There Is No Hard Problem of Consciousness

Kieron O'Hara
Tom Scutt

Artificial Intelligence Group
Dept. of Psychology
University of Nottingham
University Park
Nottingham NG7 2RD
U.K.

Abstract: The paper attempts to establish the importance of addressing what Chalmers (1995) calls the 'easy problems' of consciousness, at the expense of the 'hard problem'. One pragmatic argument and two philosophical arguments are presented to defend this approach to consciousness, and three major theories of consciousness are criticised in this light. Finally, it is shown that concentration on the easy problems does not lead to eliminativism with respect to consciousness.

1. Introduction

Cognitive psychology has made many important advances over the past couple of decades. Conceived as a response to the hollow dogma of behaviourism, as a discipline it has clearly shown itself to possess greater explanatory power, and has been widely accepted as an important paradigm for the description and modelling of a number of psychological phenomena.

For example the work of Marr and Hildreth (1980), not only outlined the higher-level functions (such as edge detection) that the visual system would need to be capable of, but also predicted many of the lower-level neural mechanisms for computing such functions, which have since been revealed by neurophysiological and psychophysical research.

One reason for the success of cognitive psychology is its technological aspect. Its methodology is distinctive for its concentration on computational and formal aspects of psychological phenomena. When confronted with some problematic behaviour, the typical response of a cognitive psychologist is to ask: *How is that done?* He or she typically lives in a world of simulations and models. This concentration on computation provides

cognitive psychology with the important links to the engineering discipline of artificial intelligence. In AI, the engineer asks: *How could that be done?* And we can see that the explanations in cognitive psychology will find a ready audience in AI, ready to test the psychological hypotheses to the limit.

Because of the important links between cognitive psychology and engineering, however, many thinkers find its theories unsatisfactory as *general* explanations of human psychological phenomena. The cognitive psychologist, by concentrating on computational aspects of such phenomena, will, according to these critics, miss out on precisely those aspects which are *not* computational. The cognitive psychologist's reply is that anything that is literally true about psychological phenomena will be fully accounted for by a concentration on their computational aspects, and his or her research programme will be an attempt to delve for satisfactory computational or formal explanations.

The result is a stand-off, with the critic emphasising the aspects of the phenomena which are *prima facie* independent of any possible computational account, and the cognitive psychologist standing by the record of the discipline in past investigations. In any particular area, the matter can only be resolved convincingly by the success or failure of the cognitive psychologist in providing successful theories.

One such area is that of consciousness. The question is how any straightforward physicalist account, such as would be provided by cognitive psychology, can show how such apparently non-physical phenomena as awareness are possible. The explananda of such an account are not often described as exactly as one might hope, as we shall discuss in section 3, but there is a general consensus that, however they are described, they are problematic. In Nagel's formulation of the problem (1974), if something has consciousness, then it is *like* something to be that thing (it has a subjective aspect). But surely — the critic adds — it isn't like *anything* to be a computer program, or formal theory. How can the cognitive psychologist ever succeed in reducing the *sui generis* phenomenon of consciousness to a series of physical phenomena formally described?

This leads to distinction drawn by the critic between the *easy problems* and the *hard problem* of consciousness. The easy problems are the problems that are clearly amenable to the methodology of cognitive science: Chalmers (1995) suggests as examples the ability to discriminate, categorize and react to environmental stimuli; the integration of information by a cognitive system; the reportability of mental states; the ability of a system to access its own internal states; the focus of attention; the deliberate control of behaviour; and the difference between wakefulness and sleep. All these phenomena — and doubtless others — will be susceptible to the methodology of cognitive psychology.

Contrasted with these problems is the hard problem, which is that of explaining consciousness itself, how consciousness is possible at all. Answers to the easy problems would fail to address the hard problem, which is the subjective aspect of consciousness, because the easy problems are

clearly approachable using only objective terms. The sum of all the (objective) answers to all the easy problems would never be an explanation of subjectivity itself.

In this paper, we will defend the claims of cognitive psychology against the critic we have described, by arguing that the hard problem (or HP as we shall call it) is not a serious problem. For a problem to be a genuine problem, some sort of idea of a solution must be available (e.g. some way of recognising a solution when one stumbles across one); whereas all discussion of the HP seems to preclude any sort of answer being given. Our chief argument has two components. The first component is a pragmatic claim, to the effect that, since we know how to address the easy questions (at least in rough), and we haven't the foggiest idea of what to do about HP, we would be wasting our time if we didn't explore the easy problems first; we establish this in section 2. The second component is a philosophical claim that we do not know what HP is as yet, and that the only possible way of approaching HP is via the easy problems; this is the business of section 3. It may be thought that the existence of attempted answers to HP precludes the possibility that HP cannot sensibly be addressed; in section 4 we discuss some of these attempted answers to HP (the theories of Edelman, Crick and Chalmers), to show why they fail to address any actual problem. Finally, in section 5, we will defend our position against charges that we have eliminated, or would like to eliminate, the notion of consciousness from scientific discourse.

2. Methodological Reasons to Ignore HP

In the venerable joke, little Johnny is peering hard at the lawn. "Why Johnny," says his mother, "what are you doing?" Johnny points across the garden, and says "I lost my pocket money over there." "If you lost your money over by the garage, why are you looking by the pond?" asks his mother, and little Johnny duly replies: "'Cos the grass is shorter over here."

One should never analyse humour, but since this joke isn't actually very funny, we will, just this once. Why is it a joke that Johnny is looking where the grass is short, and not where the grass is long? Well, he thinks that he will save effort by looking where the grass is short, because such a search is easier in short grass than in long, but we know that he will expend greater effort, because the money is almost certainly concealed in the long grass, and furthermore, little Johnny has told us that himself. So, although a smaller amount of effort would be expended in the short grass per blade (or whatever unit is appropriate here), all that effort is guaranteed to be in vain, since the money is somewhere else entirely. And since he is aware of all that, and has simply failed to draw the obvious conclusion, little Johnny is just being stupid.

The analogy to consciousness is evident. The critic we discerned in the opening section argues that the answer to HP is to be found in the long grass over by the garage, and that cognitive psychologists who mess about with

the easy questions by the pond will — to the extent that they see themselves as thereby addressing HP — be wasting their time.

However, approaching consciousness via the easy problems is only open to the same criticisms as little Johnny's approach to retrieving his money if two conditions hold, to render the analogy suitably watertight. In the first place, in the joke, we know that the money is in the long grass. For Johnny to be *really* stupid for the purposes of the joke, he must know that the money is in the long grass too, but the main point is that there is a laugh to be had as long as we in the audience know roughly where the money is. And in the second place, it must be the case that Johnny *could* look in the long grass, that he has a procedure for finding the money such that, if he employed it in the long grass, he would achieve his goal.

When we come to examine the extent to which the joke is analogous to the situation in the study of consciousness, we find that neither of these two conditions holds. We do not know where to look for the answer to HP, and in particular, we don't know that we are necessarily wrong to approach the question via the easy problems. The joke would not be funny if little Johnny knew only that he had lost his money *somewhere* in the garden; in that event, he would be rational, not stupid, to begin to look where least effort would be expended. For the second condition to hold, it would have to be the case that there was an alternative approach to consciousness to the one provided by the cognitive psychological investigation of the easy problems. But as it is, no-one really has any well-worked-out plan for approaching the problem any other way. If, in the joke, the long grass was completely and necessarily impenetrable, Johnny could not be blamed for not looking there.

We have, in the study of consciousness, an obvious problem which is that addressing the easy problems is the only game in town. We do not know that addressing the easy problems will automatically fail to provide insight into HP. We do not know any other way of addressing HP. We do know how to address the easy problems — indeed, Chalmers (1995) seems to hold this as definitive of the easy problems. Pragmatically, it does look as if looking into the easy problems is as likely to be successful as any other way of approaching consciousness, and therefore it would seem to be as rational a way of proceeding as any.

Consider the alternative; the only alternative to the easy way is to develop a new paradigm of psychological investigation. For example — and this is an extreme, if genuine, example — take Penrose's (1994) suggestion that a new branch of physics is required to explain consciousness. Pragmatically, is there any reason at all why one should take the Penrose route to a solution to HP? Addressing the easy problems is hard enough, and work in those cognitive psychological fields is likely to produce interesting and important results (as no-one denies), even if the end result is not an answer to HP. The advantage of such an approach is that we know where to start, there is a body of work to build on, there are surrounding theories in which such work can be embedded and located. But developing a new branch of physics is the sort of paradigm-shifting, epoch-making task that, unless one was Nobel

Prize material, one should probably avoid. Worse, one is not even guaranteed success with respect to the HP, since any such suggestion remains tentative, though the Nobel Prize ought to be consolation enough. It is difficult to see why anyone who is neither seriously irrational nor seriously clever, and who wishes to contribute to a solution to HP, would not take the pedestrian route and go for the easy problems.

We call this argument the *pragmatic argument*.

3. Philosophical Reasons to Ignore HP

Surprisingly, in this debate, pragmatic considerations carry little weight. Although the pragmatic argument seems to us to be sufficient to consign HP to the dustbin reserved for historical curiosities, the critic of cognitive psychological approaches to consciousness will certainly not agree. More *a priori* reasons are required here, since all the arguments to be met are firmly unencumbered with empirical detail.

Therefore, in addition to the pragmatic argument, we will now add two compelling *a priori* arguments to leave HP alone. Both are concerned with the lack of understanding of the nature of HP that precedes satisfactory answers to the easy problems.

The first *a priori* argument, which we shall refer to as the *context argument*, simply makes the obvious point that an understanding of a concept may be facilitated, at the very minimum, by the investigation of contexts in which that concept is important. Take the example of the concept of *life*. The analogy of HP in this context is the question of what life is (call this LHP). The analogy of the easy problems would be various biological questions to do with the processes that go on in things that are alive. Many major philosophers addressed the LHP, but very few addressed the easy problems. Indeed, between Aristotle and Harvey it is not obvious that *anyone* did. We can also agree that, remarkable as they were, Aristotle's discussions of the easy problems were inadequate. Yet as the discipline of biology grew up in the modern era, more and more easy problems fell to the advancing forces. Now, what happened to LHP? Since the vitalists, no-one has seriously attempted to answer the question; indeed, very few thinkers think the question is worth asking at all. Those who do still wonder about LHP are increasingly likely to assimilate it with one of the easy problems. The answers to the easy problems gradually undermined the air of mystery that underlay LHP, and suddenly LHP didn't really seem like a problem any more.

Consciousness may or may not be analogous to life in this way. It may be the case that answering all the easy problems completely and satisfactorily would still not get rid of HP. But that cannot be decided *a priori*. Clearly, answers to the easy problems are bound to lead to adjustments in the concept of consciousness. For example, there is no doubt that the research on epileptics who had had their corpus callosum severed (Gazzaniga, 1970; Nagel, 1971), so-called 'split-brain' patients, has led to a serious readjustment of the notion of the unity of consciousness, and no philosopher

or psychologist investigating in this field can afford to ignore such results. But surely no-one would claim that the phenomena of consciousness thereby revealed would have been discovered *a priori*, without empirical research into this problem (presumably an 'easy' problem). Hence, had this easy problem been neglected in favour of HP, any putative solution to HP would quite simply have been wrong, because it would have failed to account for important properties of consciousness (in which case it couldn't have described consciousness fully). Furthermore, the only way to provide criticisms of the putative solutions to HP would be to perform the research into the easy problems to discover the properties of consciousness overlooked by HP-research.

So the very notion of consciousness we are dealing with will change as more answers to easy problems roll in (right *and* wrong answers). Older notions of consciousness will be seen to have been mistaken, or simply wither on the vine as their role in scientific research diminishes. Hence, addressing the easy problems will seriously affect any answers to HP that we get, via their effect on the notion of consciousness whose nature is the topic of HP. The lesson is clear: we cannot be confident of any putative solution to HP until we have a reliable picture of the topography of the ground via the less ambitious research into the easy problems. Not all easy problems would have to be solved, but a well-supported psychological consensus will, without a shadow of a doubt, precede any serious assault on HP. That consensus is not yet with us, and will only follow a long period of dedicated research into the easy problems.

The second *a priori* argument, which we shall call the *epistemological argument*, is an adaptation of a recent argument from Colin McGinn (1989). McGinn's argument has been widely criticised, and with good reason. But we feel that there is a kernel of truth in it, which is highly germane to the issue at hand. McGinn argues that we will never understand consciousness (i.e. we will never solve HP), and argues innovatively on epistemological grounds.

His argument goes as follows. We are all good physicalists, so any explanation of consciousness (solution to HP) will consist of some physical predicate, P, satisfaction of which by an entity will entail that that entity has consciousness. But, argues McGinn, it is inconceivable that we would ever recognise that a mere physical predicate confers consciousness in this way; we just do not understand consciousness in the right terms. So even if we discovered P, we would never recognise that it was the solution to HP.

There are many dubious points in McGinn's argument (which in addition we have necessarily caricatured). It is not obvious to us which side of our particular fence McGinn sits; he would probably advocate the pursuit of the easy problems on the ground that HP will be forever out of our reach, and therefore would probably see the point of our methodological argument in section 2. But he would also discount the possibility of HP fading away, or changing rapidly, as a result of the investigation of the easy problems, and

in this he agrees with the critic of cognitive psychology we discerned in section 1.

What McGinn has done is highlighted the importance of our recognition that a solution to HP *is* a solution to HP. We do not agree that it is necessarily the case that we will never make such a recognition. But clearly some prerequisites for such a recognition must obtain before McGinn can be shown to be wrong. One such prerequisite is that we must know what we are talking about; until the nature of consciousness is more or less well-understood, we will never recognise that a theory solves HP. Note that this is independent of whether or not the theory actually *does* solve HP. Hence until HP is well-understood, it won't be rational to address HP.

As a parallel, consider the work of ancient philosophers about the nature of matter. Their version of HP was MHP, what the nature of matter is. The easy problems in this realm are the way that matter reacts in various circumstances (in effect, the problems of chemistry). Now it turned out that Democritus, and later Lucretius, hit the nail on the head with their solution to MHP, which was an atomic theory. Although their theories weren't particularly detailed in the way in which modern atomic theories are detailed, they were more nearly right than Anaxagoras or Empedocles.

But, because these authors were addressing MHP before the easy problems had been elucidated to any significant degree, it is fair to say that — from the point of view of chemistry at least — their answers to MHP are effectively worthless. However wonderful *De Rerum Natura* is as a piece of philosophical poetry, it is clear that its value is quite independent of the discovery that the right answer to MHP is that matter is atomic, as evinced by the poem's high reputation during periods when it was thought that the answer to MHP was that matter is *not* atomic. And it is also clear that Democritus and Lucretius weren't right for the right reasons. They were neither cleverer, nor more rational, nor more well-informed than their rivals; it is simply that their shot in the dark happened to hit the target, while their rivals' didn't.

It is our contention that the thinkers who approach the notion of consciousness via HP are going to be subject to the same problems. It may indeed be the case that a new branch of physics is required to solve HP, as Penrose claims. It may be the case that information has an irreducible phenomenal aspect, as Chalmers suggests. But since neither of these authors has any idea of the detailed consequences of their solutions to HP, then neither of their answers can be considered as authoritative. And the reason that neither author is able to be authoritative is that work on the easy problems has barely commenced. A research programme, the programme of cognitive psychology, is underway to address the easy problems. That programme may fail completely, or may gain partial success, or may claim to provide complete solutions to the easy problems. When the programme has been carried through sufficiently to make an estimate of the extent of the success of the research, anyone who makes a stab at HP will be on much more solid ground. And even if one of the currently competing claims —

all shots in the dark, as we claim — *does* hit the target, that will be coincidence, and will be rejected as such.

4. Edelman's, Crick's and Chalmers' Attempts to Address HP

The obvious reply to the arguments adduced in sections 2 and 3 above is that, although we have claimed that addressing HP will be at best very difficult and at worst impossible, there has been quite a number of serious attempts at solving HP. In turn, the obvious rejoinder to *that* would be that these attempts fail in the ways that we have suggested that they will fail.

We have already discussed Penrose's (1994) claim that a new branch of physics is required. This does not even look like an approach to HP, at least until the putative branch of physics is exhibited. Similarly, McGinn's (1989) idea that a physical predicate P, true of the brain, will explain "the mind", is hardly of use in the context of HP, given that he doesn't suggest what this predicate may be (because he thinks that we can never understand how the predicate will explain the mind), which makes his argument more like a testament to his faith in materialism than anything else. Any solution to HP must give us a handle on the problem. Other attempts simply do not have the right form. Eccles' theory (Popper and Eccles, 1977) that random quantum effects will undermine physical determinism and 'let in' the mind will ultimately fail, because the effects of the mind are specifically *not* random.

Our aim in this section will be to take a number of influential attempts on HP, and show that they succumb to the three arguments given above. Our aim is to criticise theories that are plausible, in the sense that Penrose's, McGinn's and Eccles' are not, and furthermore which have gathered attention and support from the community. We hope they are sufficiently representative. Of course, criticising existing theories can never be *decisive* — since another, better, theory may always come along later — but we expect that the exercise will demonstrate the plausibility of our arguments. To that end we consider the theories of consciousness of Gerald Edelman (1992), Francis Crick (Crick, 1994; Crick and Koch, 1990) and David Chalmers (1995).

4.1 Edelman and the Pragmatic Argument

Edelman distinguishes between Primary and Higher-Order Consciousness. Animals with primary consciousness have phenomenal experience, but are not aware that they have it. They have no notion of self, and therefore nothing to 'attach' the experience to. Primary consciousness consists of a re-entrant loop connecting value-category memory to current perceptual categorization (1992, p.120). There would be a fleeting phenomenal aspect, but

> Creatures with primary consciousness, while possessing mental images, have no capacity to view those images from the vantage point of a socially constructed self. (Edelman, 1992, p. 124)

Edelman is very concerned — rightly, of course — to make sure that this notion of consciousness has a place in neuropsychology, and is not merely a metaphysical conceit.

> What is the evolutionary value of such a system? Obviously, primary consciousness must be efficacious if this biological account is correct. Consciousness is not merely an epiphenomenon. (Edelman, 1992, p. 121)

The question then is how this promise can be cashed out. And this is where what Chalmers (1995) terms the bait-and-switch takes place. Edelman lists the properties that make consciousness useful (such as relating the creature's present inputs to previous dangers or rewards as experienced). But what actually happens is that Edelman lists the benefits of the re-entrant loop, the connection between value-category memory and current perceptual categorization, *i.e. the thing he attached the consciousness label to in the first place*. This is a common tactic: take some neuro-architectural item that seems to be doing the sort of thing we want consciousness to do; attach the consciousness label to it; and then show the benefits of consciousness by pointing to the assets of our newly labelled neuronal structure.

We can see that Edelman has fallen foul of the pragmatic considerations we discerned in section 2. The only way we can make progress is to tackle the easy problems; concentration on HP will always fail to provide a lead. Hence, when attempting to tackle HP, the temptation will always be to move covertly towards the easy problems. In Edelman's case, the re-entrant loop is a phenomenon that is eminently open to study, and he applies his considerable expertise to that end, with interesting results. But the methodological difficulties involved in gaining a vantage point for HP will always lead the researcher to move via the easy problems, and we see why a common class of theories emerges, of which Edelman's is an exemplar, in which a neuronal phenomenon is assimilated to the phenomena of consciousness. Chalmers is correct to say that such a strategy will always be unsatisfying for someone interested in HP.

Incidentally, Edelman also exemplifies another methodological pitfall in assaulting HP. As Dennett puts it, Edelman's theory is

> ... an instructive failure. It shows in great detail just how many different sorts of question must be answered before we can claim to have secured a complete theory of consciousness, but it also shows that no one theorist can appreciate all the subtleties of the problems addressed by the different fields. (Dennett, 1991, p.268, n.1)

There are so many easy problems to be understood, in so many fields, that the researcher who would take a crack at HP will be hard pressed to find (and argue for) a point of view that transcends them all. But such a point of view is vital if HP is to be regarded as a problem essentially distinct from the (sum of the) easy problems. Again, pragmatically, there are few who can move so easily between all the various discipline boundaries that they

can sift through the evidence provided by the various answers to the various easy problems and make correct judgments about them.

4.2 Crick and the Context Argument

Francis Crick, rather than give an exact *location* for consciousness (he does give locations for certain phenomena, for example, free will, in his (1994), but these speculations are probably better overlooked; indeed, perhaps his tongue is in his cheek), points instead to an observable phenomenon — certain 35-75 hertz neural oscillations in the cerebral cortex (Crick and Koch, 1990). These oscillations appear to be correlated with the binding of information from various sense modalities, but it remains unclear whether Crick and Koch think that these oscillations form the basis for consciousness, are merely correlated with consciousness, or form an observable aspect of consciousness. Chalmers (1995) is right to criticise Crick and Koch on the grounds that suggesting a correlation between two phenomena in no way explains how one gives rise to the other (or indeed, which gives rise to which).

Of course, concentration on particular observables amounts yet again to an assault on an easy problem. Once more, HP has been sidestepped, as Chalmers points out. But Crick and Koch are prepared to get excited over the properties of the easy problem they have concentrated on.

> We have suggested that one of the functions of consciousness is to present the result of various underlying computations and that this involves an attentional mechanism that temporarily binds the relevant neurons together by synchronizing their spikes in 40 hz oscillations. (Crick and Koch, 1990, p.272)

Let us assume that they are right, for the sake of argument. Now, consciousness as we presently understand it is not like this at all. If Crick and Koch have got it right, then our view of consciousness will have to change, to conform to their results. But note that it is their concentration on a particular *easy* problem that has given them this perspective on HP. In that event, the form that HP has finally taken would have depended on the results of some investigation into an easy problem.

On the other hand, if Crick and Koch have not provided us with any sort of handle on HP, as Chalmers says, then it is clear that they have uncovered some interesting results concerning the easy problems. No view of HP could *fail* to take these results into account. Good work on the easy problems will always have to be incorporated into HP. And such results are unlikely to be predictable in advance; recall our remarks concerning brain bisection and the unity of consciousness.

In other words, as we claimed in section 3, the easy problems, when solved, alter HP; hence HP could not be considered as reaching any sort of 'final form' until the easy problems were beginning to be cracked. They don't all have to be solved, of course, but we need to be comfortable enough in this area so that we do not expect to be surprised by any new discoveries. 'Comfortable enough' is a vague term, of course, but its application seems

clear. We are pretty familiar with the functions of the lungs, for example, and do not really expect to have our understanding of them turned upside-down (though of course that understanding is always liable to be adjusted in the light of new discoveries). On the other hand, Big Bang theory is so new and untested that any speculation on the basis of its current form is probably going to be made redundant by massive changes in the theory itself. Our claim is that consciousness theory is more like Big Bang theory than lung theory, and that we are not yet comfortable enough with consciousness theory for it to be worth the risk in wasted resources to start speculating.

4.3 Edelman, Crick, Chalmers and the Epistemological Argument

The theories we have so far considered also provide evidence for our epistemological argument. Both Edelman and Crick are distinguished neurobiologists who have claimed to have located consciousness in the brain. In neither case is this fully explained, although in both cases it is work on the 'easy' problems which has led to a suggested solution of HP. We do not share Chalmers' pessimism that work on the former can never help with our search for the latter. However, at present it is too early (indeed, it often looks plain ridiculous) to point to a part of the brain and say 'this is where consciousness lies'. Diagrams of cognitive architecture in modern texts on the nature of mind often look like medieval maps with labels saying "Here be Dragons". Until we have a much clearer idea of what we are looking for, it is far too early to say where we have found it.

Chalmers, on the other hand, takes a very different route, but will still fall foul of the epistemological argument. His position is founded on an *a priori* assumption that there can be no scientific account of consciousness (at least in physical terms as currently understood). He claims that there needs to be a non-reductive explanation of the hard problem of consciousness. Reductive explanation, according to Chalmers, won't get us anywhere. He gives an example of where such examples succeed and fail:

> If someone says "I can see that you have explained how DNA stores and transmits hereditary information from one generation to the next, but you have not explained how it is a *gene*", then they are making a conceptual mistake... But if someone says "I can see that you have explained how information is discriminated, integrated, and reported, but you have not explained how it is *experienced*", they are not making a conceptual mistake. This is a nontrivial further question. (Chalmers, 1995)

Indeed. But that is because at the moment it is impossible to see how experience *could* be explained in terms of mere (sic) information processing — we lack the necessary bridging concepts to see one in terms of the other. In fact we are in much the same position as a geneticist would have been if such a term existed two hundred years ago.

From our present viewpoint, a statement such as 'I see how you have explained the nature and transmission of electromagnetic radiation in the range 400 - 700 nanometres, but you have not explained how it is light', seems somewhat naïve. But before these scientific facts about the electromagnetic nature of light were carefully laid out by that painstaking and haphazard endeavour we call science, the validity or otherwise of this sentence would have been open to debate (and a very ill-informed debate at that). At the end of the next millennium, a statement that consciousness can be explained in terms of X, Y and Z may be widely accepted as being beyond doubt. The same statement from our present viewpoint would be viewed with some scepticism, because we haven't yet done the work (on the easy problems) which would justify this statement. Chalmers himself admits that 'there is a direct correspondence between consciousness and awareness'. Wouldn't it make sense, this being the case, to research awareness (the easy problem) to better understand consciousness? Is Chalmers *really* saying that understanding the former doesn't help us *at all* in understanding the latter? The naïve question which began this paragraph is naïve only so long as it is understood in a particular way (e.g. from the standpoint of the twentieth century). Before, in Kuhnian terms, the paradigm shifted, it is a perfectly good question. For the paradigm with respect to consciousness to shift, however, a good deal of evidence has to be built up with no obvious home in any general theory — this is what provides the impetus for a paradigm shift. And this evidence will only be built up by addressing the easy problems.

So what are we to make of Chalmers' alternative to a physical account of consciousness? In the 'dual-aspect theory of information', it appears to us that 'Information Space' — the place where 'experience might have a subtle kind of causal relevance in virtue of its status as the intrinsic nature of the physical' (Chalmers, 1995) — is the pineal gland of Chalmers' account. However overstated some may have been in their heralding of Crick and Koch's neural oscillations as an aspect of consciousness, at least such work on the easy problems may get us nearer to a solution of HP (or nearer to a *reasoned* rejection of a physical account). It is not so clear that the non-reductive account of Chalmers offers any such reward. It is, in our current epistemological state, impossible to understand how information can have a phenomenal aspect. It may turn out to be true, but we, in the twentieth century, have no idea what the implications of such a claim are. It is as unexplanatory to us as atomism was to the Ancient Greeks.

5. Eliminativism With Respect to Consciousness and HP

There is one final point that we would like to make, and that is to rebut the charge that our insistence that concentration on the easy problems and consequent neglect of HP amounts to eliminativism with respect to consciousness. HP, such a charge would maintain, is *the* fundamental, definitive problem of consciousness. Anything else, such as the ability to discriminate environmental stimuli, is mere effect. Hence to avoid or reject

HP amounts to rejecting the scientific investigation of consciousness. Such a strategy results either in the elimination of consciousness from our discourse, or the cheerful acceptance of its irreducibly occult nature.

A sophisticated version of this charge would attempt to line us up with an eliminativist argument such as that of Kathy Wilkes (1988), where sceptical play is made with the notion that consciousness is not used to any great effect in natural science, is not rooted 'in observation or experiment' (1988, p.38), cannot be seen as a natural kind, and has little or no explanatory power. The idea then is that, since the idea of consciousness fails to have any scientific *use* (although of course it may be the case that scientists have as a *goal* the elucidation of consciousness — i.e. the investigation of HP), nothing would be lost were consciousness removed from scientific discourse altogether.

First, we note that this argument will not support the elimination of the notion of consciousness altogether, in the way that Churchland, for example, wants to eliminate mental terms from discourse generally. Wilkes' argument will only apply to *scientific* discourse, as she herself is more than willing to concede (1991). Hence she is advocating treating the term 'consciousness' rather like the term 'dirt', admitting that it has a ready usage in ordinary everyday talk, while keeping it out of more technical scientific discourse.

Second, we note that if those who would support the investigation of HP are right in saying that HP is the only way of 'getting at' consciousness, then concentration on the easy problems and the removal of HP from the realm of scientific discourse *would* constitute an elimination of consciousness from science. As we argued in section 3, we do not accept that they *are* right in so saying.

Our main reply to this charge, however, is to side-step it. We are quite happy to agree with Wilkes that her *conditional* premiss is correct (with a small modification). That is, we accept that *if consciousness has no use in any practical context, then it should be eliminated from such contexts*. We have substituted the term 'practical' for 'scientific'. We do not want to define 'practical' here, but it certainly includes the scientific. Whether all practical contexts are scientific is a question we would like to leave open (many famous philosophers could be lined up both for and against that proposition).

It follows from this that we would accept the eliminativist charge if we accepted Wilkes' non-conditional premiss that *consciousness has no use in any practical context*. However, we don't. There are a number of areas where some aspects of the notion of consciousness are used ineliminably. We think that the easy problems are examples of these (although this begs the question against our pro-HP critic). Perhaps the clearest case of a scientific (and therefore practical) context in which the notion of consciousness is centrally involved is the discipline of anaesthesiology, which is the science of making people lose conscious awareness. This is not the place to go into detail on this matter, but see (Nikolinakos, 1994) for a

recent discussion of scepticism about consciousness in the context of anaesthesiology.

We think that the elimination or otherwise of the notion of consciousness depends entirely on that notion being used practically, and would not wish our attachment to it to be merely sentimental. Fortunately, we can get serious work out of the concept, in a number of psychological and/or medical fields, and therefore, despite our rejection of HP, we are very happy to accept that the notion of consciousness has an important role to play in our scientific and practical discourse. Of course, the nature of that notion would depend on its use in those practical contexts — many aspects of the 'common-sense' or 'folk' notion of consciousness would turn out to be of little or no use in any practical context. But we think that this is a good discipline, and should help prevent metaphysical obscurities muddying the waters of practical discourse.

6. Let's Take It Easy

Our aim in this paper has been to establish that the easy problems of consciousness, as addressed by cognitive psychology, remain the most promising route to success. We adduced three arguments against approaching HP directly:

- the *pragmatic* argument — we know how to address the easy problems, whereas we can only speculate with respect to HP;

- the *context* argument — the elucidation of the easy problems will probably change the nature of our understanding of HP, and so any attempt to solve HP currently would be premature; and

- the *epistemological* argument — a solution to HP would only be of value if its success could be recognised because the problem was well-enough understood.

We illustrated our claim with discussions of three prominent attempts to solve HP. Edelman, trying to avoid vacuousness, ends up conflating his answer to the easy problems with his answer to HP, showing the perils of ignoring the pragmatic argument. Crick, like the researchers into the split-brain phenomena before him, doesn't solve HP, but provides important input into HP, subtly changing its nature, as we would expect from the context argument. Because our understanding of HP is so limited, Edelman, Crick and Chalmers all fall foul of the epistemological argument. Our final point was to say that our preference for the easy problems and consequent neglect of HP does not amount to eliminativism with respect to consciousness.

In this concluding section, we would like to expand a little on the nature of the easy problems and the hard problem. Part of the difficulty here is caused by the names given to the two groupings. Chalmers says

> The easy problems of consciousness are those that seem directly susceptible to the standard methods of cognitive science, whereby a phenomenon is explained in terms of computational

or neural mechanisms. The hard problems are those that seem to resist those methods. (Chalmers, 1995)

Of course, rhetorically, easy problems don't seem to be worthwhile, and an assault on the hard problem appears more heroic. No prominent player in the game, least of all Chalmers, would make the crass error of saying that the easy problems (e.g. explaining the ability to discriminate, categorize and react to environmental stimuli, the integration of information by a cognitive system, the reportability of mental states, the ability of a system to access its own internal states, the focus of attention, the deliberate control of behaviour, and the difference between wakefulness and sleep) are actually *easy*. Equally, no-one would suppose that HP were not hard.

But it is interesting to explore the reasoning behind calling the easy problems easy and HP hard. The easy problems are not easy, but we do have a handle on how to approach them. We have no guarantee of success, but at least we have a rough idea of where to go from here. What sets HP apart is the fact that we haven't the foggiest idea of how to begin. Compare, for example, Penrose's solution with Chalmers' or Crick's. They are not even remotely in the same ball park. We have three answers by three leading thinkers in three central disciplines, and they point in three mutually exclusive directions.

Consider an analogy with bridge building. Easy bridge building involves constructing structures to cross small spans. Hard bridge building involves constructions over more challenging spans, in unfavourable conditions. These notions of 'easy' and 'hard' do not map onto the ones used in this debate. Suppose we termed bridge building 'easy' just in case we had a rough idea of how to start, no matter how problematic the project, and 'hard' just in case we did not know how to start. Then an easy project might be the construction of a bridge across the Atlantic, and a hard project might be the construction of a bridge into the fourteenth century. A hard bridge builder then might echo Penrose's solution, that a new branch of physics needs to be developed before hard bridge building might begin in earnest. But the obvious response would be to say that, at our current state of ignorance, hard bridge building ought to be shelved, in favour of the easy sort, at least for the time being.

Science is an incremental activity. For every Einstein there are thousands of Scutts and O'Haras, adding their two penn'orth to the current state of knowledge. Ambition plays a role in science, but ambition needs to be tempered by a firm grasp of the possibilities available at any point. In particular, our understanding of the large questions often crucially depends on our answers to all the small and relatively insignificant questions. When enough of the small questions have been answered, then an attempt to answer the hard questions becomes a possibility (unless the hard questions have been shown to be wrong-headed).

We all want to understand consciousness. But the way to achieve that understanding is not to divert intellectual resources toward HP. Because the small questions relating to consciousness remain unanswered, we simply do

not understand HP sufficiently well to determine whether there is a serious question to answer (and if so, precisely what that question is), or whether the whole thing is some ghastly metaphysical error best avoided. We can have opinions on these questions, sure, but those opinions are fundamentally ill-informed, and will remain so until the lessons from the easy problems have been learned and absorbed.

References

Chalmers, David (1995) 'Facing up to the Problem of Consciousness', *Journal of Consciousness Studies*, II

Crick, F. (1994), *The Astonishing Hypothesis: The Scientific Search for the Soul* (New York: Scribners)

Crick, F. and Koch, C. (1990) 'Towards a Neurobiological Theory of Consciousness', *Seminars in the Neurosciences*, II, pp.263-275

Dennett, Daniel C. (1991), *Consciousness Explained* (Harmondsworth, Penguin)

Edelman, Gerald (1992), *Bright Air, Brilliant Fire* (Harmondsworth: Penguin)

Gazzaniga, Michael S. (1970), *The Bisected Brain* (New York: Appleton-Century-Crofts)

Marr, D. and Hildreth, E. (1980), 'Theory of Edge Detection', *Proceedings of the Royal Society of London, Series B*, CCVII, pp.187-216

McGinn, Colin (1989), 'Can We Solve the Mind-Body Problem?', *Mind*, XCVIII, pp.349-366

Nagel, Thomas (1971), 'Brain Bisection and the Unity of Consciousness', *Synthese*, XXII, pp.396-413

Nagel, Thomas (1974), 'What is it Like to be a Bat?', *Philosophical Review*, LXXXIII, pp.435-450

Nikolinakos, Drakon (1994), 'General Anesthesia, Consciousness and the Skeptical Challenge', *Journal of Philosophy*, XCI, pp.88-104

Penrose, Roger (1994), *Shadows of the Mind* (Oxford: OUP)

Popper, Karl and Eccles, John (1977), *The Self and Its Brain* (Springer International: Berlin)

Wilkes, Kathleen V. (1988), '_____, yìshì, duh, um, and Consciousness', A.J. Marcel and E. Bisiach (eds.) *Consciousness in Contemporary Science* (New York: OUP) pp.16-41

Wilkes, Kathleen V. (1991), 'The Relationship Between Scientific Psychology and Common-Sense Psychology', *Synthese* LXXXIX, pp.15-39