic and confident then the world is an exciting place full of challenges to grapple with. If one is pessimistic and disempowered, then the world is an awful place full of threats to one's happiness and survival. It is not true that any outlook is as good as any other, nor that if one has one outlook then one is locked into it forever, but to a certain extent one's outlook does shape the world and make it impossible for one to objectively assess the outlook itself. Questions concerning the objectivity of these things will be raised in part IV.

II

REALITY, LANGUAGE AND EXPERIENCE

'But if our freedom in empirical thinking is total, in particular if it is not constrained from outside the conceptual sphere, that can seem to threaten the very possibility that judgements of experience might be grounded in a way that relates them to a reality external to thought.'

-JOHN MCDOWELL, *MIND AND WORLD*

In part I I developed a view of the objective world as containing all kinds of things, most things being seen to have objective existence. Objective existence was shown to be not a matter of being part of an unknowable underlying reality but, for things in the world around us, a matter of being independent of our perceptions and beliefs: something is objective if it has a life of its own and has the capacity to surprise us. The things which depend on us in the sense that we create them in the first place - our artefacts - are to be thought of as just as objective as those things which existed before we did. The relationship between ourselves and our artefacts is like the relationship between a parent and a child - once the child has been born its continued existence becomes independent of its parents and it becomes more than just the child of these parents.

But we would not be able to experience an objective world at all without language: objectivity requires the possibility of conceiving of things existing when we are not experiencing them and this requires the possibility of referring to objects. Further, the possibility of one's experiencing certain kinds of things, the possibility of one's experiencing anything as meaningful, for example, relies on one's possession of particular concepts, the concept *meaningful*, for example. Does this kind of dependency threaten the reality of properties like

meaningfulness? There are those who argue that truth is relative to conceptual scheme, and that whilst it might be appropriate within a particular conceptual scheme or language game to apply the term 'meaningful' to certain things, that is not to say that those things really are meaningful. We could have developed a conceptual scheme without the concept of *meaningful* and if we had, it is suggested, then meaningfulness wouldn't have existed. Or we could have developed the concept only found different things to be meaningful, and this casts doubt on whether the things we take to be meaningful really are meaningful.

1. The relationship between language and experience

Language is in some ways similar to experience - it presents reality as being thus and so. But the way experience presents reality and the way language represents reality are fundamentally different. Experience presents things from a viewpoint whereas language does not. The way experience tells us about reality is as a result of being caused by reality, and, if an experience is reliable, its content and the content of reality are structurally related, each composed of elements in a certain arrangement. Language, on the other hand, can present things as being thus and so without being caused by things being thus and so (this must be true, since we can utter falsehoods using language), and the structure of what we say need not mirror the structure of that we are talking about.

Experience is *of* the world, language is *about* the world. If someone experiences something, it must exist. Experience is tied to reality and we have a problem explaining how reality and experience come apart when they do. With language the opposite seems true - we can talk about non-existent things and about

how things could have been just as easily as we can describe what we see. In this case the problem seems to be: how can language and reality be brought together? If experience relates to reality as a footprint relates to a foot, which it seems to, then it is relatively easy to see how the structure of experience can yield content about reality - the structure of reality can, to a certain extent, be read off from the structure of experience. With language things are different: whereas an experience of a dog, like a painting of a dog, can be said to show us what a dog is like, the word 'dog' cannot.

1.1 The emergence of objectivity

Conceptualisation is a necessary condition for experience of an objective world. This claim and the idea that having concepts is equivalent to having an objective world is discussed by Adrian Cussins in 'The Connectionist Construction of Concepts'.¹ Cussins starts off by stipulating that conceptual content be 'content which presents the world to a subject as the objective, human world about which one can form true or false judgements.'² He later argues for this view of objectivity using the ideas of P. F. Strawson and Gareth Evans.

To have the concept *a* of an object it is not enough that one be presented with that object, that one experience *a* in a certain context, that one experience *a* as *F*, for example. Thought is essentially structured: to be able to think that *a* is *F*, one must also be able to entertain the thought that *a* is *G*, *H*, and so on. This is Evans' generality constraint.³ We can think of the difference between non-conceptual and conceptual content as the difference between an unstructured,

¹ In Margaret Boden ed., *The Philosophy of Artificial Intelligence*.

² *Ibid.*, p. 381.

unarticulated representation and structured representation. An animal that doesn't possess concepts can see that-*a-is-F*, predator-here-now, e.g., but not that *a* (identified as a particular) has the property *F*. To an animal, what is seen is a one-off situation, not something with separable parts which can be recombined on other occasions. The animal will not be able to recognise *a* as the same thing if it sees *a* again and thus *a* does not have continued, objective existence for the animal.⁴

Experience of an objective world is experience of things which exist independently of our experience of them. In order to experience a thing as objective we must be able to think of it as existing when it is not presented to us and to re-identify it as the same thing when it is presented to us again. In order to think of a thing in this way we must be able to refer to it, so, in general, we must have a language. The objective world is the world of thought, imagination and flexibility, of escape from the actual into the possible. In order to be able to think about things, to consider possibilities, we must be able to separate what we see is the case into component parts and then recombine the parts in different ways in thought. An animal is trapped in the actual, being unable to make the jump from *a*'s being *F*, to the possibility of *a*'s being not-*F*. Being unable to separate *a* from *F*, it is as though, to the animal, everything it registers is necessarily the case - the

³ See *The Varieties of Reference*, pp. 100 ff.

⁴ Someone might object that animals often seem to be able to refer to particulars - a dog recognises its owner, for example. The debate about whether animals can be said to possess concepts or at least 'proto-concepts' is not one I am interested in entering here: I am not primarily interested in which creatures have access to an objective world, but am just taking on a view about what having access to an objective world entails.

animal's representations are not of a world with a fixed past and open future, not of a world at all.⁵

Cussins says: 'the emergence of objectivity is the transition from mere experience to experience *of* the world.'⁶ This raises two questions: firstly, how is experience transformed by conceptualisation in this way? Secondly, if different conceptual schemes are possible, which they seem to be, then is it not misleading to talk of objectivity? Shouldn't we say instead that the emergent world is one of many possible conceptual worlds, that the transition is from mere experience to experience of *a* (conceptual) world? I turn now to the first question. The second question will be raised in section 2.

1.2 From mere experience to experience of the world

Experience is essentially from a viewpoint or perspective. Language is viewpoint-free. This is a fundamental difference. When I assert that *p*, what I am doing makes no reference to my perspective. If *p* is a fact then it is an objective fact. If this were not the case then sentences would not stand on their own but would need a reference to perspective; my assertion would take the form '*p* (relative to *SK*)', for example. But this could not work. The assertion '*p* (relative to *SK*)' can be translated into a pure assertion *q* where *q* is 'From *SK*'s perspective it seems true that *p*' and then the question 'From what perspective is *q* seen to be true?' is left unanswered. It is impossible to attach a viewpoint to a sentence in a way that then traps it into being relative, i.e. non-objective.⁷ Whatever the content of a report,

⁵ For a fuller discussion of this idea see John McDowell's exploration of the difference between possessing a world and inhabiting an environment (*Mind and World*, pp. 114-119).

⁶ *Ibid.*, p. 410.

⁷ Someone might think that the answer to the question about *q* would be '*q* is true relative to our practices', but then this new claim doesn't appear to be relative.

however many viewpoints are described, the outermost framing is not viewpoint-relative; the reporter's viewpoint drops out. If one says 'It seems to *A* that it appears to *B* that *C* thinks that *D* is under the illusion that *p*.' then the question 'From what perspective is this being seen from?' still goes unanswered. Whenever anyone makes an assertion, they do so from a point of view, but in assertion, the point of view is discarded. That there are different ways of coming to see that *p* doesn't seem to affect the objective nature of *p*.

There are two ways of dealing with this. The first is to say that the structure of language is misleading, that even though assertions appear to be viewpoint-independent, they are not. So when someone makes an assertion, we should interpret the assertion in the light of who makes it and from what perspective, not take its content to be viewpoint-transcendent.⁸ A problem with this approach is the question of what to say about assertions which aren't tied to the way things seem: as well as being able to make assertions about how things seem to us, we can make assertions which don't seem true to us. I can assert that the sun and the earth revolve around the moon, or even that all three revolve around a huge chocolate cake if I feel like it, and those assertions would be perfectly valid moves in the language game, answerable to the facts, without needing to be related to the way things seem to me. In some cases we might hit upon the truth by asserting something contrary to the way things seem to us. The second approach is to take the content, *p*, of an assertion to be viewpoint-independent but to say that there is an important relationship between the fact that

⁸ It might be the case that our assertions are viewpoint-transcendent but relative to conceptual scheme. The two questions are separate.

p and the different ways of seeing that p . This is the approach I will take and I will offer an account of the relationship.

The possibility of transcending one's particular viewpoint, transcending the actual, is built into language. How is this possible? How does conceptualisation facilitate objectivity? How is it possible to jump from a 's seeming F from one's perspective to being able to assert that a is F ? How does one make the jump from mere seeing or seeing *as* to seeing *that*?⁹ I will refer to the jump from seeing *as* to judging *that* as 'the objective move'. When we learn the seems/is distinction, what exactly are we learning? And how do we learn it? There are more valid ways of seeing *as* than of judging *that* - if someone does see something *as* ϕ , then in general seeing that thing *as* ϕ is a valid way of seeing *as*. But someone's judging *that* ϕ does not guarantee that ϕ is valid, i.e. true. This is not to say that there are no constraints on seeing *as*. There are some ways of seeing *as* that are ruled out - there is no possibility of someone seeing the moon as a pencil, for example.¹⁰ But although we can dismiss such contrived examples, we can't completely dismiss actual ways of seeing *as*. Judging *that* requires we take into account more than the way things seem to us, though. When we see or judge *that* something is ϕ , we see or judge that it is objectively ϕ . Seeing or judging *that* is what assertion is about. When we make assertions what we say is detached from our viewpoint, we say something that stands independently, that other people can agree or disagree with (though they might not be in a position to learn its truth the way we are).

⁹ The jump is from having experience with non-conceptual content to having experience with conceptual content. I take it that one can see *as* without possessing concepts - an animal can see another animal as threatening, for example - whereas seeing or judging *that* is a conceptual ability.

¹⁰ Or if there is, I have not been able to come up with a possible context for such a way of seeing. The important thing is not that there are some ways of seeing *as* which are absolutely impossible, but that within a certain context there are some valid and some invalid ways of seeing *as*. One

If there are more valid ways of seeing *as* than judging *that*, then how is seeing *as* related to judging *that*? We cannot have a one-to-one mapping from seeing *as* to judging *that*, so either we must weed out some of the ways of seeing *as*, establishing a one-to-one function between certain ways of seeing *as* and ways of judging *that* (like the square root function over the Real numbers which maps each positive number onto a single value but doesn't assign values to negative numbers at all), or we must construct a many-to-one function mapping the many ways of seeing *as* onto ways of judging *that*.

One view of what we learn when we learn to distinguish between the way things seem and the way they are is that we learn to divide appearances into true appearances and deceiving appearances. The true appearances are the appearances which present things as they really are, i.e. the ways of seeing *as* which are to be thought of as ways of seeing *that*, and so when we are presented with a true appearance we are justified in asserting the truth of its content. When we are presented with an illusory appearance then all we can say is that things seem to be a certain way even though they are not really that way: we have a way of seeing *as* which doesn't map onto a way of judging *that*. Extreme examples of such illusory appearances are optical illusions, like the Muller-Lyer figure, and hallucinations of the kind Macbeth is in the grip of when he sees the infamous dagger. But these deceptions are just the tip of the iceberg: according to this kind of view, most of our everyday perceptions are in fact deceptive.

In the first chapter of *The Problems of Philosophy*, Bertrand Russell considers the variability of appearances. A table first thought to be a certain shade

aspect of what we value as creativity is expansion of the boundaries of valid seeing *as*, but this is expansion along intelligible lines.

of brown, on reflection is seen to have no inherent colour, since it appears to be different colours to different perceivers under different conditions. He says ‘When, in ordinary life, we speak of *the* colour of the table, we only mean the sort of colour which it will seem to have to a normal spectator from an ordinary point of view under usual conditions of light.’¹¹ He argues that in order to avoid favouritism, ‘we are compelled to deny that, in itself, the table has any one particular colour.’¹² The same applies to shape, and indeed all properties: the (surface of the) table might be thought to be rectangular but it only looks rectangular when seen from directly above (or below).

According to the view of the objective move we are considering, it must be possible to single out *the* colour and shape of the table. Perhaps Russell is wrong to describe the singling out of certain conditions as favouritism, though. Maybe appeal to normal conditions is legitimate. Maybe we can strip away the inessential, the distortive, leaving certain right ways of seeing *as*, certain normal conditions which are truth-conducive. The other ways of seeing *as* are then to be rejected as distortions of reality or figments of the imagination. The table is brown and rectangular and only the ways of seeing *as* which present it as being so are valid.

In order to make the move from appearance to truth we must be able to differentiate between normal, i.e. truth-conducive, conditions and other conditions. If it isn’t possible to do so, then we are left with certain conditions being given preference over others in an arbitrary way, the way the world is is given by the way the world appears under normal conditions, and normal conditions are just

¹¹ Ibid., p. 2.

¹² Ibid., p. 3.

those conditions which usually obtain, or those conditions which are average or universal. This leads to a view of objectivity as based on human agreement of the worst kind, appealing to the majority, and forgetting that there are objective truths only viewable from certain perspectives or by certain perceivers who are more perceptive than average.

Further, the often proposed equation between something's being red and its looking red under normal conditions is a mistake. Let us suppose that we have specified the normal conditions in some way and then let us ask whether something would necessarily be red if it looked red under those conditions. Suppose an object looked red under those conditions but then glowed yellow when we took it into non-normal conditions. Would it be correct to call the object 'red'? Surely not. In order for something to be red it must look the right way under *all* conditions - looking red under normal conditions is not enough.

Even if we could make a valid distinction between normal and abnormal conditions, there would be a problem accounting for most of our perceptual knowledge, according to this view. Recall that tables only appear strictly rectangular when viewed from directly above or below. If only those perceptions of tables from directly above or below are true appearances, then since none of us usually view tables from such awkward positions, it is difficult to explain how we do come to know anything about the shapes of tables. Russell would say that we infer the real rectangular shape from the non-rectangular appearance, that although perception is deceptive, we have a way of compensating for it.

But this cannot be right: the implication is that if we could exchange a deceptive appearance for a true appearance in every instance we'd be better off

epistemically, if we always saw tables from directly above or below we would never be deceived about their shape, for example. But if we only ever saw tables from directly above or below, we wouldn't have enough information to know what their real shape was - we would only have two two-dimensional views and would be in ignorance about the height dimension. We wouldn't know how long a table's legs were, and we wouldn't even know whether its main work area was of a suitable thickness, i.e. relatively thin. We wouldn't be able to distinguish between tables and thick blocks of wood on legs. With some shapes, there would be no non-distorting position from which to view them. Cubes, for example, would be completely unknowable, since if we were to view a cube from a non-distorting position, i.e. to look directly at one of its faces, then all we would see would be a square. Any other position which would give us a view of more than one face would be a distorting view, since it would not present the faces as square. With spheres, on the other hand, any view would yield a true appearance, since a sphere looks the same shape from all angles.

Russell gives us two options: either we accept some appearances and reject others on the basis of prejudice, or we reject all appearances, concluding that the senses do not give us truths about the world. There is a third option, though: we can accept all appearances; we can reject the first view of the relationship between ways of seeing *as* and judging *that* in favour of a many-to-one relationship. We can see the assertion that *a* is *F* as arising out of the totality of different ways that *a* can appear whilst being *F*. Russell is wrong when he equates the colour of the table with 'the sort of colour which it will seem to have to a normal spectator from an ordinary point of view under usual conditions of light.' It is wrong to say that a

brown table only seems brown in normal light: to someone who has a fuller understanding of brownness, a brown table will seem brown to them in all sorts of conditions. Likewise, it seems wrong to say that a coin only seems round when viewed from straight on, that the rest of the time it seems elliptical. This is wrong, even if in a sense we can describe the appearance as elliptical (e.g. when describing the shape we would paint if we were painting the coin from this angle): it doesn't *seem* elliptical, it seems round, because part of what it is to be round is to appear that way to perceivers like us when viewed from that angle. The proposition '*a* is *F*' is not to be understood as equivalent to '*a* is such that it always appears the same, i.e. strictly *F*': this characterises *F*-ness as one-dimensional, making it impossible for a coin ever to look 'round in an elliptical way' but having to present itself to the eye as a perfect circle from every angle. Nor is it to be taken as equivalent to '*a* seems *F* to normal perceivers under standard conditions': this fails to account for the fact that *F*-ness explains all the diverse appearances of *a*. '*a* is *F*' is to be thought of as irreducible, for all objective properties *F*, since to fully translate *F*-ness one would have to say something like '*a* seems *F*₁ in condition *c*₁, *F*₂ in *c*₂, ...' And the only thing tying *F*₁, *F*₂, ... together is *F*-ness.

To understand the concept *F* is not to only know how to recognise *F*'s in normal conditions: this is just the start. Deeper understanding of the nature of *F*-ness, requires knowledge of how *F*-ness manifests itself in many different conditions. Someone who can only identify a red book in daylight and flounders

when asked what its colour is in twilight or ultraviolet light has not yet grasped the concept *red* fully.¹³

1.3 Essence and naming

We have a many-to-one relationship, many possible ways of seeing *as* being mapped onto any one way of seeing or judging *that*. Redness is the property responsible for *all* of these appearances; redness is greater than all the actual appearances we have experienced; redness has a life of its own, being independent of any actual appearance. In naming something 'redness', all we have done is labelled *whatever is responsible for all of these appearances*. The thing we have picked out can still surprise us at any point in the future and we may never completely grasp its essence.

This view of the relationship between a concept and its instantiations denies the possibility of giving an analysis of the meaning of a concept by giving necessary and sufficient conditions for its application. Essence is to be thought of as multifaceted, not as an essential core. This is in opposition to a view such as Kripke's, according to which discovering the essence of something is finding just those aspects of it which are necessarily so, unchangeable, the essence of water being taken to be its chemical composition, for example. According to a Kripke, once we have found out that water is H₂O we have found out the essential nature of water.¹⁴ According to my view, the essence of water is not just this feature,¹⁵ the essence of water is much too big to be tied down to one definition.¹⁶

¹³ In the background here is the idea that conceptual matters such as knowledge of how to apply a concept ϕ blend into empirical matters concerning the nature of ϕ 's. This will seem problematic to philosophers who remain unconvinced by Quine's arguments against the analytic/synthetic distinction.

¹⁴ See his discussion of natural kinds in *Naming and Necessity*, pp. 116-134.

But then the essence of redness might seem to be unknowable, to be part of the unknowable underlying reality previously rejected. Are we not back to postulating an unknowable underlying reality? No: what I have in mind is to be thought of as imaginary construction rather than underlying structure. The essence of redness is not to be thought of as something more real, something somehow redder than any of the red appearances we are presented with, but as a logical entity analogous to a mathematical construction like a geometric point or the limit of an infinite series, it is the point of convergence of all red appearances. The essence of redness is given by the totality of possible red appearances (and this includes the scientific as well as the perceptual appearances). Just as the limit of a series may never be reached but is given or pointed to by the way the function approaches it, redness is known by the nature of its appearances.

Whenever we pick out an object or property and say anything about it we pick out *that* object or property. All names of things with objective existence, whether physical objects, properties, feelings, values, songs, etc., act as demonstratives. It is only when the object being picked out doesn't have objectivity, doesn't have a life of its own, that it can be completely pinned down by description. To attach one description to an object essentially, as Kripke does,¹⁷

¹⁵ As Hilary Putnam has suggested, it is only true that water is necessarily H₂O if we rule out possible worlds in which the laws of nature are different. In a chemically different possible world, we might not want to say that water was H₂O (see *Realism with a Human Face*, pp. 69-70).

¹⁶ Too big to be tied to a cluster of definitions either: essences are uncodifiable.

¹⁷ I am sure Kripke would protest against my characterising his view in this way - he denies that proper names and substance terms have a sense, which is, as he sees things, equivalent to denying that they are shorthand for a description or cluster of descriptions. Kripke rejects the idea that the meaning of proper names or substance terms can be given in a description, ruling out the possibility of truths like 'Aristotle was the pupil of Plato' being true necessarily in virtue of the meaning of the name. So this one dimension of Aristotle is not forced. However, his stipulation that truths of constitution are necessary means that the physical or chemical dimension of an object or substance is taken as being its essence, everything else about it being seen as inessential. This is to make essence one-dimensional, and that is what I object to. I am not against the idea of distinguishing between necessary and non-necessary truths in general, but if 'possible worlds are

is to refuse to recognise the object's autonomy, its objectivity. The only time we can attach a description to something essentially is when a thing is originally defined using a description. Suppose we stipulate that the name *N* refer to just whatever fits description *d*. Then one of two things will happen: either *N* will be seen to have various other properties, will be observable or describable in many ways - it will be seen to have more than one dimension - or *N* will not come to be seen as anything more than just the thing that fits description *d* - either because there is no one thing picked out by description *d*, or because whatever is picked out by description *d* is wholly described by the description *d*. In this case we have not succeeded in naming anything at all.¹⁸

Imagine a group of underground explorers who come to a cave, see a flash of strangely coloured light and decide to investigate the cause of the light. They use the name 'X' to refer to whatever caused the flash of light. Then, if the cause of the flash of light is found to be a substance previously undiscovered, then the name 'X' will no longer be tied to the description 'the cause of the flash of light', it will apply to *that* substance and the description 'the cause of the flash of light' will come to be seen as just one of the many true descriptions of X. If nothing which accounts for the flash of light is discovered in the cave and the strange-coloured light is never seen again, then the name 'X' will die out.¹⁹

stipulated, not *discovered* by powerful telescopes' as Kripke tells us (ibid., p. 44), then the question is, why does he deny the validity of stipulating possible worlds in which the laws of physics are different?

¹⁸ This is true of the things we invent as well as the things we discover. If we invent a game and it takes off, then the fact that we invented it in the first place will do nothing to detract from its objectivity.

¹⁹ This second possibility is what happens when the existence of something is denied, e.g. the case of Phlogiston. (It might be the case that more than one description applies to a proposed thing and it still fail, e.g. there might have been a set of descriptions that were thought to apply to Phlogiston. The point is that if a name is to pick out something real there must be an indefinite number of possible descriptions that apply to it.)

So either a named thing outgrows the description originally used to pinpoint it, or it fails to get off the ground as an objective thing at all. If something has objective reality then it can't be pinned down to one description or way of seeming. Thus being objective is being two things: firstly, as we have discussed, it is essential to the objectivity of a thing that its nature is knowable from many different viewpoints, not just from one privileged view. This means that the objectivity of something grounds many ways of seeing *as*. Secondly, as well as a thing's objectivity being a guarantee of the validity of different ways of seeing it, its objectivity is what makes it possible for us to make mistakes about its nature. It is because something has objectivity that it can seem ϕ to us without actually being ϕ . This can seem paradoxical: objectivity is what allows for all ways of seeing *a* to be ways of seeing that *a* is *F* (whilst *a* remains *F*) and also what allows for the possibility of saying that although *a* seemed *F* it wasn't really *F*. The paradox dissolves when we realise that even if all ways of seeing *as* lead to ways of seeing *that*, not all ways of seeing *as* ϕ lead to ways of seeing that ϕ . The objective move translates ways of seeing *as* into ways of seeing *that*, but the rules of translation do not always leave the content of experience unchanged.

Whenever things surprise us, whenever we discover a situation in which the way things seem and the way things are comes apart, we look for an explanation of why our experience was misleading and thus shed light on the nature of the thing experienced. For example, the discovery that straight sticks looked bent when dipped into water led not to a dismissal of our perception of sticks in these circumstances, but to an account of the nature of refraction. Then our perception is altered by understanding: the stick no longer seems mysteriously

bent once we understand why its straightness and the difference in density between air and water causes the appearance we are presented with. With understanding, we come to see the stick as straight when we see it standing in a glass of water, just as we come to see coins as round from all angles.

How do we differentiate between discovering a different way of seeing that a is F and discovering that a has changed and is no longer F ? What makes us group together certain appearances under the concept F and reject others? What stops us from concluding that just as when a straw is dipped in a glass of water it remains straight, that when we bend a straw it remains straight? What makes this a change in the object and not just another way of it seeming straight? Someone might say that the difference is that when we dip a straw in a glass of water and it looks bent, we can then remove it from the water and see that it is still straight. But this is no good - after bending the straw I can bend it back too. There are reversible changes, so we can't differentiate between a change and a new way of seeing that something is the same in this way.

Whether a has changed or merely shown us another side of its nature is not something that can be determined empirically, because it is what determines what we say about our experiences. A concept F groups together certain things - the F 's - but not because there is something absolute that dictates that we group things in this way: given any selection of things, a , b , c , ... there are any number of possible concepts these things could be instantiations of. Suppose we are given a collection of round green objects to take as examples of F 's: are we to take F to pick out round things or green things? We might ask for further examples of F 's; we might ask whether a square green object was F . But if we found out that it was, would

that prove that F meant 'green'? No: it's entirely possible that triangular green things wouldn't count as F 's. There are all sorts of ways in which things can be grouped together as similar. It is a fact of human nature that we do group things in certain ways rather than others. Whatever the reason we group things together in the ways we do, i.e. develop the concepts we do, it can't be only because of the way the world presents itself to us.²⁰

1.4 Different ways of seeing the same thing

So we have lots of different ways of seeing that a is F , different routes to the same fact. How are we to characterise the differences? If S comes to believe that a is F on the basis of seeing a in one situation and S' comes to believe the same thing on the basis of seeing a in a completely different setting, or on the basis of being told that a is F by someone, then although S and S' have acquired the same fact, is there not some epistemological difference between them? Suppose that S knows that the earth is round because she has seen it from space and S' knows the earth is round because she has been told so. Surely what S knows surpasses what S' knows; S knows that the earth is round in a way that is deeper or richer than the way S' knows it, and this is either to be thought of in terms of there being more facts connected with the roundness of the earth that S knows, or in terms of there being more to knowledge than just the resultant fact, i.e. there being differences in knowledge which are finer grained than differences in facts acquired. In either case, we will need to refer to non-factual differences, I will argue. The differences between S and S' will be characterised as differences in non-conceptual content.

²⁰ This summarises the argument Wittgenstein presents against what Colin McGinn calls 'self-identifying signs'. The conclusion of the argument is that a rule such as 'use the concept F to apply

Suppose we try to characterise the difference between what *S* knows and what *S'* knows as a difference due to the fact that *S* not only knows that the earth is round, but also knows the truth of a whole set of related propositions about the earth that *S'* doesn't know. But then let us list these related truths and then pass on the list to *S'*, then *S'* will come to know all the related things *S* knows. Would we then say that there's no epistemological difference between them? It seems clear that there is still a difference. Although *S* and *S'* both possess the concept *earth* and both understand the proposition 'The earth is round', *S*, being an astronaut, seems to have a better, richer understanding of the concept *earth* and so of the proposition. Part of what makes *S*'s understanding better can be put down to the fact that *S* possesses more astronomy-related concepts and knows more facts of astronomy than *S'*. But the problem is not solved in terms of the number of concepts and propositions understood, because there will still be a possible difference between *S* and someone else who understands all these extra concepts and propositions but has a different level or depth of understanding.

The possibility of epistemic difference without a difference in facts known is important as a counter to over-enthusiastic reductionism - being able to reduce one set of propositions to another is not necessarily being able to do away with the need for the first set. Talk of different ways of knowing the same thing shades into talk of different things known - if *S* knows *p* in one way and *S'* knows *p* in another way, then in a sense they know different things. But we cannot express the difference easily - we cannot pinpoint one thing, *p*, known by *S* and another thing, *p'*, known by *S'*.

to things like these' cannot logically rule out all but one course of action (cf. *Philosophical Investigations* §§. 85-201).

The Fregean notion of sense as mode of presentation might be thought to be helpful here. The notion of sense, as put forward by Frege, is the solution to the problem of how ' $a = a$ ' and ' $a = b$ ' can share truth conditions and yet express different thoughts.²¹ This will not stretch quite far enough for our purposes, though. According to Frege the sense (mode of presentation) of ' a ' is objective, i.e. all those who understand ' a ' grasp the same sense.²² Frege only allows there to be as many ways of asserting the fact that a is F as there are combinations of different terms which refer to a and F . (If $a = b$ and $F = G$ then there are 4 ways.) What the many-to-one mapping from ways of seeing as to ways of seeing *that* requires is a way of characterising the difference between different people's ways of coming to assert ' a is F '.

Our understanding of all but the simplest, thinnest²³ concepts changes and improves as we encounter new situations in which concepts apply (or fail to apply). As well as there being a distinction between understanding and lack of understanding, there are degrees of understanding. A child who has a basic grasp of a concept ϕ is capable of understanding some truths about ϕ 's, but does not yet have a full understanding of the nature of ϕ 's. How do we characterise the difference between the child's understanding of the concept ϕ and a more complete adult understanding of it? Perhaps the difference can be given in terms of truths

²¹ See 'On Sense and Meaning' (in Peter Geach and Max Black eds. *Translations from the Philosophical Writings of Gottlob Frege*).

²² Although he has problems with this in 'The Thought: A Logical Inquiry' when considering the possibility that two people utter different thoughts if one of them says 'I have been wounded' and another says 'Dr. Gustav Lauben has been wounded'. This seems to imply that two other people who think of Dr. Lauben differently but who both use the name 'Dr. Lauben' to refer to him utter different thoughts when they say 'Dr. Lauben has been wounded'. Frege seems to completely avoid dealing with this problem, saying 'as far as the proper name 'Dr. Gustav Lauben' is concerned, [they] do not speak the same language, since, although they do in fact refer to the same man with this name, they do not know that they do so.' (p. 25).

known about ϕ 's: the adult knows more truths about ϕ 's than the child does. So then what the child needs to grasp is the truth of more propositions. Suppose the child knows that a certain book is red but doesn't yet know what the book would look like when lit by a coloured street light. When she first sees the book in this way, she does not acquire a new truth - that the book is red is already known to her - so how can the increase in her knowledge be characterised in terms of facts?

Perhaps the answer is that in addition to coming to see the truth about the colour of the book in a new way, the child acquires a truth with subjective content; she acquires a truth about the way the book seems to her under certain conditions. But how are we to specify the content of this truth without being misleading? We shouldn't say that she acquires the truth that the book looks orange when lit by orange street lights - the book doesn't look orange any more than a coin looks elliptical when viewed from an angle. The truth she acquires is a truth about the way red things look in orange light. But spelling out 'the way red things look in orange light' is the difficult part. The child is learning a new way of looking red, an orangey way of looking red, we might say (an elliptical way of looking round). If we want to say that the book seems red to her whilst seeming orange to her, then we must say that there is one sense in which the book seems red and another sense in which it seems orange if we are to avoid paradox. The natural thought is that the seeming red is to be taken as conceptual and the orangey aspect is to be taken as non-conceptual.²⁴

²³ Bernard Williams uses the term 'thick concepts' to refer to specifically ethical concepts (see *Ethics and the Limits of Philosophy*, pp. 140 ff.). I see thickness as a general property of concepts - some concepts are thicker than others. It is not surprising that certain ethical concepts are thick.

²⁴ My argument is derived from the argument Christopher Peacocke presents in chapter 1 of *Sense and Content* to the conclusion that experience has representational and sensational properties.

1.5 Conceptual experience and the dawning of aspects

What is the change that takes place when we learn to experience ϕ 's as ϕ 's, to see and hear words as words, for example? When we learn concepts we come to experience the things falling under them in a different way, we come to experience them as doing so. We come to see that something falls under a concept and what happens is the same kind of thing as what happens when we discover a previously unnoticed aspect of something. What happens when one looks at a figure such as the duck-rabbit figure²⁵ and notices an aspect not previously seen? The figure seems to change before one's eyes, and yet it clearly does not *really* change. One sees the figure as a picture of a duck, then as a picture of a rabbit, and this change is conjoined with a continuity in seeing - you see that the figure remains unchanged. Wittgenstein reports this with clarity:

What is incomprehensible is that *nothing*, and yet *everything*, has changed, after all. That is the only way to put it. Surely *this* way is wrong: it has not changed in one respect, but has in another. There would be nothing strange about that. But "Nothing has changed" means: Although I have no right to change my report about what I saw, since I see the same things now as before - still, I am incomprehensibly compelled to report completely different things, one after the other.²⁶

What happens when one of the aspects of the duck-rabbit figure dawns upon one also happens when one comes to see anything in a new way, whenever one gains a new way of seeing *as*. On grasping a concept one's experiences often change in this way: one comes to see things falling under that concept as doing so; one learns

²⁵ This figure, which can be seen as a duck or as a rabbit, is brought up by Wittgenstein in the *Investigations* p. 194.

²⁶ *Remarks on the Philosophy of Psychology Volume II*, §.474. The core of Wittgenstein's treatment of aspectival perception is to be found in *Philosophical Investigations* II xi and *Remarks on the Philosophy of Psychology Volume I*, §§.506-546 and *Volume II* §§.473-549.

that what one is seeing is a case of seeing that *a* is *F*. One's experience becomes experience *that*, i.e. it becomes conceptual.²⁷

There are some experiences which only seem open to those with certain concepts: seeing a triangle as an isosceles triangle doesn't seem open to someone without an understanding of what makes a triangle an isosceles triangle, without the concept *isosceles triangle*, for example. As Wittgenstein puts it, 'The substratum of this experience is the mastery of a technique.'²⁸ A better example, also given by Wittgenstein, is the way our visual and aural perceptions of words change as a result of learning a language. Once we have learnt a language it is near-impossible to listen to someone speaking that language and hear their words as meaningless noises, to avoid hearing their words as words. The same happens if we try to look at written text and see it as a series of meaningless squiggles or shapes. On the other hand, there are some aspects of experience which can be separated from the conceptual, which can be experienced by those lacking concepts and which don't seem to be qualitatively affected by our learning the relevant concepts. Examples are pain - a baby can experience pain and it's implausible that the pain of a pre-linguistic baby differs in quality from that of a child who has learnt the concept 'pain'. Similarly, to say that the taste of sugar is different once one has learnt that it is sugar one is tasting, that the concept *sugar* is required to taste it *as* sugar, is also far-fetched.

²⁷ There is much argument about whether or not experience is conceptual. Someone who disputed that experience was conceptual would say that we should draw the line between perception and judgement, with seeing on one side and seeing *that* on the other. Then 'seeing *that*' would be thought of as a misleading phrase, better interpreted as 'judging *that* (on the basis of what one sees)'. I take the arguments McDowell presents against the possibility of appealing to an 'extra-conceptual Given' for justification to show that experience must be conceptual. I disagree with McDowell that this makes experience completely conceptual.

²⁸ *Philosophical Investigations*, p. 208.

I shall draw a distinction between conceptual and non-conceptual objects of experience in this way: something, ϕ , is conceptual if one must have certain concepts, e.g. the concept ϕ , in order to experience something as ϕ . According to this definition, once we have learnt a language all of our experiences will have some conceptual and some non-conceptual elements. What happens when one notices an aspect or comes to see something as falling under a new concept is that at least some of the non-conceptual elements remain unchanged whilst the new aspect or conceptual element appears.

Sometimes learning a concept leads to the ability to see things one would not have been able to see before, to differentiate between things which were previously indistinguishable. One learns to see things which were under one's nose but which one was previously blind to: in learning a language one learns to hear words where previously one had only heard noises, for example; in learning emotional concepts, e.g. in learning the difference between jealousy and envy, one comes to notice aspects of one's emotional life one previously passed over, one comes to be able to differentiate between two previously indiscernible kinds of feelings, for example. The development of emotional concepts seems to facilitate development in the emotions themselves: the ability to refer to an emotion draws that emotion out into the foreground and one can study it more attentively. Similarly with the acquisition of the concepts involved in understanding music, or in wine tasting: the difference between a good vintage and a bad one might be undetectable to those who haven't mastered the right conceptual framework. Experience of meaningfulness is conceptual: a minimal condition for experiencing something as meaningful is that one has the concept *meaningful*, so if there is a

problem about the objective reality of conceptual properties then there is a problem concerning meaningfulness.

One important facet of the phenomenon of the dawning of an aspect or the grasping of a concept is the sense of discovery - the duck in the duck-rabbit figure was there all along, there was something common to chestnut trees before one learnt to group them together, to see them as the same kind of tree. And the sense in which the thing discovered was 'there all along' is independent of whether anyone had previously noticed it. Suppose the duck-rabbit figure had been originally drawn by someone aiming to sketch the outline of a rabbit and then someone else, not knowing what was being sketched had seen it as a duck, we would say that the duck aspect had been there even though its existence had been an accident. With concepts, what happens when a child learns to see something as a ϕ for the first time seems to parallel what happens when the concept ϕ is first used, when ϕ 's are first noticed. In both cases there seems to be a discovery of something that was there all along, even though in the second case the thing that was there all along had not been noticed by anyone.

If there was not a sense in which things remained constant before and after the dawning of an aspect there wouldn't be the possibility of coming to discover something new about a familiar thing. If there was no continuity between the thing before and the thing after the aspect shift, then what we would have would be a situation in which the thing changes, in which one thing disappears and another appears in its place: we would be observing a causal chain of events, not noticing something about the way things were all along. What is the difference between the duck-rabbit figure and a toy that would present a picture of a duck when a button

was depressed and then flick to a picture of a rabbit when the button was released? In what sense is the noticing of an aspect of the duck-rabbit figure noticing something that was there all along? What makes it true to say that the object hasn't changed but only one's perceptions of it?

When we notice something different, we ask: has a change occurred or have we discovered something that was there all along? If a change has occurred then we will conclude that *a* is no longer *F*. If not, we will conclude that *a* is still *F* and the notion of *F*-ness will expand to include the new kind of case. In the first case the change occurs in the object and requires a causal explanation. In the second case, the change occurs in our understanding of the meaning of the term and requires no causal explanation - we have discovered a new fact (or a new way of seeing the same fact). But what about when we notice something metaphorical in an object?

When we notice the duck aspect of the duck-rabbit figure we notice that the figure is of a duck. What about when we look at the clouds in the sky and see one of them as a duck? In general, there are ways of seeing *as* which don't seem to yield literal truths or ways of seeing that, but yield metaphorical truths. What is the difference between a typical observation statement like 'That bird is a sparrow' and a metaphorical one like 'The orphan was a little sparrow'? One answer is that the first kind of statement latches onto something objective whereas the second doesn't. The first expresses a truth because the property ascribed to the object is actually possessed by that object whereas the second is not literally true, it is creative; it links a property not actually possessed by a thing with it and somehow conveys information. This answer seems mistaken: metaphors can become truths

and when they do it doesn't seem that an objective property is acquired by the object. When an apt metaphor catches on, the metaphorical use of a word becomes part of the word's literal meaning. Examples of the metaphorical expansion of words are the use of the word 'in' in non-spatial contexts like 'It's in your mind - it's a figment of your imagination!' and the words 'crow', 'blue', 'thread' as in 'She's such an old crow', 'I'm feeling blue', 'I've lost the thread of your argument'. If the orphan does not have the property of being a sparrow then how can the old woman have the property of being a crow? In what relevant way does the woman I now quite literally call a crow differ from the old woman first called a crow metaphorically? If one is inclined to say that even now the old woman may be called a crow but she is not *really* a crow, one should think of other examples such as when we talk about someone's being in a position of power. It is literally true that some people are in positions of power - the fact that we are not locating someone spatially when we describe them as being in a position of power doesn't imply that we are using the phrase 'in a position' metaphorically. Something can start off as a metaphorical description and then, because it is insightful, catch on and become taken as literally true.²⁹

We use words in new contexts to grasp at truths about the intangible; we use metaphors to express truths we don't have the words to express literally - we call music sad, conversational, violent, for example. As a result we develop some understanding of things which are subtle and difficult to understand. These ascriptions, when appropriate, are surely true: some music really is passionate and someone who objected "But music can't be passionate: in order to be passionate

²⁹ Cf. Nietzsche: 'Truths are ... metaphors which are worn out and without sensuous power ... to be truthful means using the customary metaphors' ('On Truth and Lie in the Extra-Moral Sense' in

something must be a being capable of experiencing passions” would be making the mistake of attaching a necessary condition to the concept *passionate*. We should reply to such a person “No, that is just one aspect of the concept *passionate*, a central aspect admittedly, but there are others - look, can’t you see that this music is also passionate?”

What is it for something to really have a property ϕ , then? Either the facts shift, falsehoods becoming truths, or the facts are independent of our expression of them, sameness of sentence structure needn’t imply sameness of fact expressed. What is the fact expressed by ‘*a is F*’? We cannot describe it as the fact that *a is F*, because the point in question is whether the same fact is being expressed now by ‘*a is F*’ as before.³⁰ Can we describe it at all? I suggest that we cannot (although we can refer to it as “the fact now expressed by ‘*a is F*’”).

A number of questions have been brought up only to remain unanswered. I have been unable to present more than the bare bones of a view of the relationship between experience and language. It has been important to show how language, whilst being viewpoint-transcendent in that the truths we utter are true independently of the viewpoints which led us to them, draws on experience. It was emphasised that understanding is more finely-grained than knowledge of truth values: understanding is more than just knowing the facts and this means that the different viewpoints and ways of coming to see truths are important in ways that do not stand out if we view knowledge acquisition as nothing more than fact-collecting.

The Portable Nietzsche, p. 47).

³⁰ As well as the idea that the same sentence can express different facts, it is true that different sentences can express the same fact. David Bell argues that the fact that different sentences can express the same (Fregean) thought forces us ‘to deny that thoughts *have* a determinate, intrinsic structure.’ (See ‘The Formation of Concepts and the Structure of Thought’, p. 596.)

2. The relationship between language and reality

We have seen that truth is viewpoint-independent, that language gives us an objective world. But is truth independent of language? If we had developed different concepts would we have been presented with a different objective world? The possibility of experiencing certain things requires the possession of concepts but it is sometimes suggested that it is not just experience of ϕ 's that is dependent on the concept ϕ ; that ϕ 's wouldn't have existed had we not developed the concept ϕ . It seems to us that when we notice an aspect or come to see something as a ϕ as a result of learning the concept ϕ that we *discover* that things are a certain way. Is this right? I now turn to the relationship between reality and conceptualisation, asking whether we discover or create essence.

The things our words pick out might seem to be dependent on our particular conceptual scheme. The things our concepts pick out, the property of greenness, for example, seem real/objective but what if we had developed different concepts, would they still have existed? If we had developed different concepts, such as *grue* and *bleen*³¹ instead of *blue* and *green*, then we would have 'discovered' different essences, grueness and bleeness instead of blueness and greenness. Doesn't this threaten the reality of the things our words do pick out? Are the objects of our discourse created or discovered by us? If greenness is a real property then surely we must say that had we developed the concept *grue* instead then grueness would've been a real property? Either we find a reason for giving preferential treatment to greenness, saying that greenness is real but grueness isn't,

³¹ These 'bent' concepts are introduced by Nelson Goodman in *Fact, Fiction and Forecast*. The predicate 'grue' 'applies to all things examined before [time] *t* just in case they are green but to other things just in case they are blue.' (p. 74). 'Bleen' is the complementary predicate, covering blue things examined before *t* and green things examined after *t*.

or we say that both greenness and grueness are real, or we say that what is real is relative to scheme, that greenness is real relative to our conceptual scheme and grueness is (or would be) real relative to a scheme with the concept *grue*.

In what sense could objects in the world be dependent on conceptual scheme? What can the idea that existence is relative to conceptual scheme mean? If certain objects have objective existence, then their existence can't be relative to conceptual scheme. If the facts are objective then they can't depend on conceptual scheme in the sense of holding relative to scheme. Is the relationship between scheme and the truths expressible in it parallel to the relationship between viewpoint and the truths ascertainable from it, as discussed in part I? We can allow that there are some facts only expressible in certain schemes: if a scheme has no concept similar enough to the concept *round*, for example, then the fact that the earth is round cannot be expressed in it. But if no scheme with the concept *round* existed, then would it be accurate to say that there were facts about roundness which couldn't be expressed or would this be for there to be no facts about roundness, for nothing to be round?

I will argue against the idea that existence is relative to conceptual scheme, aiming to develop a view of the relationship between our concepts and things in the world which parallels the view I developed of the relationship between our experience and the world. In part I I emphasised the importance of the subjective whilst claiming that truths ascertained subjectively are objectively true. In this section I will emphasise the importance of conceptualisation whilst claiming that the truths we discover are true objectively. I hope to develop a realism that does justice to the intuitions behind conceptualism.

2.1 Scheme-independent or preconceptual reality

It might be thought that the two questions ‘Do the things we refer to exist independently of our conceptual scheme?’ and ‘Are the truths expressible from within our conceptual scheme true independently of that scheme?’ stand and fall together. Someone might think that if existence was scheme-dependent, i.e. if certain things wouldn’t have existed had we not developed our conceptual scheme, then truth must be scheme-dependent too - there wouldn’t have been truths about those things had we not developed our conceptual scheme, since they wouldn’t have existed. This is wrong, though: there are scheme-independent truths about even those things we clearly do invent, such as chess. If it is true that there have been no games of chess played on the moon, then it is scheme-independently true, for example. Even if it turns out that we must concede that *existence* is relative to scheme, we might still be able to avoid the conclusion that *truth* is relative to scheme. On the other hand, if existence is scheme-independent, at least some truths will also be: if certain objects, properties, etc., have scheme-independent existence, then the truths about their existence will be scheme-independent. I will concentrate on the possibility of scheme-independent existence.

There are three views of the possibility of scheme-independent or preconceptual existence: the conceptualist view according to which nothing exists independently of conceptualisation,³² the metaphysical realist view according to which some of the things we talk about exist independently of conceptualisation

³² Hilary Putnam’s ‘internal realism’ is an example of conceptualism: according to his view ‘*what objects does the world consist of?*’ is a question that it only makes sense to ask within a theory or description.’ (*Reason, Truth and History*, p. 49).

and others don't,³³ and the more egalitarian realist view according to which nothing depends on conceptualisation for existence.

The conceptualist claim seems obviously false if the questions we are asking about the priority of essence or conceptualisation are taken to be questions about what existed before human beings started conceptualising. If we say that essence is created by language not uncovered by it, then we can be thought to be telling a story of how things came about as a result of language, and saying that prior to the existence of language nothing existed. If this were true then words like 'dinosaur' would have no meaning. If nothing existed prior to language, then evolutionary theory would be false and its talk would either be nonsensical or empty. This way of seeing the conceptualist thesis makes it seem to be competing with science (and religion) in offering an account of how things came into existence.³⁴

If the conceptualist thesis is not a thesis about the temporal order of things, then how is it to be understood? Maybe it is to be understood as saying that although the stuff of the world existed prior to conceptualisation, it took concepts to organise it or group it together into things like dinosaurs. The metaphor of noumenal dough that our concepts cut into cookies could be invoked at this point, but that metaphor is misleading: if the noumenal dough is thought to be uniform,

³³ Michael Ayers' realism is a clear example of metaphysical realism: according to his view the individuality of physical objects is 'peculiarly "absolute" and independent of human concepts or ways of looking at the world.', and unlike events, properties, etc., physical objects 'are natural unities or natural structures which come into existence, continue to exist and cease to exist quite independently of any conceptualising on our part.' ('Individuals without Sortals', p. 114).

³⁴ Cf. Elizabeth Anscombe: 'if we say the essence expressed by the grammar of the word 'red' is itself the creation or product of that grammar aren't we saying that nothing would have been red if there had not been human language? ... And *that* we do not want to say.' (in 'The Question of Linguistic Idealism', p. 113). To Anscombe the project of describing the relationship between essence and grammar is the steering of a course between 'the falsehoods of idealism and the stupidities of empiricist realism' (p. 115). The realist claim that essence determines grammar must

like normal dough, then it misses the point that before concepts there were differences between the stuff we cut dinosaurs out of and the stuff we cut days out of.³⁵ The dough cannot be uniform and if it is not, then the metaphor fails to be appropriate, we must resort to a metaphor such as the metaphor of a spectrum to be divided into discrete groups. But then we are forced to accept the existence of preconceptual structure, even if we deny explicit preconceptual distinctions.

And, further, it seems wrong to see everything according to this metaphor: it seems wrong to say that before we developed the concept *dinosaur* there were no dinosaurs, that there was only dinosaur stuff. Consider the difference between dinosaurs and weeks. Dinosaurs seem to be natural unities but weeks are a contrivance: nothing divided time into weeks before we did, until we created the concept *week*. It might be thought that dinosaurs, planets, days, etc., have preconceptual existence whereas weeks, marriages, values, etc., do not. What is the substance of this distinction though?

Metaphysical realists such as Michael Ayers draw a distinction between the things that exist as natural unities and those things our concepts unify. In 'Individuals Without Sortals'³⁶ Ayers argues for a realist theory of identity, i.e. he argues that identity is not relative to conceptualisation in the case of physical objects. This is because 'physical objects are natural unities or natural structures

also be false, since taken in this way it too implies that we would be able to do natural science just by studying our concepts.

³⁵ There are other metaphors on offer: Davidson considers the metaphor of conceptual schemes organising reality and the metaphor of scheme fitting reality (see 'On the Very Idea of a Conceptual Scheme' pp. 191 ff.). His argument against the coherence of the idea of organising a reality which doesn't already contain objects to be organised seems valid but the idea of concepts carving up reality remains unchallenged: we can carve up anything as long as it exists and our knife is sharp enough!

³⁶ In *Canadian Journal of Philosophy*, Volume IV, No. 1.

which come into existence, continue to exist and cease to exist quite independently of any conceptualising on our part.³⁷ Things are different for modes, though.

Substances, i.e. physical objects, are contrasted with modes, i.e. attributes, properties, events, etc. Ayers argues that substances are to be seen as natural or real individuals, i.e. individuals which are ““given” not “constructed””³⁸, whilst modes are mind-dependent in that they are individuated by a ‘distinction of reason’ not by a real distinction, their individuality being ‘posterior to their individuation by us.’³⁹ When we consider episodes, he says, it is our concepts that decide whether something counts as an effect or a part of the episode. He says ‘The only reason why the breeze now blowing over England is a result, rather than an element of the hurricane out to sea may be that it is too gentle to count for us as the latter.’⁴⁰ Substances are natural wholes, though. He asserts: ‘To say that a horse is a natural whole is to say that it is not a part of anything in the way in which its natural parts are parts of it.’⁴¹

Ayers starts off with natural unities like horses and then widens the field to include artefacts and mere lumps of matter. The idea of natural unity is too stringent to include these things so Ayers waters it down from the biological notion to the idea of ‘simple cohesion’,⁴² the overriding idea being that of causal unity. Presumably what Ayers means by ‘causal unity’ is something like ‘the kind of unit which appears in causal laws’, but then it is hard to see why events are not natural unities.

³⁷ Ibid., p. 114.

³⁸ ‘Substance: Prolegomena to a Realist Theory of Identity’ in *Journal of Philosophy*, 1991, p. 69.

³⁹ Ibid., p. 71.

⁴⁰ Ibid., p. 74.

⁴¹ Ibid., p. 74.

⁴² ‘Individuals Without Sortals’, p. 115.

Ayers says that things or substances are ‘objectively discrete unities’⁴³ which can be individuated without reference to concepts. He argues that whilst Frege’s claim that the question ‘how many?’ cannot be answered without knowing the answer to the question ‘how many *what?*’ holds for some objects of thought, it does not hold for substances. Frege gives the example of being asked to give the number of a pile of playing cards: the answers ‘52’, ‘4’ and ‘one’ are all answers to the question as it stands; there are 52 cards, 4 suits and one pack.⁴⁴ Ayers wants to say that in this situation it is objectively true that there are 52 things. This is because of the difference between objects like cards and groups like packs: groups are insubstantial, merely being brought together artificially by the mind, whereas objects are objective naturally occurring unities.

Ayers means to separate out one set of individuals as basic objects and this will not work. It is important to him that natural unities like horses are not themselves parts of anything else in the way that their parts are parts of them, since if they were parts of some greater natural unity then it would be up to us whether to take a horse or the greater natural unity it was part of as the natural individual. What tells most strongly against his position is the fact that many natural unities do have definite parts which are themselves natural unities. Ayers considers a chair made of seventeen pieces of wood and also the eyes of flies, flies having composite eyes, and says ‘the possibility of such numerable parts may itself be thought to cast doubt on the possibility of definitely numerable objects: should the chair count for one or for seventeen?’⁴⁵ Ayers seems not to see the force of this objection. His reply is that this is not a worry about objects like chairs since the

⁴³ Ibid., p. 128.

⁴⁴ See *The Foundations of Arithmetic*, pp. 28 ff.

same possibility applies to biological or strictly natural objects: the question ‘How many eyes does a fly have?’ can be answered ‘ten simple and two compound’,⁴⁶ for example. He is just missing the point, though: the fact that the number of eyes of a fly cannot be given independently of stating whether we mean compound or simple eyes is not something he can allow any more than the fact that the chair could be taken as one thing or seventeen things is. The most that Ayers can hang on to is the importance of natural unities but he must come to see that there are many different natural principles of unity, and this means that he must give in to the idea he wants to argue against: the idea ‘that no object is “really” or absolutely or utterly discrete’.⁴⁷ There are many real principles of unity, so there are many objects which are ‘really discrete’, but there is not one ultimate such principle, so we cannot single out some of these objects as ‘absolutely or utterly discrete’. Nature sees not just physical objects as unities: natural selection individuates species, genes, properties, and so on; causation seems to individuate properties and events as well as physical objects; physics individuates subatomic particles, waves, forces, and so on.

Ayers’ distinction between substances as scheme-independent individuals and modes as conceptually individuated will not work but the discussion of natural principles of unity suggests an argument against conceptualism. Without falling back on the idea of *the* absolute structure of reality, we can acknowledge the existence of preconceptual structure: there were natural principles of unity prior to conceptualisation. We can distinguish between natural or non-conceptual unities and conceptual unities in the following way: something is a natural unity if it is

⁴⁵ Ibid., p. 140.

⁴⁶ Ibid., p. 139.

picked out by a natural principle of unity; something is a conceptual unity if it is not, if its existence is essentially tied to the existence of a concept or group of concepts.⁴⁸ So, returning to dinosaurs and weeks: we see that dinosaurs are natural unities and weeks are conceptual unities - it might seem that weeks wouldn't have existed if we hadn't developed the concept. Days are natural unities too: a day is the time it takes the earth to make a revolution on its axis; colour and shape in general seem to be natural, but particular colours and shapes, like burnt sienna and regular heptagons, might be conceptual; numbers and other mathematical entities seem to be conceptual too. This distinction, although in need of tightening up, does seem valid and the existence of natural unities means that conceptualism as a general thesis is to be rejected.

But does this mean that the metaphysical realist thesis that there is an absolute preconceptual structure is true? No, for two reasons: firstly, we have argued that there is not one monolithic structure of reality, that all distinctions are distinctions made by something, whether a person, a non-human animal, 'mother nature' or gravity, etc. The structure of reality *sub specie aeternitatis*, whatever that is, is thus not the only real structure. We can think of the question 'Is the world organised into ϕ 's independently of our conceptualisation?' as equivalent to the question 'Is the world organised into ϕ 's *by something* prior to our doing so?'. This is closer in spirit to the conceptualist view than to that of the metaphysical realist who claims that there is a fundamental structure to reality that is independent of the actual structuring made by us and other intentional or causal

⁴⁷ Ibid., p. 141.

⁴⁸ I don't mean to suggest that our conceptualisation is not natural: the distinction I am making is best thought of as a distinction between a certain kind of natural principle of unity and all other natural principles of unity.

systems. The second reason for thinking that metaphysical realism is false is that conceptual unities are no less objective than natural unities and further, just as we can talk about prehistoric natural unities, we can talk about conceptual unities in the same way: we can talk about the number of weeks a particular dinosaur lived for, for example. Once we have the concept *week* we can apply it to the past. Doesn't this show that in a sense weeks *did* exist prior to our concept? Also, given any contrived concept we can apply that to preconceptual reality too: we can ask whether brontosaurus had grue skin, for example. Doesn't the distinction between conceptual unities and non-conceptual or natural unities turn out to be unhelpful: shouldn't we reject it and account for the difference in the way we view things like dinosaurs, days, planets, etc., and things like weeks, values, numbers, etc., in terms of the distinction made in section 1.5 between conceptual and non-conceptual objects of *experience*?

2.2 *Carving nature at the joints*

According to a metaphysical realist, there is an intrinsic structure to the world and this structure is independent of our conceptualisation. If our concepts are to describe reality they must match this structure, they must 'carve nature at the joints'. I have argued against the idea of there being only one set of natural joints. But perhaps we can still make sense of the idea of our conceptual scheme being answerable to reality: perhaps we need to place some restrictions on acceptable conceptual schemes and reject those conceptual schemes which don't match up to any of the various sets of natural joints but actually cut across them. Is there any

sense in which our concepts *blue* and *green* are better, more correct, than *grue* and *bleen*, for example?

Ralph Walker thinks so: he describes the difference between squareness and ‘squircularity’ (which holds of a thing at time *t* if it is square and *t* is before noon on 15th September 2003 or if the thing is circular and *t* is at or after noon on that date⁴⁹) as follows: ‘Squareness constitutes a genuine similarity among things; squircularity does not. Certainly the squircular things form a class, and there could be nothing wrong with classifying them together if one felt inclined to do so, but there would be no respect in which all the squircular things would be alike. ... There is no reason why [squircularity] *should* be less of a similarity: it just is.’⁵⁰

Even if some concepts are more natural or objective than others in this sense, is that important? Walker talks about the possibility of finding out that the concepts one uses are the wrong ones. He says that if on the relevant date square things carried on being square, the imagined people using the concept *squircular* would come to realise that we were right and they were wrong; ‘their understanding of the way the world works after the critical time would have received a massive blow’.⁵¹ Presumably if things had turned out the other way round, the square things becoming circular at that time, we would have been the ones to suffer from the realisation that our concepts were flawed. But is this really what would happen in such a situation? Wouldn’t we hold on to our concepts *square* and *circular* and try to find an explanation for the change in shape of the previously square things? Surely we would hold on to our concepts and ask what had caused square things to change shape and become circular? We would have

⁴⁹ See *The Coherence Theory of Truth*, pp. 152ff.

⁵⁰ *Ibid.*, pp. 152-3.

noticed a change that begged for a causal explanation, whereas for the *squircular*-users there would have been no need for explanation at all: things would have been carrying on in the normal way. The concepts we have are the tools we use to talk about the world and if something strange happens in the world we try to find an explanation using the tools we have rather than downing tools and giving up.⁵² The difference between using the concept *square* and using the concept *squircular* is that the question of why all the square things became circular on September 15th 2003, if they did, would be an interesting question to people with the concept *square*, but would not be to people with the concept *squircular*.

The idea that our concepts should pick out objective similarities is misguided - we can say things about the world regardless of what our concepts group together. Anything we can say using the concepts *blue* and *green* can be said using the concepts *grue* and *bleen*. We can translate 'grue' as follows: something is grue if and only if it is green and observed before time *t* or if it is blue and observed at or after time *t*. Where translation is possible we have two ways of expressing the same facts.⁵³ So when the imagined *grue*-user says that the grass is grue and I say that it is green we are both picking out the same fact - that the grass is that particular colour. If two conceptual schemes are translatable, i.e. if they express the same facts, then neither one is better than the other with respect to the

⁵¹ Ibid., p. 155.

⁵² This is the idea behind one answer to the problem of induction: if it were to turn out that our regularities became or revealed themselves to be irregular what we should do is not to give up on induction but to find out why the regularities broke down and work out new ways of inducing.

⁵³ Although it is less clear that we would want to say this if the concepts of another scheme were severely kinked with respect to ours: does 'The cabbiphant is on the mabltain' (where something is a cabbiphant if it is a cat and observed in Europe, or a rabbit and observed in the Americas, or an elephant and observed elsewhere and something is a mabltain if it is a mat and observed in Britain, or a table and observed in the Americas, or a mountain and observed elsewhere) really express the same fact as 'The cat is on the mat'?

structure of reality, if two schemes express different facts, then the question doesn't arise.⁵⁴

What about conceptual schemes which pick out different individuals, though? Imagine people who not only had bent colour and other property concepts but also referred to bent individuals, imagine, for example, that they had the concept *tabair* described by Kripke: 'a 'tabair' is anything that is a table not found at the base of the Eiffel Tower, or a chair found there'.⁵⁵ It is the possibility of schemes with different individuals that leads philosophers to conclude that existence must be relative to scheme. For if it isn't, if all the possible things that schemes could refer to exist independently of those schemes, then we must conclude that the world is populated with an infinity of strange things, tabairs, tabeds, tabanos, tabawers, and so on. If the concept-independent reality of the objects individuated by each actual distinction implies the concept-independent reality of the objects individuated by all of these strange distinctions and more, then we surely have an absurdity.

Perhaps we can privilege the objects of actual distinctions over the objects of merely possible, contrived, unreal distinctions: perhaps we can say that an object is real only if it is picked out by an actual distinction made. This would mean that the difference between blue and not-blue exists because humans have

⁵⁴ Davidson argues that since we can't make sense of translation failure between conceptual schemes talk of different schemes is mistaken. I disagree: the fact that two schemes are translatable is not enough to show that they are the same - there are important differences between schemes that express the same facts, differences in which facts are taken as basic, for example. We take *green* and *blue* as basic and *grue* and *bleen* as parasitic on them whereas the imagined *grue*- and *bleen*-users' scheme puts things the other way round. Claiming that translatability yields identity of schemes is like saying that the metric and imperial systems of measurement are one and the same. Cf. Goodman: 'To project "grue" and "bleen" instead of "blue" and "green" would be to make and live in a different world.' (*Ways of Worldmaking*, p. 101.)

⁵⁵ See *Wittgenstein on Rules and Private Language*, p. 19.

the concept *blue*.⁵⁶ Also this distinction existed prior to our developing the concept *blue* if our non-conceptual behaviour distinguished between blue and grey, which seems likely, in predicting the day's weather on the basis of the colour of the sky, for example. Any difference the laws of nature noticed would be real: the colours that nature differentiated between, the colours of an animal's coat being relevant to its survival, for example, would have an earlier reality than those only distinguished by us. If we don't separate real objects and properties from the contrived ones in this way then how can we possibly do so? Must we do so, though? Hartry Field suggests not.

Hartry Field pinpoints the realisation 'that our conceptual schemes might have been sufficiently different that we didn't think in terms of entities like dinosaurs, but thought instead in terms of entities of quite different sorts (undetached dinosaur parts, for example)⁵⁷ as the motivation for the claim that objects' existence is dependent on conceptual scheme. He then offers an ingenious block to the conclusion that if we had thought in terms of undetached dinosaur parts that would have been for undetached dinosaur parts to exist instead of dinosaurs. He says 'Most of us recognize that more exists than one need assert the existence of: it is rarely to the point to assert the existence of undetached rabbit parts as well as of rabbits; this does not mean that we are denying the existence of

⁵⁶ There is a sceptical line of thought attributed to Wittgenstein by Kripke which questions whether we do actually have the concept *blue*. How do we know, Kripke's sceptic asks, that we don't really have the concept *bleen*? My answer to this form of scepticism about meaning is to acknowledge that our concepts are capable of taking surprising turns but to say that our lack of certainty about what will happen to our concept in the future, at time *t*, for example, is due to their objectivity not due to there being 'no fact of the matter about [what] I mean [by a word]' (*Wittgenstein on Rules and Private Language*, p. 71). If Kripke's sceptic had asked the people in the past who used the word 'crow' to refer solely to birds of the genus *Corvus* how they knew that 'crow' didn't also refer to ugly old women, then they would perhaps have found that possibility as unlikely as we find the possibility that 'blue' does, or will at some point, refer to bleen things. We do of course know a lot about what 'blue' *doesn't* refer to: it doesn't refer to addition, for example.

⁵⁷ 'Realism and Relativity' p. 557.

the undetached rabbit parts. If this obvious fact is borne in mind, there is clearly no way to draw anti-realist consequences merely from the fact that two equally good theories could differ in their existence claims.’⁵⁸ Surely Field is right: the objective existence of undetached rabbit parts is not threatened by our failure to single them out. So we should resist the claim that existence is relative to scheme.

What does seem relative to scheme is basicness of objects: relative to our conceptual scheme the colours blue and green are basic and grue and bleen are defined in terms of them. Relative to the scheme of the imagined *grue*- and *bleen*-users, things are the other way round. In §.19. of the *Investigations* Wittgenstein gets us to imagine a language consisting only of orders; the speakers shout out things like “Slab!” in order to get others to pass them things like slabs of brick:

But what about this: is the call “Slab!” in example (2) a sentence or a word? - If a word, surely it has not the same meaning as the like-sounding word of our ordinary language, for in §.2. it is a call. But if a sentence, it is surely not the elliptical sentence: “Slab!” of our language. - As far as the first question goes you can call “Slab!” a word and also a sentence; perhaps it could appropriately be called a ‘degenerate sentence’ ... in fact it is our ‘elliptical’ sentence. - But that is surely only a shortened form of the sentence “Bring me a slab”, and there is no such sentence in example (2). - But why should I not on the contrary have called the sentence “Bring me a slab” a lengthening of the sentence “Slab!”? - Because if you shout “Slab!” you really mean: “Bring me a slab”. - But how do you do this: how do you mean that while you say “Slab!”? Do you say the unshortened sentence to yourself? And why should I translate the call “Slab!” into a different expression in order to say what someone means by it? And if he means the same thing - why should I not say: “When he says ‘Slab!’ he means ‘Slab!’”? Again, if you can mean “Bring me a slab”, why should you not be able to mean “Slab!”? - But when I call “Slab!” then what I want is, that he should bring me a slab! - Certainly, but does ‘wanting this’ consist in thinking in some form or other a different sentence from the one you utter?⁵⁹

This passage highlights the temptation to treat our own language and its structure as basic. The key to the view Wittgenstein is putting forward is the answer he gives to the suggestion that “Slab!” is really a shortened form of “Bring me a slab”; the answer being that this is only true in so far as the converse is also true -

⁵⁸ Ibid., pp. 555-6.

that “Bring me a slab” is a lengthened form of “Slab!”, neither form is more basic in itself.

2.3 Creating systems and worlds

We are free to use our imagination and create systems within the world. From non-objective starting points, from the stipulation of concepts, axioms and rules, objective entities arise. I will briefly consider two systems - mathematics and chess - and will ask what makes them capable of objectivity. I will then turn to value systems, especially ethical systems. It seems true to say that $2 + 2$ would not have equalled 4 had we not made it so. Or if this seems contentious then consider instead a claim about chess: it surely would not have been the case that bishops could only move diagonally if we had not invented chess, its pieces and their moves. With the invention of chess (and mathematics), we create entities - bishops, knights, kings and queens (numbers, operations, etc.) - and there is no reason to think that the entities we created were the only ones that could have been created - we were not forced by reality to create one game rather than another. If someone were to ask about chess why the king can only move one square whilst the queen can move any number of squares in one direction, we would have no real answer: we would say that that is just the way the game is. There may be a historical explanation for the nature of chess, there may be reasons why chess evolved into the game it did, but there are no justifications, no reasons why chess should have evolved one way rather than another. We could have created a different game in which king and queen had the same moves - a more egalitarian

⁵⁹ *Philosophical Investigations*, pp. 8-9.

game than chess! - but we didn't. Similarly if someone were to ask why the square root of -1 is i then we would not be able to give them any justification.

Some systems are successful and others are not. What makes a system successful? Given that there was nothing about reality prior to the invention of chess that the rules of chess had to answer to, what makes chess a good game? Chess will probably still be played long after real-life kings and queens have become a thing of the past: its relevance or success is independent of the historical conditions that led to its invention. Chess is a successful game because it is incredibly rich, deep and challenging; much more so than one would expect such a simple system to be. When we invented chess we created a world that flourished. Likewise with mathematics. Chess was not 'out there' waiting for our discovery - we invented it from scratch - so the important question, at least concerning systems like chess, is not how true a system is to a reality prior to it, but how real the world of a system can become. Some systems work and others don't, and even if we can't pinpoint the qualities a successful system must have, in practice we can easily differentiate between good and bad systems.⁶⁰

What is the status of the truths of a system? We created chess and stipulated its rules, and so its rules are not independent of us. And yet there are objective truths to be found about the game in general, and about particular games. Once we have created the game, objective truths about it will emerge. If one fails to notice a move one could have made in a game, one fails to notice an objective fact - that a certain move was available. Even though facts about chess couldn't

⁶⁰ There are parallels between the evolution of systems and the evolution of life forms and in both cases the things that survive are the best examples we have of survivability. We know that these things have whatever it takes, then the question 'what is it that enables them to survive?' can be asked.

have arisen before we invented the game; once the game has been invented, facts about chess are just as objective as facts about things that weren't created by us.

I hope to have shown that the relationship between conceptual schemes or systems and reality does not threaten the objectivity of the objects of conceptualisation, then. Even the things which only arise because we develop a particular system, the objects of chess, for example, are to be thought of as objective. I now move to consider value systems.

Would the existence of objective truths that arise once we set up a system be enough concerning value systems? Surely not: the truths we can discover about chess are objective, i.e. they hold true independently of whether or not we take them to be true, but they only apply to the game of chess. Suppose someone decided to use a chess set to play a new game, then suppose that someone else came along and told them that they were not making the right moves. The first person could just reply "I'm not playing chess, I've invented a new game" and then it would be silly if the second person refused to accept that, saying "You can't do that, it's a chess set, you have to play chess with it!". We are not bound to play chess; it's up to us to choose which games we want to play.

But consider an ethical situation analogous to this one: suppose one came upon a person torturing a cat and told them that what they were doing was wrong. The reply "I'm not playing morality; I'm playing a game with different rules" is unacceptable. If ethical systems are like games in that their truths only arise once one has entered into the system, then ethical truths are not as wide-ranging as we would hope; they only apply to people who aim to be ethical.

A truth of chess such as ‘The best move open to John was to take his opponent’s queen’ can be thought of as more or less equivalent to ‘Given the rules of chess, the best move open to John was to take his opponent’s queen’ but an ethical truth such as ‘The right thing for John to do was to own up to his crime’ doesn’t seem to translate into ‘Given our moral rules or conventions, the right thing for John to do was to own up to his crime’. If it is right for John to own up, we want to say, it would still be right for him to do so in a society that failed to value honesty. Similarly, the claim ‘John lived a meaningful life’ doesn’t seem equivalent to ‘Given the norms of his time, John lived a meaningful life’.

There are two ways in which ethics could be different to games like chess: it could be that ethics is answerable to reality in a way that chess is not, or that ethics is no more open to outside evaluation than chess but that ethics is a game we must all play. Are ethical systems and their concepts open to evaluation in ways that other systems are not? Are our value concepts assessable as right or wrong, unlike the rules of chess or colour concepts?

The thought that we could have developed different games than the ones we did doesn’t worry us, the thought that we might not have developed any games at all is not horrifying. But the thought that we could’ve developed different ethical systems than the one(s) we have is worrying, and even worse is the thought that we might not have developed any ethical system(s) at all. Is this thought worrying because we would have been missing features of the world if we had not developed an ethical system or because even though the ethical world emerges as we develop an ethical system it is a world of fundamental value?

Ethics, values, meaningfulness, etc., do not seem to be parts of the world in the way even things like colours and weeks are: it doesn't seem true that the ethical world was there all along any more than it seems that the world of chess was there all along. Creatures capable of valuing evolved and so the world of value emerged. Such creatures might never have come about. This fact might seem to show that value systems aren't answerable to an earlier reality but it doesn't: all it suggests is that value systems aren't answerable to prehuman reality. Given that human beings did evolve and were capable of suffering, of causing or preventing suffering in others, of deliberation, etc., it seems right to say that any way of life which developed to meet our needs would have to have had an ethical element.

We learn the concepts of our culture, we learn what counts as good, just, meaningful, etc., but are these things *really* good, just, meaningful, etc.? Couldn't our culture just be mistaken about things; couldn't one strive to live a meaningful life on the basis of one's culture's norms and all the time be striving in the wrong direction entirely? This sceptical worry seems to have particular force regarding values: why is this? Is it because there is evidence of real disagreement among people about values, because we can really consider the option that our values are wrong and those of another culture right? In part III I will argue that evaluations are underpinned by emotion and that this is why scepticism about values is a real threat.

It is hard to accept that there is no ultimate viewpoint from which to decide matters of meaningfulness: the questions 'What makes my life meaningful?', 'How should I live?' can seem unanswerable. There are so many conflicting

answers: the good life is the contemplative life, or the Christian life, or the life of happiness, or the life of luxury, etc. There is no such ultimate viewpoint: to be meaningful is to be meaningful *to* someone. This idea will be looked at in part IV and I will address the problem of whether meaningfulness can be relative in this way and yet be objective.

III

INVOLVEMENT, OBJECTIVITY AND EMOTION

‘When people are feeling friendly and placable, they think one sort of thing; when they are feeling angry or hostile, they think either something totally different or the same thing with a different intensity’

-ARISTOTLE, *RHETORIC*

I have looked at two main areas of scepticism so far: in part I the possibility of our experience of the external world being completely erroneous was rejected, and in part II the possibility of the real world being out of reach because of conceptualisation was broken down. These forms of scepticism might be thought of as purely philosophical or academic, not real threats to our knowledge of the world. I now turn to more empirical, more disturbing, sceptical questions; questions which arise as a result of the role played by feeling and the subconscious in our involvement with the world.

I will be looking at the role of feeling in matters of meaningfulness, arguing that meaningfulness arises out of feelings of meaningfulness. But one's feelings aren't always trustworthy and unless there is a way of separating reliable feelings from non-reliable feelings, then feelings of meaningfulness will never be sufficient for meaningfulness and scepticism about meaningfulness will be a real threat.

I will develop a view of feeling as the cause of belief in certain situations - the feeling of meaningfulness being the cause and the source of one's belief that something is meaningful, for example. This will undermine the rationalist separation between rational thought and irrational feeling/desire. I will not draw a firm distinction between feeling and emotional experience, although the idea of an

unconscious feeling will be important. Feelings are known as experienced and some feelings and experiences are inextricable. Meaningfulness, love, and other things we cherish are always experienced emotionally (i.e. felt). This view of the rational as based on the so-called irrational will lead to a sceptical problem about justification: if (some of) our beliefs are caused by feelings and if these beliefs seem true to us because they match the way we feel about things, then unless our feelings are rationally constrained, then there is no rationality. So we must ask the question: how are feelings caused and can we separate valid from invalid feelings?

1. The importance of emotional involvement

In part I I rejected the equivalence made by Nagel between detachment and objectivity, arguing that not all detachment is good, i.e. conducive to finding the truth, and that what makes one viewpoint more objective than another is that it is more inclusive, that it contains the second viewpoint, not that it is more detached. In the following I will be arguing for the importance of involvement as taken in contrast to detachment: in some cases, emotional involvement will be seen to be necessary for knowledge.

In human beings there is an important link between knowledge and feeling: involvement is important in the acquisition of knowledge, even in the acquisition of the most abstract knowledge, in solving mathematical problems, for example. There has to be something about the solution to a problem or the answer to a question that makes it *better* than non-solutions or wrong answers. If one got the same satisfaction out of a non-solution as one got out of a solution, then one wouldn't aim to find a solution. Are solutions and truths satisfying to find because

they are good in themselves or are they good because they make us feel good, give us a satisfying feeling? The phrase 'resolving a tension' is delicately balanced between two things, between the emotional and the objective. We resolve tensions in philosophy, wrestling with contradictions, hoping to find a solution. Resolving a tension also describes the move from an unsatisfactory state of mind or being to a more comfortable one. However we answer the question, we will not be able to understand the quest for truth, enquiry, without understanding the emotional side of these things.

Imagine a God who viewed the world from a detached point of view. I have already argued that even though such a God might be able to know the truth values of all statements, he would not be able to *understand* those truths only ascertainable subjectively. I now wish to suggest that such a God would not understand *anything*; that understanding requires involvement. The detached view of the world according to which only physical objects with primary qualities exist is not the view of the world given by complete detachment but just the view of the world given by purely physical involvement: the reason this view is thought to be the detached view is that we take it that everything in the world is physical, so physical involvement is universal. An immaterial God would lack even this form of involvement and a God without emotions would be unswayed by any of the things that seem important to us, unswayed in favour even of one set of axioms over another and so would be unable to accept a system at all. Like Buridan's ass, such a being would wait forever undecided.

1.1 Feelings and objectivity

Given that some involvement is necessary to understand most (if not all) things, we can ask how much involvement, how emotional an involvement, is good. If emotional involvement and the emotional experience it yields are necessary for meaningfulness, as I will argue, then there is a problem if emotional experience is biased and incompatible with objectivity. It would seem that we would be forced to choose between living a life that seemed meaningful and living with an objective view of life.

An example that illustrates the problem is the following: when one falls in love one gains something good, one's life is enhanced, but often what also happens is that one loses one's objectivity, seeing the object of one's emotions as the most perfect being on earth, seeing even their most bland utterances as profound and their most annoying habits as endearing. It is quite clear that emotional involvement *can* lead to loss of objectivity, and, as examples such as this remind us, loss of objectivity is a common feature of emotional involvement. But the question to be considered is: what is the tension between emotional involvement and objectivity; must emotional involvement lead to loss of objectivity? Must emotional involvement lead one away from truth?

The idea that our emotional attachment to the world clouds our vision needs questioning. The first question is: what does attachment cloud or block our vision of? When we say love is blind we don't mean that a person in love will fail to see things going on around them, we just mean that they will fail to notice faults in the person they love, faults that seem all too obvious to the rest of us. They lose

their detachment, their objectivity about the object of their attachment, but may still be as good a scientist or politician as they were previously.

Someone may fall in love and thereby lose interest in doing the things that were previously important to them. We can imagine a brilliant scientist falling in love with someone and neglecting her work as a result, perhaps deciding that she would rather just get married and have children. Should we think of this kind of phenomenon as loss of detached objectivity, though? If we do then we seem to mischaracterize the scientist's earlier state of mind: it is not that she is neutrally going about her work and until she is swept off her feet in a blaze of passion, or at least if this were the best way of describing her work, if she were doing it unenthusiastically, as nothing more than a way to pay the bills, then we wouldn't think of her giving it up as a bad thing. We must assume that she is doing research for its own sake if we are to see her giving up work as a bad thing, and this means that she cannot be seen as detached: someone who is involved in her work is just as attached or partial as someone who is involved in a relationship.¹

What is the difference between someone who is blind to their loved one's faults and someone who is able to be objective about their loved one, seeing some bad points? If the relevant difference is that the first person is more deeply involved than the second then this gives us a reason for thinking that involvement ought to be limited in order for objectivity to be preserved. If deep involvement must be in opposition to objectivity, then there is a fundamental problem: we want deep involvement - this is what gives our lives meaning, I will argue - but we also want objectivity - we must know that our involvement is real.

¹ There is a question of whether any involvement is as good as any other and this has not been addressed.

Suppose the person one falls in love with really is wonderful and so there is no discrepancy between one's own view of them and the views of others: even people who hardly know this person find them loveable and one's friends approve of one's involvement. In this case, one's emotional involvement doesn't seem to threaten one's objectivity: one has fallen in love with someone and as a result come to see them as being wonderful, and they truly are wonderful.

But did one's emotional involvement lead one to come to see them as wonderful or were things the other way round, i.e. was it their being wonderful that led to one's emotional involvement? If the former, then, someone might say, one's being right about them is sheer coincidence. Perhaps one fell in love with a person because of something else entirely, perhaps one fell in love with the first person one met who shared one's interest in train-spotting, for example. If so then one's emotional involvement is not truth-conducive, one is just lucky. If the latter, then involvement is the rational thing given what is the case, one's emotions are governed by reason rather than themselves leading the way.

The possibility of falling in love with people who are not wonderful, of being misled by one's emotions, makes it impossible to take all emotional involvement as valid.² We need to be able to separate truth-conducive or true involvement, from deceptive involvement. We want to know which forms of

² There are two questions that could be asked about such involvement, one pragmatic and one philosophical. The pragmatic question is the question of whether one benefits from the involvement, whether it's a healthy involvement. This leaves to one side the philosophical question of whether one's involvement gives one access to truths. Is falling in love with this person good for one? Is falling in love with this person leading one to live a true life or to live a deluded life? There are ways of relating the two questions and giving one explanatory priority over the other: one could hold the view that truth is in essence good/healthy, that if one is living a deluded life then one is by definition emotionally unhealthy. Or one could hold a more modest view which claimed that there is more likelihood of disappointment and hurt if one is deluded, but that it is perfectly possible for someone to thrive even though deluded. Starting from the pragmatic side of things, one could take truth to be what works and say that if someone lives a healthy life on

involvement lead one to (subjectively ascertainable) objective truths and which forms lead one to falsity? This might be thought of as a purely epistemic question, and it might be thought that whatever difficulties there are in distinguishing between true and false involvement, there is no metaphysical problem, but that would be wrong: unless we have some criterion for distinguishing between the two then it won't make sense to assume that there is any underlying difference.³

One way of distinguishing might be the following: the things a form of involvement leads to are objective/true if and only if they can be confirmed by others who aren't involved. If this is right then in the case of falling in love, if one's friends are completely at a loss to see what is good about the person one has fallen in love with, if even though they try to like the person, they just can't, if they end up thinking 'I just don't know what she sees in him', etc., then we should say that one's involvement has led to a loss of one's objectivity. Often this is what we say in this kind of situation, but are we right to do so? Sometimes when one falls in love with a person one comes to see a side of them which remains hidden to others.⁴ Involvement might be necessary in order to discover or bring out the best in someone. If one finds oneself in a situation in which the person one loves is rejected by one's friends, should one take their view of things as having any bearing on one's feelings and choices? Or should one, on the contrary, take a critical look at one's friends instead? There doesn't seem to be a clear-cut answer

the basis of certain emotions then those emotions are justified. I am here only concerned with the philosophical question.

³ The problem can be applied to a practice as well as to an individual: we can ask of a shared practice whether it is the result of true involvement or not. (Of course there are some who think we can't make sense of the question concerning the practice.)

⁴ An extreme illustration of this possibility is found in the story of the marriage of Gawain - Gawain is bound to marry an ugly old hag who, he finds out, is really a beautiful maiden under an enchantment. She asks him to choose whether she is to take her real form during the day, when everyone will see her, and the form of a hag during the night, when only he will see her, or the

to this - it depends on circumstances. So it seems that we cannot accept a general rule to the effect that if the view one's involvement yields doesn't tally with the views of others then one's objectivity has been lost.

Maybe the important thing is that one still sees as valid the observations of one's involvement when one ceases to be involved. Maybe we should say that if when one's love wanes one looks back and still sees the person one loved as wonderful, then one's involvement was valid and didn't threaten one's objectivity. If one looks back and finds it strange to think that one had been so besotted by that person, then one's involvement had led to a loss of one's objectivity.

But if appeal to the views of others doesn't work, then neither will appeal to the views of one's later self, for one's later self may be at fault: one's memory might be lacking in objectivity, especially if one's involvement ended traumatically. One might have forced oneself to think badly of the person one loved in order to put them out of one's mind and get on with life. If one is able to convince oneself that the relationship wouldn't have worked out, that they were not the right person for one, maybe even that one's involvement was a mistake, then one is less likely to yearn for them and this might be the only way one can get on with things and ever be able to come to love anyone else. (And what if one's love never wanes? Should this be taken as proof of the validity of one's involvement or as proof of the opposite?)

Both of these suggestions are made on the assumption that the fruits of involvement are answerable to other views, either to a detached view, or to all views. But we have denied the link between objectivity and detachment and

other way round. If he chooses the first option his friends will feel only pity for him and will be unaware of the beautiful side of his apparently ugly wife.

between objectivity and availability to all viewpoints and argued for the possibility of things which are objective but only available from quite narrow viewpoints. So if the beliefs one acquires through involvement are not answerable to more detached views, what are they answerable to; how are we to determine the objectivity of the fruits of involvement? The test for objectivity that the view of objectivity as detachment yields - in order to find out whether x is ϕ , see if every viewer sees x as being ϕ - has been rejected, but now we need an alternative test.

Someone who was hostile to the possibility of truth-conducive emotional involvement would offer the following dilemma: either emotional involvement leads one to believe things others don't believe or it doesn't. If it does then it is deceptive and leads one into error. If it doesn't, if it only presents things the way they are anyway, then there is nothing to be gained from it. Either emotional involvement comes into conflict with dry detachment, in which case we should reject it, or it doesn't, in which case we should forget about it. Such a way of putting things is wrong, though: it leaves out the most important possibility - that what is found out through emotional involvement doesn't conflict with or merely duplicate what one already knew from a more detached point of view but adds to it.

The acceptance of the idea of objective truths which are only ascertainable subjectively leads to a problem: if there are some things which are true but only available to those who are involved in the right way and other things which seem true to the people involved with them but which are actually false, then how could we possibly make sense of the difference between these two sorts of cases? We cannot appeal to an external perspective, since both the truths and the falsehoods

can only be fully appreciated from the inside. And likewise we cannot appeal to the internal perspective, since in both cases it seems to the involved parties that they are presented with truths. There are two areas in which the possibility of a solution may be found: first of all, it is not the case that all of the truths internal to a perspective or practice are known to those on the inside - there is room for error within a perspective or practice. Secondly, there are often people involved in a practice who also hold an external view.

I suggest that an alternative test can be formulated, a test which takes on board what is important about the availability to all viewpoints test: something is objective if it plays the role of an objective item in the best explanation of how it seems to those involved with it and to those who aren't involved; something is objective if it appears in the most objective (i.e. containing) view. A view of objectivity as detachment makes something objective only if it is available from a detached point of view, whereas a view of objectivity as containment makes something objective only if it is available from a reconciling, containing point of view.⁵

On the containment model of objectivity, when there is conflict between views, the objective view must be *constructed*. If we have a conflict between things seeming ϕ and things seeming not- ϕ , then there is a problem to be solved: we don't have an automatic procedure for dealing with such conflicts but must exercise skill and imagination in finding a way to reconcile the two. Working on the assumption that objective reality is non-contradictory (which I take as read,

⁵ There will often be a problem deciding what is the most objective view - one person's view of another's religious claims could claim objectivity and explain religious claims as metaphors rather than factual claims about God, whilst the religious person could claim objectivity and explain the other's hostility to religion as scientism, for example. The problem is deciding which containing view is right - we then need a view which contains both of them!

since it is something we are unable to reject and is also necessary for the idea of conflict between views) we must find the hidden nature of things that yields the apparent contradiction.⁶ This form of invention or construction of viewpoint should not be thought of as constructivist in the sense of it being the case that we construct objectivity. We construct viewpoints so as to get at, find the truth about, objective reality.

Involvement itself doesn't lead to a loss of objectivity but deep involvement in one thing at the expense of everything else does. Often when one becomes involved in something new one becomes blind to everything else. It is possible to remain objective whilst involved in something - one can sometimes have a perfectly accurate idea of just how one's involvement appears to others. One's involvement leads to a loss of objectivity only if it leads to a loss of the ability to understand how things are to those who remain uninvolved. If one is still able to take an external view of oneself and one's involvement then one's view is more complete than either the views of the 'outsiders' or the views of the 'pure insiders'.

One can sometimes see quite accurately what is happening to one, one might be able to see the wider picture as well as feeling strongly about the person one is emotionally involved with, for example. One might say things like 'for some reason I have fallen head over heels in love with this person, they have come to be the most important thing to me', acknowledging the possibility that the strong feelings one has might not be the appropriate way of responding to that person. Emotional involvement doesn't necessarily preclude objectivity as I

⁶ This idea is to be found in the writings of F. H. Bradley: Bradley argues that the principle of non-contradiction must be taken as an 'absolute criterion' and builds his view of absolute reality as a

characterise it on the containment model even if it does preclude detachment. But there is a tension: to see things in this way is to doubt the cognitive significance of one's emotions. If one trusts one's emotions then one will say that it is the person's goodness and beauty that caused one to fall in love with them. The more objective approach seems to bring with it the idea that emotions are not necessarily trustworthy, that even one's current emotions might be misleading.

There might seem to be an irresolvable tension between complete involvement and objectivity. Someone might say that real involvement must be complete, that being open to the possibility that one's love is unfounded is not really being in love, for example. The idea of religious faith is the idea of unquestioning dedication to one way of life at the expense of other ways. Is it inevitable that full involvement rules out the possibility of questioning and objectivity? I think not: the idea that having a real commitment to one way of seeing things or one way of life rules out the possibility of appreciating other ways of seeing or living arises as a result of the view that in order for something to be real or valuable it must be universal. Further, the idea that deep involvement precludes detached criticism must be wrong: surely the things we involve ourselves in the most deeply ought to be strong enough to withstand criticism, not be in need of protection from inquiry. The things we care about most and are most captivated by ought to be things we want to scrutinise from all angles not things we want to put away carefully in glass cases.

unified whole consisting of experience upon this modest foundation (see *Appearance and Reality*, Book II).

1.2 Bare goods and emotion

If life is to be meaningful at all then it must contain some non-instrumental goods (or at the very least it must seem to) - for if it didn't, if all goods were instrumental, then there'd be no way of getting them off the ground. As Aristotle puts it, we cannot 'choose everything for the sake of something else for at that rate the process would go on to infinity, so that our desires would be empty and vain'.⁷ Metal is good because it can be used to make arrowheads, arrowheads are good because they contribute to the making of arrows, arrows are good because they can be used to kill animals and killing animals is good because it provides food, food is good because it enables survival. But survival, staying alive, is just a good thing, the end all of these instrumental goods get their value from.

An instrumental good can be sought without one's caring one way or another about it itself because of what one does care about. One only seeks to bring about instrumental goods because they lead to non-instrumental goods. The non-instrumental goods are what give an activity a point. With instrumental goods we can offer reasons for their being desirable; we answer the question 'why do you bother making arrows?' by pointing to the end result of the process, for example. The end justifies the means (when it does!). We refer to non-instrumental goods in order to justify instrumental goods. What then can we refer to in order to justify these non-instrumental goods? If someone asks of the arrowhead-maker 'But what is so good about staying alive?' they will draw a blank. Being alive is good in itself - it is a bare good, its goodness is just given: being alive is better than being

⁷ *Nicomachean Ethics*, 1094^a19-21. Aristotle sees this as just obvious, needing no argument. Although I don't see it as being completely obvious - it's possible that a meaningful way of life could emerge without there being any ultimate goods, that a system of goods could support each other holistically, for example - I see ultimate goods as the most natural way of supporting meaningfulness.

dead, pleasure is better than pain. I will argue that the possibility of bare goods is tied up with the possibility of emotional experience, that only because some things just are satisfying to us or loved by us when we encounter them can there be bare goods. (This is not to say that all the things we find satisfying are bare goods, but the converse; that bare goods are things which have the capacity to satisfy us.)

Bare goods are either to be thought of as goods independent of us, part of ultimate reality or as goods which are unequivocally good to us, part of the world as it appears to us. Either bare goods are given by detachment or by involvement then. I have argued against the idea of ultimate reality but even if those arguments fail, the idea of appealing to ultimate reality as given by the most detached view for goods of any sort is hopeless: from a detached point of view nothing matters, nothing is good or bad. So bare goods must be arrived at via involvement; they are the things that present themselves as being good to us.

The existence of things good in themselves is necessary for meaningfulness of life: someone who has lost their interest in life is someone who has lost such goods and so only has instrumental desires: in order to survive I should eat; in order to keep my relationship together I should show an interest in my partner; in order to do well in my exams I should study, and there's a big 'if' hanging over everything - if I thought there was anything good about life, my relationship, my success, etc., I'd have real reasons to do these things, but I just don't. Such a person lacking real reasons for doing anything doesn't really lack *reasons*. What they lack is beyond the reach of reason, they lack the foundation or source of reason. If someone doesn't feel that life is worth living then what they lack is a feeling. We cannot offer them reasons - we cannot say 'life is worth living for the

sake of ... or in order to ...' All we can offer them is the possibility that the way they feel might change, that they might come to see life as worth living.

There are many feelings gathered under what I will call 'the feeling of meaningfulness': the euphoria of love, the feeling of having hit upon the right answer - 'eureka!', the feeling of awe at the world or God, the feeling of being totally alive at the end of a strenuous climb, etc. These feelings all seem to have the power to give life meaning in spite of everything else. One can be broken, in pain and despair, with nothing, and on hearing a piece of music be so uplifted by it as to feel that although life is full of misery, whilst there is music like this then there is something to live for. These feelings are necessary for the meaningfulness of life and further they sometimes seem to guarantee it. If they do so, then is that because they are reliable indicators of meaningfulness or because they constitute meaningfulness? What is the status of feelings of meaningfulness as the bare goods required to live a good life?

Some things make us feel good, others make us feel bad, what is the relationship between the way things make us feel and the way things are? Are these feelings indicators of meaningfulness or are they constitutive of it? Is the satisfaction, the good feeling, certain things give us a sign that those things are good or is it because they make us feel good that these things count as good? We take it that the former is true, that the things we feel are good, meaningful, right, true, etc., are objectively so, that they would be so even if we failed to feel the way we do, since our feelings can be and often are mistaken. But it might well be that the latter is in fact true. We seem to take some feelings as indicators of goods but as infallible indicators - we seem to want feelings to be responses to objective

reality but without the possibility that they could be wrong. We take the love we feel for a person to be an indication of a good relationship not to be constitutive of it - the meaningfulness of a relationship is not guaranteed by the quality of one's feelings, since one's feelings might not be indicative of the way things really are, the person one loves might be deceiving one, for example. But we still seem to want to say that the feeling itself is a guarantee of a bare good. The possibility of a solid basis - of knowledge about the meaningfulness of life - seems to require our feelings to be infallible, whereas the possibility of objectivity seems to require our feelings to be guides to an independent reality.⁸

1.3 The possibility of error about the content of one's emotions

Feelings of meaningfulness can result from non-meaningful causes. Suppose one had a mystical experience during which one felt oneself to be communing with the divine. Suppose that one then developed a whole religious system on the basis of that experience. Then suppose that one found out that the experience had been engineered - perhaps one had been the subject of a psychology experiment without knowing it and one's consciousness and emotions had been manipulated using hallucinogens and subliminal cues. When one finds out that the causes of one's 'profound' experience are not divine at all, the meaningfulness one's life had gained from thinking they were will likely be lost again. Or suppose that one is happy and sees life as meaningful. and that the main cause of one's happiness is the belief that the person one loves, who matters more than anything else in the world, loves one too. Then suppose one found out that it had all been a lie, that

⁸ There seems to be a parallel between certain emotions and the axioms of a system - the axioms of mathematics, for example: we build our lives on the basis of certain emotions we can't doubt.

one's dear one had just been playing a game, deceiving one into thinking one's love was reciprocated. Suppose the whole thing had been plotted as an elaborate deception. On finding out that the fundamental basis of one's life is a deception, that what one experienced as meaningful is actually devoid of meaning, one is thrown into meaninglessness. If one is an optimist and has other resources then one will eventually recover and one's life will become meaningful again. If not, if the love of one's life had been everything to one, then the deception and the fact that what seemed more certain than anything else was a sham will lead to hopelessness. If an experience that seemed so true was false, then one comes to feel that nothing can be true.

These imagined situations are only more extreme versions of situations most of us have found ourselves in at some point - the disappointment of betrayal or unrequited love, the disappointment of being let down by things we felt sure about. If love counts for anything, then when one falls in love with a person, then surely that person ought to feel the same way, otherwise one's love is pointless, misguided and the whole idea that feelings of love are somehow significant, that they are anything real, is wrong. At least one's feelings ought to be a response to a person's goodness, not just a chemical rush caused by something arbitrary like the way they smell, their pheromones, etc.! If a feeling is completely overwhelming, the strongest feeling one has ever had, then its turning out to be unfounded suggests that feelings in general are not to be trusted. How are we to trust our feelings if the best of them can deceive us? In order for life to be meaningful we must trust (some of) our feelings, but experience seems to show us that we cannot do so.

Experiences in general cannot carry guarantees about their causes. Whatever we experience could be other than the way it seems. If feeling is to be sufficient for meaningfulness then either feeling must be transparent, it must present its causes truly, or feelings of meaningfulness are such that they make life meaningful regardless of their causes. The second alternative is a possibility in some cases - if a piece of music affects one strongly and one feels filled with joy as a result, then the way one feels, that life is worthwhile after all, seems safe. It is the music as one hears it that matters, not what caused the music. And if one were to find out that there was no 'real' music - that one was the only person to hear anything, or if one was to find out that the cause was completely random - that it wasn't composed or played by a sentient being but was the noise of the wind and rain on certain surfaces, its existence would be just as wonderful. When this kind of thing happens it seems that meaningfulness emerges out of randomness and so is independent of the wider context. Is this because music, unlike love, doesn't bring beliefs with it? Meaningfulness is in these cases independent of whatever causal relations underpin things; it is not threatened by the fact that we can give an explanation of what brought about the meaningful experience.⁹

But even if this is true of some feelings of meaningfulness it is clearly not true of most of them. Is it essential that one have at least some completely secure feelings of meaningfulness? Or is it only important that one has reason to believe that such feelings are available to one even if one is yet to experience them, like someone believing that the possibility of falling in love is there and that should the

⁹ Although even with music there are *some* things which can threaten the feelings of meaningfulness we experience: finding out that a piece of music was originally composed as a Fascist anthem, for example, could call into question the profundity of one's experience of the music.

right person come along it will happen, or the religious believer who has yet to experience the Divine for herself?

There are two aspects of emotion that together might be able to dissolve this worry: their internal properties and their relation to the world, to their objects/causes. Emotions are themselves things with certain features; strange things, admittedly, but no stranger than ideas, tunes, lies, etc. It has been one of my themes that things such as these, which seem strange from a physical point of view, are nonetheless objective. The internal or intrinsic properties of an emotion are thus not to be thought of as 'merely subjective' and dismissed as unreal - if an emotion is overwhelming then there is something in the world that is overwhelming (i.e. that emotion). If an emotion is joyous, then there is (real) joy in the world (even if there's nothing to be joyous about). The questions that then arise about how to understand the joy one feels, how to see the rest of the world in the light of one's feelings, await answers, but not answers which deny the properties the emotion itself possesses, which deny the existence of joyousness, for example. If someone is overwhelmed by joy and wonder at the world then the internal properties of the feeling they experience are enough to give them the belief that there is *something* good. If they then interpret their feeling as of the natural world, or of the work of God, or of their own spiritual nature, or whatever, they might make a mistake; they might misinterpret things. But whatever the best interpretation, they will not be wrong about there being some good that the feeling gives them: some feelings just *are* good.

The internal properties of an emotional experience, its feel, intensity, etc., are as important as its relational properties, i.e. what the experience tells the

subject about the rest of the world. (The emphasis being on ‘the rest of’ - the internal properties are not devoid of content albeit non-conceptual.) Can meaningfulness be supplied by the internal properties of an emotional experience or is the possibility of feelings of meaningfulness about something which is objectively good that is the possibility of meaningfulness? In order to answer these and other questions it is necessary to look at the nature of emotion.

2. Emotion and its relation to evaluative belief

In this section I will develop the core of a theory of emotion as the basis of the meaningfulness of life. I will then go on to examine further sceptical worries about emotion due to the fundamental role it plays in the formation of some beliefs. The idea that emotion is irrational or arational is problematic: if, as I argue, emotion infects evaluative belief, then the irrationality of emotion, if emotion is indeed irrational, threatens rationality itself. I will start by looking at the causal-evaluative theory of emotion given by William Lyons in *Emotion* and will develop an alternative though related view. The view I outline will take emotion as causally prior to (evaluative) belief whilst remaining a cognitive theory of emotion, where cognitive theories are thought of as those which make ‘some aspect of thought, usually a belief, central to the concept of emotion’.¹⁰ Cognitive views can be contrasted with phenomenological, physiological or behaviourist views which don’t take as central to emotion its links with rationality. The key difference between the view I will develop and the views of cognitivists such as Lyons is that I will present the cognitive and evaluative aspects as arising as a

result of feeling, not vice versa. It is not, as Lyons suggests, that we come to believe that things are thus and so and as a result come to feel a certain way, e.g. ‘a bird lover out birdwatching sees a golden eagle in full flight. She thinks it marvellous ... and this strongly affects her physiologically.’¹¹ This puts the cart before the horse - it is because the birdwatcher is stirred emotionally at the sight of a golden eagle that she acquires the evaluative beliefs she does (and that is why she became a birdwatcher in the first place).

2.1 Are emotions rational?

Beliefs are judged as true or false, not as rational or irrational. Beliefs are firmly placed within the realm of reason and so the idea of a belief itself being anything other than rational is ruled out. When we talk about irrational beliefs what we are referring to are beliefs it is irrational to hold. What makes it irrational to hold a certain belief *B* is nothing more than that given the evidence, it is extremely unlikely that (the content of) *B* is true. So the only room for irrationality is in the believer, not in the belief itself. With emotions (and desires) it seems that there is room for a different kind of irrationality - we sometimes call emotions irrational and mean more than just that it is irrational for someone to feel a particular emotion. What is the relevant difference between beliefs and emotions?¹²

¹⁰ *Emotion*, p. 33. An alternative characterisation is given by Lyons on the following page, according to which my view would not count as cognitive: he says ‘For Aristotle, and cognitive theories in general, the primary cause of the feeling and physiological aspects of emotion is belief.’

¹¹ *Ibid.*, p. 58.

¹² I view desires as feelings: desiring *x* is to be thought of as feeling *x* to be good. This is significantly different to equating desire with belief, i.e. to saying that desiring *x* is *believing* *x* to be good. The difference between feeling something to be good and believing it to be is important: the problem of weakness of will is best thought of as a tension between feeling something to be good whilst believing it not to be. Arguing for this view of desires is beyond the scope of this thesis but I take it to have intuitive appeal: Hume for one groups desires and emotions together under the heading of ‘the passions’.

There are two ways of denying emotions rationality: firstly, we could deny that they have cognitive content, which would make them arational (although they could still be thought of as reasonable, i.e. appropriate, urges). Secondly, within cognitivism about emotions, we could allow that emotions represent the world as beliefs do, but claim that they do so badly, distorting the way things are or projecting onto the world, or that they are not constrained by consistency and the other standards of rationality that beliefs are constrained by, and so are likely to present a false view of the world. The first option will not be explored: I will assume a broadly cognitivist framework.

The fact that emotions can be inappropriate leads some people to view the emotions as irrational. The fact that we can feel frightened by things which aren't (objectively) frightening, i.e. dangerous, is taken to imply that emotions such as fear are subjective and so not open to assessment; there are thought to be no rational standards to which emotions must answer. A particular emotion (or desire) can be shown to be irrational and yet it remains, unlike a mistaken belief. One may remain scared of spiders even though one realises that one's fear is irrational, one may carry on loving someone even though one has come to believe them to be a despicable person, one may find oneself unable to avoid feeling angry at a certain person even though one believes that the person means no harm and doesn't deserve such a harsh response, etc. If one sees these cases as typical of emotion then one will view emotion as irrational. Hume's view of the passions is of this kind: 'Tis not contrary to reason to prefer the destruction of the whole world to the scratching of my finger. 'Tis not contrary to reason for me to chuse

my total ruin, to prevent the least uneasiness of an Indian or person wholly unknown to me.’¹³

But a view of emotion as irrational in this way cannot account for the fact that emotion can be constrained rationally. The cases in which our emotions and our beliefs come apart account for only a part of our emotional life. In many cases, our emotions seem rational: our emotions often do evaporate when we come to see them as inappropriate. If one feels angry at a friend for letting one down and then finds out that the friend has been involved in a car accident, then one’s anger is likely to disappear instantly and be replaced with concern. Links like this between emotion and belief suggest that we must take a cognitivist approach to emotion.¹⁴

This leads to a twofold rationality constraint on a theory of emotion: a theory of emotion must allow for the possibility of emotion being rational, being answerable to norms of rationality, but not insist that all emotion is rational. I move now to consider Lyons’ cognitivist account of emotion. Lyons starts with the following definition:

X is to be deemed an emotional state if and only if it is a physiologically abnormal state caused by the subject of that state’s evaluation of his or her situation.¹⁵

One worry about this definition is whether the idea of a physiologically abnormal state can be understood in any way other than as equivalent to an emotional state, for if not, the definition is vacuous. Although we may be able to individuate

¹³ *A Treatise of Human Nature*, p. 416.

¹⁴ This dissipating of emotions only seems to work with factual discoveries, though, when one discovers that the facts one’s emotions are a response to are other than one thought. When one’s emotions come to seem irrational not in the light of further evidence but in the light of other feelings or evaluative beliefs, then which feeling prevails seems less to do with which feeling (or evaluation) is right or rational but more to do with which is the stronger in intensity.

¹⁵ *Ibid.*, pp. 57-8.

physiological states, it is certainly not clear that we can talk about them as being normal or abnormal without reference to a prior distinction between emotional and non-emotional states of mind. I will leave this worry, though, and concentrate on a much more important issue: the question of whether it can be true that emotions are generally caused by evaluations. Phenomenologically this seems wrong in two ways: firstly, emotion often seems more immediate than or prior to evaluation or belief acquisition, and, secondly, there is often a conflict between a person's evaluation of a situation and the way they feel about it.

We can wake up in the morning feeling full of joy, wonder why we feel that way and then notice that the sun is shining for the first time in a while and it's a beautiful day; we can meet someone and immediately feel drawn to them and later find out, after getting to know them, that there are good reasons for feeling like that; we can feel fear in a situation without knowing what we are scared of, and then turn around and realise that there's a menacing figure following, and so on. In all of these situations feeling precedes evaluation, feeling precedes conscious awareness altogether. We don't experience any evaluation prior to the feeling, so if there is any evaluative element preceding emotion in these cases, then it must be subconscious. This may turn out to be the best way of thinking of emotion, but if so then Lyons' kind of cognitivism will not turn out to be correct: subconscious evaluations cannot be thought of as evaluations made by the subject of the emotion but as subpersonal evaluations, evaluations made by the body, not the (rational) mind. Then emotion is caused by evaluations the *body* makes. I will suggest that we do think of emotion as caused by subconscious evaluations; as

forming a bridge between the subconscious and the conscious. But this is to diverge from the kind of cognitivism Lyons endorses.

Babies respond emotionally from birth, babies cry, frown, smile, shows fear, etc., but surely they do not evaluate the world prior to doing so. How is Lyons to account for this? He might say that babies do not have emotions, that it would be to make the mistake of behaviourism to assume that we can ascribe emotions to them purely on the basis of their behaviour.¹⁶ Lyons distinguishes between emotions and feelings, reserving the term 'emotion' for feelings that fit his definition. He might want to say that babies have feelings but not emotions, but this is not plausible and just serves to show that his definition is not quite right: babies and infants feel frightened, happy, angry, and these feelings are surely not significantly different from the adult versions; not so different as to be thought of as more similar to purely physiological feelings like twinges, pangs, and so on. Alternatively, Lyons might respond that babies do evaluate the world, but this is not plausible.

The second difficulty is that sometimes we find that our feelings and evaluations don't match. Sometimes we find ourselves responding emotionally, getting angry with someone, for example, whilst believing that the response is wrong, that the person doesn't deserve our anger, for example. Or we find ourselves feeling tempted by things we know full well are bad for us, cigarettes, alcohol, fatty food, an affair, etc. But feeling cannot be caused by evaluative belief in these cases, since then there would be no room for such a conflict.¹⁷ If one

¹⁶ His brief discussion of behaviourist theories of emotion suggests that this would be his response. See *ibid.*, pp. 17-25.

¹⁷ Unless it is possible for one to evaluate something in two conflicting ways, to evaluate something as good and as not good, for example. And then there would be a problem of the rationality of belief.

evaluates something as bad and yet feels irresistibly drawn towards it, then one's feeling cannot be thought of as determined by one's evaluation. Weakness of will is one outcome of such a conflict between a feeling and an evaluative belief - weakness of will is what happens when the feeling wins.¹⁸ Sometimes what happens is that in order for one to resolve the conflict one accepts the feeling and then looks for a rationalisation to explain away the evaluation: an alcoholic might know full well that he would risk falling back into alcoholism if he has 'just this one drink', but under the grip of his addiction he might feel like having the drink so much that he will contrive a new evaluation of the situation, generate a new belief to replace the belief which is getting in the way of what he wants then and there; he will convince himself that it will be okay, that one drink will be fine and that that will be the end of it. Someone might, to alleviate their feeling of guilt at constantly losing their temper with someone, convince themselves that they aren't reacting irrationally, but that the person who is the object of their anger is selfish, or stupid, or just plain annoying! Sometimes feelings generate evaluations that are appropriate, but often they generate evaluations that just rationalise their irrationality. A problem, which will be looked at later, is that there are beliefs which seem rational and autonomous but which are really just the slaves of the passions that generate them. How widespread we take these beliefs to be ranges from thinking of them as quite rare, only common in people who are particularly irrational, to thinking of them as the norm, at least regarding evaluative or normative beliefs.

¹⁸ Continnence in the Aristotelian sense is also a problem since for the continent man there is the same conflict, it's just that in his case the evaluative belief wins. On the other hand, the virtuous man, whose feelings are in harmony with his beliefs about what is virtuous, and in the vicious man, whose feelings and beliefs both support a life of vice, present no such problem.

We have seen that Lyons' account of emotion has problems dealing with certain emotions, then. Let us look at his treatment of everyday emotion. Maybe we can separate emotions into those which are rational and so susceptible to his account and those which are irrational and so not explained by him. To narrow the account to be concerned only with rational emotions is not too damaging, since these are probably the most important, in the same way as true beliefs are the most important beliefs. If, however, Lyons fails to give a plausible account of even ordinary or rational cases of emotion, then his account will be doomed.

One example Lyons gives is of a man who loves a woman for her shyness: 'John loves Frieda because he believes that she is shy, and that shy people are generally speaking more kind and affectionate than extroverts'.¹⁹ Another is the accounting for the grief someone feels at the death of their daughter by citing the evaluation 'my daughter's death is a grave loss and misfortune to me'.²⁰ These seem completely artificial - surely Frieda would be most insulted if John told her that this was why he loved her! Surely the dead daughter would be turning in her grave! Consider also the following Lyons-style explanation: 'John appreciates the music of Beethoven because he believes that it is Classical and that Classical music is generally more refined than non-Classical music' Surely if someone liked Beethoven's music for this reason, then we'd say that they didn't truly appreciate it at all, that they didn't appreciate the music for itself. This is also what we'd say about the account of John's love for Frieda and other such examples.

¹⁹ Ibid., p. 73.

²⁰ Ibid., p. 96.

What is wrong with these accounts? Someone might say that Lyons is attempting the impossible, that he is attempting to break down emotion into components and that we can't do this because emotion is primitive. Someone who held such a view would say that there are no reasons why we feel the way we do, that we just feel the way we do: we love someone as a whole, not because of any of their particular characteristics, for example. But this is not the only possible objection we can make: even if we accept the possibility of analysing emotion, as I do, these analyses seem faulty. The beliefs Lyons puts forward as the causes of emotions don't even seem plausible as accompaniments to or effects of the emotions he considers. The kind of reasoning from general principles to an emotion as conclusion that Lyons offers doesn't match the phenomenology of emotion. Our appreciation of individuals is often prior to our appreciation of the general categories the individuals fall in: whatever the reason for loving someone who is shy, someone's liking for shyness in general, seems to be reducible to their liking of shy individuals, not the other way around. This seems even clearer in the case of music appreciation: we don't form beliefs about different genres of music and then judge individual pieces according to which genre they fall into, we respond to individual pieces of music and so, as a result, we form general beliefs about musical genres.²¹

But there are other analyses with more plausibility: John is scared of the tiger because he believes it to be dangerous, John is angry at Jack because he believed that Jack had forgotten to lock the tiger's cage, John admires Jane because he believes she is brave to have rounded up the tiger, etc. These examples

²¹ Sometimes people are taught that a certain kind of music is good, but then, even though they will have acquired evaluative beliefs, i.e. even though they will know that they ought to appreciate that

seem to be examples of emotions that can be accurately characterised in terms of cognitions and even if it is wrong to think of the emotions as caused by the beliefs in these cases, there must be some non-accidental relationship between the content of the beliefs cited and the emotions. And even if we reject Lyons' proffered accounts, of the examples given above, there might be more plausible cognitive accounts to be given. We could explain someone's appreciation of a piece of music in terms of their evaluating it as passionate, well-crafted, subtle, etc. We could explain someone's loving a person in terms of their evaluating them as generous, intelligent, attractive, we could explain someone's grief at the death of a loved one in terms of them seeing that person's death as devastating, etc. The difference between these plausible accounts and Lyons' contorted ones is that they refer to beliefs about or evaluations of particulars without the need for appeal to general principles of evaluation in order to generate them.

Lyons appears to think that evaluations must be deductions from general principles. Why is this; why does he not appeal to the goodness of a particular directly? The most likely explanation is that he doesn't think that values are objective properties inherent in particulars. An evaluative belief, such as the belief that a piece of music is majestic can either be thought to arise as a result of appreciating certain (neutral) features of the music, x , y , and z , and then applying the general principle 'pieces of music with features x , y , and z are majestic' or as a result of directly perceiving that things are thus and so, that the piece of music is majestic, for example. But the latter requires that values be there to be appreciated, that values be real properties of individuals. It is Lyons' rejection of this view of values that forces him to put forward the implausible accounts he

kind of music, they will not actually enjoy it.

does. In order for a cognitivist account of emotion as evaluation to offer plausible evaluations to accompany emotions, it must incorporate a realism about values: the values we respond to emotionally must be thought of as features of the world directly open to perception. But then there is a problem: some emotions don't seem to be responses to objective features of the world.

2.2 An alternative account of emotion

Lyons' causal-evaluative theory of emotions fails on at least two counts: it fails to match the phenomenology of emotion and it fails to satisfy the twofold rationality constraint, i.e. it allows emotions to be rational but by tying them to evaluative beliefs it makes them too rational; it doesn't allow for them to be irrational in ways that beliefs can't be.

Why does Lyons insist that emotions always be caused by evaluative beliefs? Surely not because he has evidence for this: as we have seen the phenomenology goes against him. The most likely explanation is that he holds a view of rationality according to which only beliefs and thoughts can be rational in themselves, that in order for emotions to count as rational at all they must be underpinned or caused by beliefs. He says 'A feeling does not of itself provide one with a good reason for wanting to do anything and so for action. An evaluation does, particularly the evaluation involved in emotion, for such an evaluation, being of the world, ... generates desires to change the situation *vis-à-vis* the subject, or else to prolong it.'²² Lyons gives us no argument for the claim that feelings themselves cannot be rational, he merely asserts it. In section 1.2 I

²² Ibid., p. 65.

argued that only feelings can provide ultimate grounding or reasons and if my argument is valid then Lyons' claim is plainly wrong.

I suggest that we put things the other way round. Recall Lyons' definition of an emotional state: '*X* is deemed to be an emotional state if and only if it is a physiologically abnormal state caused by the subject of that state's evaluation of his or her situation.'²³ My alternative definition is as follows:

X is an emotional state if it is a physiological(ly abnormal) state with an evaluative aspect and it causes the subject of that state to acquire evaluative beliefs about his or her situation.

Note that I have dropped the 'only if' clause: it must be possible for non-conceptualisers to have emotions and for some emotions to fail to generate beliefs. The definition doesn't specify the cause of an emotional state because I take it that it is just properties of things in the world, values, etc., that are the causes of emotion (though on another level of description it is prior physiological states that are the immediate causes of emotional states).

I argued that emotion cannot result from evaluative belief, since emotion is not as rational as evaluative belief. But then it seems that evaluative belief cannot always be caused by emotion since that would result in the opposite problem: evaluative belief wouldn't be able to be more rational than emotion and so the twofold constraint would not be met. The possibility of conflict between feelings and evaluative beliefs suggests that neither one can be wholly determined by the other. But this objection to the idea that evaluative beliefs are causally determined by emotions can be met: it is possible to feel something, ϕ , to be good at one time and bad at another time. If one acquires the belief that ϕ is good then when one later comes to feel that ϕ is bad then one will either carry on believing it to be

good, or one's new feeling will overturn the belief one holds and one will come to believe it to be bad. The important difference between my view and Lyons' view can be thought of metaphorically in the following way: my view starts with openness (emotion being less rationally constrained than belief) which gets narrowed down (belief) whereas his view starts with narrowness (belief) and so precludes the openness required (emotion). There are more possible emotions open to one than there are possible evaluative beliefs open to one, thus, at least some emotions aren't caused by evaluative beliefs.

I have talked about evaluative beliefs being acquired emotionally, but it is not the case that all evaluative beliefs are acquired emotionally: the belief that umbrellas are good things to have in heavy rain is surely not the result of emotion but is just something we learn is true. Which beliefs are to be thought of as emotionally acquired, then? Can we make a distinction between methods of belief acquisition in this way, distinguishing emotionally acquired from other beliefs? We cannot in general judge from the content of a belief whether it was acquired emotionally or not: the belief that *C* is cruel could be emotionally acquired by *A* and neutrally acquired by *B* as a result of being told so by *A*, for example. So the distinction does not coincide with the traditional fact/value distinction that would separate beliefs with factual content from those with evaluative content.²⁴ The distinction is comparable to the distinction between perceptual and non-perceptual beliefs: in this case we are distinguishing between beliefs acquired in perception and those acquired some other way. It is possible for beliefs about perception to be non-perceptual in this sense as it is possible for a particular belief *B* to be

²³ Ibid., pp. 57-8.

perceptual for one person and non-perceptual for another - if *A* sees that the sky is blue and tells *B* that it is, then the belief that the sky is blue is perceptual for *A* and non-perceptual for *B*.²⁵

The discussion in section 1.2 of bare goods suggests that the beliefs relevant to meaningfulness of life are emotionally acquired beliefs: the difference between someone whose life is meaningful and someone else whose life fails to be is not to be thought of in terms of the differing content of their beliefs about their lives, their beliefs about whether or not their lives are meaningful, for example, but in terms of a difference in emotional involvement. Someone who knows that her life is meaningful in some sense, without having acquired the belief emotionally, lacks the meaningfulness the person who has acquired the belief emotionally has, she may have the belief that her life is meaningful, but she lacks understanding of what that amounts to.

2.3 Emotional knowledge in Aristotle's Nicomachean Ethics

Myles Burnyeat, in presenting an interpretation of Aristotle's account of habituation and the role it plays in the acquisition of virtue aims 'to reconstruct Aristotle's picture of the good man's development over time, concentrating on the earlier stages'.²⁶ Burnyeat argues that not only is habituation a prerequisite for knowledge of virtue but it is actually the method of acquiring it. He points to what Aristotle says at 1098^b33-4 to support this: "The that" is a first thing and a starting point. Of starting points some are seen by induction, some by perception, some by

²⁴ Although the distinction I am drawing can perhaps help to explain the appeal of the fact/value distinction: perhaps it is the distinction those who draw a fact/value distinction are looking for.

²⁵ Although the belief is acquired perceptually by *B* in a sense - *B* hears *A*'s words. The important difference is that the content of the belief is not perceived by *B* whereas it is by *A*.

²⁶ 'Aristotle on Learning to be Good' in A. E. Rorty ed., *Essays on Aristotle's Ethics*, p. 69.

a certain habituation, and others in other ways again.’ Burnyeat says ‘we learn by induction that all men breathe, by perception that fire is hot. In ethics the appropriate mode for at least some starting points is habituation. ... The thesis is that we first learn (come to see) what is noble and just not by experience of or induction from a series of instances, nor by intuition (intellectual or perceptual), but by learning to do noble and just things, by being habituated to noble and just conduct. ... Aristotle is not simply giving us a bland reminder that virtue takes practice. Rather, practice has cognitive powers, in that it is the way we learn what is noble or just.’²⁷

Burnyeat is claiming that knowledge and understanding of the virtues arises purely out of training. The obvious question to ask is: how can one learn that something is noble or just by becoming habituated into doing it; how does knowledge arise out of habituation? Burnyeat thinks the key to the answer given by Aristotle is the link between pleasure and virtue: someone who is learning to be virtuous will be learning to take pleasure in doing virtuous things. He considers the link between the phenomenon of learning to enjoy something and that of learning that the thing is enjoyable. In a weak sense of the word ‘learn’ one can learn that something is enjoyable without coming to enjoy it, but there is a stronger sense of the word in which one can only learn that something is enjoyable by coming to enjoy it. Here, Burnyeat says, ‘The growth of enjoyment goes hand in hand with the internalisation of knowledge.’²⁸ Further, ‘Aristotle holds that to learn to do what is virtuous, to make it a habit or second nature to one, is among other things to learn to enjoy doing it, to come to take pleasure - the appropriate

²⁷ Ibid., p. 73.

²⁸ Ibid., p. 76.

pleasure - in doing it.’²⁹ Aristotle also claims that the virtuous man must choose virtuous actions for their own sake, i.e. because they are just and noble. Now if this choosing of virtuous actions for their own sake is reliant on the enjoying of virtuous acts for themselves and their intrinsic value and this is acquired through habituation then we can make sense of the way knowledge of virtue is acquired in that way: Burnyeat says ‘I may be told, and may believe, that such and such actions are just and noble, but I have not really learned for myself (taken to heart, made second nature to me) that they have this intrinsic value until I have learned to value (love) them for it ... To understand and appreciate the value that makes them enjoyable in themselves I must learn for myself to enjoy them.’³⁰ Beliefs about virtue, etc. cannot make one enjoy virtuous acts, cannot bring about the right emotions. But, on the other hand, emotions can bring about beliefs.

Burnyeat’s account of how knowledge of virtue arises out of habituation yields a picture of this kind of knowledge as active and personal. He contrasts it with the detached, academic knowledge that certain acts are virtuous. The kind of knowledge acquired in this way is a direct result of learning to feel a certain way - unlike the dry, impersonal knowledge of the same things, which needn’t even be accompanied by emotion. In order to have real knowledge of virtue, and also of meaningfulness, and value in general, one must feel it: it is not possible for someone to have this kind of knowledge without also having the right feelings and desires. To really know that one’s life is meaningful requires one to feel that one’s life is meaningful.

²⁹ Ibid., p. 77.

³⁰ Ibid., p. 78.

3. The causes of emotion and association in emotional life

In this section I will explore the ways in which emotion is thought to be less rational than other aspects of mental life. I will start off with worries about the possibility of inducing emotions artificially: if emotions can be induced independently of the way the world is, then the way emotional experience presents things to us will seem suspect. I will then turn to the associative nature of emotion.

3.1 Induced emotions and scepticism about the evaluative element of emotion

William Lyons discusses a series of psychological experiments done by Schachter and Singer in which subjects were given doses of adrenaline and then placed in contexts of euphoria or anger. Those who'd been told that they'd been injected with adrenaline and about the effects of adrenaline didn't report their state as emotional. Those who hadn't been told that they'd been injected with adrenaline reported feeling euphoric or angry. Lyons uses these experiments to emphasise that physiological changes are not enough in themselves to determine emotion, that unless a physiological change coincides with an evaluation, it won't be interpreted by the subject as an emotional change.

These experiments throw up a serious problem for Lyons' account, a problem Lyons seems unaware of. Lyons uses Schachter and Singer's results to back up his causal-evaluative theory of emotion when they actually refute it. The conclusion that one is led to on considering Schachter and Singer's results is not that emotional states are in actual fact 'abnormal physiological states caused by the subject of that state's evaluation of his or her situation', which is the definition

Lyons gives,³¹ but that emotional states are abnormal physiological states (not necessarily caused by an evaluation but) *thought by the subject to be* caused by (or in some way related to) an evaluation by the subject. What is actually happening to the experimental subjects is, it seems, that a physiological change is causing an evaluation and thus being seen as an emotional change rather than the other way round. Contra Lyons, this is what happens in general.

The way the experimental subjects interpret their reactions is problematic: we know that their responses were induced chemically and that seems to show that the emotions weren't appropriate. How should we describe what happened to them? Either we say that they were deceived, thinking that they were reacting emotionally to a situation when in actual fact they were just reacting to the drug, or we say that they weren't deceived, that they really were reacting emotionally, but then we are forced into saying that in general an emotion directed at a particular thing needn't be caused by that thing, it is only necessary for it to be thought of as the cause by the subject. The sceptic about emotion will use results like these to argue that emotions are arbitrary, that they aren't ever really about the things we take them to be about, but are just linked by us with the things that happen at the right time to be used to explain our physiological swings: the way we feel tells us nothing about the nature of the world, but in fact distorts the view we have of the world to fit in with it. Or a less radical sceptic would accept that real emotions are reactions to the way things are but would point to 'fake' emotions such as the ones experienced by the subjects of these experiments and argue that we can never know whether a particular emotion is real or not.

³¹ Ibid., p. 58.

It could be argued against the sceptic that the fact that adrenaline was introduced into the subjects' bodies isn't enough to show that their emotions weren't genuine: the situations the subjects were placed in warranted emotional reactions and so their bodies would've reacted by producing natural adrenaline had there not already been an excess of adrenaline there; the same amount of adrenaline would have resulted even if the subjects had not have been injected with it. In favour of this approach it can be pointed out that if the subjects had been placed in calm situations they wouldn't have reacted angrily. In order for the sceptical point to hold, it must be shown that emotions can be induced without the help of emotional scenarios; that subjects can be brought to react emotionally to neutral stimuli. But this experiment does show that an essential causal link between properties such as dangerousness and physiological states cannot be made: we can respond emotionally to something in the world without the contributing physiological changes having been caused by that thing.

3.2 The causes of emotion

When can knowledge of the cause of an emotion threaten its validity? Normally an emotion is caused by properties of things in the world: our fear is caused by dangerousness, for example. At the neurological level, emotional states are presumably caused by other brain states. What are we to say about the possibility of inducing emotions neurologically? If one's mood - contented or depressed - is due to the presence or absence of certain chemicals in one's brain, or other neurophysiological factors, and these can be artificially produced, then how can one's mood have cognitive significance? Should we just separate the two levels of

causal explanation and deny that they compete, or should we allow that a neurological or evolutionary explanation of what caused an emotion can threaten its validity? Further, emotions sometimes have subconscious causes which seem to threaten their validity: if someone is hypnotised to react emotionally to a certain stimulus, then surely the emotion isn't a valid emotional response to that stimulus?

William James argues that the cause or origin of a religious experience is irrelevant to assessing it, that we must judge such an experience by its fruits. He considers the 'assumption that spiritual value is undone if lowly origin be asserted',³² the assumption that pinpointing the underlying process or condition of a religious experience shows it to be 'nothing but' that process or condition: 'Alfred believes in immortality so strongly because his temperament is so emotional. Fanny's extraordinary conscientiousness is merely a matter of overinstigated nerves. William's melancholy about the universe is due to bad digestion ... Peter would be less troubled about his soul if he would take more exercise in the open air.'³³ According to some, even the visions of the saints are to be dismissed as due to epileptic seizures, hysterical fits, etc. James argues that accepting the existence of some underlying process shouldn't lead us to view religious experiences as any less valid, since 'there is not a single one of our states of mind, high or low, healthy or morbid, that has not some organic process as its condition.'³⁴ Unless we are eliminative materialists across the board, then, we have no reason to think that religious experience is threatened by materialism about the mind.

³² *The Varieties of Religious Experience*, p. 32.

³³ *Ibid.*, pp. 32-3.

³⁴ *Ibid.*, p. 36.

James' argument carries over to emotional experience in general but there are two problems with it: firstly, there are some mental states that we want to describe as invalid or abnormal and if we cannot invalidate religious experiences by reference to the underlying brain states then we won't be able to do so for any mental states. The second problem is that even if the causes at a low level, the neurophysiological causes, don't conflict with our view of religious experiences, there are other causes which do seem to - subconscious causes and evolutionary causes, for example. When the causes of our emotions don't conflict with our understanding of those emotions then James' argument can be accepted. When the cause of an emotion seems to threaten the rationality of that emotion then things are different.

The first problem is not a real problem, in fact it serves to alert us to a mistake - the mistake of thinking that it's possible to view brain states as normal or abnormal without reference to the mental level. We don't have any way of distinguishing between normal and abnormal brain states qua brain states: at the neurophysiological level there's nothing wrong or right, just causality. To explain an abnormal mental state in terms of chemical imbalances and not have a criterion for what is normal in neurochemical terms is to have already judged the abnormality at the level of intention, etc. What makes this particular state an *imbalance* in the brain is that what is experienced doesn't fit with our conception of reality.

The second problem is raised by James' discussion of post-hypnotic suggestion. James describes the process: 'You give a hypnotized subject, adequately susceptible, an order to perform some designated act - usual or

eccentric, it makes no difference - after he wakes from his hypnotic sleep. Punctually, when the signal comes or the time elapses upon which you have told him that the act must ensue, he performs it; - but in so doing he has no recollection of your suggestion, and he always trumps up an impoverished pretext for his behaviour if the act be of an eccentric kind.³⁵ In situations like this the actual cause of the subject's behaviour - the suggestion planted by the hypnotist - falsifies any rational explanation: the fact that the subject behaved as they did as a result of post-hypnotic suggestion goes against them having any *reason* for behaving in the way they did. Here, unlike with physiological causes, the cause invalidates the reason. In contrast with the point James made about physiological causes, i.e. that all mental happenings have physiological correlates or causes, we can see a difference between normal behaviour and hypnotically suggested behaviour: normally there is no underlying suggestion, there is a reason, which is why normal behaviour is rational and the behavioural responses to post-hypnotic suggestion are not.

Behaviour, perceptions, and emotional responses can result from the workings of the subconscious, and when this happens then the rationality of that behaviour, perception, or emotion is called into question. The subconscious is not completely unconstrained - a subject cannot be persuaded to do something he is vehemently opposed to - but it is less constrained than the rational mind and will output behaviour, etc., without needing a full-blown reason for it.

As well as post-hypnotic suggestion, there are other ways in which the subconscious can bring about irrational behaviour. Suppose we have innate tendencies that govern the things we respond to, the things we find attractive,

³⁵ Ibid., p. 235.

frightening, etc., as a result of evolution. Suppose our genes, like the hypnotist, have planted suggestions in our subconscious such as ‘when you see *X*’s, run; when you see *Y*’s, attack; when you see *Z*’s, mate’, then the reasons we give for running, attacking and mating will be false. Evolutionary explanations for aspects of human life seem to threaten the rationality of those aspects, then.

The idea that like animals our responses to other human beings are really governed by smells, hormones, innate programmes, etc., is disconcerting. Similarly with the idea that when we listen to someone we’re learning more from their tone of voice, body language, etc., than from the content of what they say. It might seem that we must accept the claim that rational life is epiphenomenal, the non-rational, subconscious processes being the causally relevant ones. But this is overly pessimistic: the processes underlying conscious life cannot be arational - how would they have been conducive to our survival in that case?

If the most rational of our thought processes are based on and affected by subconscious or emotional patterns and processes then it will be important to discover the irrationalities and limitations of these patterns and processes. I suggest that we reject the assumption that these patterns and processes are inherently irrational in contrast with rational thought but look closely at the particular ways they can fall short of our rational ideals. The emotions and the subconscious are the roots of meaningful life then but they are also less refined than conceptual thought: we can see some emotional responses as being in need of correction due to being governed by good (survival-conducive) rules of thumb rather than precision. ‘Run first, ask questions later, from a distance’ is presumably a better danger-response than ‘don’t run until you’re sure there is real danger’.

Some of our fears are bound to be unwarranted. Similarly with our other emotions.

3.3 Real and projected values

It is sometimes said that emotions work by association, that emotions can be triggered by something as a result of the thing being associated mentally with other emotionally charged things. A piece of music can be overwhelming not because of its own character but because it reminds one of someone one loves and misses; one can love someone because she reminds one of one's mother, and so on. With experience in general, something's appearing ϕ is good reason for thinking it really is ϕ . With emotional experience, it is suggested, this does not hold true. Some claim that values are not real but are projected by us onto the world.

There is a *prima facie* distinction between those of our emotional responses which are responses to things as they really are, to real values, and those which are responses triggered by something other than real values. It might seem that we can differentiate between real/objective values and subjective/projected ones with reference to shared reactions: something is a real value if everyone responds in the same way to it and something is a projected value if different people respond differently to it. It can then be pointed out that even the most objective of values are not actually responded to uniformly: psychopaths fail to respond to moral values, for example. Maybe the formulation can be modified to refer to what our responses ought to be rather than to what our responses actually are: something is a real value if everyone ought to respond to it in a certain way. But this is still the wrong place to draw the distinction since there are real values which we don't

expect everyone to respond to uniformly and these are to be contrasted with projected values: I can love a person and be the only one to do so, but this does not necessarily mean that my love projects rather than responds to that person's goodness. The question of whether I love them for themselves or whether I love them because they remind me of my mother, for example, can still be asked. When we are trying to understand our emotions then this is the kind of question we should ask.

Frances Berenson discusses emotion in music and argues against the 'purist' claim that there can be no true ascriptions of emotion to music, that our emotional responses to music are just in us, not in the music.³⁶ Rather than drawing a distinction between those features which are just in the music and those features which are just in us, projected by us onto the music, which would be the most obvious way of distinguishing between objective and subjective properties, Berenson argues that the important distinction is between the features picked out with essential reference to the music (and also to the self) and those features picked out with reference exclusively to the self. She relates this distinction to Roger Scruton's distinction between meaning and association and she says 'my first category embraces cases of complete involvement with the music, in the second category we use the music as a backcloth for extra-musical experiences. In the first category the reference to the self focuses on explaining how I came to understand some emotion in the music. ... That sort of reference is different in kind from reference to the self where the music is used merely as a means to some further end which can only be linked to some specific episode in my past which I

³⁶ 'Interpreting the Emotional Content of Music' in Michael Krausz ed. *The Interpretation of Music*.

then proceed to dwell on.’³⁷ Musical involvement allows one to discover things about the music. In contrast, ‘in pure association there is nothing new to be discovered, it is a calling forth of a specific past experience which we get emotionally involved with, forgetting the object which gave rise to it.’³⁸

A question that arises at this point is the following: what about when the music conjures up a certain emotion which is linked to some specific episode in my past but which I don’t consciously dwell on, which I don’t realise is due to association, etc.? The purist will say that all emotion in music is like this. The projectivist about value will say that all values are really like this, that even when we think that we’re responding to the music itself, we’re not - we’re using it as a trigger or backdrop for our own emotional projections. Berenson doesn’t see anything problematic about the fact that training or initiation into a musical tradition is often necessary for musical appreciation. But if it is only those who have been exposed to certain things who are able to hear the emotion in a piece of music then what is the difference between these ‘specific episodes in one’s past’ which don’t threaten the objectivity of what is heard and the ones Berenson wants to describe as subjective? Composers, poets, etc., rely on the web of associations shared by their audience in order to evoke atmosphere, mood and setting: if the only difference between minor chords really being sad and a certain piece of music not really being anything to do with *X*’s dead mother is that we all associate minor chords with sadness and only *X* associates the piece with his dead mother, then this is not a difference between things as they really are and things as they merely seem to us.

³⁷ Ibid., p. 69.

³⁸ Ibid., p. 71.

It might seem that we should accept a full-blown subjectivism or projectivism about values, claiming that values don't exist in the world, that we overlay them onto the world. But if this is what happens then what governs which things we project which values onto? How are we to understand the process that starts with a neutral object and ends up with an emotion and a value judgement? One view would be that there is a projecting mechanism in our brain that gets activated randomly and so colours our perceptions and causes us to react emotionally. On this view there is no correlation between things in the world and our values and so this is clearly not the right view of things. Suppose that there is a neutral world onto which we project our values in some orderly way. What is the relationship between this world and our values? How are we to gear our projections to the neutral world?

Perhaps there is subliminal perception of the neutral world which activates the projective mechanism and which overlays the neutral world with values for us to then perceive consciously. The projection cannot be prior to all perception because if it were, then it would be related randomly to the world. But then if we allow for there to be some input from the world that governs our projecting, then it is not obvious that values are purely projective. If there is a feature *F* of the neutral world which activates the projection of the property of dangerousness, for example, then what prevents us from saying that feature *F* is dangerousness?

There must be some structural relationship between things in the world and one's emotional responses, so even if one is projecting and not responding directly to things as they are, projecting dangerousness onto spiders, for example, then one's reaction is still a reaction to spiders. In other words, there is a correlation

between things in the world and our responses: our responses are structured in this way, not randomly related to the world. It cannot be that our emotional responses are completely off kilter, completely unpredictable; things must be as follows: some kind of thing in the world, X , triggers an emotion E or triggers an associational chain which leads to emotional response E' . In both cases we have a regular emotional response, although in the case of E' our response is not really a response to X : we take E' to be a response to X , but it is really a response to (a memory or innate representation of) something else, A , which is the thing at the end of the associational chain which prompts E' .

Because our emotional responses are regular we can come to rely on them even if they aren't appropriate; we can compensate for them. Sometimes we can even rework the emotional responses themselves and come to respond in more appropriate ways. If one comes to realise that certain harmless characteristics, such as someone's pinching the bridge of their nose when thinking, or laughing nervously, which irritate one do so because they remind one of one's father, for example, then even if one can't help responding in this way, one can at least reinterpret one's irritation and refrain from making judgements on the basis of it.

3.4 Association

Freud distinguished between 'primary process', which is the mental functioning of the archaic mind, and 'secondary process', which is conceptual functioning. According to Jonathan Lear, primary process 'tends to be expressed in concrete images, rather than in concepts, and it tends to proceed by loose associations'.³⁹ Further, it is not constrained by logic. He describes the relationship between

archaic association and conceptual thought as follows: 'The loose associations of archaic mind represent the infant's primitive sense of similarity and relevance. This sense of similarity is itself honed and refined, so that by the time a person is able to think conceptually, the associations of primary process appear to him to be loose.'⁴⁰ Lear suggests a mechanism for the development of associative pleasure: an infant associates its mother's breast with the satisfaction of hunger (which is to be thought of as a basic pleasure) and as a result comes to hallucinate its mother's breast when hungry, this hallucination or visualisation forms the basis of desire, which can later be conceptualised, and also becomes pleasurable in itself, independently of its origin in the satisfaction of hunger. So not only will the infant hallucinate its mothers breast when hungry, there will be an opposite association set up - it will come to find breasts pleasant.

Consider an example from later life: one finds a certain piece of music pleasurable, then, from this starting point, one is made to associate the piece of music with a certain product - a brand of jeans, cigarettes, alcohol, or a type of car - as a result of a successful advertisement campaign. What then happens is firstly that the music prompts one to think of the product when one hears it again, and, secondly, that one comes to enjoy or appreciate the product itself, the look of a particular brand of jeans, for example. The product becomes imbued with the pleasure one gets from the music, but no longer depends on actually being accompanied by the music once the association has been set up. In this way pleasures can be multiplied - new pleasures come about as a result of being associatively connected with old ones. Like King Midas turning everything he

³⁹ *Love and its Place in Nature: A Philosophical Interpretation of Freudian Psychoanalysis*, p. 76.

⁴⁰ *Ibid.*, p. 85.

touches into gold, pleasures turn everything associated with them into pleasures. Likewise with pains. Thus things that are not really good come to be loved, things which are not really scary come to be feared, etc. A sceptic would argue that the existence of this associative spreading of pleasures, pains, etc., is enough to show that emotions are unreliable in general, that if we rely on them to tell us what is good and what is bad we will be misinformed more often than not.

One response to this kind of scepticism is to deny that association is all that widespread as a cause of emotion. This is not a convincing response though: associational elements seem to be pervasive; it's plausible that all but the most basic of our emotional responses were set up associatively. A different approach would be to argue that although our emotional responses result from a network of associations, this is unproblematic: there is no reason for assuming that association is irrational, that the different network of links is still a real network of properties. My response will be related to this.

The unscrupulous use commercial advertising makes of the associative nature of the subconscious suggests that the associations set up yield emotional responses which are purely negative, deceptive, but this is not the case. Associations of the same kind underlie some of our most valued emotional experiences: good music, poetry, myth, seeing a casual gesture as meaningful, etc., all rely on the power of the subconscious to make associative links. As well as relying on the meaning of words, poetry relies on features such as colour, tone, rhythm, which are associative or evocative. Association enriches our experiences, giving us in the experience of a particular thing echoes of countless other things. We listen to a piece of music and find ourselves transported vividly to another

time and place, we return to a particular place and find ourselves accompanied by the voices of those we spent time with there, we love the subtlest gesture or movement because we see in it the person we love, etc. Those who lack the power of association seem to live dull lives, seeing just what's immediately given. If there is something good about our ability to associate, to build mental connections between things which aren't essentially connected in reality, then what sort of good is it? We enter the realm of fantasy, leaving behind objective reality. But how can fantastical projection be better than bare objective perception?

If the world around one is rich then perceiving is good, if the world around one is sparse, then imaginative projection may be better. If one finds oneself in a bare prison then the facts of one's existence may be grim and not worth dwelling on. One might do better to retreat to the realm of the imagination, to imagine the outside world, to relive conversations with other people, to develop imaginary worlds, etc. Sometimes it's better to just forget about how things are and imagine things being better. Sticking with the objective facts is not enough for us: our reality is more than just the external world. The things we discover as a result of creative thinking, imagining, etc., can be equally as real and important as the things we discover as a result of perceiving things in the world. Whatever differences there are between the external world and the world of the mind, we shouldn't think of the external world as real and denigrate the realm of imagination.

Consider the following scenario: someone goes to an ancient tomb and hears voices around them. There are three kinds of explanation: either they are imagining things, their imagination is conjuring up the voices, or someone has set

up a tape player or other device to deceive them, or there are voices of ghosts or echoes of the past or something in the tomb itself. If the important question is whether something is part of a person's imagination or whether it has an external source, then if the first explanation is the correct one, the experience is to be dismissed, whereas if either the second or third explanations are true then the experience is to be considered legitimate. This doesn't seem right at all - it seems to be the second explanation that is uninteresting. If we find out that a strange experience we had was just a deception then we would dismiss it. (Maybe we would become curious about why someone wanted to deceive us, but this would be to go off at a tangent.) Whereas if either of the other two explanations is the truth then we would have an interesting phenomenon. If lots of people heard voices in the tomb, or if one person heard them on many occasions then would that show that the source of the voices is something in the tomb rather than in the imagination? If not then how are we to distinguish between imaginative projection of voices and perception of voices?

3.5 Emotional seepage and transference

If one is open to the possibility that one's emotional responses to some things are the result of associations rather than responses to the things themselves, then one can start to look for associative explanations for all of one's emotional responses. One can ask of things one previously took for granted: why do I feel the way I do? If one notices that all the men one falls in love with happen to have the same

colour hair as one's father then one should consider the possibility that this is no coincidence, for example.⁴¹

As well as one's emotional responses being open to question in this way, one's outlook and evaluative beliefs are also susceptible to the same treatment: if an evaluative belief is acquired emotionally then any distortion in our emotional reaction will be transmitted to the belief. Once we start digging, we might discover all sorts of strange explanations for our beliefs. Consider the following example: I have always believed that locking doors and shutting windows is a bad thing and that in general it is much more pleasant to have an open house than a closed one. I took it that this belief was related to my belief in the importance of trust and the idea that it's worse to worry too much about one's belongings than to actually risk them getting stolen. Then I remembered that as a child I had once been locked in our outside toilet: could this experience have brought about my belief? Possibly. It seems that if one finds plausible candidates for the causes of one's attitudes then one must take them seriously and consider whether or not they threaten the validity of one's attitudes. In this example the best explanation for my belief that it's better to leave doors unlocked is not that I'd considered options and decided that locking doors and being security conscious was a bad way to be but the reverse - my preference for keeping doors unlocked generated my beliefs about this being more relaxed and trusting as a kind of rationalisation. The preference itself is plausibly the result of my associating locked doors with being trapped.

⁴¹ There are other explanations for such patterns in a person's attractions - evolutionary explanations, for example. If elements of attraction are inherited, then an explanation for why a woman is attracted to men like her father is that she has inherited her mother's genes and thus inherited an attraction to the same sorts of men her mother is attracted to. The existence of explanations such as this don't in any way conflict with my general point, which is that some emotional responses, although responses to things in the world, are not true responses to things as they really are, so we cannot always take our emotional responses at face value.

If associational links can be set up easily, just by presenting a subject with two things together often enough, then can we be conditioned to feel *anything* is good? If one can be persuaded to buy a certain brand of jeans by being presented with images of that brand of jeans conjoined with things one already likes, then can one be persuaded to support Fascism, or to approve of murder, in the same way? The escalating amounts of money spent by political parties on advertising suggests that our political attitudes are malleable to some extent. But presumably even the most successful advertising campaign will not persuade the most staunch Tory to vote Labour. Is it only those who have yet to make up their minds who can be swayed by the power of association? And if so, then what is it that makes others immune, their firm belief or the fact that a strong association which conflicts with the one on offer has already been set up?

Aristotle discusses the emotions in the *Rhetoric* and the main conclusion he draws concerning them is pragmatic: emotions can be used as a tool to sway people. He says ‘The Emotions are all those feelings that so change men as to affect their judgements’⁴² and goes on to give a strategy for arousing anger which would be useful to an orator trying to persuade his audience to go to war. The mechanism seems to be a kind of seepage or spread: a strong emotional response to one thing leads one to feel strongly about other things that one is then presented with. If emotions can be easily manipulated and if one’s emotional state affects one’s judgement, then judgement can also be manipulated. A sceptic could argue that a view of the world as good or meaningful is not a true view of the world but the result of emotional seepage - when one feels happy about things in one’s life

⁴² *Rhetoric*, 1378^a20-21.

this happiness seeps out into one's view of the world as a whole, similarly when one is unhappy about specific things.

Aristotle's claim, that one's feelings affect one's judgements, is supported by a number of psychological experiments. One is reported by Gerald Clore as follows: 'Subjects initially engaged in a task in which emotional thoughts and feelings were activated. Later, they read an evaluatively ambiguous description of a person and rated how much they liked him. The results showed that the subjects' perceptions of the ambiguous person were biased by their prior affective experience. Subjects whose happy thoughts and feelings had been activated saw him more positively than did subjects entertaining sad thoughts.'⁴³

It might be countered that there is no real problem, that the different judgements one makes when one is under the influence of different feelings are not conflicting - they each focus on, although perhaps magnify, different aspects of a situation. When one is happy and comes to see the world as good then one is able to see the real goodness of the world. When one is miserable and sees the world as bad then one sees the real bad in the world. Unless we think that the world must either be totally good or totally bad, this is unproblematic. Different feelings aid judgement, it could be claimed, by showing one the different sides of things and there is only a problem if one fails to realise that the judgements one's feelings lead one to make are just part of the truth of the matter.

If one becomes depressed, all the things in one's life, things one loved and valued previously, just lose their shine. Everything becomes dead to one: one can see that according to some standards one no longer identifies with, one has a good,

⁴³ 'What is the Function of Emotions?' in Paul Ekman and Richard Davidson eds. *The Nature of Emotion*, p. 104.

meaningful life, that one should be happy. But one just isn't. That there are good things in one's life can make things worse: one is unhappy and has no reason for being so, so one cannot even *understand* one's unhappiness. One may try to work out why one feels the way one does, to find out what has changed: could it be due to difficulties with one's work or one's relationship, or has one come to see the truth - that there really is no meaning, no point to anything? If one recovers, then one may come to see that there was no *reason*, that a black cloud had just descended on one and eventually one had somehow managed to leave it behind. When there are no reasons one can find for one's depression, one point to chemical and neurophysiological *causes*. And if one's feeling one's life to be meaningful can seem so shaky then how can feelings ever be thought of as firm enough to give one the truth? Seeing the world as meaningful can come to seem a naïve and shallow way of seeing: meaningfulness is shown to be like the surface of a calm sea, smooth and stable until turbulence breaks it into many pieces. Once meaning has been broken into little pieces in this way, can it ever be repaired? William James captures the worry poignantly: 'Of course the music can commence again;- and again and again - at intervals. But with this the healthy-minded consciousness is left with an irremediable sense of precariousness. It is a bell with a crack; it draws its breath on sufferance and by accident.'⁴⁴

If we take feelings of meaningfulness to have the potential of showing us things about the world, i.e. that it is meaningful, then we must surely allow that feelings of depression also have the potential to show us things about the world. And the things these feelings seem to show are things like that everything is meaningless, everything good comes to an end, everything we venerate has a base

⁴⁴ Ibid., p. 145.

source, etc. Unless we are able to single out depressive feelings and show that they are false, that they are abnormal and deceiving, which doesn't seem possible (we can single out abnormal states of mind but only pragmatically - the abnormal is that which prevents healthy life) we are compelled to take into account the way the world seems to the depressive, to the pessimist.

The sceptical questions raised here have not been fully answered: I have only been able to suggest ways in which the worries might best be tackled. I argued that emotion and the subconscious are fundamental to the possibility of meaningfulness, that for one's life to be meaningful one's life must feel meaningful, that only emotionally acquired beliefs have the power to make one's life meaningful, and showed that although there are many real ways in which emotions can fail to be reliable this should not lead us to give up hope completely. There are real values, objective things we are right to respond emotionally to, and the fact that we so often find ourselves to be mistaken in our emotional responses should not lead to radical scepticism about them, to the conclusion that there are no real values. A healthy wariness about one's emotions can only aid one's understanding: the questioning of an emotion will either lead one to reject that emotion as unreliable or will give one good reason for think that it is to be trusted after all. Disassociation from one's emotions, on the other hand, cannot be good, no matter how unreliable one's emotions are: without an emotional foundation nothing meaningful can be built.

IV

ELEMENTS OF MEANINGFULNESS

‘Most people feel on occasion that life is absurd, and some feel it vividly and continually.’

-THOMAS NAGEL, *THE ABSURD*

‘It is a terrible thing to try to live a life without believing in *anything*. But surely that doesn’t mean that just any set of concerns and beliefs will do ... Surely, if any old set would do, that is the same as life’s being meaningless?’

-DAVID WIGGINS, *TRUTH, INVENTION, AND THE MEANING OF LIFE*

‘Everything is indeed permitted if God does not exist, and man is in consequence forlorn, for he cannot find anything to depend upon either within or outside himself.’

-JEAN-PAUL SARTRE, *EXISTENTIALISM AND HUMANISM*

‘What we are doing here, *that* is the question. And we are blessed in this, that we happen to know the answer. Yes, in this immense confusion one thing alone is clear. We are waiting for Godot to come’

-SAMUEL BECKETT, *WAITING FOR GODOT*

I have looked at three areas of scepticism which seemed to threaten the possibility of meaningfulness. I will now deal directly with the question of what makes life meaningful and what is necessary in order for it to be possible for life to be meaningful. In section 1 I will draw out links between meaningfulness, the way one sees one’s life, optimism and unity. The existence of a strong link between one’s life being meaningful and seeing one’s life as meaningful will lead to worries about the objectivity of meaningfulness, which I will try to answer. Then an essential link between one’s life being meaningful and one’s being optimistic about the world will be argued for.

In section 2 the question of the substance of a meaningful life will be addressed: I will look at worries that suggest that the substance of one’s life is arbitrary. Then the idea that meaningfulness requires more than just that one’s life

can be seen as being good, that meaningfulness requires one's life being *right* in some sense. I will explore the idea of destiny in this context.

1. Meaningfulness, objectivity and optimism

There are many objections to the idea of meaningfulness as objective, some of which I have already answered. That meaningfulness can only be discerned from certain viewpoints, by those with the necessary concepts and the right feelings or temperaments does not fundamentally threaten the objectivity of meaningfulness, if the work done so far has been successful. The link between meaningfulness of life and life's seeming meaningful might suggest that meaningfulness cannot be objective - that one's life's being meaningful is reducible to one's life's seeming meaningful to one, and then meaningfulness is merely an appearance. I will argue that this is wrong.

1.1 Meaningfulness and absurdity

Nagel had been grappling with what he later located as the tension between subjective and objective viewpoints long before he wrote *The View From Nowhere*. In an early paper entitled 'The Absurd'¹ he argued that our lives are absurd in a general philosophical sense over and above the particular absurdity of certain situations and that this absurdity arises due to 'the collision between the seriousness with which we take our lives and the perpetual possibility of regarding everything about which we are serious as arbitrary, or open to doubt.'² It is the existence of a viewpoint 'outside the particular form of our lives, from which the

¹ First published in *Journal of Philosophy*, LXVIII, no. 20 (1971) and reprinted in his *Mortal Questions*.

seriousness appears gratuitous'³ that is the root of the problem. Once we come to see that there is such a viewpoint, Nagel assumes that we must come to see that 'a sense of the absurd is a way of perceiving our *true* situation.'⁴ The fact that we *can* view our lives as absurd, coupled with the fact that this view is more detached, i.e. more objective on Nagel's view, than the immersed views which enable us to see life as serious and important, leads to absurdity being the real truth of the matter. Life is absurd and all we can do is approach it with irony, realising that nothing really matters, i.e. nothing matters *sub specie aeternitatis*. As a consolation, he reminds us that if nothing really matters then the fact that nothing matters doesn't even matter!

Nagel examines possible ways of avoiding the conclusion of absurdity and bravely rejects them all as false comforts. The idea that one's life concerns could matter as a result of them being shared by the rest of humankind, by God, or by the cosmos, is shown to be irrelevant - we can raise doubt about the meaningfulness at the next level too - 'a role in some larger enterprise cannot confer significance unless that enterprise is itself significant.'⁵ There is always the possibility of a more detached view from which even the largest of enterprises seems meaningless and this is the problem.

As we saw in part I, Nagel's view of objectivity as detachment is unsatisfactory. The idea that those things only perceivable or understandable through involvement are less real, less objective, than those things available to all perceivers was rejected. The view from nowhere was rejected in favour of the most all-encompassing view. I will first argue that on Nagel's detachment model

² *Mortal Questions*, p. 13.

³ *Ibid.*, p. 14.

of objectivity he is mistaken that absurdity is objective and will then go on to describe the relationship between meaningfulness, meaninglessness and absurdity according to the containment model of objectivity.

Nagel says that absurdity is a specifically human phenomenon - a mouse's life is not absurd, for example, since the mouse 'lacks the capacities for self-consciousness and self-transcendence that would enable him to see that he is only a mouse. If that *did* happen, his life would become absurd'.⁶ But this means that absurdity is not viewable from the most detached view, only viewable by those who have access to the illusion of meaningfulness. From the viewpoint of physics, for example, nothing is meaningful and nothing is absurd - everything is lacking in meaning. Nagel concludes that life really can be absurd or pointless but not that it really can be meaningful. If it is true that '[a] situation is not absurd until the absurdity is perceived'⁷, but that once the absurdity has been perceived things are objectively absurd, then why is the equivalent claim concerning meaningfulness not also true?

A mouse's life cannot be absurd, but it cannot be meaningful either: self-consciousness, far from converting meaningfulness to absurdity is necessary for meaningfulness too. Absurdity and meaningfulness are on a par in this respect; neither can arise without self-consciousness. Absurdity may be more objective than meaningfulness, in the way that an ethical viewpoint is more objective than a personal viewpoint according to Nagel, but neither will remain as we acquire a more detached viewpoint. It will be sheer meaninglessness that remains: everything is pointless, since nothing matters *sub specie aeternitatis*. Nagel's

⁴ Ibid., p. 23, my italics.

⁵ Ibid., p. 16.

conclusion, that ‘a sense of the absurd is a way of perceiving our true situation’, that we should approach life with a sense of irony not despair, does not hold: the ironic approach that enables us to see life as absurd but bearable is not objective enough.

Nagel is clear that there can be no objective meaningfulness. There are two reasons for this, one is that his faulty model of objectivity blinds him to the possibility of objective phenomena not viewable from a detached point of view, but there is another valid reason, due to the recognition of the fact that one’s concerns, values and projects are not shared by all and that if one were to attempt to give a justification for them, one would have a difficult task. The realisation that it is possible to view things differently, that there are indeed people who have fundamentally different values to one’s own, and that there are viewpoints from which one’s values are invisible, is a realisation liable to shake one’s convictions at least temporarily.⁸ This issue will be dealt with in section 2.

Let us now consider the relationship between meaningfulness, absurdity and meaninglessness as given by the containment model of objectivity. We start off with two conflicting views: the view of life as meaningful, of certain things as being important, that we get as a result of involvement, and the view of life as meaningless that we get as a result of detachment. In order to find a view more objective than either of these two views of things we look to find a way of understanding the relationship between them, to find a third view which resolves

⁶ Ibid., p. 21.

⁷ Ibid., p. 23.

⁸ There are some people who seem immune to the threat posed by the existence of people with different value systems. Those people who are so single-minded and sure of their convictions as the only right ones do not seem to understand what on earth Nagel is talking about. The inability to question one’s own values seems to be a sign of intellectual immaturity, or blindness, a lack of objectivity, a failure to perceive a real problem.

the apparent conflict between them. We cannot disregard either view: we must take account of the meaningfulness our lives seem to have and also of the meaninglessness they seem to have. We are led to reject the possibility of God-given meaning that is plain to all, since such a view of meaningfulness would have to dismiss the insights we gain through detachment. But if we are not to dismiss the insights we gain through involvement, then a modest meaningfulness must be insisted upon: it must be true that life can be meaningful but that meaningfulness is not to be thought of as reliant on things being meaningful *sub specie aeternitatis*.

1.2 The relationship between meaningfulness of life and life's seeming meaningful

The meaningfulness of one's life is bound up with the way one sees one's life, the decisions one makes, etc. It might seem that meaningfulness is lacking in objectivity, that there is no epistemic gap between life's seeming meaningful to one and its being meaningful, that one cannot be wrong about the meaningfulness of one's life, that truths about meaningfulness are 'up to us', that one makes life meaningful by believing it to be, by refusing to give in to meaninglessness. I will argue that this is mistaken, that truths about the meaningfulness of life are perfectly objective.

Whether one's life is meaningful or not depends on how one lives: it is not the case that human life is essentially meaningful, that *any* human life would be meaningful. There are all sorts of things that can happen to one or that one can decide to do that have bearing on whether one's life is meaningful or not. One thing that seems essential for meaningfulness is understanding: in order for one's life to be meaningful, one must have some kind of understanding of it. The

clearest examples of meaninglessness are examples of people living without understanding why, without seeing a reason for any of the things that happen to them or for doing any one thing rather than another. We see human beings, unlike other animals, as capable of living meaningful lives because humans have the ability to make sense of their lives. Meaningfulness of one's life requires understanding and one's life seeming meaningful to one. Understanding and awareness of meaningfulness are partly constitutive of meaningfulness, then.

Someone could disagree with this saying that it's possible for one's life to seem meaningful and important to everyone but oneself, and if so, surely one's life is meaningful despite its seeming meaningless to one? If one sees one's life as futile and meaningless, but the people around one see great value in one's life, is one's life meaningless, or is one blind to its meaningfulness? I want to say that unless one sees one's life as meaningful, then one's life is meaningless. But what justifies this privileging of one's viewpoint over the viewpoints of others? The fact that one must understand and value one's life in order for it to be meaningful. I am not denying that one's life can have meaning for others in this situation: one's life can give someone else's life meaning in the same way as other things in the world, some of them inanimate, can give someone's life meaning. This can be accepted and it still be true that one's life is not meaningful *qua* life. (Cf. someone's finding meaning in the stars: it is not that the stars *themselves* have meaningful existences but that their existence, beauty, patterns, etc., can inspire.)

Suppose that what would have led one to see one's life as meaningful is that one's work be appreciated, and suppose that one dies unrecognised, seeing one's life as meaningless. Then suppose that one's work is appreciated

posthumously. Is one's life still meaningless? I want to say 'yes', that without self-understanding life cannot be meaningful. In situations like these, rather than saying that someone's life acquires meaning, we should say that their life could have become meaningful. I have no argument for viewing things in this way; my intuition is that self-understanding and seeing one's life as meaningful are necessary for meaningfulness, that the difference between human beings and animals incapable of living meaningful lives is in our ability see significance, meaning and coherence.

Meaningfulness of life implies life's seeming meaningful: does the opposite implication also hold, i.e. does life's seeming meaningful imply that it is meaningful? The answer is clearly 'no': as well as there being things which had they happened *would have* affected the meaningfulness of one's life, there are other things which *do* affect the meaningfulness of one's life without one knowing about them. If there are things which exist unbeknownst to one which undermine the meaningfulness of one's life, then one's life is meaningless even though it seems meaningful to one. If such things exist but one does not in fact find them, then the meaningfulness of one's life is still threatened: one is living a lie. (I am not supposing that anyone would necessarily be in a position to know whether such things exist, but am taking a realist line about undiscovered facts.) There is an asymmetry between meaningfulness and meaninglessness, then: life's being meaningful implies life's seeming meaningful, but life's being meaningless doesn't imply life's seeming meaningless.

If this is right, then whether or not one's life is meaningful is an objective matter: there is an objective truth about the meaningfulness of one's life. Now if

someone wanted to answer the question of whether their life was meaningful, how should they go about it? Suppose they thought that what they should do was to find out the objective truth. Suppose that they had access to the Book of Truths, which contained all truths, past, present and future, and that they could look up the truth about the meaningfulness of their life in it. Suppose that they found out that their life was meaningful and then thought ‘oh that’s alright then, I needn’t worry about doing any of the things that seem important, or about working on my relationship, or in fact about anything at all. Since the meaningfulness of my life is secure, none of this matters.’ What would be wrong with taking this attitude?

The question ‘is my life meaningful?’ is like the question ‘will I ϕ at t_2 ?’. Imagine someone who on realising that she didn’t yet have enough evidence to know whether or not she would ϕ at that time, decided to sit tight and watch herself in order to find out. Then when the time came, she could look at her watch and say to herself “well now I can answer the question ‘will I ϕ at t_2 ?’ in the negative, since I am not ϕ -ing!” This is silly - we want to say that if a person spends all her time observing herself in this way, she will not get a *truer* account of her actions, but will be interfering with them. Similarly with the question of whether one’s life is meaningful: we want to say ‘if you sit around wondering whether your life is meaningful rather than doing meaning-conducive things and finding a way to fit them into a coherent whole, then your life jolly well won’t be meaningful!’

There are timeless objective truths about the meaningfulness of life, but what makes it true that one’s life is meaningful is what one does with one’s life, that one understands one’s life, sees it as meaningful, makes the right decisions

along the way, etc. The relationship between the truth about the meaningfulness of one's life and what one does with one's life is an example of the relationship between truths about future contingents involving our actions and what we do, and the truth is 'up to us' in the case of meaningfulness only in the way that it is in the case of other future contingents.

There is the traditional problem of future contingents: if 'X will ϕ at t_2 ' is true timelessly, whenever asserted, then it is true at t_1 (where t_1 is earlier than t_2) and so X is not free at t_1 to decide to refrain from ϕ -ing. If this argument is valid then fatalism holds.⁹ The problem is ingeniously dissolved by Gilbert Ryle¹⁰: X's ϕ -ing is not necessitated by the truth of 'X will ϕ ', even though the truth, being timelessly true, precedes the action in some sense. On the contrary, it is because X will choose to ϕ that 'X will ϕ ' is true. The fatalist worry that if 'X will ϕ ' is already true then X is not free to choose whether or not to ϕ is answered by pointing out that the connection between the truth about X and X's action is not causal, the prior truth of 'X will ϕ ' does not cause X to ϕ . Rather, the relationship is logical, truths such as this are 'shadows cast by the events'.¹¹

The conclusion one should draw concerning the meaningfulness of one's life is that even though there is an objective truth about whether or not one's life is meaningful, one's knowing it, if one could know it,¹² is no help to one in deciding how to live: consulting this truth would be like watching one's shadow in order to find out what movements one was making.

⁹ The problem is originally brought up by Aristotle in *De Interpretatione*, 9.

¹⁰ In the lecture called 'It was to Be' in *Dilemmas*.

¹¹ *Dilemmas*, p. 21.

¹² There is a line of thought that suggests that it cannot be possible for one to know truths like this: that if one could know truths about one's own future then, if fatalism is not true, one could act to avoid unpleasant consequences and so render the truths one knew false.

Most people's lives will seem to comprise of periods of meaningfulness and periods of meaninglessness. In a sense, the meaningfulness of life does come and go. This is consistent with the claim that the meaningfulness of one's life as a whole is an objective, timeless matter. The meaningfulness of one's life as a whole cannot imply that one's life always seem meaningful to one, then. I suggest that the meaningfulness of one's life be taken to imply that one's life seems meaningful to one according to the most objective view of one's life open to one, and that if one's life seems meaningful to one according to that view, and there are no unknown factors which cancel out meaning, then one's life is meaningful.

When one comes to see one's life as meaningful, does one discover that one's life is meaningful, that it was meaningful all along, or does one's life become meaningful? In coming to see meaning in one's life, one doesn't necessarily come to see that one's life was meaningful all along: a new element might have been necessary to add meaningfulness. But sometimes what happens is that one comes to see meaning that was there all along. The same with meaninglessness: there are some things one finds out which seem to show one that one's life had been meaningless all along, that one had been wrong. And there are other things which destroy the meaning that had existed in one's life: if one suffers a tragedy then one's life may become meaningless.

Regardless of whether one discovers meaning in one's life or one's life becomes meaningful, when one's life seems meaningful at all then it is one's life as a whole that seems meaningful: it is not that we come to see a part of life as being a meaningful part in a meaningless whole, but that we come to see how things fit together, the good and the bad, to make the whole of life meaningful. If

one's life comes to seem meaningful at time t , then the whole of one's life will come to seem meaningful, even if one had suffered greatly and felt that life was pointless previously. When one's life comes to seem meaningful, things slot into place, one acquires the missing pieces and comes to see one's struggles and suffering as worthwhile or at least bearable. Conversely, if one loses the things that matter to one and one's life becomes meaningless, it comes to seem to one that one's life as a whole is meaningless, that the good in one's life has come to nothing, that there is not a coherent picture at all.

If we make sense of our lives as we go along, then it's always possible for something to come along in the future and undermine the meaningfulness of our lives. The possibility of such things means that meaningfulness is never secure, that it is something we must nurture and work at, and that it is something we are fortunate to have, whenever we have it. Someone might think that the possibility of loss of meaningfulness threatens our claims to meaningfulness in the same way the possibility of discovery that one's life is not really meaningful does: they might ask 'if the meaningfulness of my life relies on chance, on my not coming into contact with things capable of destroying me, on my meeting the right people, not being thwarted in my endeavours, etc., then how can it be real?' We can answer this question by saying that meaningfulness is no different to anything else in this respect, anything with contingent rather than necessary existence. The fact that I could burn all of my books doesn't bring the actual existence of my books into question, for example. Meaningfulness is more precious and also more precarious than things like books: the worry that meaningfulness can easily be lost is very real.

1.3 Meaningfulness and optimism

When one comes to see that one's life is meaningful, even though one doesn't necessarily come to see that one's life was meaningful all along, one does come to see that the world is meaningful and has been all along, i.e. one comes to see the world as the kind of place meaningfulness can exist in. This parallels what happens when one comes to see a book as being red: it comes to seem red to one, and one realises that it was red all along. If one's life is ever truly meaningful then the world must be. And what happens when one comes to see that one's life is meaningful is one finds a way to live which enables one to see the world in this way. So the truth that the world is meaningful, if it is in fact a truth, is a straightforward example of a truth only ascertainable subjectively.

In order for one's life to be meaningful one must see it as being meaningful. And to see one's own life as meaningful one must see the world as meaningful, since to not see the world as meaningful would be to treat the meaning one experiences oneself as an illusion. It is not necessary that one see *everything* in the world as meaningful, but just that one sees the world as the sort of place meaningfulness can exist in, that meaningfulness really exists in the world. This is to be optimistic about the world in a broad sense, to see the world as containing something worthwhile, good. So one's life's being meaningful requires one to have some kind of optimism about the world. Is any optimism about the world valid though? If we must believe that the world is good in order to live meaningful lives, then what if the world isn't good? A pessimist would say that one's life can only seem meaningful if one believes a falsehood, so meaningfulness is an illusion.

I will define optimism as the belief that there is enough real good in the world to make life worth living. This is a minimal form of optimism to be contrasted with the kind of pessimism that denies the existence of any real good, claiming that all good is an illusion, or that acknowledges the existence of good which is always trumped by greater evil. The difference between someone who is unable to be happy no matter how lucky they are and someone who is unable to remain unhappy no matter how unlucky they are - between the pessimist and the optimist - is due to one possessing this attitude and the other lacking it. If one treats the meaningfulness of life as an independent matter and tries to *find out* whether one's life is meaningful, or whether the world is a good place, one will hit a brick wall: the facts about the world don't seem to decide things. To an optimist, the existence of one small grain of goodness in an otherwise evil world is enough to make life worth living, whilst to a pessimist, nothing is really good; some things might appear to be, but this is just an illusion. Our experience of the world doesn't show us which is true - a neutral appraisal of the world seems to show that there is a mix of good and bad - and so the difference between optimism and pessimism cannot be given in terms of facts about the world independent of our perceptions of and attitudes and feelings towards the world. (The same fact is expressed by 'The glass is half-full' and 'The glass is half-empty' but the two utterances result from different ways of seeing the glass. Cf. part II, section 1.4)

In general, someone who claims that ϕ 's do not exist, for any ϕ , must convince us that the fact that he has not found any ϕ 's is good evidence for there being no ϕ 's, especially if someone else claims to have found ϕ 's. If someone who had never left England claimed that there were no decent Scotsmen, we would not

be convinced by them - we would tell them to go to Scotland before making such a sweeping claim. Neither would we be convinced by a neo-Nazi who claimed that there were no decent Jews - we would conclude that such a person would not recognise a decent Jew if she saw one. I suggest that the claim made by the pessimist that there is no real meaning or goodness in the world is to be taken in the same way: we should respond by saying "It is not surprising that you have not found anything to be meaningful, since you are convinced that nothing can be meaningful."

But the pessimist can say the same kind of thing to the optimist about bad or evil - the pessimist can challenge the optimist as follows: "It is not surprising that you haven't encountered the kind of evil so great that it renders life meaningless since you have closed your eyes to the possibility of its existence. What you call the meaningfulness of life is an illusion you sustain by refusing to see that all good is underpinned by bad." Another example of this kind of exchange is the exchange between those who think altruism or non-selfish action is possible and those who think that all acts are selfish: the optimist about human motivation may point to examples of sheer kindness and the pessimist then offers an interpretation of the motivation for that kindness in terms of the agent's selfish desire to promote co-operation, or the enjoyment they themselves get out of doing things for others. The optimist can then point to heroic acts which not only don't benefit the agent but actually harm or endanger her. The pessimist can then attempt to deconstruct these acts in terms of the agents needs and desires, her need to prove something or impress others, for example, or the fact that she has been morally conditioned.

Maybe the right thing to say is that neither optimism nor pessimism is accurate - that the world is neither good nor bad but neutral. This can be meant in two different ways: first, one could hold a view according to which good and bad were thought of as projected onto the world, and so not really part of the world itself. Such a view is to be rejected, as we have seen. A second way of fleshing out the claim that the world is neutral is by saying that it is clear that there is (real) good and bad in the world, but that there is no reason at all for thinking the balance must go one way or the other. There is absolutely no reason for thinking that there is more good than bad, or that all bad ultimately leads to good, etc. If we reject the idea of a benevolent God who has set things up so as to make the world good then we lose any reason we might have had for thinking that the world is ultimately good.

This is surely right: there is nothing we could point to in order to decide things in favour of the world being good or bad. Then the question is: what could justify optimism over neutralism? Everyone, whether optimist, pessimist or neutralist, accepts that there is some good and some bad in the world, the difference is in the relationship between good and bad. It is as if there are two different final views - the optimist's overall or final view is of a good whole or a happy ending whereas the pessimist's overall view is of unresolved and meaningless suffering.

If an optimistic approach is not dictated by the facts, by the way the world is prior to our taking it, then optimism is more like a stipulation than a discovery - we stipulate or take it as axiomatic that the world is good rather than discover that it is good. Once we have made this stipulation, though, certain things follow. As

the results of part II show, the fact that the meaningfulness of one's life depends on one's taking a certain attitude to life doesn't show that there is no objective fact of the matter about whether one's life is meaningful. The non-objectivity of the optimistic approach does not lead to the non-objectivity of the results of taking that approach to life any more than the non-objectivity of the rules of chess leads to there being no objective facts within chess. If we take an optimistic view then what things are good will still be an objective matter.

1.4 Is optimism justified?

The difference between one's taking an optimistic view of the world and one's taking a pessimistic view might seem to be a question of how one's life has gone. It might be true that the world is good to some people and not to others. Someone who is lucky enough to have had a good life, to have had all their basic needs satisfied and to have been given the opportunity to fulfil their potential, etc., will naturally have reason for optimism whereas someone who has had to struggle for food and shelter and has had no help along the way, whose desires have always been thwarted, etc., will see the world as a cruel place. Whilst it is undoubtedly true that one's circumstances are influential on the development of one's attitude to life, that one is more likely to be optimistic if one's life has been good and more likely to be pessimistic if one has suffered, the question I want to address is independent of this issue. The difference between the optimist and the pessimist is not best thought of as a difference in facts about each one's life: one's attitude, whether optimism, pessimism or neutralism, is not to be thought of as relative or as a truth with wholly subjective content. The optimist claims that the world is a

good place, whereas the pessimist claims that this is false, that the world is inherently bad. Even if their claims are both based on the way the world has treated them and so are inferences from the known to the unknown, their claims are not to be thought of as relative: the optimist is not just claiming that her little bit of the world is good, but that the world as a whole is to be thought of as good, and the pessimist is claiming that there is no real good in *any* of the world.

Being happy with one's life and with the world is only a good thing if one's life and the world are good. If one's life and the world are in need of improvement then optimism or acceptance blinds one to what one ought to see. A paradoxical situation arises once one starts to see that things in the world are not what they should be. If one is oblivious to or shuts oneself off from poverty, suffering, war and murder then one may live a happy life. But that kind of happiness would be gained at considerable loss: if seeing the world as a good place requires one to ignore those things in the world which are terrible, then one's optimism is shallow and false. On the other hand, once one opens oneself up to suffering and pain, it seems one can no longer believe that the world is a good place.

The awful and inevitable-seeming conclusion is that optimism, happiness and meaning are delusions which can only be sustained if one is blind to the way things really are. I will argue against this conclusion by describing real constraints on optimism, and so limiting the optimism I want to defend, and then showing that limited optimism is justified: to be optimistic that one will be able to find good/meaning in whatever life throws at one is important.

One defence of optimism might be to argue pragmatically, to say that an optimistic approach is more likely to lead us to find good than a pessimistic approach: applying a principle such as ‘always look for good in things’ is the best way to find good in the world. This kind of pragmatism is often seen as supplementing a neutralist view, as giving us a way of coming to find good in the world without having to be committed to believing optimism to be true. But if it were not the case that there was good to be found in all or most things, then such a pragmatic principle *couldn't* be retained: if in actual fact good was only to be found in a small number of things, one would soon have to reject such a principle, unless by some coincidence one happened to come across just those things. In order for this kind of pragmatic optimism to be a useful outlook, it must be true, or at least be a close approximation to the truth, and so not be merely pragmatic. So in order for us to use a principle of optimism like this, we must believe it to be *true*, or at least believe that if we believe it, it will be true.

There are two kinds of optimism: naïve optimism and deep optimism. Naïve optimism is the way of the privileged and sheltered; it is blind to the depths of human suffering because it has never been confronted with real suffering. Naïve optimism is an immature and incomplete, i.e. non-objective, approach: a naïve optimist fails to see many things about the world, assuming that because things in her part of the world are rosy, things everywhere are rosy. This kind of optimism is rightly attacked by those who take suffering seriously. What happens to a naïve optimist when awful things happen to her; when her optimism is put to the test? Either her optimism will prevail and pull her through or it will be lost. There are two ways in which optimism can prevail, either by maturing into deep

optimism or by shielding one from the horror of the situation. Deep optimism is optimism that has seen trauma, been overwhelmed by it, and eventually resurfaced. There are ways of sealing oneself off from trauma, of skating over the surface of things, numbing oneself to the terrible things that happen. By refusing to enter into traumatic situations, either by avoiding them altogether or by avoiding engaging psychologically, one can remain a naïve optimist. If one does confront trauma one of two things will happen: either one's optimism will be lost or deep optimism will emerge.¹³

The belief that the world is completely good, that there is no bad in the world, is naïve and just plain false. No one, not even the firmest believer in an omnipotent and benevolent God, can deny the existence of suffering. Then the question is, for theists and optimists alike, what should we say about the existence of suffering? We can account for the existence of bad as a necessary consequence of the existence of good: good and bad are interdependent - if something is good then its lack or loss is often bad and so if there is any good, some bad is inevitable. But over and above the bad we can account for in this way, there are also some things much worse, things which seem not just bad but evil. Some suffering is necessary for and cancelled out by greater good and if we suffer and then rise above that suffering, or learn from it, then we can look back on our suffering as something worth going through, something which benefited us in the long run. Some suffering is impossible to bear, though.

¹³ A confirmed pessimist would say that all optimism is shallow and naïve, that the survival of any optimism just shows that the situation was not traumatic enough and that given a traumatic enough experience any optimist can be reduced to despair. I accept the force of this - I believe that given enough of the kind of pressure one is susceptible to, one's optimism can be destroyed and one can be broken. But this fails to show that optimism is false and misguided, just that it is precarious.

As well as the possibility of one's rising above suffering, there is the possibility of one's being crushed by it. Just as one may look back on suffering and view it in a positive light, one may not manage to convert one's suffering into something worthwhile at all; one may suffer so badly as to come to feel that life is not worth living, to feel that one's life is meaningless. I suggest that we characterise evil as this kind of suffering. This can be the realisation of what one has done to others, the horror one has experienced at the hands of others, the overwhelming sense that things will never get any better, the realisation of the mindlessness of everything, and so on.

If we characterise evil as that which is so bad as to make life no longer worth living, then evil is relative. Whether something is evil relative to someone depends on how they see it, on whether they are overwhelmed by it or find a way of accepting and dealing with it. This is not to say that nothing is really evil: if a thing is evil to someone, i.e. if it breaks them, then it really is evil. If one is an extreme pessimist then everything is evil to one - the world is an evil place - whereas if one is a firm optimist then almost nothing is. An argument for optimism can be made on the basis that it reduces evil and so makes life easier to live: optimism is necessary for survival in the face of suffering, in order to make sense of suffering.

But how this works is important: the pessimist can counter that optimism may be necessary for coping with great suffering, but that the way it enables one to cope is by deception and by blinding one to the way things are. Even if given the choice between seeing the way things really are and living a life of misery as a result, and seeing things in a limited or distorted way and thus living a life with

some happiness, most of us would choose the latter, this is not really a justification of optimism. This would just be acknowledging how weak we are and saying that for human beings a life of deception is the only chance of happiness. If we were stronger, optimism would not be necessary to shield us from reality. If this is the way the justification for optimism goes, then there is no real meaningfulness to life and it would be more honest to just face up to meaninglessness than to live by false optimism. But this is wrong; this is not the way the justification for optimism must go: optimism makes life easier, not by blinding one to how bad things really are but by actually making things better.

Imagine creatures who felt no pain and were virtually indestructible. A lot of the things which constitute evil for us - death, endless pain, sadism, etc. - just wouldn't be evil to such creatures. Now imagine creatures who felt unbearable pain at the slightest touch. The physical world would be an evil place to these creatures; they could have no sense of being at home in the world, since at every turn the world would assault them. I suggest that we see optimists as armoured in a way that is analogous to the way the first creatures are physically armoured and that we see pessimists as vulnerable in the same way as the second creatures. The physical natures of these two sorts of creatures makes a difference to whether the world is good or evil (relative to them). Similarly with the outlooks of the optimist and the pessimist.

The nature of the world is not independent of our attitudes and beliefs any more than it is independent of our physical makeup, not because everything is subjective, there being no objective truths about the world, but because the whole world contains our beliefs, attitudes, actions, etc., and so its nature cannot be

independent of these things any more than the flavour of a soup can be independent of any of its ingredients. There is a tendency to forget that as well as being about the world or directed towards the world, our beliefs, attitudes, actions, etc., are themselves things *in* the world; things which contribute to the nature of the world constitutively and also have causal power to affect other things in the world. A belief's content represents the world as being a certain way and so stands back from the world. But a belief also has a place in the world - a causal and constitutive one, like anything else. So to ask whether the world is good independently of our beliefs, attitudes, etc. is to ask whether part of the world is good, not to ask whether the whole world is objectively good. This misunderstanding must be cleared up if we are to understand the relationship of an attitude like optimism to the nature of the world.

The general beliefs a person has about the nature of the world will affect the world around them. If one believes that there is good in everyone then one will be more likely to find good in anyone one comes across, good that was already there. One's beliefs about the nature of the world colour the world around one and reveal aspects of the world to one. But as well as revealing the good that's there anyway, one's optimism also affects the world around one via one's actions, expressions, etc. This does not just mean that the optimist's subjective world is different from the pessimist's. The way the world appears to one governs the way one reacts, and one's actions extend outside of the subjective, effecting other people and things. Then these people and things react and have further effects, some of them on oneself, and so on. The pessimist who frowns at people as she walks past them may cause them to think she hates them and then they will frown

back. She might not realise that she is frowning and so will think that people generally frown for no reason. Frowns and hatred create more frowns and hatred; smiles and happiness create more smiles and happiness. It isn't just that the world is seen differently by the optimist and the pessimist; the world is changed by the optimist and by the pessimist. The area of the world around or affected by the optimist is actually better than the area around or affected by the pessimist.

The difference between those things which are the way they are independently of what we believe about them and those things which aren't is thought to coincide exactly with the difference between those things which are objective and those things which are merely subjective. With things like the nature of the world as a whole - whether the world is good, for example, this doesn't seem quite right. It is not up to us to just believe that the world contains real good and so make it true that the world contains real good. If the world doesn't contain real good then our belief would be false.

Believing that there is real good in the world is not necessarily believing that there has already been real good in the world. The optimist may be optimistic of the future in spite of seeing the present as bad. Beliefs this general about the world are not to be thought of as temporally located: one sees the world as a whole as being a certain way. One is justified in becoming an optimist on the basis of potential for real good as well as on the basis of actual good, then.

With ourselves and other human beings, things are even more malleable: our natures are less determinate than the natures of other things. If a glass is half-full of water (or half-empty) then it is so - it isn't completely full or completely empty - and nothing one believes about it being will make any difference. With a

human being things are not so fixed: one can change, sometimes as a result of wanting to change, and one of the causes of change is what is believed about one. One can help someone to become a better person by believing them to be good, for example. As well as one's beliefs about one's nature and the natures of others being caused by these natures, they also shape these natures. When one asks oneself 'what kind of person am I?' one is not just asking for the facts as they stand, one is also asking 'what kind of person could I be?' with the aim of becoming that kind of person.

Optimism and pessimism are views about the world as a whole, which includes the past and the future as well as the present. Whereas one can't change the past, unless fatalism is true, one can change the future. Or, more accurately, one *makes* the future. One imagines how things could be and then works towards bringing about the best possibility. A pessimist is unable to imagine good futures, and so is unable to work towards them. Someone who saw their current situation as miserable could still be an optimist, believing that a good future is possible.

There are three things involved in successful creation: first, a dissatisfaction with the way the world is, second, a vision of the way the world could be, and third, seeing oneself as the link between the two. If one is only able to have a utopian vision, and is unable to relate that vision to the actual world, then one only succeeds in creating a fantasy world which will make the actual world seem even less satisfactory. If one sees the actual world as the seed of a utopia, which one can bring to fruit, then one's dissatisfaction is relieved. Even whilst one is in the process of developing, the world will be seen as good, because of what one sees in it as potential, in the way an expectant mother may view the child

she carries as a precious thing, even though it is not yet precious in actuality. To be inspired is to be able to see the future in the present, to be lifted out of the actual, to see the actual not as solid and fixed. Then one can love something - one's life, for example, for what it will become and not just hate it for what it currently is. It is only by being an optimist that one can make the future a good one rather than a bad one.

2. The substance of meaningfulness

I have been considering the question of whether life can really be meaningful, whether meaningfulness is objective. Now I move to the question of what makes life meaningful. We have seen that in order for one to live a meaningful life one must be optimistic about the world and see one's life as good, capable of being meaningful. But a good life needn't be a meaningful life: it's possible that a life be good and yet seem meaningless. Consider the following scenario: someone is given a choice between two or more different ways of life. Let us suppose that they are guaranteed happiness whichever life they choose and that there is nothing inherently better about any one of the possible lives. How are they to make their choice? If their choice is made arbitrarily then is the life they choose a meaningful life? It seems that the fact that a life is a good life to live is not enough: we want to be able to say that the life a person lives is the *right* life for them to live. I will argue that this is right and will aim to show that we can make sense of the idea of rightness of a life without referring to the idea of God-given purpose; that we can free the ideas of rightness and destiny from the theistic baggage usually associated with them.

2.1 Arbitrariness of values, concerns and projects

From a detached point of view one's life seems meaningless. If one engages with the world, involves oneself in projects and relationships, then one's life can become meaningful. From a detached point of view all possible projects and relationships are on a par; none of them seem better, or more worthwhile, than others. So from that point of view one's choices are arbitrary: there's no reason for one to choose one way of life rather than another. If we reject the priority of a detached point of view, if we reject the idea of there being a privileged viewpoint which presents the ultimate truth of the matter, then do we lose the chance of finding a concrete answer to the question 'how should we live??' Granted we must find something to live for, surely not just any old thing will do. As David Wiggins puts it, 'It is a terrible thing to try to live a life without believing in *anything*. But surely that doesn't mean that just any set of concerns and beliefs will do ... Surely, if any old set would do, that is the same as life's being meaningless?'¹⁴

The worry Wiggins expresses can be converted into an argument to the conclusion that our values, concerns and projects, and so our lives, are meaningless. We start from the premise that our values, concerns and projects are arbitrary or are chosen arbitrarily, and combined with the further premise that things which are arbitrary or chosen arbitrarily are meaningless, we are led to conclude that our values, projects and concerns are meaningless. Then, the further step to the conclusion that our lives are meaningless seems straightforward. This argument is not unproblematic: it is certainly not obvious either that our values,

¹⁴ 'Truth, Invention and the Meaning of Life' in *Proceedings of the British Academy*, LXII (1976), p. 333.

concerns and projects are arbitrary or chosen arbitrarily, or that if this were the case, they would thus be rendered meaningless. I will argue against both premises: I will show that the first premise is false, since at least some of our values, concerns and projects are non-arbitrary. Further, I will argue that a thing's being chosen arbitrarily needn't make it meaningless.

The realisation that there is no ultimate, universal or God-given grounding for our lives, our values, concerns and projects, is what generates this form of scepticism about the meaningfulness of life.¹⁵ The realisation that what is important and meaningful to us may seem arbitrary to other people, to God, to the universe, can seem like a realisation of the truth of the meaninglessness of the world and everything in it. I have acknowledged the validity of this kind of realisation: it is an important truth about life that one's life seems meaningless from certain viewpoints. But in denying that these viewpoints are more objective than those which present life as meaningful, that they have the last say on the matter, we are able to conceive of a way out of despair.

David Cooper, describing the view of Jean-Paul Sartre, says 'To exist at all, I must be engaged, for example through my having values, without which nothing could appear more worth doing than anything else ... Yet these values, which justify my particular everyday choices, are themselves the outcome of a 'fundamental choice' for which no justification is available.'¹⁶ For an existentialist like Sartre, we start from nothing and so anything we do is free in the sense of lacking in justification.

¹⁵ Cf. Sartre: 'Everything is indeed permitted if God does not exist, and man is in consequence forlorn, for he cannot find anything to depend upon either within or outside himself.' (*Existentialism and Humanism*, pp. 33-34).

¹⁶ *Existentialism*, p. 143.

The most radical existentialists and postmodernists, whilst advocating this kind of complete freedom or foundationlessness will still appeal to something: Rosi Braidotti, a postmodern feminist, argues for a ‘new nomadism’¹⁷ in contrast with the traditional view of the subject as an unchanging unity. Sartre himself puts forward the claim ‘Man is nothing else but that which he makes of himself’ as ‘the first principle of existentialism.’¹⁸ From the discovery that there are no fixed, universal, absolute foundations it is natural to move to the conclusion that there are no foundations, that everything is up for grabs. But this is an impossible way to live: there must be something, some starting point. If not then everything is just random movement. Whether the starting point is the belief that openness to change is good - Braidotti’s starting point - or the belief that human beings are completely free - Sartre’s starting point - there must be something. And this starting point is non-arbitrary.

We must distinguish between arbitrary things and arbitrarily chosen things. One is born with certain characteristics, traits, etc., into a certain environment and these things are arbitrary - there is no *reason* for one to be born with certain characteristics, traits, etc., rather than others (although there are genetic and other *causes*). But it is not true to say that *one* received those things arbitrarily, since they are what make one who one is: if one hadn’t have had those characteristics one wouldn’t have existed; a different person would have existed instead of one.¹⁹ The question ‘why was I born with *this* body and in *this* time and place?’ is

¹⁷ *Patterns of Dissonance*, pp. 277 ff.

¹⁸ *Existentialism and Humanism*, p. 28.

¹⁹ I am assuming that persons are physical and so am failing to consider views such as the view of persons as souls who receive bodies to be incarnated into. According to such a view there would be reasons why one was born with the body one was born with, etc. But they would be *God’s* reasons, not ours. I will have more to say about the idea of God’s reasons supplying our lives with meaningfulness in section 2.3.

meaningless: one is who one is because one was born with this body in this time and place.²⁰ Who one is is an arbitrary matter, then.

Just as the physical characteristics one is born with can affect one's suitability to a particular life, a life of athletics, for example, other characteristics may make certain choices, values and projects suitable for one. If I am born with a set of characteristics, *X*, and that set makes me perfectly suited to a certain way of life, then there's nothing arbitrary about my choice to live that life. It is wrong to think of a person as a featureless being who simply exists and must make arbitrary initial choices before they can be said to have any nature. Sartre is just wrong when he says 'to begin with [man] is nothing. He will not be anything until later, and then he will be what he makes of himself.'²¹ We all have our own starting points, and even if it is true that these starting points are arbitrary in some general sense, each one is necessary with respect to its owner. The fact that our starting points, our axioms, are all different is at the root of the worry: the thought is that unless there is one right starting point or set of values then none are really right at all.

Someone who thought that unless something was universal, fixed and absolute it would be arbitrary would be blind to the view I am putting forward; the view that there are non-arbitrary but individual starting points. Clearly what makes life meaningful for one person need not make life meaningful for another person. One cannot decide how to live one's life by appealing to human nature in general: if one is trying to decide whether to become a philosopher, a pharmacist

²⁰ Or, more accurately, this body, time and place gave birth to one.

²¹ Ibid., p. 28.

or a postman, one cannot appeal to anything general, since all three professions are acceptable in general. One must appeal to the specifics of one's situation.

We can put the non-universality of what makes life meaningful to one side, then. Is it true that for any individual there may be radically different options? We often have choices to make: we choose one thing over another; we choose one set of choices over another, and so on. Some would say that anything is possible, that the only limits to what a person can choose to do or be are those set by a limited imagination; that the ultimate in human development is the realisation that nothing is fixed, that we control and create our own realities; that nothing is true and everything is permitted. We become creators, re-creating ourselves in whatever image we desire: just as we are able to alter our appearances with surgery, we are able to alter our personalities with therapy. There is no reason to settle for second best, for what nature has given us.

Decisions, decisions! But how to decide? If nothing is fixed, then why should one decide to do one thing rather than another, to be one way rather than another? "It is your choice - you are completely free: there is nothing either in the world or in you that makes one choice better than another." But this would not be freedom, it would be pointlessness. This complete openness is paralysing and we stand, like Buridan's ass, unable to see any reason for deciding one thing rather than another. To make a decision one must think there are real differences between the available options.

Or one must think that one's choices don't matter, that any way one chooses, any form of involvement, is meaningful. But this answer is no good: it just isn't true that any form of involvement is meaningful and the only reason one

could have for thinking it was true is that one thinks that either it must be true that different forms of involvement can be assessed from a detached point of view or it must be true that all are on a par. There is another option: we are able to assess forms of involvement from within; in involving ourselves in one way of life we don't necessarily accept wholesale the values and outlook of that form of involvement; as things develop we assess whether our involvement has led us to achieve what we hoped to achieve, etc. There are things we can appeal to.

What can we appeal to? We have rejected God's guidance, and being individualists we reject the dictates of society too. All that remains is one's own feelings and intuitions. But one's feelings often prove themselves untrustworthy. It seems that there is nothing left. Unless there is a foolproof guide to what is right, to our destiny, we will be forced to make arbitrary choices, and this will leave life meaningless.

I will argue that this is wrong in two ways: I will argue that even if we can't locate one source which is always to be relied upon, an ultimate and foolproof guide, this is not to say that there are not different things we can rely on in different situations: it might be that in some cases our feelings are to be trusted, whereas in others a judgement that a particular feeling is untrustworthy is to be trusted, and so on. This will be the subject of section 2.3. Also, I will argue that even if some of our choices are made without complete certainty, they may still lead to meaningfulness.

If one makes a decision to commit oneself to one way of life, to a relationship with one person, for example, realising that there are other options, then is one's way of life, one's relationship, meaningless? Not necessarily:

although there may be no decisive reason for choosing to commit oneself to this one person rather than another, the fact that one's choice is to a certain extent arbitrary does not threaten the meaningfulness of the life one builds on the basis of having made that choice. By making arbitrary leaps into the dark we sometimes manage to hit upon something meaningful.

There being more than one possible way of life does not remove objectivity completely: the decisions one has to make in order to live are similar to the decisions a mathematician makes when extending mathematics beyond what is already determined, when mathematicians stipulated an imaginary number - i - to be the square root of -1, for example. This was not a discovery - mathematicians didn't discover that the square root of -1 was i . But from this non-objective stipulation all kinds of things developed. Extensions in mathematics are judged on the basis of how fruitful they become - if a possible extension quickly leads to a dead-end then it is rejected. We should view the decisions we make in life in this way: where there seems to be nothing we can appeal to in order to decide the matter for us, where we have to make an unjustified leap, we do so rightly without thereby losing the possibility of objectivity. The decisions we make bring about one future and this future, although brought about by us, is objective.

If one had to make a new decision at every point, then objectivity would be threatened. Imagine that the world never presented things as being one way rather than another, that one always had to decide which way to take things. One's relationship with the world would be like having a conversation with someone who never said anything but just repeated one's words back to one, or someone who agreed with whatever one said. The things in the world that we interact with

are all unpredictable in that they do things independently of what we do and believe, but they are also predictable enough to be interpretable. We live in the space between randomness/chaos and control: we like to be surprised by things, but not completely surprised.

The kind of involvement that makes life meaningful is involvement with objective reality. This explains why the thought that one must make a decision at every point would be threatening. If one had to make a decision at every point then this would show that what one was involved with did not have objective existence. We sometimes have to make choices where there does not seem to be one obvious right choice and when we make a decision to go one way rather than another without there being a right choice for us to discover, we truly make the decision. What we do is make the future; we close off all but one possible future. There may be no objective fact about which future is best - in effect we decide which future to call the best in our endorsing it. Once we have done so, there is room for objectivity: we bring about a new dimension to reality which contains things independent of our having made the decision we did.

If one only finds out what works by engaging with it, then life is risky. If, like in mathematics, there are some things which become meaningful ways of life for one and other things which lead to dead-ends, then one's decisions are important. Sometimes things are made easy for one: sometimes one is just plunged headfirst into a form of involvement, other times one has a strong hunch that something is going to be worth involving oneself in. Often decisions are difficult - one must choose between two or more ways of life which seem equally likely to be enriching. (Or one seems forced to choose between options none of

which seem satisfactory.) If one believes in destiny then one believes that there are some right decisions, decisions one must make if one's life is to go well, if one is to live the life one is destined to live.

Once one engages with the world, certain things which seemed indistinguishable from a detached vantage point will become distinct, certain options will present themselves as better than others, like the way blurred shapes in the distance become clear and distinct as one draws closer. It might not be possible to predict what the right decision would be prior to one's reaching the right time to make it, it might be impossible to tell that the person one is talking to is the right person for one to spend the rest of one's life with, although as time goes by and one get to know them, this becomes obvious, for example. There are constraints on the decisions one makes, but these constraints only emerge as one travels through life.

There is a tendency to desire things to fall into place, to be able to make decisions early on in life that will then hold for the rest of one's life. What I call the 'happy ever after' syndrome is the desire for an ultimate solution, the hope that once one has found one's true love, or one's vocation, or has achieved one's goals, etc., that meaningfulness will then be guaranteed for evermore. This kind of desire is self-defeating, though: if one can only be said to be living a meaningful life if one is involved, then a 'happy ever after' ending which would allow one to make one's choices once and for all, to decide on a way of life and to then be able to just live without having to be perceptive and responsive to what happens, is an impossibility. Retaining involvement and making choices at every stage, as one goes along, is an essential part of living a meaningful life. Meaningfulness

without involvement, given to one once and for all by God on high, is impossible. The things which are meaningful, which are worth involving oneself with, are objective, and so not completely given by what we know of them, and to sustain involvement with such things requires continued attention. The 'happy ever after' endings which one sometimes yearns for are impossibilities - if one were to capture the world completely once and for all then one would yearn for a new unknown. The things one values and which make life meaningfulness are those things one imagine will never be completely known, and so are worthy of one's continued devotion.

The 'happy ever after' endings are endings which require that the things one yearns for can be pinned down by one, but if this were to happen then they would cease to be objective. If the things one values and seeks involvement with are those things which are objective - other people, the natural world, God, etc. - then such yearnings are destructive; one desires something only because it is out of one's grasp, and then one grasps it and one's desire disappears. In contrast, the things one truly desires for themselves will remain desirable as one becomes closer to them.

2.2 Meaningfulness and unity of life

A life filled with good or meaningful things, or made up of meaningful elements, need not be meaningful. If one were to spend a period of time engrossed meaningfully in one activity and then emerge and find another completely different absorbing activity, then another, and so on, one might on reflection see

one's life as meaningless, there being no connection between its parts, no direction.

Suppose one wants to do many things, maybe there are even different things one wants to be. One might have a strong desire to be a philosopher but also have just as strong a desire to be a musician, for example. If one's needs and desires are not one-dimensional, if one is pulled in different directions, then one's life will not have direction and unity, and when looked at as a whole will seem meaningless. This might not seem important: it might be thought that just as it is unimportant that the meaningfulness of a life cannot be discerned *sub specie aeternitatis*, it is unimportant that one's activities and projects seem meaningful when one disengages from them: all that matters is that one gets satisfaction out of doing them whilst one is doing them. But if one loses the sense of meaningfulness something had when one finishes doing it, then one's projects will all end up seeming meaningless unless uncompleted, and one's life when viewed as a whole will also seem meaningless.

After one has completed a project that had previously absorbed one, one is likely to be left feeling empty, the focus of one's life having been removed. A common phenomenon is the frantic filling of gaps in one's life as one finds them: a relationship breaks down and the gap left by the departure of one's partner is hurriedly filled with a new person or project; one loses one's job and finds the prospect of not having anything to do with one's days more frightening than the prospect of poverty, and so on. This constant need for activity and busyness doesn't seem conducive to meaningfulness, rather it seems to signal meaninglessness: one desperately fills up of one's time in order to avoid standing

still and seeing how empty one's life really is. If the meaningfulness of one's relationships, projects, etc., remains with one after they come to an end, though, one will not be thrown into this kind of panic.

Is the meaningfulness of a project a question of successfully completing it or of finding satisfaction in the process of doing so? Earlier I suggested that part of the satisfaction one gets from creating is due to one's ability to grasp the end as a possibility before it becomes actualised. So whilst one is in the midst of creation one gains whatever can be gained from both the activity itself and the end one is working towards. But when one finishes, achieves one's end, what is left? If the product of one's activity is taken away from one - a child grown up, a painting sold, a manuscript sent to the publishers, etc. - one seems to be left with nothing.

But if one can grasp the future whilst involved, whilst creating, if one's experiences are not all best thought of as located in an extensionless present, then one should have the ability to hang onto what is no longer present just as vividly. When one is in love, one's beloved appears in one's thoughts and dreams, one re-experiences times spent together over and over; the distance in time and space is collapsed in one's mind. This can be as important as time spent together. The ability to think of someone who is absent and feel happy as a result of doing so is an example of the ability to value something one is removed from. Similarly once something has come to an end.

It might seem that the problem of conflicting aims and directions is only a problem because we are mortal, that if we were immortal and not restricted to a relatively short lifetime, there wouldn't be a problem, we would be able to do all the things we wanted. But this is wrong: the problem is not just our wanting to do

different things but our wanting to *be* different things. An eternal life divided into fifty-year chunks, each chunk dedicated to one project, does not seem meaningful. The meaningfulness of life seems to depend on our being able to answer the question ‘Who are you?’: it is not just a battle for time, but a battle for supremacy - one’s projects battle it out for the honour of being one’s defining label, the others being relegated to the role of leisure activities, hobbies.

The problem is one of how to see one’s life as a meaningful whole, how to reconcile the parts of a life. If one devotes one’s life to one clear project, one escapes the problem. If one has one overriding inclination, talent, and view of one’s life, one will not find oneself having to decide between different projects, and neither will one find oneself having to choose between ways of seeing one’s life: one’s overriding project will define one - ‘philosopher’, ‘poet’, etc. If one has many diverse abilities and aims it might seem impossible that one will be able to find a way of life which harmonises them and brings them all together. There are two acceptable ways for things to turn out in this case: either some of one’s aims and desires will dissolve as one’s life unfolds and a coherent set will remain, or a new and imaginative way of harmonising the most unlikely combination of desires and aims will present itself to one. What if neither of these happens? Must one just live in the hope that one of them does? Or must one decide to leave behind what seem to be central and important parts of oneself?

2.3 Meaningfulness, destiny and the right way of life

The death of God has been taken to mean the death of morality, meaningfulness and purpose. Dostoevsky’s slogan ‘If God did not exist, everything would be

permitted.’ is accepted by theists and existentialists alike. I will argue that this is wrong, that meaningfulness, destiny and morality should not be seen as relics of religious or superstitious ways of life any more than the objects of science should be seen as leftovers from more primitive views of the world. The opposite is true: the earlier religious or superstitious ‘theories’ or ‘metaphysics’ were attempts to understand and explain the source of meaningfulness. I will argue that the idea of God as creator fails to capture and explain what is important about meaningfulness and destiny and so the rejection of God-given answers is not a rejection of the possibility of meaningfulness or destiny. This view is in direct contrast with Sartre’s view that without God there can be no values and Alasdair MacIntyre’s view that ‘the deontological character of moral judgements is the ghost of conceptions of divine law which are quite alien to the metaphysics of modernity’.²²

My main concern is with meaningfulness, with what makes a way of life the right way for someone to live, not with strictly moral questions of right and wrong. I could concede the point that deontological questions do not make sense without a God whom it is our duty to obey, without losing anything important, so I will just put this question aside. I will argue that although God is appealed to in order to justify our actions, that such appeal is unsatisfactory; that neither God, nor a conception of human beings as rational agents, nor any inflexible code, can help us decide how to live. This is for two reasons, one being that God’s law is too abstract and detached to help us in personal matters, and the other being that, as Sartre rightly points out, even in following God one is still making a choice to do so, a choice that God cannot help us to make.

²² *After Virtue*, p. 111.

Sartre tells the story of a pupil of his who was having difficulty deciding whether to stay at home with his beloved mother or to join the Free French Forces and fight the Germans. Sartre asks ‘what could help him choose?’²³ and argues that the Christian doctrine could not, since it preaches love and respect both for one’s parents and for one’s countrymen. Similarly, the Kantian demand that we treat human beings as ends rather than means is of no help, since ‘if I remain with my mother, I shall be regarding her as an end and not as a means: but by the same token I am in danger of treating as means those who are fighting on my behalf’.²⁴

What this should tell us is that even if God does exist, there are important decisions we must make without His guidance. Moving out of the sphere of morality to one’s decisions concerning how to live a meaningful life, we ask: can God help one decide what is the right way of life? The answer is: surely not. God cannot help one decide between two acceptable ways of life: if one wants to know whether to devote one’s life to having a family or to painting, then even the dictum ‘use the talents God has given you’ leaves one on one’s own in working out which of one’s talents are the important ones, what is the best use one can put them to, etc.

The view that value must come from an infallible authority is mistaken: it separates us from value, making us reliant on God in order to access value, and requiring us to value things only as long as God values them. If something is good because God approves of it, then we are not required to value it for itself (we may even be incapable of doing so). Consider the following analogy: a child who has been told to do something, to clean her bedroom, for example, may be inclined to

²³ Ibid., p. 36.

²⁴ Ibid., p. 36.

work only whilst supervised and then fall back into playing with her toys when her parents are no longer watching, keeping an eye open for them and quickly picking up a duster again when one of them comes to see how she's getting on. In this kind of situation, the watchful parent is necessary to ensure the job gets done. But this is precisely because the child herself sees no value in doing the job. If she had been asked instead to devise a new game involving her favourite toys then she might have taken this aim, made it her own, and worked on it excitedly without the need for supervision. It is only when we don't see a value as real, when we fail to identify with it ourselves, that we need outside, e.g. divine, enforcement. God cannot supply us with real values, then, he can only be used to force us to go through the motions of valuing what we don't in fact value.²⁵

If God cannot be the source of meaningfulness, then what on earth can be? I do not hope to answer this question, but will aim to present a brief account of the phenomenon in need of explanation. This will serve as a challenge to those such as Sartre who would claim that the phenomenon is not real.

When we say someone is destined to become a great pianist, what do we mean? We don't just mean that playing the piano would be a good thing for them to do. But on the other hand, we don't mean that there is no way they can avoid becoming a great pianist. We mean something stronger than that a certain way of life would be good, but we mean something weaker than out-and-out fatalism, that that way of life is literally the only option. We mean that that way of life is right for them. But how are we to understand this kind of rightness? Some couples seem made for each other and others seem a total mismatch. What do we mean

²⁵ This is not to say that God cannot be an important part of one's life: if one loves God, then appeal to God will be fruitful. But this is a different situation: here one has already made the

when we say that two people are ‘soul mates’ or ‘made for each other’? We do not mean that someone or something actually made them for each other, but neither do we just mean that they get on well and that their relationship is successful. There is a difference in kind between relationships which seem (merely) workable and those which seem ‘right’.

Sometimes one dreams of doing certain things, or of certain things happening to one: one dreams of sitting at a beautiful piano, of concert performances, for example. If one’s dreams are vivid, detailed and more sustained than childish daydreams and if they pull one forwards, give one a definite end to work towards, then shouldn’t we describe them as premonitions, as indications of one’s destiny? They are insights into the future which enable one to bring about that future. Whatever the relationship between dreams like these and one’s future, and we can be sure enough that it’s not mysterious, the match between them is central to the notion of destiny. It is either that someone who has such vivid and overpowering dreams will be more inclined to work hard towards realising them, or it is that the clarity of the dreams makes it easier for one to know what one is working towards. The likelihood of one’s actualising one’s dreams is directly proportional to the vividness, reality and force of one’s dreams: if one can ‘see’ a possible future completely, then one is not working towards something vague and intangible, trying to find out what to do as one goes along. On the contrary, the future is so real in one’s imagination that it is as if one has already succeeded; as if one has already been through the struggle and the battle has been won.²⁶ This

decision to value God. The point is that God cannot order us to love Him, but if we do in fact love Him, then He can give us guidance.

²⁶ Cf. Beethoven’s description of composing: ‘I see and hear the picture in all its extent and dimensions stand before my mind like a cast, and there remains for me nothing but the labour of

means one is in a position to guide oneself through it in reality, working with the certainty of someone travelling a route for the second time.

In these cases we cannot distinguish between something's happening because one strongly wanted it to happen and one's clearly seeing that it had to happen. Seeing a future and wanting that future come together in the way they do in smaller ways in everyday life, one's wanting to eat the apple on the table and one's knowledge that that is what one will do, for example. The only limitation seems to be a limit of the imagination: once one can imagine a way, there's nothing to stop one. (That's not to say that inspiration comes easily or that it happens as often as one would like.)

There are times when things seem to click, to slot into place. These times are to be contrasted with the times when one is confused about what to do. Sometimes one feels things to be clearly right, sometimes one feels that one has found one's path, and the confusion just drops away and one's options narrow down. When one makes an important decision, makes a choice to do one thing rather than another, even though one was not forced into making that decision, one may still say "I couldn't have done otherwise". This means something like "I would've been a different person had I done otherwise" or "I wouldn't have been able to live with myself had I done otherwise". This is partly relief at having done what one considers is the right thing, but it is also partly the realisation of having been constrained. Whenever one makes the right choice, one says "In the same situation I would do the same thing again" and this is the same as the description

writing it down' (in *A Dictionary of Musical Quotations*, compiled by Ian Crofton and Donald Fraser).

we give of a deterministic system - given the same initial conditions, the same things will follow.

One's destiny is not just based on one's personal characteristics - the individual should not be thought of as independent of the environment. The miracle of things coming together, of being in the right place, meeting the right person, etc., only seems miraculous and requiring divine help if we see the individual as separate from the rest of the world. If we don't make such a separation then the coming together of two compatible creatures seems no more miraculous than the way the shape of a footprint matches the shape of the foot that made it. In living the right life one comes to see oneself as the right shape for *this* space in the picture. One's ideas and desires don't spring from nowhere, they are caused by all sorts of things around one, and so one is the way one is because of the way things around one are, and vice versa.

The possibility of inspiration, of finding direction and purpose is an amazing one. That meaning can arise in the world seems miraculous. Maybe once we acquire an understanding of the nature of meaningfulness this sense of wonder will evaporate and we will take meaningfulness for granted. Maybe with such understanding will emerge a procedure that if followed would automatically lead to one's life's being meaningfulness. Until then, meaningfulness is something we can only strive towards in the hope that 'through grace'²⁷ we will encounter it.

²⁷ Cf. Martin Buber's account of our relationship to the *Thou* (*I and Thou*, pp. 24 ff.).

BIBLIOGRAPHY

Anscombe, G.E.M., 'The Question of Linguistic Idealism' in her *Collected Papers Volume I: From Parmenides to Wittgenstein*, Basil Blackwell, 1981.

Aristotle, *Nicomachean Ethics*, in *The Basic Works of Aristotle*, Richard McKeon ed., Random House, 1941.

— *De Interpretatione*, in *The Basic Works of Aristotle*.

Ayers, Michael R., 'Individuals without Sortals' in *Canadian Journal of Philosophy*, Volume IV, Number 1, 1974.

— 'The Ideas of Power and Substance in Locke's Philosophy' in I.C. Tipton ed. *Locke on Human Understanding*, Oxford University Press, 1977.

— 'Substance: A Prolegomena to a Realist Theory of Identity' in *Journal of Philosophy*, 1991.

Bell, David, 'The Formation of Concepts and the Structure of Thoughts' in *Philosophy and Phenomenological Research* Volume LVI, Number 3 (September 1996).

Berenson, Frances M., 'Interpreting the Emotional Content of Music' in Michael Krausz ed. *The Interpretation of Music*.

Blackburn, Simon, *Spreading the Word*, Oxford University Press, 1984.

— 'The Individual Strikes Back' in *Synthese* 58.

— 'Reply: Rule-Following and Moral Realism' in S. Holtzman and C. Leich eds., *Wittgenstein: To Follow a Rule*.

Bradley, F.H., *Appearance and Reality*, Swan Sonnenschein & Co, second edition, 1908.

— *Essays on Truth and Reality*, Clarendon Press, 1914.

Braidotti, Rosi, *Patterns of Dissonance: A Study of Women in Contemporary Philosophy*, Elizabeth Guild tr., Polity Press, 1991.

Boden, Margaret, ed., *The Philosophy of Artificial Intelligence*, Oxford University Press, 1990.

Buber, Martin, *I and Thou*, Ronald Gregor Smith tr., T.&T. Clark, 1937.

Budd, Malcolm, *Wittgenstein's Philosophy of Psychology*, Routledge and Kegan Paul, 1989.

Burnyeat, Myles F., 'Aristotle on Learning to be Good' in A. O. Rorty ed. *Essays on Aristotle's Ethics*.

Burroughs, William S. and Gysin, Brion, *The Third Mind*, The Viking Press, 1978.

Camus, Albert, *The Myth of Sisyphus*, Justin O'Brien tr., Penguin Books, 1955 [first published as *Le Mythe de Sisyphe*, 1942].

- Cartwright, Nancy, *How the Laws of Physics Lie*, Clarendon Press, 1983.
- Cooper, David E., *Existentialism: A Reconstruction*, Blackwell Publishers, 1990.
- Cussins, Adrian, 'The Conceptualist Construction of Concepts' in Margaret Boden ed., *The Philosophy of Artificial Intelligence*.
- Darwin, Charles, *The Origin of Species*, Penguin Books, 1985 [first published 1859].
- Dawkins, Richard, *The Extended Phenotype*, Oxford University Press, 1982.
— *The Blind Watchmaker*, Penguin Books, 1991 [first published 1986].
- Davidson, Donald, 'On the Very Idea of a Conceptual Scheme' in *Inquiries into Truth and Interpretation*, Oxford University Press, 1986.
— *Essays on Actions and Events*, Oxford University Press, 1986 [first published 1980].
— *Inquiries into Truth and Interpretation*, Oxford University Press, 1986 [first published 1984].
— 'A Coherence Theory of Truth and Knowledge' in Ernest LePore ed., *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*.
— 'Empirical Content' in Ernest LePore ed., *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*.
- Dennett, Daniel C., *Elbow Room: The Varieties of Free Will Worth Wanting*, Clarendon Press, 1984.
— *The Intentional Stance*, MIT Press, 1987.
— *Consciousness Explained*, Allen Lane: The Penguin Press, 1991.
— *Darwin's Dangerous Idea*, The Penguin Press, 1995.
- Degler, Carl N., *In Search of Human Nature: The Decline and Revival of Darwinism in American Social Thought*, Oxford University Press, 1991.
- Descartes, René, *Philosophical Writings*, G.E.M. Anscombe and P.T. Geach tr., Nelson's University Paperbacks, 1970.
- Dworkin, Ronald, 'How Law is Like Literature', in *Critical Inquiry*, September 1982 [reprinted in his *A Matter of Principle*, Clarendon Press, 1985].
- Ekman, Paul and Davidson, Richard J., eds., *The Nature of Emotion*, Oxford University Press, 1994.
- Evans, Gareth, *The Varieties of Reference*, Clarendon Press, 1982.
- Field, Hartry, 'Realism and Relativism' in *Journal of Philosophy*, 1982.
- Frege, Gottlob, *The Foundations of Arithmetic*, J.L. Austin tr., Basil Blackwell, 1950.

- ‘On Sense and Meaning’ in *Translations from the Philosophical Writings of Gottlob Frege*, Peter Geach and Max Black eds., Basil Blackwell, 1980 [first published 1952].
- ‘The Thought: A Logical Inquiry’, A.M. and Marcelle Quinton trs., in *Philosophical Logic*, P. F. Strawson ed., Oxford University Press, 1967.
- Freud, Sigmund, *Introductory Lectures on Psycho-Analysis*, George Allen & Unwin Ltd., 1922.
- Goleman, Daniel, *Emotional Intelligence*, Bloomsbury, 1996.
- Goodman, Nelson, *Fact, Fiction and Forecast*, University of London Press, 1955.
— *Ways of Worldmaking*, The Harvester Press, 1978.
- Hacker, P.M.S., *Insight and Illusion*, Oxford University Press, 1986.
- Herrigel, Eugen, *Zen in the Art of Archery*, Arkana, 1987 [first published 1953].
- Hesse, Hermann, *Narziss and Goldmund*, Geoffrey Dunlop tr., Penguin Books, 1971 [first published 1959].
- Holtzman, Steven and Leich, Christopher, eds., *Wittgenstein: To Follow a Rule*, Routledge and Kegan Paul, 1981.
- Hookway, Christopher, *Quine*, Polity Press, 1988.
- Hume, David, *A Treatise of Human Nature*, L. A. Selby-Bigge ed., Oxford University Press, 1958.
- Huxley, Aldous, *The Doors of Perception and Heaven and Hell*, Penguin Books, 1959.
- James, William, *The Varieties of Religious Experience*, Fontana, 1960.
— *Essays in Radical Empiricism*, Gloucester, Mass. 1967 [first published 1912].
- Jung, Carl Gustav, *Synchronicity: An Acausal Connecting Principle*, tr. R. F. C. Hull, Ark Paperbacks, 1955.
— *The Archetypes and the Collective Unconscious*, Routledge, 1991 [first published 1959].
- Kant, Immanuel, *Foundations of the Metaphysic of Morals*, Lewis White Beck tr., Bobbs-Merrill, 1959.
- Kaufmann, Walter, *Critique of Religion and Philosophy*, Harper & Brothers, 1958.
— *Nietzsche: Philosopher, Psychologist, Antichrist*, Princeton University Press, 1974.
- Krausz, Michael, ed., *The Interpretation of Music*, Clarendon Press, 1993.
- Kripke, Saul, *Naming and Necessity*, Basil Blackwell, 1986 [first published 1972].
— *Wittgenstein on Rules and Private Language*, Basil Blackwell, 1989.

Lakoff, George, *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*, University of Chicago Press, 1987.

Lear, Jonathan, *Aristotle: The Desire to Understand*, Cambridge University Press, 1988.

— *Love and its Place in Nature: A Philosophical Interpretation of Freudian Psychoanalysis*, The Noonday Press, 1990.

LePore, Ernest, ed., *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, Basil Blackwell, 1986.

LePore, Ernest & McLaughlin, Brian, eds., *Actions and Events: Perspectives on the Philosophy of Donald Davidson*, Basil Blackwell, 1985.

Lovibond, Sabina, 'True and False Pleasures' in *Proceedings of the Aristotelian Society*, 1989-90.

Lyons, William, *Emotion*, Cambridge University Press, 1980.

— *Approaches to Intentionality*, Clarendon Press, 1995.

McDowell, John, 'Virtue and Reason' in *The Monist*, 1979.

— 'Wittgenstein on Following a Rule' in *Synthese* 58.

— 'Non-Cognitivism and Rule-Following' in S. Holtzman and C. Leich eds., *Wittgenstein: To Follow a Rule*, Routledge and Kegan Paul, 1981.

— *Mind and World*, Harvard University Press, 1994.

McGinn, Colin, *Wittgenstein on Meaning*, Basil Blackwell, 1984.

MacIntyre, Alasdair, *After Virtue: A Study in Moral Theory*, Duckworth, 1985.

Mackie, J. L., *Ethics: Inventing Right and Wrong*, Penguin Books, 1977.

Malachowski, Alan, ed., *Reading Rorty*, Basil Blackwell, 1990.

Malcolm, Norman, *Wittgenstein: Nothing is Hidden*, Basil Blackwell 1986.

— 'Wittgenstein on Language and Rules' in *Philosophy* 64.

Mander, W. J., *An Introduction to Bradley's Metaphysics*, Clarendon Press, 1994.

Nagel, Thomas, *Mortal Questions*, Cambridge University Press, 1979.

— *The View from Nowhere*, Oxford University Press, 1986.

Nietzsche, Friedrich, 'On Truth and Lie in the Extra-Moral Sense' in W. Kaufmann ed. and tr., *The Portable Nietzsche*, Chatto and Windus, 1971.

— *Beyond Good and Evil*, R. J. Hollingdale tr., Penguin, 1973.

Nozick Robert, *Anarchy, State and Utopia*, Basil Blackwell, 1974.

— *Philosophical Explanations*, Oxford University Press, 1981.

- Peacocke, Christopher, *Sense and Content: Experience, Thought and their Relations*, Clarendon Press, 1983.
- Pears, David, *The False Prison Volume I*, Oxford University Press, 1987.
— *The False Prison Volume II*, Oxford University Press, 1988.
- Putnam, Hilary, *Mind, Language and Reality: Philosophical Papers Volume 2*, Cambridge University Press, 1975.
— *Reason, Truth and History*, Cambridge University Press, 1981.
— ‘Reply to Two Realists’ in *Journal of Philosophy*, 1982.
— *Realism and Reason: Philosophical Papers Volume 3*, Cambridge University Press, 1983.
— *Realism with a Human Face*, Harvard University Press, 1990.
- Quine, W. V. O., *From a Logical Point of View*, Harvard University Press, 1953.
— *Word and Object*, The M. I. T. Press, 1960.
- Rorty, Amélie Oksenberg, ed., *Essays on Aristotle’s Ethics*, University of California Press, 1980.
- Rorty, Richard, *Consequences of Pragmatism*, The Harvester Press, 1982.
— *Objectivity, Relativism and Truth: Philosophical Papers Volume 1*, Cambridge University Press, 1991.
- Roubiczek, Paul, *Existentialism: For and Against*, Cambridge University Press, 1964.
- Russell, Bertrand, *The Problems of Philosophy*, Oxford University Press, 1912.
- Ryle, Gilbert, *The Concept of Mind*, Penguin Books, 1949.
— *Dilemmas*, Cambridge University Press, 1954.
- Sartre, Jean-Paul, *Existentialism and Humanism*, Methuen, 1948.
- Strawson, Peter F., *Individuals: An Essay in Descriptive Metaphysics*, Methuen, 1959.
- Stroud, Barry, *The Significance of Philosophical Scepticism*, Clarendon Press, 1984.
- Varela, Francisco J., Thompson, Evan and Rosch, Eleanor, *The Embodied Mind*, The MIT Press, 1992.
- Walker, Ralph C. S., *The Coherence Theory of Truth: Realism, Anti-Realism, Idealism*, Routledge, 1989.
- Watson, Gary, ed., *Oxford Readings in Philosophy: Free Will*, Oxford University Press, 1982.
- Wiggins, David, ‘Truth, Invention, and the Meaning of Life’ in *Proceedings of the British Academy LXII*, 1976.

— ‘Moral Cognitivism, Moral Relativism and Motivating Moral Beliefs’ in *Proceedings of the Aristotelian Society*, Vol. LXXXI, 1990/91.

Williams, Bernard, *Ethics and the Limits of Philosophy*, Fontana Press, 1985.

Wittgenstein, Ludwig, *Notebooks 1914-16*, G.H. von Wright and G.E.M. Anscombe eds., G.E.M. Anscombe tr., Basil Blackwell, 1961.

— *Tractatus Logico-Philosophicus*, D.F. Pears and B.F. McGuinness trs., Routledge and Kegan Paul, 1961.

— *Philosophical Remarks*, R. Rhees ed., R. Hargreaves and R. White trs., Basil Blackwell, 1975.

— *Philosophical Grammar*, R. Rhees and A. Kenny eds., Basil Blackwell, 1974.

— *The Blue and Brown Books*, Basil Blackwell, 1958.

— *Remarks on the Foundations of Mathematics*, G. H von Wright, R. Rhees, G.E.M. Anscombe eds., G.E.M. Anscombe tr., Basil Blackwell, 1956.

— *Philosophical Investigations*, G.E.M. Anscombe tr., Basil Blackwell, 1953.

— *Zettel*, G.E.M. Anscombe and G.H. von Wright eds., G.E.M. Anscombe tr., Blackwell, 1981.

— *Remarks on the Philosophy of Psychology vol. I*, G.E.M. Anscombe and G.H. von Wright eds., G.E.M. Anscombe tr., Basil Blackwell, 1980.

— *Remarks on the Philosophy of Psychology vol. II*, G.H. von Wright and H. Nyman eds., C.G. Luckhardt and M.A.E. Aue trs., Basil Blackwell, 1980.

— *On Certainty*, G.E.M. Anscombe and G.H. von Wright eds., D. Paul and G.E.M. Anscombe trs., Basil Blackwell, 1974.

Wolheim, Richard, *F.H. Bradley*, Penguin Books, 1959.

— *The Mind and Its Depths*, Harvard University Press, 1993.

