Should We Expect To Feel As If We Understand Consciousness?

Mark C. Price

Psychology Faculty, University of Bergen, Norway

We tend to assume that progress in answering the 'hard question' of consciousness will be accompanied by a subjective *feeling* of greater understanding. However, in order to feel we understand how one state of affairs arises from another, we have to deceive ourselves into thinking we have found a type of causal link which in reality may not exist (Rosch, 1994). I draw from and expand upon Rosch's model, which specifies the conditions under which this self-deceptive kind of causal attribution arises. I argue that the mind–body relationship may not meet these conditions, especially because of its potential novelty and uniqueness. We should *not* therefore expect to subjectively *feel* we understand consciousness.

1 Introduction

We would desperately like to understand, as the philosopher Thomas Nagel has put it, how on earth a subjectively experienced first person viewpoint on the world can emerge from a few pounds of subatomic particles (Nagel, 1986). But despite having exercised many of the best minds of science, the 'hard question' of how consciousness *per se* arises in the universe sadly continues to evade an answer. At best, all that objective science seems able to throw up are correlations between neurophysiological or cognitive descriptions of brain processes on the one hand and the occurrence of conscious experience on the other. But as causal explanations of consciousness, such correlations appear fundamentally insufficient. Chalmers (1995) is correct to point out that in all theories of consciousness to date there exists what Levine (1983) has termed a blatant *explanatory gap* between our objective accounts of brains and the fact that reality seems to include the existence of first person subjective viewpoints. This situation is repeatedly stressed in discussions on consciousness. A typical example is provided by Gray (1995, p. 9) who complains that 'there is no hint of a theoretical understanding of [the link between conscious experience and brain and behaviour] that would take us beyond brute correlation towards a "transparent" theory of causal connection' (my insertion).

Opinion varies as to whether and how the explanatory gap may be bridged. Some argue there can be no bridge at all (McGinn, 1989). Some suggest that epistemological innova-

Originally published in *Journal of Consciousness Studies* 3 (4), 1996, pp. 302–312. Reproduced with permission.

tion may be needed and that the gap may somehow be a function of our whole framework of inquiry and explanation (e.g. Nagel, 1986). And some, like Chalmers (1995) suggest that the gap may require an ontological solution which involves admitting new fundamental properties of the universe. What many approaches nevertheless share is the assumption that the occurrence of an explanatory gap is in itself problematic, that what is special about the hard problem is this unbridged gap. I would like to explore a slightly different perspective. I would like to suggest (1) that explanatory gaps are in fact ubiquitous in our causal explanations of the world, (2) that we are just very good at covering up these gaps, and (3) that what is special about consciousness is not the presence of a gap, but the fact that the gap just happens to be particularly obvious and difficult to obscure.

In other words I want to turn the hard question on its head. Instead of asking why consciousness suffers from an explanatory gap, I want to take a step back and ask why other aspects of our world-view do not seem to suffer from gaps of this sort. This will require looking at the psychology of understanding, and in particular at our perception of causation, of how we come to see one thing (e.g. consciousness) as arising from another (e.g. brains); what is it that allows us to walk away from a problem, having banished any explanatory gaps, with a smile on our face and warm glow in our hearts, feeling, 'Yes, I understand that now'? Two proposals will emerge from this examination. First, we may be placing too much importance on the presence of explanatory gaps. Second, we may be placing too much expectation on an answer to the hard question; *even if a scientifically acceptable answer of some kind were staring us in the face, we still might not feel as if we had understood.*

Let me begin by emphasizing a crucial distinction between, on the one hand, psychological accounts of our perception of cause and, on the other, philosophical accounts of causation 'as it really is', causation as an extensional relationship between states of affairs, what Hume (1740) called causation 'in the objects'. Psychological accounts of our perception of cause are concerned with what I shall refer to as our *everyday conception* of causality; i.e. with our intuitive lay-person's beliefs about the kind of beast that causes are. Psychological accounts are also concerned with whether we *perceive* explanatory gaps and with our *feelings* of understanding. In contrast, philosophical accounts of causation 'as it really is' are concerned with ontological statements about the world as we aim to describe it scientifically.

The gist of the argument I want to present, which is based heavily on a paper by Rosch (1994), is as follows. In order to subjectively feel as if we understand causally how a state of affairs in the world comes about, we need to think that we have located a cause of a type that corresponds to our everyday conception of causality. This concept of cause differs from causation 'as it really is', as described by philosophical accounts. In fact under some philosophical accounts causes themselves do not really exist. There may therefore be, in reality, an explanatory gap between a state of affairs that we are trying to explain causally (henceforth an 'outcome'), and the state of affairs from which it arose (henceforth a 'ground'). Such explanatory gaps may be ubiquitous. However we happen to be very good at hiding these gaps by deceiving ourselves into thinking we have located a cause of a type that corresponds to our everyday conception of causality. The conditions under which we can do this are described by psychological accounts

of our perception of causality. If we examine these conditions then the relationship between brains and consciousness turns out to be of a type that does not easily lend itself to such deception. *Feeling* that we understand consciousness is therefore going to be difficult.

To illustrate the argument I shall employ one example of a psychological account of causal perception, and one example of a philosophical account of causation. The choices of example are designed to make the argument as clear and as extreme as possible. Obviously other choices would have been possible, but I shall postpone until later a discussion of whether the argument rests or falls on the particular framework I use to illustrate it. As an example of a psychological account of causal perception, I shall use the recent account suggested by Rosch (1994), which is partly based on traditional Buddhist psychological texts. I choose this account because it seems especially relevant to the issue of when we should expect to perceive explanatory gaps. As an example of a philosophical approach, I shall use the regularity theory of causality which has its origins in Humean scepticism, and which challenges whether causes exist at all. This approach is certainly not universally accepted, but neither is it universally rejected and there are modern versions of it which have considerable mileage (e.g. Mackie, 1974). I choose this approach for two reasons, partly in the spirit of devil's advocate. First, it contrasts nicely with the claim that we need more than correlations. Second, it contrasts maximally with our everyday conception of causality and will therefore make the argument as clear as possible.

2 Our Everyday Conception of Causality

Under our everyday conception of causality, causes are themselves really thought to exist. When we seek to understand how one state of affairs (A) gives rise to another state of affairs (B), our natural tendency is to follow our folk psychological instincts in search of an 'A makes B happen' type of explanation (Rosch, 1994). The kernel of the explanation, or causal nexus, is believed to lie in some kind of necessary connecting link between A (the ground) and B (the outcome). Our *feeling* of understanding is a state of mind which depends on thinking we have located this causal nexus.

3 The Regularity Conception of Causality

In contrast there is at least one philosophical tradition, initiated by Hume (1740), that regards the causal nexus as a bankrupt concept (see for example Mackie, 1974; Rosch, 1994; Searle, 1983; Von Wright, 1974). It teaches that causes themselves do not exist. Like the pot of gold at the end of the rainbow, the causal nexus is a mythical creature that recedes as fast as it is approached. We cannot causally explain why certain states of affairs arise from certain other states of affairs in any necessary *a priori* sense.[1] In other words there is no hidden string between ground and outcome. The idea of a caus-

[1] In fairness to Hume, he did speculate that such necessity might still arise as qualities with which we are as yet unaquainted. For discussion of some of the subtleties of Hume's arguments, see Strawson (1989).

al nexus is in principle non-sensical because ground A and outcome B cannot at the same time be different from one another *and* account for each other. We seek to link two different states of affairs, but, by virtue of the very fact that they are different, an explanatory gap must remain between them. Causation 'as it really is' consists just of regularities in the relationships between states of affairs in the world (Hume's 'constant union and conjunction of like objects').

Of course modern regularity-based accounts of causation are more sophisticated than this. The cock may regularly crow every sunrise but clearly does not cause the sun to rise. Accidental regularities must therefore be distinguished from regularities that are 'necessary' in the counterfactual sense[2] that outcome B would not have occurred *unless* ground A had occurred. In addition, one has to distinguish grounds that are merely necessary (in the counterfactual sense) from those that are sufficient for a particular outcome. 'Causes' are usually complex; i.e. if outcome B is to arise, many grounds (X,Y,Z) may be needed. In practice we are rarely interested in all of these grounds, but rather in what Mackie (1974) refers to as INUS conditions — insufficient but necessary parts of an unnecessary but sufficient condition.

Regularity theories do have their problems, especially in distinguishing accidental from non-accidental regularities, but I shall gloss over these and merely refer the interested reader to modern defences such as those of Mackie (1974) and Smart (1993). For present purposes the main point is that regularity theories are an extreme example of the idea that there is no real cause 'in the objects' and that all we really have are correlations between states of affairs, albeit of a complex kind.

But if there are no causes, what are we doing as scientists? What are our causal explanations? Popper (1980) is helpful in pointing out that we do not need causes as such to do science. He argues that to talk about causes as showing how one state of affairs follows necessarily from other, is to use 'necessary' as 'a mere word. . . a label useful for distinguishing the universality of laws from "accidental" universality' (p. 438). What is important is the search for universal laws and the 'principle of causality' can be regarded as the metaphysical version of 'the simple rule that we are not to abandon the search for universal laws and for a coherent theoretical system. . .' (p. 61). 'To give a *causal explanation* of an event means to deduce a statement which describes it, using as premises of the deduction one or more *universal laws*, together with certain statements, the *initial conditions*' (p. 59).

So what we think of as causal relationships are merely instantiations of some universal regularity. To ask what the cause of something is, turns out to be a shorthand for asking for an expansion of its lawful relationship to other things (eventually to everything in the universe although we don't usually get that far!). And at least according to regularity accounts, we can look for these universal laws without recourse to 'real causes'. Perhaps we should not therefore be preoccupied with trying to pinpoint what it is about brains that actually *causes* consciousness. To Chalmers (1995), regularities are not

[2] Note that the counterfactual sense of necessity should not be confused with *a priori* necessity. Mackie (1974) provides a good discussion of these different senses of necessity.

enough. He demands 'an account of why and how' (p. 207). But good science need not imply that any causal explanation of consciousness be transparent. Perhaps we should after all be content with an empirical programme of research which gradually fills in the jigsaw of relationships between the physical brain and consciousness.

Chalmers (1995) argues that cognitive science is unable to bridge the explanatory gap between accounts of brain function and consciousness because cognitive science is geared to explaining functional properties, and consciousness is not (or is not just) a functional property. However from the preceding argument the problem is more fundamental than this. It is a basic problem to do with the unbridgeability of causal explanatory gaps in general. The real issue, which I turn to below, is why the particular explanatory gap surrounding consciousness is not hidden from us in the way that explanatory gaps often are. The solution that Chalmers adopts is to propose that consciousness be considered a new fundamental entity, in the same way that, for example, the electromagnetic force was introduced to account for the otherwise inexplicable behaviour of certain physical systems. Consciousness would be characterized by new basic principles which Chalmers says would 'ultimately carry the explanatory burden' (p. 210). But it is difficult to see how this approach bridges the gap. New fundamentals merely beg the question. Their origin still demands explanation. We still need to explain how, for example, the electromagnetic force comes about to start with. To a certain extent Chalmers recognizes this. He states, 'Of course, by taking experience as fundamental, there is a sense in which this approach does not tell us why there is experience in the first place' (p. 210). However, given his stated aim of bridging the explanatory gap surrounding consciousness, Chalmers does not seem to acknowledge the full extent to which this qualification undermines the whole force of his arguments.

4 Our Psychological Perception of Cause

If we accept the ubiquity of explanatory gaps, what prevents us from usually noticing these gaps? What makes us sometimes *feel* as if we have understood? How do we end up thinking we have found the kind of causal nexus that philosophical analysis considers problematic. To answer this we need to turn to theories of our psychological perception of cause, and in particular to the interesting analysis recently provided by Rosch (1994). Rosch's main argument is that there are various ways in which we deceive ourselves into seeing outcomes as already contained in their grounds. If, when we focus on ground A, we find outcome B in some sense already present within A, the explanatory gap between A and B will seem to disappear. This is because A and B are no longer seen as different. Under these circumstances, getting from state of affairs A to state of affairs B will not seem to require an explicit connecting link; we do not need to look for a cause because a causal link is implicitly contained within the fact that B is already in A. The illusion that a causal link is somewhere present then supports a feeling of understanding. In general events will only be perceived as coherent if they are seen to arise from themselves in this fashion.

From a logical point of view Rosch's proposal may appear incoherent and circular. Explaining the occurrence of B by pretending it is the same thing as A in the first place is no explanation at all. A proper scientific account of how anything new arises is re-

quired to be non-circular and show how a state of affairs arises from something that is not already itself. But we have to remember that this is a psychological account, not a logical one. What we are talking about here is an aspect of the way that nature or nurture has endowed us. According to Rosch all explanations that derive events from something other than themselves only come to feel like explanations because somewhere along the line they surreptitiously accomplish the trick of introducing the outcome itself; i.e. the account is turned into one that suffers from the hidden circularity of the outcome being already contained in the ground. Rosch argues that there are four basic types of situation in which outcomes are seen in this way as contained in their grounds. These are: (1) when a property is transferred from ground to outcome; (2) when we perceive an object or intend an action; (3) when grounds and outcomes are seen as the same entity, but transformed in some way; (4) when an outcome appears to be a property of a category to which the ground belongs.

If we accept Rosch's account of where our subjective sense of causal understanding comes from, any emerging jigsaw of relationships between the brain and consciousness should only generate a *feeling* of understanding if consciousness is seen to arise from its ground in some way. Let us therefore consider whether any of Rosch's four ways of making outcomes arise from grounds could accommodate a brain-consciousness relationship. In the following discussion I make the assumption that the brain is the ground which gives rise to the outcome of consciousness.

(1) Transfer of a property from ground to outcome:

An example is provided by mechanical causal events such as one ball hitting another, where the property of motion is transferred from ball A to ball B. If the object of interest in the outcome is the motion of ball B, the aspect of the outcome to which we are attending can be said to be already present in the ground. It is of note that Hume himself said that the perception of a causal relationship in motion is due to the relation of resemblance between ground and outcome.

For the hard question of consciousness, the property that we are crucially interested in is the first-person subjective nature of consciousness. However this property, which is the very essence of consciousness, does not seem to be obviously present in what we are taking to be the ground, i.e. brains. The transfer of a property from ground to outcome will not therefore help us to hide the explanatory gap surrounding consciousness. Indeed the inapplicability of this particular method of making outcomes appear to arise from their grounds succinctly characterizes the fundamental mystery of the hard problem.[3]

(2) Perceiving an object or intending an action:

[3] The method might nonetheless be applicable to explanations of the contents of consciousness: to the extent that there is any structural isomorphism between the contents of consciousness and the way information is physically (or functionally) represented in the brain, a causal relationship between the two may be perceived.

When we perceive something, we believe our perception (an outcome) to be similar to the object that gives rise to the perception (the ground). And our actions (an outcome) usually seem similar to our preceding intention to act (the ground). In both cases there is a natural tendency to perceive a causal relationship; we see objects as directly causing perceptions, and intentions as directly causing actions. To the lay person, the explanatory gap between an object and its perception, or between an intention and its resultant action, tends to be hidden. Because of this, commonsense explanations at a mentalistic level of description are often taken as a form of causal understanding and introductory psychology courses place much emphasis on pointing out that there *is* a problem of perception and action to be explained. However as psychologically educated scientists we are unlikely to have our inquisitiveness dampened by this particular lay instinct. In any case, the hard question of consciousness is not about how particular objects of perception or particular actions map onto particular *contents* of consciousness — it is about how our internal representations are conscious at all.

(3) Seeing grounds and outcomes as the same entity, but transformed in some way:

Typical examples include objects which change their location, age, form, etc. One sense in which consciousness and its ground, the brain, can be considered as the same entity, is perhaps captured by the mind–brain 'identity thesis'. The strongest version of the thesis, which is allied to eliminative materialism, states that subjective consciousness and the physical brain are entirely identical. To most people this position seems rather meaningless and difficult to maintain. Leibniz's Law of the Identity of Indiscernables requires that all properties of identical entities are shared, whereas the crux of the mind–brain problem is exactly that the physically described brain (or the functionally described cognitive system) does *not* share the crucial property of first person subjectivity.

In its weaker and more interesting sense the identity thesis proposes that consciousness and the physical brain are two different perspectives, or levels of description — inner and outer — that derive from the same ontological base. Consciousness is what it is like to be a brain 'from the inside', but consciousness and the brain are not fundamentally different types of stuff. (For a good elaboration of this view see Globus (1976); Popper & Eccles (1984) give further examples.) This version of the identity thesis has been labelled variously as 'parallelist', 'correspondence', or 'dual-aspect' theory. Chalmers (1995) himself embraces what he refers to as a 'double-aspect theory of information' in which it is information rather than matter which has two perspectives. By locating the mind–brain divide within a difference of perspective, all these varieties of the weak identity thesis attempt to escape the potentially problematic metaphysical implications of mind–brain dualism. But as Wimsatt (1976) warned, taking the identity thesis seriously involves at least a moderate kind of dualism, 'in that no identity theory will involve a satisfactory 'fit' unless it explains. . . the substantial and absolutely central differences between first and third-person (or subjective and objective) perspectives' (p. 232). In other words it is all very well to think of consciousness and its ground as the same thing viewed from differing perspectives, but this merely begs the question of how such radically differing perspectives can come about. Accepting the ontological identity of a ground and its outcome seems insufficient to hide the explanatory gap

between them when they are separated by so novel a transformation. (The importance of novelty is something I elaborate on below.) The identity thesis is more an attempt to avoid the problems of dualism than to explain how consciousness arises. It cannot in itself throw us off the scent of a causal nexus and lull us into a sense of understanding.

(4) Seeing an outcome as a property of a category to which the ground belongs:

An example would be to say that opium puts you to sleep because it belongs to the group of substances with dormative power. One way in which this type of manoeuvre can be used to apparently bridge the gap between physical brains and consciousness is provided by panpsychism — the view that consciousness is a property of all matter and that everything is in some sense conscious. According to panpsychism, the arising of consciousness in brains is not particularly problematic because (i) brains are a type of material object, (ii) material objects have the property of consciousness, and so (iii) brains have consciousness. No specific connecting link between brains and consciousness is required since consciousness (the outcome) is already found in the ground (brains). Of course panpsychists still need to explain why matter itself has consciousness. To most modern commentators this quandary seems as problematic, if not more so, than the problem of how brains could give rise to consciousness. In our contemporary culture panpsychism is therefore usually rejected as a serious candidate theory. However there are notable exceptions, such as Globus (1976) and Chalmers (1995). Chalmers, in an albeit disguised form of panpsychism, is willing to consider that conscious experience may be 'much more widespread than we might have believed, as information [of which consciousness is a fundamental inner aspect] is everywhere' (p. 217; my insertion). For anyone who *is* happy to accept consciousness as a property of all things (and in animistic cultures this is the norm), it becomes easier to take the mind–body relationship for granted. Accepting consciousness as a property of a general category, to which the brain belongs, generates at least a temporary feeling of having dealt with how brains give rise to consciousness. The problem has been pushed back to the ground and as long as we can keep our limited attentional capacity engaged on the mind–body relationship, we fail to notice the explanatory gap.

5 Similarity and Familiarity

So far it does not seem that any of Rosch's four ways of making events *seem* to arise from themselves will, for the majority of us, easily succeed in hiding the explanatory gap between consciousness and brain. Of course the account could be (and is likely to be) incomplete. But we can do better than arguing just that Rosch's prescriptions for hiding explanatory gaps do not apply too well to consciousness. We can use her framework to predict that consciousness will be a *particularly* difficult customer.

Consider again Rosch's last way of hiding explanatory gaps, namely seeing an outcome as a property of a category to which the ground belongs. This has two particular subclasses which Rosch does not mention. These occur very commonly, are capable of breeding the same impression of causal coherence, and are of particular relevance to the relationship between brain and consciousness. They are (1) situations which are in some way *similar* to other types of situation, and (2) situations which are in themselves

just very *familiar*.

First take similarity. Suppose that the relationship between a ground A and its outcome B has a degree of structural isomorphism with another ground–outcome situation, say the relationship between X and Y. Ground A can then be seen as a member of a more general category of states of affairs which have B-like outcomes. If the causal link between X and Y is felt to be understood, then this sense of understanding may transfer. A good example is provided by the computational metaphor of cognitive psychology: we think we know how computers work, brains are like computers, so we think that we know broadly how brains work. What we are talking about is the power of analogy, not just to provide a structural template for the shape of an explanation, but to obscure explanatory gaps and promote a *feeling* of understanding. Lakoff (1988) has argued that understanding is generally based on a group of structural templates which derive from the structure of our experience of our own bodies and our experience of social interaction. If this is so, these fundamental templates may themselves be prone to hiding explanatory gaps in the way that, according to Rosch, characterizes our folk understanding of perception and intentional action. Therefore transfer of a feeling of understanding, via the process of analogy, could be very general.

Next consider familiarity. Even if a relationship between A and B is *dissimilar* to other relationships we know of, it may still be *familiar* — perhaps the occurrence of the relationship is a very common event. If so, there is a sense in which B is a property of the category which is just the 'category of instances of A leading to B'. To the extent that instances of this category are familiar, and the category therefore also highly familiar, we may tend to see outcome B as a property of the category, not of the particular instance. Again there is a sense in which outcome B is contained in ground A, and we may be less likely to notice the absence of an explicit connecting link between A and B. Familiarity may therefore be another important way of obscuring explanatory gaps. If I cycle over broken glass, I really feel as if I understand why I get a puncture. This is because I am happy to let my explanation rest at a certain level, characterized by relationships I accept because they are common regularities and fit snugly into my world picture of what goes with what (knowledge about air in tyres, sharp objects, holes, pressure release, etc.). In this way every-day explanations leave us with the impression of getting to the causal nexus of events.

If similarity and familiarity help to obscure explanatory gaps, then when we encounter *unfamiliar* relationships which are also *unlike* anything else in our experience, the opposite will hold true. Explanatory gaps will be particularly obvious. Modern physics provides an excellent example of an area where this is commonplace, and where feelings of locating a causal nexus and of 'transparent' understanding consequently break down. All we are left with are mathematical regularities which defy direct visualization and expose a reality that seems fundamentally alien to the human mind (Davies & Gribbin, 1991).

Consciousness may find itself in a similar boat. It seems not improbable that the relationships between a physical system and consciousness may turn out to be so unique and unlike other types of relationship we have so far encountered, that neither analogy nor familiarity wil be able to lull us away from the search for a causal nexus. The expla-

natory gap will therefore continue to be glaringly obvious. Perhaps in time, as we become used to the relationships we have mapped, familiarity may breed a certain sense of understanding. But we should *not*, at least initially, expect to *feel* as if we understand.

6 Conclusion

In conclusion, whether we see explanatory gaps is at least partly to do with psychology, with what goes on in our head, and not just with the way the world *is* objectively. We must not ignore the psychology of the hard problem. Our ability to gauge the success of our scientific attempts to understand consciousness will be clouded if we do not consider the factors that govern our psychological perception of causal relationships, and, especially, the question of whether these factors prescribe that the relationship between brain and consciousness is likely to be of the type that lends itself to a psychological perception of causality.

We should not necessarily expect that progress in answering the hard question of consciousness will be accompanied by a *feeling* of understanding. As we search for an answer, we would therefore be unwise to depend on any intuitive sense of understanding as a guiding heuristic. In particular, those of us who are concerned with an experimentally based exploration of the relationships between consciousness and brain processes (whether at a neurophysiological or cognitive level of description) should not worry whether we are getting nearer to a feeling of understanding. Rather, we should be content to get on with the job of mapping the terrain.

Discussions of consciousness are often apologetic that their theoretical contributions do not appear to make inroads on the hard question. While I would not want to suggest that some humility in the face of the hard question is amiss, I think that this particular apology is unnecessary. For the reasons I have discussed, we should not necessarily expect a moment of blinding insight when we suddenly manage to get our heads round the problem. As far as the hard question is concerned, there is nothing second best about just chipping away at the relationships between brain and consciousness, and we do not need to apologize for doing it.

In order to avoid confusion, let me stress that I am not trying to argue any of the following. First, I am not trying to say anything about what form scientific theories should take, or denying any of the usual criteria for their success such as predictive power, falsifiability, coherence with other theories, parsimony, etc. Second, I do not wish to deny the existence of a hard question; even if we are 'only' looking for lawful regularities, our search must address the existence of conscious experience itself. Third, I am not denying that there is an answer to the hard question, or that we should be trying to find one. Finally, I do not want to belittle the very great problems presented by the hard question. It is one thing to plot the relationships between neurophysiology, cognition and consciousness experience, but for reasons that I discuss elsewhere (Price, 2001), it is an empirically non-trivial task to discover which brain processes are *necessary and sufficient* for consciousness.

My central point has been that the apparent adequacy of our progress in answering the hard question depends partly on psychological processes which mediate our percep-

tion of causality. In presenting this argument I have used the particular framework of Rosch's account of our psychological perception of cause on the one hand, and a regularity account of causation on the other. However I believe that the argument transcends this framework. Rosch's account is likely to be at best incomplete, but it importantly illustrates that our psychological reaction to what scientific exploration offers us can vary. It therefore cautions us to be wary about how we judge the fruits and direction of our scientific endeavour, and warns us that the hard question is not just the science, but also our reaction to the science. Similarly, a regularity account of causation may be contentious of causality, and factual accounts of causality 'in the objects'. I have taken an extreme account in order to make the contrast as clear as possible, but I suspect that other accounts will still not be able to offer us anything that completely matches our everyday concept of causation. This is because our everyday concept is not really a coherent one: how *can* two things that are different lead to each other? Like the notion of free will it is something that we are endowed with rather than a logical concept. And as with free will it is rather difficult to leave our everyday concept of causation behind. To climb out of our heads and adapt our everyday concept to fit our scientific or philosophical concept is far from easy.*

References

Davies, P C W, and Gribbin, J (1991). *The Matter Myth: Towards 21st Century Science*. London: Viking Press.

Chalmers, D J (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies* 2 (3), pp. 200–219.

Globus G G (1976). Mind, structure and contradiction. In *Consciousness and the Brain*, edited by G G Globus, G Maxwell, and I Savodnik. New York: Plenum Press.

Gray, J A (1995). Consciousness — What is the problem and how should it be addressed? *Journal of Consciousness Studies* 2 (1), pp. 5–9.

Hume, D (1740/1964). *A Treatise on Human Nature*, edited by L A Selby-Bigge. Oxford: Clarendon Press.

Lakoff, G (1988). Cognitive semantics. In *Meaning and Mental Representations*, edited by Umberto Eco *et al*. Bloomington: Indiana University Press.

Levine, J (1983). Materialism and qualia: the explanatory gap. *Pacific Philosophical Quarterly* 64, pp. 354–361.

Mackie, J L (1974). *The Cement of the Universe*. Oxford: Oxford University Press.

McGinn, C (1989). Can we solve the mind–body problem? *Mind* 98, pp. 349–366.

Nagel, T (1986). *The View from Nowhere*. New York: Oxford University Press.

Popper, K R (1980). *The Logic of Scientific Discovery*. New York: Routledge.

Popper, K R, and Eccles, J C (1984). *The Self and its Brain: An Argument for Interactionism.* London: Routledge & Kegan Paul.

Price, M C (2001). Now you see it, now you don't. Preventing consciousness with visual masking.

* Author's Note. I would like to thank Henry Howlett and especially Peter Grossenbacher for their extremely helpful comments on early drafts of this paper. I am also very grateful to Jonathan Shear and to the two anonymous referees whose critical comments helped to clarify my arguments.

In *Finding Consciousness in the Brain: A Neurocognitive Approach*. Advances in Consciousness Research 8, edited by P G Grossenbacher. Amsterdam: John Benjamins.

Rosch, E (1994). Is causality circular? Event structure in folk psychology, cognitive science and Buddhist logic. *Journal of Consciousness Studies* 1 (1), pp. 50–65.

Searle, J R (1983). *Intentionality*. Cambridge: Cambridge University Press.

Smart, J J C (1993). Laws of nature as a species of regularities. In *Ontology, Causality and Mind*, edited by J Bacon, K Campbell, and L Reinhardt. Cambridge: Cambridge University Press.

Strawson, G (1989). *The Secret Connexion: Causation, Realism and David Hume*. Oxford: Clarendon Press.

Von Wright, G H (1974). *Causality and Determinism*. New York and London: Columbia University Press.

Wimsatt, C W (1976). Reductionism, levels of organisation, and the mind–body problem. In *Consciousness and the Brain*, edited by G G Globus, G Maxwell, and I Savodnik. New York: Plenum Press.