

PHILOSOPHY A2 UNIT 3 SPECIMEN PAPER SUGGESTED ANSWERS TO SELECTED QUESTIONS

In writing the answers to past exam questions, I have referred to AQA's mark schemes (available on their website) as far as possible. However, my answers should not be treated as carrying AQA's 'stamp of approval'. They represent my considered opinion as to how these questions might reasonably be tackled. The BPA takes no responsibility for the examination performance of students who have used these answers.

There are of course a range of valid ways of tackling most of these questions and mine represent, in each case, just one possible approach.

Question 2

'The peculiarity of the soul is not that it is visible to none but that it is visible only to one.' To what extent do we enjoy privileged access to, and certain knowledge of, our own minds? (50 marks)

The starting point of Descartes' constructive philosophy is his 'cogito'. He maintained that even if there was nothing else that he could be certain of, he could at any rate be certain that he himself existed. As well as the idea of certain knowledge, this also involved the notion of privileged access, since according to this way of thinking, I can be certain that I exist, but I cannot be certain (at least initially) that you or anyone else does. Thus each of us has privileged access to his own 'self'—assuming that others actually have selves and are not mere automatons!

But what is this 'self' that Descartes believed himself to have access to? Hume complained that when he tried to perceive his self or 'I', he could only catch some particular impression or idea. It has been suggested in view of this that what Descartes should have said was not 'I think, therefore I am', but just 'There is a thinking'. A fleeting perception of a 'thinking' going on is not in itself sufficient to establish that there is a self that is doing this thinking and which survives essentially unchanged from moment to moment.

However, the question was about the extent to which we enjoy privileged access to, and certain knowledge of, our own *minds*. Thinking and similar processes are obviously mental in some sense, even if they do not take place within a substantive self. Perhaps the thesis of privileged access and certain knowledge can be maintained without reference to substantive selves, but simply in relation to the contents of our mental states. It is true that even this goes against the physicalist approach to the mind that is currently dominant. If the mind has no independent reality, then all knowledge is ultimately of physical facts and to these, it would be argued, there is no privileged, infallible access. But though currently in vogue, the truth of physicalism can be persuasively disputed, precisely because it leaves out the 'felt', subjective quality of experience.

So let us consider the claim that we possess privileged access to, and certain knowledge of, the contents of our own mental states. I shall begin with the issue of certain knowledge.

It appears at first sight very difficult to suppose that the judgements I make about what I am experiencing at this very moment could be mistaken. For surely since they are *my* experiences, and I am having them right *now*, I am ideally placed to know what they are. However, A.J. Ayer argues convincingly that doubt is indeed possible and coherent. Suppose I say, concerning one of my sense-contents: 'This is white'. Then at least part of what I am saying is that the sense-content is of a sort that not only I, but also most other people, would call 'white'. But if I had defective colour vision, then this might very well not be true, in which case I would have to regard myself as mistaken in this perceptual judgement. Furthermore, Ayer continues, the reference to other people is unnecessary. Suppose I discovered that whenever I had a sense-content that I was inclined to call white, I made a certain overt bodily movement. But then suppose on one occasion, I thought one of my sense-contents was white, but did not make the movement. I would very likely conclude that I was wrong to think that I made the movement *every* time I had a white sense-content. However, Ayer points out, it would be logically open to me to say instead that I had actually been mistaken in judging this particular sense-content to be white. Thus, Ayer suggests, nothing empirical, not even the content of our own current experiences, is beyond the possibility of coherent doubt. It might seem as if the judgement 'There is now thinking going on' must be infallible when I make it regarding my own mental state. But how can I be certain that 'thinking' really is the correct word for this sort of experience, that it represents the correct way of classifying it?

Now let us turn to the issue of privileged access to the contents of our own minds. What exactly is meant by this phrase? I interpret it to mean a particular way of acquiring beliefs about the contents of our mental states which only we, as the subjects of those states, can have. This way of acquiring beliefs is normally reliable, but, if our previous findings are correct, may not *always* result in true beliefs. It is essentially the mechanism of *attending* to our mental states. Clearly I can attend to my own mental states, but you cannot do so—at least not in the sense intended here. But is this a necessary truth? Is it inconceivable that someone could have this kind of access to someone else's mental states? Suppose we came across someone who was able to give detailed descriptions of other people's thoughts and feelings which turned out to match exactly the descriptions given by those people themselves, without there being any 'natural' explanation of how he does this. Wouldn't we be forced to conclude that this person was able to 'attend' to the mental states of others in the same way that they themselves can? To this it might be objected that we would only be entitled to conclude that this person had *some* (perhaps telepathic) way of knowing about others' experiences, not that he was able to stand in relation to them in the same way that each of us stands to our own experiences when we are said to attend to them. But the identity conditions for mental states are relevant here. It can be argued that if mental states *a* and *b* are qualitatively identical, then they must also be numerically identical. Suppose we are both having the experience of staring at a white sheet of paper, and suppose that just for an instant there is no difference at all in the qualities of our mental states. Are there two mental states involved, yours and mine, or just one? One reason for giving the former answer is that we could distinguish between yours and mine by reference to the fact that they are preceded and followed by different

mental states. But now consider the following thought experiment. You have a doppelganger on the other side of the universe whose mental states, throughout his entire life, are identical to yours. Isn't it at least arguable that there is only one mind involved here, which is correlated with the states of two different physical bodies, rather than two distinct minds? If that is so, you have just as much 'access' to your doppelganger's mind as you do to the one we normally regard as your own—for it is the same mind!

Of course, none of this alters the fact that in the normal run of things, we do have a kind of privileged access to our minds (and near-certain knowledge of our current mental states). What it does suggest is that these facts may be more contingent than we might have expected.

Question 4

To what extent, if at all, could the possession of rights be based on utility?

(50 marks)

At first sight, utility appears to be an unpromising basis for rights. It is easy to construct examples which apparently show that for utilitarians, rights are meaningless. Suppose there are three hospital patients. The first needs a new kidney, the second a new liver and the third a new heart. They can all be saved if they get these organs. Doctors can kill me and extract these organs to give to them, thereby saving three lives with the loss of only one. The balance of utility seems to favour this, but at the expense of violating my right not to be killed. If utilitarianism can't even defend such a fundamental right as this, what hope is there for extracting from it any serious contribution to rights theory?

However, the concept of a right itself should not be thought immune from problems. Legal rights are perhaps the most straightforward. For example, the legal right to the free disposal of one's property is (subject to certain constraints) enshrined in law, and the content of the law is empirically ascertainable. However, the notion of a legal right is too thin to ground the rights that are typically appealed to in morality. For such rights are often invoked when assessing and criticising legal systems themselves and therefore have to be grounded independently of them. These are so-called 'natural rights'. It is these rights that utilitarianism might be thought inadequate to handle. But exactly what natural rights do we have? The right to life would presumably be beyond dispute. But it was recently claimed that all human beings have a right to 'internet access'. Is this correct? How can we tell which natural rights people have and which they are only imagined to have by various interest groups? A theory such as utilitarianism at least holds out the promise of being able to distinguish systematically between real rights and merely imagined ones: the real ones, according to the utilitarian, are those whose recognition and protection would maximise social utility. Of course, there are other moral theories besides utilitarianism. Deontology, for example, represents a very different sort of approach from the utilitarian one, in which the notion of obligation is treated as logically prior to that of value. It has often been thought particularly well placed to explain and determine natural rights. (After all, it seems that every right a person has corresponds to a duty or obligation on the part of others.) However, the question arises of what moral *obligations* we have and here

deontology tends to rely on intuition alone. While we can expect some agreement between different people about broad, imprecisely stated, obligations (e.g. to tell the truth, not to harm others), the details tend to provoke disagreements, and the reliance on intuition is of no help here. (There is also the problem of what to do when different moral obligations conflict, e.g., when telling the truth would involve harming another.) At least utilitarianism appears to provide a way of adjudicating between conflicting intuitions, of deciding which ones are to overrule others in particular cases—by reference, that is, to the principle of utility.

But if utilitarianism is indeed to secure a viable concept of rights, we will need to have some answer to critics who cite cases like the one with which I started this essay, in which it seems that the utilitarian is obliged to overrule fundamental human rights. At first this seems impossible, for it appears to be part of the nature of rights that they ‘trump utility’ (in the phrase of Ronald Dworkin). My rights would appear to be meaningless if they can be overridden as soon as someone has established that society would be marginally better off as a result. But how can a utilitarian accept a concept of rights that has them trumping utility when his own theory construes utility as morally supreme?

The answer, as a number of utilitarian writers have pointed out, is to adopt what is called *indirect utilitarianism*, in which the principle of utility is not (normally) applied to each possible action individually, but is used instead to underpin and justify *systems* of rights and duties, and it is these rights and duties that determine which actions are right or wrong. Roughly speaking, the indirect utilitarian would say that the correct system of rights and duties for a society is the one whose adoption would maximise utility within that society. In relation to the forced transplant example, we would need to ask this: would a society in which people had effectively no right to life, in which they could at any time be killed for the sake of actions that appear to maximise utility in the short term, be likely to lead to more overall utility than one in which a much more robust right to life was recognised, one for which hardly any exceptions were permitted? When one considers the dangers involved in the former, not least the general fear and insecurity that it would engender, it is hard to dispute that the latter would be better in terms of maximising utility.

A utilitarian justification can also be found for other rights, such as the right to free speech. Society benefits from the wide availability of knowledge; indeed without accurate knowledge both of human nature and of the non-human world, it cannot hope to make significant progress in the betterment of people’s lives. But much knowledge, including some that is critical for the improvement of society, is difficult to come by. It requires subtle and demanding methods of investigation, and provokes disagreement and controversy. It also requires a constant flow of new ideas and a process of testing those ideas through rational discussion. It certainly cannot thrive in a situation where the state proclaims that it has all the necessary knowledge and that any questioning of its received wisdom is not to be tolerated. Hence (as Mill pointed out) it is essential that society tolerate the expression of the widest possible variety of opinions, including those that appear at first sight manifestly false, as there is no way of telling in advance from what directions truth may eventually emerge. Thus a right to free speech is seen as a necessary consequence of the moral imperative (as utilitarians see it) to maximise general well-being. Reasonably plausible utilitarian arguments can also be offered on behalf of other widely recognised human rights,

such as rights of association, of democratic representation and of freedom from arbitrary interference by the state.

However, defenders of natural rights are still unlikely to be convinced. They may point out that fundamental rights seem too easily overridden in a utilitarian scheme. Consider the right not to be punished by the state for a crime that one has not committed. Granted, the utilitarian can probably justify the rejection of a criminal justice system in which random individuals are habitually framed by the police. For if such a system were permitted, it would undermine the very security that it is the purpose of a criminal justice system to preserve. (Police would effectively become morally indistinguishable from 'regular' criminals.) However, it is harder for the utilitarian to justify the rejection of a *one-off* proposal to frame an innocent person in a situation where the very survival of society is at stake. (Perhaps we are dealing with a small society that is being torn apart by a particular sort of violent crime; the authorities urgently need to 'set an example', and none of the real culprits can be caught.)

However, the problem with this sort of example as an argument against the ability of utilitarianism to accommodate rights is that it assumes that rights must be absolute, that they admit of no possible circumstances in which it would be morally legitimate to override them. But it is not at all clear that this is so. Even the most basic rights, such as the right to life, are considered to be overridable in some situations (as when killing an attacker is the only way to save one's own life). Since the example above involves a situation in which the very survival of the society is threatened, then it is at least plausible to suppose that morality would require the violation of an *almost* inviolable right not to be punished for a crime that one has not committed, repellent though such a violation would be. The fact that in this scheme, the justifiable punishing of the innocent would be a very rare thing makes it legitimate to talk about a *right* here, as contrasted with a situation where individuals can be treated in any way for the sake of some supposed short-term gain in social utility. From this standpoint, it seems that the normal concept of a natural right can indeed be incorporated into indirect utilitarianism.

Question 5

Assess the view that there are such things as universals.

(50 marks)

The term 'universal' refers to the general qualities, properties and relations that are instantiated in particular things. Many philosophers have been interested in knowing whether universals actually exist over and above the particular things that exemplify them. Realism about universals was developed in perhaps its most extreme form by Plato, who postulated a realm of permanent and perfect 'forms' to which the fleeting and imperfect individuals of the familiar world approximate. But it is not necessary to adhere to such an elaborate metaphysical view as this to suppose that universals have some kind of genuine reality.

So-called *nominalists* would prefer to avoid this. They may be influenced by 'Occam's Razor': 'Do not multiply entities beyond necessity'. Occam's razor urges us

to make do with as few fundamental types of entity as we can. But can we dispense with universals? Here there is a point that needs clarifying. Are we expected to show how a language could dispense with predicates altogether, or is it sufficient to devise a special way of explaining the function of predicates, one that does not commit us to the reality of independently real universals?

Certainly it is hard to imagine a language without predicates. It might be suggested that '*a* is red' should be taken to mean something like '*a* resembles *p* in colour', where *p* is some paradigm case of redness. This approach could be used not only for simple predicates, but also complex ones, even those that have no instances, such as 'winged horse'. (We have paradigm cases of being winged and paradigm cases of being a horse—we don't actually need paradigm cases of things that come into both categories to make sense, for example, of the statement '*a* is not a winged horse'.) But there are two problems with this proposal. One is that it is not immediately clear what we are to do about uninstantiated *simple* properties. Perhaps with some ingenuity this difficulty could be overcome, but there is a more fundamental one, which was pointed out by Russell. This is that resemblance-relations are themselves universals, and if we are to admit one type of universal, why not the whole lot?

But arguably the program of showing that we can have a 'predicate-less' language is not necessary anyway. Just because we make statements like '*a* is red', it does not follow, according to this view, that we are committing ourselves to the reality of *things* that 'red' and other predicates refer to. The suggestion is that it is only the names or other singular terms of the language that should be thought of as referring to things. So as long as we can avoid making statements like 'Redness is an interesting property', in which our colour-word occurs as a *noun*, we are not supposing universals to have any reality.

Can this be done? For each meaningful statement in which a property or relation word occurs in noun position, we would have to be able to give a plausible paraphrase in which it occurs only as a predicate. For example, 'Red is more similar to orange than it is to blue' can be translated as 'Red things are more similar in colour to orange things than they are to blue things'. However, easy cases like this one should not mislead us into thinking that the nominalist program, even when interpreted in the proposed way, can be completed. Many philosophers of science tell us that mathematics, with its wide range of universals (mostly numbers of some sort) is indispensable to physics and thus indispensable to a correct description of the world. The issue is not finally resolved, but until someone shows otherwise, it is reasonable to assume that universal terms cannot be excised from our descriptions of reality.

In view of this, it may be desirable to address a more fundamental question: even if we are committed to using universal terms in noun position in order to say all the things we need to say about the world, does this actually mean that we are 'committed to the reality' of universals? Does the fact that we can use a word like 'redness' imply that there is such a *thing* as redness? This leads us to the issue of what ontological commitment really amounts to. What is it to say that some given category of objects is *real*? Of course 'real' has a use in everyday language. David Cameron is real; Sherlock Holmes isn't. We have reasonably clear criteria for distinguishing the real from the unreal, giving definite answers in a wide range of cases. When a philosopher asks whether universals (or physical objects or qualia etc.) are real, is she using the

word 'real' in this familiar sense? One reason for doubting this is that no philosopher has made a *discovery* about reality that has made an impact outside the philosophical community. No-one has ever seen a headline reading 'World exclusive! Reality of universals confirmed!' So what *do* philosophers mean when they inquire about the ontological status of different sorts of things? In fact, it is arguable that divorced from their normal criteria of application in everyday life, words like 'exist' and 'real' are somewhat indeterminate in meaning. To see this, note that if nouns representing universals are indeed indispensable, this fact can be interpreted in two different ways. The standard interpretation is that universals exist. An alternative is that since universals do not exist, we must refer to non-existent things. The former interpretation ties existence necessarily to reference; the latter requires a stronger criterion for existence. (One that has been suggested is that existent things must be spatiotemporally located, which would rule out universals.) There seems to be no reason for adjudicating the matter one way or the other. In fact, the issue seems to be a purely verbal one that does not reveal any substantive disagreement. Both ways of talking seem entirely legitimate. The former *might* appear suspect because of its apparent metaphysical character. But as long as philosophers who talk this way do not attribute to universals any properties that we would not attribute to them in everyday life, there need be no objection. True, some philosophers—such as Plato with his Theory of Forms—do not adhere to this. But we need not be tempted by such excesses.

Question 7

Is utilitarianism compatible with having integrity in one's moral opinions and behaviour? (50 marks)

Utilitarianism is the theory that morality is based on *consequences*. In its simplest form, the theory states that an act is right in any given situation just in case it has total consequences at least as good as any alternative available to the agent; it is wrong just in case it has worse consequences than some alternative. This is known as *act utilitarianism*.

Act utilitarianism does seem to conflict with the ideals of moral integrity. Take, for example, its approach to the moral obligation to keep one's promises. Suppose an act utilitarian makes a promise and later faces the decision of whether to keep it. In order to decide this she must look purely and simply at the likely consequences of keeping the promise as compared with those of breaking it. If after an objectively sound assessment of these consequences, she decides that the act of breaking the promise scores even marginally better than that of keeping it, she will feel herself obliged to break the promise. But this is hardly a position that gives weight to the idea of integrity, for it seems *contrary* to integrity to break a promise for such a reason as this. To underline the point, note that the extra benefit associated with breaking the promise might not even accrue to the person to whom the promise was made. A person with such a fragile sense of commitment could hardly be described as having 'integrity'. The same point applies to the example of lying. According to the act utilitarian, we should only tell the truth when the overall consequences of doing so would be no worse than those of telling a lie. Granted, the person who follows this is not as bad as someone who tells a lie merely to escape unpleasant repercussions for

herself, since utilitarianism requires that the consequences be assessed from an impartial standpoint. Even so, such a policy does not seem correctly characterisable as being conducive to moral integrity.

In this connection, some writers have emphasised the damage that act utilitarianism apparently does to the integrity of the *self*. If each of my actions has to be decided according to an entirely new utilitarian calculation based on the situation facing me at the time, how can I remain committed to any particular set of aims or ideals? And if I cannot, how can I retain a genuine sense of *myself* as a person with a particular kind of trajectory through life?

However, act utilitarianism is not the only kind of utilitarian theory and there are others that are arguably better positioned to accommodate moral integrity. The best known is *rule utilitarianism*, which requires that we assess the consequences, not of individual acts, but of general rules. According to this view, the right act in a given situation is the one that would be required by the rule that would produce best consequences if it were generally conformed to. On this view, principles that exemplify moral integrity, such as promise-keeping and truth-telling, have independent significance—they cannot be replaced by a practice of assessing each situation directly in terms of its likely consequences. Another (closely related) approach is the two-tier utilitarian theory of R.M. Hare. Hare distinguishes between the *intuitive* and the *critical* level of moral reflection. The former is the most usual kind of moral thinking, designed for normal everyday use. It is based on relatively simple moral rules (such as ‘Tell the truth’, ‘Do not injure others’). The critical level is designed for cases where the intuitive rules conflict or for choosing between possible rules in the first place. The critical level is utilitarian, but not the intuitive one. (Hare argues that direct utilitarian calculation would be too complex and too vulnerable to distortion by personal bias for use in everyday situations.) It is clear that precepts expressing moral integrity would be applicable at the intuitive moral level in Hare’s scheme, even if they have no application at the critical level.

Thus the adoption of some form of indirect utilitarianism, such as rule utilitarianism or Hare’s two-tier theory goes a long way to answering the objection based on integrity. But does it go far enough? It might be pointed out that in theories such as these the type of thinking that is alleged to be most fundamental, that which justifies the intuitive rules in the first place, will make no appeal to moral integrity. In fact, if it happened to be the case that rules flouting moral integrity had greater expected utility than rules respecting it, then the former would have to be chosen by the conscientious utilitarian. It could be argued that this makes the basis of moral integrity too *contingent*: integrity needs to be respected under *all* possible circumstances.

Is this last claim correct? Suppose it could be shown that conformity to a particular rule of moral integrity would lead to consequences involving a very significant reduction in human well-being, as compared to the consequences of not adopting such a rule. Then it seems at the very least an open question whether we should, on the one hand, go with the rule or, on the other, reject it because we prefer not to have this loss in well-being. What we have here is surely two conflicting conceptions of morality which have an equal claim to respect. In fact, if the difference in well-being is great enough—if it amounts to a situation in which adoption of the rule would lead to a social disaster—surely it is the utilitarian conception that needs to win out. In view of

this, it seems that a certain form of deontology might be considered to be most convincing here. I have in mind W.D. Ross's approach, in which pursuit of good consequences is considered to be one of several *prima facie* moral duties. (Two caveats: (a) what Ross actually thought we should try to promote is the best possible fit between happiness and desert and the validity of this depends on the acceptability of the contested notion of moral responsibility; and (b): endorsement of Ross's deontology does not imply endorsement of his intuitionist meta-ethics.) Such an approach seems to give both integrity and the desirability of good consequences their proper weight.

Question 9

Evaluate the view that the expression of religious belief involves a commitment to a way of life rather than a claim about what exists. (50 marks)

Certainly, for many adherents, the expression of religious belief involves a commitment to a way of life. One thinks here of the kind of believer typically described as 'devout', whose life is dominated by daily rituals of both private and communal worship and for whom all decisions of significance need to be made by reference to his religious beliefs. But the view in question here is not simply the obvious point that such believers exist. There is an additional claim being made, which is that to be a true religious believer one *must* live a life of this kind and that it is not sufficient simply to believe that God exists, say, or that there will be a last judgment. One might also argue that a religious believer does not typically *claim* that God exists in the sense of putting this forward as a thesis that can be empirically confirmed. Rather, he lives a life which is dominated by a certain 'picture' of the world—a *blik*, as it is sometimes called—of which God is an essential part.

This is a view of religious commitment that is associated with some ideas of the later Wittgenstein. Consider Wittgenstein's view of what a religious belief in the Last Judgment consists in. Wittgenstein points out that there could be good empirical evidence for a last judgment. We might come across people who could reliably predict the future and these individuals might tell us that such an event will occur. But Wittgenstein goes on to point out that even if we accepted this, it would not amount to a religious belief on our part. For the latter, it is necessary that we take the right *attitude* to this future event, that we let it dominate our behaviour and our decisions about how to live our lives.

Wittgenstein also addresses the question of the *reasonableness* or otherwise of religious belief. Suppose a person says he believes in the Last Judgment because he had a dream that told him it would occur. Religious sceptics would say that it was unreasonable to base a belief like this on a dream. Wittgenstein agrees that this would not be reasonable. But he does not want to say that it is 'unreasonable', with the latter word's sense of rebuke. Instead, he wants to say that it does not *aim* at reasonableness—it does not involve *trying* to conform to the sorts of rational standards applicable in the sciences, for example.

Wittgenstein's attitude to religion seems to be linked to his views about the criticism of moves made within a given language game. A move within the game of chess can

only be criticised by reference to the rules of chess. There is no external standard by which its validity can be determined. So too it is with language games and with 'forms of life' such as religions. To complain that religious statements are not empirically verifiable, for example, is to attempt to make an 'external' evaluation of them, to fail to assess them in the only way that it can be legitimate to do so: by appeal to the 'internal rules' of that particular form of life.

Is Wittgenstein's account of religion convincing? Can it be used to provide an adequate defence against the charge that religious believers are 'unreasonable' in view of their failure to conform to ordinary standards of confirmation such as those applicable in the sciences? It is surely correct to note, first of all, that adherence to certain beliefs is not sufficient for 'being religious'. To think that the Last Judgment will occur, yet be unmoved to let this affect the way that one lives one's life is not the attitude of any genuine adherent of a religion, such as Christianity, in which a last judgment figures. However, we should be careful not to infer too much from this. True, religions have an attitudinal or evaluative dimension. But that does not imply that they do *not* also essentially have a cognitive dimension. Indeed this cognitive dimension, consisting in the acceptance of certain beliefs about what exists (God, Heaven, Hell etc.) over and above the familiar world of common sense is apparently what forms the basis for the attitudes and commitments. There is no reason to devote one's life to God if God does not actually exist.

On the other hand, the above argument might be thought to betray a misunderstanding of what Wittgenstein was saying. He was not *denying* that religious adherence consists at least partly in the acceptance of certain beliefs. His main point was that these beliefs need to be evaluated according to the internal standards appropriate to religion and not according to those of other forms of life such as science. However, the evidential standards which sceptics (and some religious believers themselves) attempt to apply to religious claims are not peculiar to science. Rather, the methods used in the various sciences are themselves extrapolations from the evidential standards that people use in their everyday lives the world over. They are not just a particular 'form of life'. They govern the formation of beliefs in all forms of life that have some sort of cognitive content (which may indeed mean all without exception). Such a view might seem to propose a kind of rational absolutism that (it could be claimed) runs counter to the facts. People do not apply the same evidential standards the world over, it might be said. For example, shamans and those cultures that accept their claims, base their beliefs about illness and its possible remedies on very different principles from those of western science. They do not engage in the controlled trials familiar to western 'science based' medicine. That is true. But it is no reason to think that shamans do not ultimately accept the same fundamental normative principles of belief acquisition as westerners. If a shaman's remedy does not work on a particular occasion, he will normally try to give *some* sort of explanation for this. This shows his acceptance of at least the following standard norm of rationality: an unexplained case of X's failing to have an effect on one occasion provides some reason for thinking that it might also fail to have this effect on future occasions.

So can we say, then, that on the basis of universal standards of rationality, anyone who believes that there is a last judgment merely because he has learned this in a dream is irrational or unreasonable? Perhaps we should not be so quick to conclude this. 'Standard rationality' is indeed something known and practised, with varying

degrees of competence, in every human culture. But many, if not most cultures, depart from this standard with respect to *certain* subject-matters. For many, if not most, cultures include superstitions and religious beliefs and, I have been arguing (or rather assuming for the purposes of this essay), these beliefs do not conform to standard rationality. But the religious believer might say ‘So what?’ Why should there not be occasional exceptions to standard rationality? The sceptic is likely at this point to appeal to the connection between truth and rationality. Beliefs conforming to standard rational principles are more likely to be true than those not conforming to it. Hence the superstitious and the religious risk believing things that are not true. Now the most self-conscious members of any culture—even one as deeply immersed in the most extreme religious beliefs and practices that one could imagine—can readily acknowledge this to be true *in general*. But what could the sceptic say to convince them that it is true in *every* case, i.e., even in relation to religious matters? It seems that any argument to this effect would itself have to appeal to principles of rationality about which the very same question of universal applicability could be raised.

In view of all this, it seems right to say that the person who believes in the Last Judgment merely because of a dream has not made a mere blunder and indeed takes a stance that is difficult to refute in any absolute sense. If we think of the burden of proof as being on the sceptic, then this creates a problem. On the other hand, if (as is more normal these days) we think of the burden of proof as being on the religious believer, then it is a problem for *him* that his religious beliefs do not conform to standard rationality. Of course this makes it pertinent to ask: where *does* the burden of proof lie? Perhaps there is no objectively correct answer to this question. If you are a sceptic and you are trying to persuade a religious person to abandon some of her non-rational beliefs, the burden of proof lies with you and the problem is yours. If, on the other hand, you are the sort of religious person who seeks to persuade others to adopt your religion, then it is a problem for you if you cannot convincingly appeal to standard rational principles in order to do so.