Philosophical Studies (2006) 127:59–85
DOI 10.1007/s11098-005-1730-x
© Springer 2006

PHILIP ROBBINS and ANTHONY I. JACK

THE PHENOMENAL STANCE

1. INTRODUCTION

Cognitive science is shamelessly materialistic. It maintains that human beings are nothing more than complex physical systems, ultimately and completely explicable in mechanistic terms. But this conception of humanity does not fit well with common sense. To think of the creatures we spend much of our day loving, hating, admiring, resenting, comparing ourselves to, trying to understand, blaming, and thanking – to think of them as mere mechanisms seems at best counterintuitive and unhelpful. More often it may strike us as ludicrous, or even abhorrent. We are 'intuitive dualists' (Bloom, 2004). We view people and objects as fundamentally different sorts of things, to such an extent that the mere reminder of our physical nature tends to trigger such extreme reactions as laughter and disgust. The intuition that no mere physical mechanism, however intricate and sophisticated, could display the essential marks of personhood, has figured prominently in philosophical discussion of the 'mind–body problem.' A familiar form of reasoning runs like this: persons have property *P*; *P* has no merely physical instances; hence, persons are not merely physical beings. Pick your favorite value for *P* – consciousness, free will, rationality, what have you – and you're on your way to turning the water of intuition into the wine of an argument for dualism.

Arguments from intuition to the metaphysics of personhood, and the metaphysics of consciousness in particular, require a leap of faith, it is generally agreed. Yet what common sense tells us about the epistemology of personhood is

harder to resist: attempts to explain what it is to be human in purely mechanistic terms appear completely inadequate. Many philosophers approach this explanatory gap (Levine, 1983) between the physical and the personal in one of two ways: either they accept the force of the intuition and endorse dualism (Chalmers, 1996), or they reject the intuition and dismiss the apparent special difficulty of explanation in this area as an illusion (Dennett, 1991).

Philosophers sympathetic with the illusion approach have tried to explain away the gap intuition in various ways.[1] Churchland (1996) argues that the intuition is simply a by-product of the immature state of our knowledge about the brain mechanisms underlying conscious awareness and that, as this knowledge matures, the intuition will fade. In her view, the sense of mystery surrounding consciousness is of the same kind, and comes from the same source, as that surrounding any phenomenon that is not yet well understood by science. Other philosophers who adopt the illusion view disagree, suggesting that the sense of mystery surrounding consciousness is largely *sui generis*. For example, Tye (1999) argues on *a priori* grounds that it stems from the distinctive epistemic profile of phenomenal concepts (roughly, experience-dependent concepts of conscious mental states) in relation to physical concepts of the sort employed in neuroscience.

We favor a different approach. We agree with proponents of the illusion strategy that the intuition behind the explanatory gap should not be taken at face value. We agree with the proponents of this strategy who trace the source of the intuition to a distinctive feature of the explanatory domain at issue (à la Tye), as opposed to those who point to a perfectly domain-general factor like scientific ignorance (à la Churchland). On our view, the gap intuition is psychologically real *and* deep, and it calls for an explanation of its own. What we propose to offer is an empirically grounded account of how the intuition arises – an account that goes well beyond the idea that phenomenal concepts are epistemically irreducible to physical ones. We think that recent work in cognitive

neuroscience, and the neuroscience of social cognition in particular, may hold the key to such an account.

Our discussion divides into three main parts. In the first part, we 'naturalize' Dennett's distinction between the intentional stance and the physical stance by mapping it onto the distinction between mindreading and folk physics. Next, we update Dennett's scheme by adding a new twist: the phenomenal stance. The phenomenal stance corresponds to a component of social cognition that, like folk physics, is relatively independent of mindreading – namely, moral cognition. Next, we discuss how our evolved capacities for thinking about the physical and social worlds interact and help to generate intuitive dualism, and we specifically identify moral thinking as generating intuitions about the non-physical nature of consciousness. The opposition between the perspectives associated with these different capacities helps to explain the explanatory gap; yet it may be possible and useful to have these capacities work in concert to generate a more unified world view. We close by pointing to cultural influences on the opposition and touching on some implications of this work.

2. TAKING A STANCE

Humans are deeply social creatures – more so, perhaps, than any other species. The peculiar depth of human sociality stems at least partly from the fact that we coordinate our behavior with others by construing them as intentional systems, that is, as creatures whose behavior is governed by their beliefs, desires, intentions, emotions, and other contentful mental states. When dealing with our fellows, we adopt what Dennett calls the 'intentional stance':

> Here is how it works: first you decide to treat the object whose behavior is to be predicted as a rational agent; then you figure out what beliefs that agent ought to have, given its place in the world and its purpose. Then you figure out what desires it ought to have, on the same considerations, and finally you predict that this rational agent will act to further its goals in the light of its beliefs. A little practical reasoning from the chosen set of beliefs and desires will in many – but not all – instances

yield a decision about what the agent ought to do; that is what you predict the agent *will* do. (Dennett, 1987, p. 17)

In psychology, this mode of understanding goes by various names: 'mindreading' (Baron-Cohen, 1995), 'mentalizing' (Frith, 2003), and 'theory of mind' (Premack and Woodruff, 1978). All of these terms denote the capacity to ascribe intentional mental states and to predict and explain behavior on the basis of those ascriptions.[2]

To get clearer about what the intentional stance comes to, it helps to consider the empirical tests used to gauge subjects' mindreading ability. The best known of these tests involve false-belief tasks, whose successful performance requires distinguishing one's own beliefs from the beliefs of another. A classic example is the Sally-Anne task, in which subjects observe a puppet show featuring two characters, Sally and Anne (Frith, 2003). The plot goes like this. Sally has a basket and Anne has a box; Sally puts her marble into the basket and leaves; while Sally is away, Anne removes the marble from the basket and puts it into the box. The subject is then asked where Sally will look for her marble when she returns. In order to generate the right prediction, the subject must refrain from assimilating Sally's epistemic state to her own. This is a first-order false-belief task, in that it involves thinking about another person's mistaken belief about the world. In second-order tasks of the same type, subjects are queried about stories in which one character has a mistaken belief about another character's mental state. In both cases, what's at issue is the capacity to represent another person's intentional states and to use that information in making sense of their behavior.

Though successful performance on false-belief tasks is widely viewed as a litmus test of mindreading (Singer et al., 2003), tasks of this sort comprise only a smallish subset of the full battery of tests used by clinicians to gauge this ability. Other examples include tests of understanding that seeing leads to knowing; tests of the production and comprehension of deception; tests of understanding metaphor, sarcasm, and irony; tests of conversational pragmatics; tests of the production of

spontaneous pretend play; tests of the ability to recognize facial expressions of emotions and other intentional states; tests of understanding complex causes of emotion, such as beliefs; and tests of understanding stories involving bluff and double bluff (Baron-Cohen, 2000b). This diversity supports the idea that competence at mindreading rests upon possession of a wide range of more specialized cognitive abilities, all of which involve the attribution of intentional states.

In introducing the intentional stance, Dennett emphasizes the contrast between it and a fundamentally different way of thinking about the world, which he dubs the 'physical stance':

> From this stance our predictions are based on the actual physical state of the particular object, and are worked out by applying whatever knowledge we have of the laws of nature. Instances of prediction from the physical stance are common enough: 'if you turn on that switch you'll get a nasty shock,' and 'when the snows come that branch will break right off.' (Dennett, 1981a, p. 45)

Although less in vogue than the intentional stance, the physical stance has also been an object of study in psychology. It has been studied directly in investigations of 'folk physics' (Baron-Cohen, 2000a), 'intuitive physics' (Boyer and Barrett, 2005), and 'theory of body' (Leslie, 1994).[3] It also been studied indirectly, in the context of investigations of the intentional stance. The comparison or control tasks for many studies of mindreading have involved physical-causal reasoning, such as: understanding stories and picture sequences depicting mechanical events (Baron-Cohen et al., 1986); predicting the behavior of gears, levers, pulleys, and other mechanical systems (Baron-Cohen, 2000a; 2002); and trying to figure out how a computer might be calculating its responses in the children's game Rock-Paper-Scissors (Gallagher et al., 2002). Another example is the false-photograph task, an analogue of the Sally-Anne task that requires understanding of how a camera works. Here the role of Sally is played by a mechanical device: a photograph is taken of the initial location of the marble, prior to its displacement, and the subject is asked to specify the location of the marble as depicted in the photograph (Leslie and Thaiss, 1992).

Folk physics also covers implicit knowledge of core principles governing the dynamics of physical objects, such as the principle of continuity, which says that every moving object traces exactly one spatiotemporally connected path; and the principle of contact, which says that two objects influence each other's movement only if they touch (Spelke, 1988). Unlike more sophisticated forms of physical knowledge, knowledge of these principles seems to emerge very early in development – as early as 4 months, judging from the results of looking-time experiments by Spelke and others. In these experiments, young infants habituated to an initial test display look longer at a variant display that violates some physical principle (say, the principle of continuity) than they do at a different variant that conforms to the principle. This suggests that infants conceive of physical objects as the sort of thing that exhibits certain regularities, and they are surprised when they encounter something that looks initially like a physical object but behaves otherwise.

Dennett's philosophical distinction between the physical and intentional stances has a lot going for it. What's more, there's firm empirical support for the idea that folk physics and mindreading – the scientific analogues of these two stances – are relatively autonomous capacities. For one thing, folk physics and mindreading appear to be subserved by different bits of neural circuitry. Preliminary evidence from functional neuroimaging (fMRI) studies of subjects observing social and mechanical vignettes suggest that causal-mechanical concepts are localized in the medial fusiform gyrus and middle temporal gyrus, whereas social-intentional concepts are localized in a disjoint set of regions: lateral fusiform gyrus, superior temporal sulcus, ventromedial prefrontal cortex, and the amygdala (Martin and Weisberg, 2003). A second point is that folk physics and mindreading are dissociable in at least one direction, and probably both. This is suggested by studies of autism, a pervasive neurodevelopmental disorder marked by poor performance on the standard battery of mindreading tests together with systematic losses of social and communicative functioning. Despite their

difficulties in the social realm, many autistics excel at tests of folk physics, including the false-photograph task (Baron-Cohen, 1997, 2000a, 2002, 2003; Baron-Cohen et al., 1999). In addition, autism is strongly heritable, and the parents of autistic children are overrepresented in professions like engineering, where superior folk physics is an asset, if not a necessity. The clinical profile of autism, in which social cognition is impaired but causal-mechanical cognition is intact, makes it likely that the physical stance is dissociable from its intentional counterpart. Research on Williams syndrome, a disorder in which visuo-spatial cognition is severely impaired but social, communicative, and verbal skills are relatively preserved, suggests a dissociation in the opposite direction as well (Bellugi et al., 2000).[4]

These points add empirical substance to the claim that the intentional and the physical stances correspond to very different, even disjoint, modes of construing the world – an idea that lies at the heart of intuitive dualism. Human beings are, after all, intentional systems *par excellence*. As such, they do not always behave like physical objects of the sort described by folk physics. For example, in Michotte's (1963) 'billiard ball launching' event, one ball rolls toward a second ball, strikes it sending it off in the opposite direction, and stops at the point of impact. This is a classic illustration of how physical objects are supposed to behave: causation on contact and no action at a distance. But people don't play by that rule. People can and very often do influence each other's movement without touching, for example, when tracking, chasing, or playing a game of tag. Since the brain categorizes something as a physical object partly in virtue of the fact that the principle of contact applies to it, and the principle does not apply to persons, this makes it natural for common sense to exclude persons from the category of physical objects. Developmental evidence supports this hypothesis. Looking-time studies of 7-month-old infants, for example, suggest that they do not apply the contact principle when reasoning about events involving people, though they do apply it to a parallel scenario involving inanimate objects (Spelke et al., 1995).

A similar story can be told about the principle of continuity. A recent looking-time study of 5-month-old infants hints at the possibility that they do not expect people to conform to the principle of continuity, though they do expect this of quotidian physical objects (Kuhlmeier et al., 2004).[5] In this case, unlike the case of contact, subsequent development builds in the parallel expectation about people. Yet its apparent absence in early infancy illustrates early developmental pressure for these two perspectives to develop as mutually exclusive – a developmental trajectory that may mature into a steadfast belief in dualism.

3. BEYOND THE INTENTIONAL STANCE

Mindreading matters a lot to the success of creatures like us. Without it, our ability to interact socially with others, whether for purposes of competition or cooperation, would be gravely compromised. This is clear from case studies of autistic individuals, who experience social isolation or marginalization as a result of their disorder (Frith, 2003). Minimal competence at mindreading appears to be essential for successfully negotiating the human social world – as well as for grasping the essence of what it is to be human. But it is not sufficient. Some individuals who are very good at mindreading can nonetheless alienate themselves from society in a quite spectacular manner, apparently because of a different sort of failure to understand human nature.

Psychopathy is a personality disorder marked by an array of cognitive and affective deficits, including lack of concern for others' distress, lack of remorse, shallow affect, impulsivity, pathological lying, and poor behavioral controls (Hare, 1991). We will focus on the first of these markers: lack of concern. This response can be measured physiologically. Normal subjects presented with images of crying and frightened babies, or watching actors feigning sorrow or fear, show clear bodily signs of autonomic arousal, such as increased skin conductance; psychopaths show significantly less autonomic arousal when shown the same images (Blair et al., 1996).

The lack of concern shown by psychopaths suggests a specific deficit in their empathetic responses to others. But we must be careful before we make any blanket claim that psychopaths lack empathy. The concept of empathy is tricky to define (Preston and de Waal, 2002). Some of the difficulty may stem from the fact that the term 'empathy' is ambiguous. Frith (2003) usefully distinguishes between two kinds of empathy, 'instinctive' and 'intentional.'[6] Instinctive empathy is "a basic emotional response that just spills out" (Frith, 2003, p. 11), and it does not involve any sophisticated mindreading. Intentional empathy, by contrast, is affectively neutral and does depend on mindreading: it involves understanding the source of the other person's distress and producing an appropriate behavioral response based on that understanding. Though both kinds of empathy – 'hot' and 'cold', as it were – are present in normal individuals, they appear to be functionally distinct capacities.

Frith's distinction bears directly on the issue of lack of empathy in psychopathy. As its name suggests, intentional empathy is a component of the intentional stance. And there is considerable evidence that psychopaths excel at intentional-stance-taking. In a study testing advanced (second-order) mindreading, subjects were presented with stories like the following:

> During the war, the Red army capture a member of the Blue army. They want him to tell them where his army's tanks are; they know that the tanks are either by the sea or in the mountains. They know that the prisoner will not want to tell them, he will want to save his army, and so he will certainly lie to them. The prisoner is very brave and very clever, he will not let them find his tanks. The tanks are really in the mountains. When the other side ask where his tanks are, he says, "They are in the mountains."

Subjects were then asked to explain why the prisoner gave this reply.[7] In this study, psychopaths performed marginally *better* than non-psychopathic controls (Blair et al., 1996). And in a more recent study using the 'Reading the Mind in the Eyes' test, in which subjects attribute intentional states to people based solely on photographs of the eye region, psychopaths did as well as non-psychopaths (Richell et al.,

2003). There is also abundant clinical and anecdotal evidence that psychopaths are skilled at psychologically manipulating others for their own benefit – the sort of Machiavellian skill that rests on sophisticated mindreading (Hare, 1993). Thus, while 'hot' empathy is impaired in psychopathy, 'cold' empathy may be more or less intact. The opposite situation obtains in autism. Since autistics are poor mindreaders – for example, they do badly on both of the mindreading tasks just mentioned (Happé, 1994; Baron-Cohen et al., 1997) – they tend to lack intentional empathy; but they exhibit normal physiological arousal to others' distress (Blair, 1999). As far as lack of the two kinds of empathy is concerned, autism and psychopathy are complementary disorders.

But it isn't just that psychopaths appear to lack a certain sort of feeling, and the motivation that feeling brings along with it. It isn't just that their affective structure leads them to desire, and thus to choose, more manipulative courses of action and less compassionate acts. There is also something they fail to *understand*, namely, the difference between right and wrong. Indeed, one of the most consistent findings is that psychopaths are at sea in the domain of morality. A standard test of competence in this domain consists in subjects' capacity to distinguish between moral and conventional transgressions (Turiel, 1983). Moral transgressions, such as punching someone or picking their pocket, violate rules protecting the rights and welfare of others. Conventional transgressions, such as cross-dressing or jaywalking, violate rules protecting more or less arbitrary features of the social order. The judgments of normal subjects regarding these types of transgressions exhibit the following features (Blair, 1995):

Moral transgressions are judged to be more serious than conventional transgressions.

Moral transgressions are judged to be impermissible even if there is no rule prohibiting them, whereas conventional transgressions are generally judged to be permissible in the absence of a prohibition.

Moral transgressions are judged to be impermissible even if a relevant authority figure suspends the prohibition against them, whereas conventional transgressions are generally judged to be permissible under those circumstances.

This is the pattern observed cross-culturally in normal children (as of 39 months) and normal adults. Likewise, autistics reliably distinguish between moral and conventional transgressions (Blair, 1996). Psychopaths, on the other hand, fail to make this distinction; they are morally incompetent (Blair, 1995).

Preliminary results from neuroimaging studies of normal subjects are consistent with the idea that mindreading and moral judgment tasks engage distinct, albeit overlapping, brain networks. In particular, the amygdala and orbitofrontal cortex (BA 10/11) may be particularly involved in processing emotional information relevant to moral judgment (Moll et al., 2001; 2002; 2003; Greene and Haidt, 2002), although full-blown moral reasoning would also require medial frontal and paracingulate areas specifically associated with the intentional stance (Gallagher et al., 2002).

These are telling results. They point to the existence of a core component of the normal understanding of what it is to be human – spared in autism but absent in psychopathy, and subserved to some extent by special-purpose neural circuitry – that is relatively independent of mindreading and the intentional stance. Tipping our hats to Dennett, we'll call it the 'phenomenal stance.'

4. FROM PHENOMENALITY TO MORALITY

The concept of the phenomenal stance is best approached by analogy with its intentional counterpart. To adopt the intentional stance toward *X* is to understand *X* as an intentional system, that is, to regard *X* as a locus of intentionality. This entails ascribing intentional states (beliefs, desires, intentions, etc.) to *X* and using those ascriptions to make sense of *X*'s behavior. Analogously, to adopt the phenomenal stance toward *X* is to understand *X* as a 'phenomenal system,' that is, to regard *X* as a locus of phenomenal experience.

Part of what it is to regard something as a locus of experience is ascribing phenomenal states (emotions, moods, pains, visual sensations, etc.) to it. But this involves more than mere

rote ascription of phenomenal states; it requires a felt appreciation of their qualitative character.[8] For example, if you don't know what it's like to feel sad, you can't understand what it is to feel sad. And if you can't understand what it is to feel sad, you can't regard something as feeling sad – at least, not in the full-blooded way that the phenomenal stance requires.

Another part of regarding something as a locus of experience is a kind of emotional sensitivity. Phenomenal states typically have some hedonic value for the bearer, either positive (e.g. pleasure, joy) or negative (e.g. pain, sadness), and when we detect them in others, we tend to react accordingly. It is usually pleasant to observe another's pleasure, and distressing to observe their distress. The capacity for such responses, what we earlier discussed under the heading of 'hot' (i.e., instinctive) empathy, is essential to the phenomenal stance.

Third, regarding something as a locus of experience involves regarding it as a potential target of moral concern. Phenomenal states, and phenomenal consciousness in general, have moral value for the observer. We care about what others are experiencing, and we care about it for its own sake. For example, part of what it is to regard something as a locus of experience is to treat it as something that one ought to shield from harm. Viewed from the phenomenal stance, a creature in distress demands our attention and mandates our best efforts to alleviate that distress. A creature facing the threat of injury (hence, pain) or death (hence, the cessation of experience altogether) engages our attention in much the same way. Aversive phenomenal states like pain, sadness, and fear, as well as phenomenal experience more generally, are primitively morally compelling. As Singer (1976) puts it, following the lead of Bentham (1789), it is axiomatic that "if a being suffers, there can be no moral justification for refusing to take that suffering into consideration." The point can be strengthened: if a being has phenomenal experience of *any* sort (i.e. if there is something that it is like to be that being), we are morally obliged to take that experience into

consideration. In this sense, phenomenality and morality are tightly linked. Grasping this linkage – not just intellectually, but emotionally (in one's bones, as it were) – is an integral component of the phenomenal stance.

In this respect, our concept of the phenomenal stance picks up on an idea in one of Dennett's early writings on the intentional stance:

> One is guilty of no monstrosities if one dismembers the computer with whom one plays chess, or even the robot with whom one has long conversations. One adopts the intentional stance toward any system one assumes to be (roughly) rational – to adopt a truly moral stance toward the system (thus viewing it as a person) might often turn out to be psychologically irresistible given the first choice, but it is logically distinct. We might, then, distinguish a fourth stance, above the intentional stance, called the *personal stance*. The personal stance presupposes the intentional stance – and seems, to cursory view at least, to be just the annexation of moral commitment to the intentional. (Dennett, 1981b, p. 240)

We applaud Dennett's focus on the moral dimension of personhood, and his distinction between the personal and the intentional – a distinction neatly illustrated by the case of psychopaths, who combine expert mindreading with moral blindness. But the differences between the personal stance and the phenomenal stance are noteworthy. First, the concept of the phenomenal stance is explicitly tied up with the concept of consciousness, whereas the concept of the personal stance is not. Phenomenal systems, unlike Dennettian persons, are essentially subjects of conscious experience. Second, the concept of the personal stance is tied up with the concept of intentionality, whereas the concept of the phenomenal stance is not. The personal stance "presupposes" the intentional stance, in the sense that one cannot regard something from the personal stance without also regarding it as an intentional system. This is because, for Dennett, "the solicitude of a gardener for his flowers, or for that matter, of a miser for his coins, cannot amount to moral commitment, because of the absence of the intentional" (1981b, p. 241). The phenomenal stance, by contrast, does not depend on the intentional stance. This is suggested by studies of autism,

where mindreading is impaired but core moral understanding is intact.

A final point about the phenomenal stance concerns its special role in mediating social interaction. The phenomenal stance is geared largely toward affiliation, the primary motor of which is instinctive empathy. The experience of empathy – what Suzanne Langer described as an "involuntary breach of selfhood" (Dolan, 2002) – affords a kind of primitive fusion of perspectives, a breaking down of the boundaries between persons. The feeling of being intimately linked to a consciousness distinct from one's own is something we value for its own sake, not as a means to some further end. The intentional stance, by contrast, is primarily pragmatic and strategic. It enables one to better compete with other individuals for limited resources, and its primary value is instrumental. This contrast between the two stances helps to explain the profile of psychopaths as 'social predators,' incapable of friendship or love (Hare, 1993).

By this point it should be passably clear what the phenomenal stance is, and how it differs from the intentional stance. We'll now turn to consider the implications of this concept for thinking about the psychological origins of the explanatory gap.

5. THE EXPLANATORY GAP AND THE BALKANIZED BRAIN

We began this paper by considering a classic philosophical question: why is it so hard to imagine the possibility of a mechanistically adequate account of persons? Relatedly, why does there seem to be an explanatory gap where mentality, especially conscious mentality, is concerned? Some philosophers have found their answers to these questions in metaphysical speculation: it is hard to imagine a mechanistically adequate account of persons because persons are not just mechanisms. By corollary, there seems to be an explanatory gap because mechanistic accounts really can't explain everything that happens in the mind. These answers are intuitive,

perhaps even 'natural.' Yet, as many have realized, there are powerful reasons to resist them. Here we make no attempt to rehearse those reasons, or to add to them. Rather our aim is to help arm the resistance by supplying a fuller and more complete – and empirically grounded – account of the gap intuition, or at least the outline thereof.

Let's begin by recapping the major themes of the discussion so far. The most general theme is that human beings are equipped with a range of disparate modes of construing, or stances toward, the world and its objects. Among the most familiar and best studied of these stances are the physical stance, which parses the world in spatiotemporal and causal-mechanical terms, and the intentional stance, which parses the world in terms of beliefs, desires, intentions, and other contentful mental states. The cognitive capacities corresponding to these stances – folk physics and mindreading – are distinct in at least three dimensions. First, they are *semantically* distinct, insofar as they involve the ascription of different, and occasionally inconsistent, sets of properties to objects. Second, they are *functionally* distinct, as suggested by evidence of double dissociation drawn from studies of autism (physical > intentional) and Williams syndrome (intentional > physical). Third, preliminary functional imaging studies suggest that these stances are *neurally* distinct, that is, subserved by different bits of neural circuitry. These differences jointly give rise to a problem of cognitive integration: they make it cognitively difficult and demanding – unnatural, to a certain extent – for us to bind information from the physical and intentional domains.[9] The result is a kind of Balkanization, with competition between the two systems making it difficult for us to grasp at an intuitive level how one and the same thing could be both a bearer of beliefs and desires and a purely physical mechanism. This goes some distance toward accounting for the widespread intuition of an explanatory gap concerning the intentional mind – and with it, the intuitive attraction of dualism in this area.

Of course, the more prominent gap intuition concerns the phenomenal mind, not the intentional one. Or, to put it another way, it does not seem that our intuitive difficulty

with seeing persons as mere mechanisms can be accounted for just by the tension that arises between the physical and the intentional stances. To make explanatory headway here, we suggested that humans are equipped with a third mode of construing the world and its objects: the phenomenal stance.

The phenomenal stance parses the world in terms of emotions, pains, pleasures, and other conscious mental states. Like the physical and intentional stances, the phenomenal stance is made up of a variety of subcapacities, chief among them the capacity for empathy and moral understanding. The latter capacities are implicated in phenomenal-stance-taking because they are automatically triggered upon detection of phenomenal states in another agent. In other words, conscious experience is primitively morally compelling.

We noted above that the intentional stance is semantically, functionally, and neurally distinct from the physical stance. The phenomenal stance is distinct from its intentional and physical counterparts in all of these ways, though it lies at a greater remove from the physical than the intentional. At the semantic level, the phenomenal stance differs from the physical and intentional stances in virtue of the difference between phenomenal properties, on the one hand, and physical and intentional properties, on the other. Unlike physical properties, however, phenomenal and intentional properties constitute contiguous subdomains of the mental, so the semantic distance between phenomenal and physical is greater than that between phenomenal and intentional. At the cognitive level, the phenomenal stance is doubly dissociable from the intentional stance, as shown by studies of psychopathy (intentional > phenomenal) and autism (phenomenal > intentional). Though the parallel hypothesis regarding the phenomenal vis-à-vis the physical has not been directly tested, there is no evidence, and no other reason to suspect, that folk physics is impaired in psychopathy (physical > phenomenal). Likewise, albeit somewhat more tentatively, there is no evidence that core moral understanding is impaired in Williams syndrome (phenomenal > physical).[10] This gives us *prima facie* reason to suppose that the phenomenal and physical stances are also

dissociable. Finally, at the neural level, each of the three stances is associated with a distinct network of brain areas. In the case of the phenomenal and intentional stances the networks overlap substantially, but this is not true of either of these stances relative to the physical stance. Hence, the neural distance between the phenomenal and the physical exceeds that between the phenomenal and the intentional.

What this all adds up to is a further Balkanization of the brain, with all three stances – physical, intentional, and phenomenal – competing amongst each other for resources. We suggested earlier that in the case of the physical and intentional stances, this competition impedes the effective integration of information across the physical and intentional domains, thereby fostering the intuition of an explanatory gap concerning intentionality. What we're suggesting now is that the same point applies, *mutatis mutandis*, to the physical and phenomenal stances.[11] Like the intuition of a gap concerning intentionality, the intuition of a gap concerning phenomenality stems at least in part from the fact that our brains are configured in such a way as to naturally promote a dualist view of consciousness.

At the heart of our account of why consciousness seems to defy physical explanation is the idea that thinking about consciousness is essentially linked to, indeed partly constituted by, the capacity for moral cognition. Making this link enables us to provide empirical justification, from clinical psychopathology and functional imaging of normal subjects, for the claim that the human brain contains a special-purpose mechanism for thinking about consciousness, alongside special-purpose mechanisms dedicated to intentionality, physics, and other cognitive domains. But it also helps to motivate our story about the explanatory gap at a commonsense level.

Here's how the argument goes. First, to think about something as conscious is to think about it as a locus of moral concern. Second, thinking about something as a locus of moral concern is naturally opposed to thinking about it as a mere physical object (however complex in structure), since a mere physical object is the wrong sort of thing to attract

moral concern. As Bloom (2004) notes, to regard a person merely as a physical object – for instance, as a mere object of physical disgust – is to exclude that person from "the moral circle."[12] Why is that? The sources of this opposition are undoubtedly multifarious and complex, but part of the story seems to be that moral concern is a non-instrumental attitude *par excellence*, whereas our attitude to mere physical objects is essentially an instrumental one. In Kantian parlance, to think of something as a moral patient is to think of it as an 'end in itself,' not merely as something to be manipulated for one's own purposes; to think of something as a mere physical object, on the other hand, is to think of it 'merely as a means' to an end. However one explains it, the conceptual tension between moral patiency and mere physicality, combined with the conceptual affinity between consciousness and moral patiency, contributes to the tension between thinking about something as conscious and thinking about it as a physical mechanism. And this further tension is what drives the intuition of an explanatory gap regarding consciousness.

An objection to this argument is: couldn't a psychopath have the gap intuition – and if so, wouldn't that show that the intuition has nothing to do with moral competence? First, it is important to realize that it is an open empirical question whether psychopaths would intuit the gap. It's far from obvious that they would. This is not a matter to be settled from the armchair. Unless you are a psychopath, your intuitions will not be informative: we simply don't have introspective access to our own cognitive architecture (Jack and Roepstorff, 2002). Second, we never sought to claim that moral capacity represents the whole of the phenomenal stance, only that it constitutes a significant part of it. At bottom, to adopt the phenomenal stance towards something is just to regard that thing as a locus of phenomenal experience, full stop. In the normal case, to understand something as a locus of experience is to respond empathically to it in relevant situations (e.g. distress displays), and to regard it as worthy of moral concern. In pathological cases, such as psychopathy, both of these components are absent. But that is

not to say that psychopaths are incapable of adopting the phenomenal stance; it's just to say that they can adopt only an attenuated version of it. So the story on offer predicts that psychopaths would feel the pull of the gap intuition less keenly than normals, not that they would be immune to it.

To see why, note that we have not argued that the opposition between the phenomenal and physical stances is due entirely to the conceptual affinity between phenomenality and moral patiency, on the one hand, and the tension between moral patiency and mere mechanism, on the other. What we've argued is that the relations among these concepts play an important – and hitherto neglected – causal and explanatory role vis-à-vis the gap intuition. This is not to say that the opposition between the phenomenal and physical stances has no other sources. On the contrary, it seems likely that the full story of this opposition will be complicated. It probably includes, for example, the fact that we think of consciousness as a 'subjective' phenomenon, in the sense that it cannot be publicly observed, whereas we think of physical objects and properties as 'objective,' or in the epistemic public domain (Nagel, 1974). It may also include the fact that we think of phenomenal states, but not physical ones, as introspectively knowable, ineffable, and non-spatial. And the list of differentiating features goes on.

A key point is to realize that the phenomenal stance is a cognitive capacity whose nature and form have been shaped by evolutionary pressures. The phenomenal stance evolved for the purpose of aiding and maintaining affiliations. In other words, our capacity to understand the experience of others, in the particular way that we 'normals' do, was shaped by an evolutionary pressure for moral behavior and understanding. The fact that some individuals lack moral behavior and understanding does little to impugn the critical role these played in determining the nature of that capacity. Compare: the structure of chicken wings can be best explained by the fact that wings evolved to fly, even though chickens don't.

So, in short: are psychopaths capable of thinking about something as a subject of experience? Probably so. Are they capable of thinking about something as a subject of experience in the same way normals do? Probably not. But our thesis does not entail or assume otherwise.

6. CLOSING THE GAPS

In characterizing the gap intuition as the product of a Balkanized brain, we have been stressing the point that it results from more or less hard-wired restrictions on information flow across competing neural networks. But we do not mean to imply that this neural configuration, or the cognitive compartmentalization that goes with it, is impervious to change. On the contrary, culture and education play a vital shaping role here. For example, many psychologists find it hard to even understand the gap intuition, perhaps because they are so familiar with thinking of persons interchangeably in personal and causal-mechanical terms.[13] Many naturalistic philosophers also say they cannot feel the intuition, though probably for a different reason – they cannot feel it because their ability to detect the intuition is swamped by their commitment to naturalism. Clearly cultural factors, such as schooling in religious or supernaturally oriented world views, or in naturalistic philosophy, can greatly influence our willingness to accept different metaphysical positions. Nonetheless it does not seem that mere adherence to the naturalistic idiom renders us immune to the intuitive pull of the explanatory gap. Many scientists appear to be talking about an explanatory gap when they write about the problems that a theory of consciousness should solve. Thus, when Chalmers (1996) described many scientific accounts of consciousness as engaging in a bait and switch, the charge struck many as accurate.

By contrast, the intuition concerning a gap between the intentional and the physical (also known as 'Brentano's problem') seems to have lost much of its popular appeal. We would suggest that this is because we frequently interact with

a class of objects that we think of in both physical and intentional terms, namely, computers.[14] The gap intuition regarding phenomenality seems more robust. Yet this too may be starting to change, given the growing popularity of general-audience books on the neuroscience of emotion (Ledoux, 1998), pain (Wall, 2002), love (Fisher, 2004), and other sectors of consciousness. We predict that another factor contributing to the ongoing revision of dualistic thinking about consciousness will be neuroscientific research on moral cognition (Damasio, 1994; Greene and Haidt, 2002). The basic principle at work here is simple: on the naturalization front, morality and consciousness march together. The more we accustom ourselves to thinking of moral cognition in physical terms, the more comfortable we are likely to be with thinking of morality – hence, with thinking of consciousness – as part of the natural order. This is one of the more important implications of our proposal concerning the phenomenal stance.

The main implication of our proposal, however, concerns how researchers should think about the explanatory gap. Despite its deep roots in human cognitive structure, the conceptual tension between personhood and mechanism need not have any echo in reality.[15] Some may regard our thesis as troubling – the latest installment in the continuing disenchantment of the world, fraught with dehumanizing perils (Wolfe, 2000). We see the situation in a more positive light. For us, the scientific enterprise promises to yield a richer, more unified, and more robust conception of our selves and our place in the world (Jack and Robbins, 2004). Progress in our self-understanding as a species will be best served not by mystery mongering, but by identifying and disarming whatever epistemological and metaphysical prejudices the brain might throw in the way of that understanding. In this paper we hope to have advanced that cause a bit further.[16]

NOTES

[1] It is important to distinguish the intuition of an explanatory gap from the explanatory gap itself. That we intuit a gap in explanatory space where

consciousness is concerned does not show that explanatory space actually contains a gap there; all it shows is that we think of this space as containing such a gap. (Consider the intuition that physical space is Euclidean.)

[2] The correspondence between Dennett's intentional stance and the capacity studied by psychologists is only a loose and approximate one. For example, rationality assumptions tend to figure prominently in theorizing about the intentional stance but only minimally in discussions of mindreading (with some exceptions, e.g. Gergeley and Csibra, 2003). For our purposes, however, the differences between the intentional stance (traditionally construed) and mindreading won't matter, so we will use these terms interchangeably.

[3] As with the mapping between the intentional stance and its empirical counterparts, we do not mean to suggest any exact correspondence between the physical stance and its counterparts (see note 2, above). In fact, the correspondence is relatively loose. For example, while Dennett tends to assume that commonsense physical knowledge is typically declarative in form, accessible to consciousness, and cognitively penetrable, psychologists studying this domain typically avoid making such assumptions. We thank Pascal Boyer for this point.

[4] So far as we are aware, this second dissociation hypothesis has not been directly tested. But it flows naturally from what is known about the neurocognitive profile of Williams syndrome, on the plausible assumption that spatial reasoning is implicated in causal-mechanical cognition. See Hegarty (2004) for a review of evidence for this assumption.

[5] Kuhlmeier et al. (2004) reports a difference in looking times between a test event involving boxes that flouts the principle of continuity and a similar event that conforms to this principle, and no difference in the case of parallel test events in which boxes are replaced with people. The authors of the study conclude from this that very young infants don't think of people as tracing continuous spatiotemporal paths, whereas they do think this of inanimate objects. This is a very surprising conclusion (the first part, anyway), and as such it should be approached with caution. It's open to question whether the key result – a null effect in the test events involving people – adequately supports the hypothesis on offer. We thank Shaun Nichols for this point.

[6] Baron-Cohen (2003) distinguishes between 'affective' and 'cognitive' components of empathy, along similar lines.

[7] Answer: since the prisoner knew that his captors expected him to lie, he intended to mislead them by telling the truth.

[8] Cf. Tye's (1999) view of phenomenal concepts as experience-dependent.

[9] The problem is essentially one of mental file-keeping, familiar to philosophers under the heading of 'Frege's puzzle' (Salmon, 1986). In the physical stance, we think about people *qua* physical objects, tracking their physical properties and adding this information to file F1; in the inten-

tional stance, we think about people *qua* intentional agents, tracking their intentional properties and adding the information to a different file, F2. F1 and F2 contain disparate sorts of information, and there is minimal sharing of information between the domains in which the two files are stored. This predisposes us to think that F1 and F2 are about different things. (Cf. Jack and Shallice's (2001) thought-experiment involving the artificial agent *Rene.)

[10] We need to hedge our bets a bit here because the array of cognitive deficits in Williams syndrome is so diverse (Bellugi et al., 2000). This contrasts sharply with the case of psychopathy, where the level of cognitive functioning tends to be relatively high across the board. On the other hand, the deep sociability and personal warmth of individuals with Williams syndrome, often remarked upon in the clinical literature, militates against the idea that they are deficient in empathy or moral sensitivity.

[11] The integration problem also applies to the phenomenal stance vis-à-vis its intentional counterpart – but to a lesser extent, due to the greater proximity between these stances in semantic and neural space.

[12] "Disgust is a response to people's bodies, not to their souls. If you see people as souls, they have moral worth … but if you see them solely as bodies, they lose any moral weight. Empathy does not extend to them. And so dictators and warmongers have come across the insight, over and over again, that you can get people to commit the most terrible atrocities using the tool of disgust" (Bloom, 2004, p. 177).

[13] Thus psychologists socially interact with the participants in their studies, as well as think about their perception and actions in mechanistic terms. Perhaps even more critically, they will usually experience the experimental procedure for themselves, thereby gaining a personal perspective on the processes under investigation (Jack and Roepstorff, 2002).

[14] Indeed, when Gallagher et al. (2002) questioned participants afterwards about their experience of playing Rock-Paper-Scissors against a person and against a computer, they used intentional language to describe their opponent in both cases.

[15] Nonetheless, this facet of our cognitive architecture may have important practical consequences for psychology as a discipline. For one thing, it makes it harder for us to integrate mechanistic accounts of the mind with our personal-level understanding of the mind. We believe that it is important for psychology to find ways to bridge between mechanistic accounts of brain function and our personal-level understanding. Specifically, we believe the most direct route to achieving this is by integrating introspective reports back into psychology (Jack and Shallice, 2001; Jack and Roepstorff, 2002).

[16] We are grateful to Paul Bloom, Pascal Boyer, David Chalmers, John Doris, Alvin Goldman, Joshua Knobe, Eddy Nahmias, Shaun Nichols, and Wayne Wright for helpful comments on earlier drafts of this paper.

Special thanks to Heidi Maibom and Rob Rupert for their contributions to the Phenomenal Stance symposium at the Pacific APA meetings in San Francisco in April 2005.

REFERENCES

Baron-Cohen, S. (1995): *Mindblindness: An Essay on Autism and Theory of Mind*, Cambridge: MIT Press.

Baron-Cohen, S. (1997): 'Are Children with Autism Superior at Folk Physics?', in H. Wellman and K. Inagaki (eds.), *The Emergence of Core Domains of Thought: Children's Reasoning about Physical, Psychological, and Biological Phenomena* (pp. 45–54), San Francisco: Jossey-Bass.

Baron-Cohen, S. (2000a): 'Autism: Deficits in Folk Psychology Exist Alongside Superiority in Folk Physics', in S. Baron-Cohen, H. Tager-Flusberg and D.J. Cohen (eds.), *Understanding Other Minds: Perspectives from Developmental Cognitive Neuroscience* (pp. 73–82), 2nd edition, Oxford: Oxford University Press.

Baron-Cohen, S. (2000b): 'Theory of Mind and Autism: A Fifteen Year Review', in S. Baron-Cohen, H. Tager-Flusberg and D.J. Cohen (eds.), *Understanding Other Minds: Perspectives from Developmental Cognitive Neuroscience* (pp. 3–20), 2nd edition, Oxford: Oxford University Press.

Baron-Cohen, S. (2002): 'The Extreme Male Brain Theory of Autism', *Trends in Cognitive Sciences* 6, 249–254.

Baron-Cohen, S. (2003): *The Essential Difference: Male and Female Brains and the Truth about Autism*, New York: Basic Books.

Baron-Cohen, S., Jolliffe, T. and Robertson, M. (1997): 'Another Advanced Test of Theory of Mind: Evidence from Very High Functioning Adults with Autism or Asperger Syndrome', *Journal of Child Psychology and Psychiatry* 38, 813–822.

Baron-Cohen, S., Leslie, A. and Frith, U. (1986): 'Mechanical, Behavioural and Intentional Understanding of Picture Stories in Autistic Children', *British Journal of Developmental Psychology* 4, 113–125.

Baron-Cohen, S., Wheelwright, S. and Stone, V. (1999): 'A Mathematician, a Physicist, and a Computer Scientist with Asperger Syndrome: Performance on Folk Psychology and Folk Physics Tests', *Neurocase* 5, 475–483.

Bellugi, U., Lichtenberger, L., Jones, W., Lai, Z. and St. George, M. (2000): 'The Neurocognitive Profile of Williams Syndrome: A Complex Pattern of Strengths and Weaknesses', *Journal of Cognitive Neuroscience* 12 (supplement), 7–29.

Bentham, J. (1789): *An Introduction to the Principles of Morals and Legislation*, Oxford: Oxford University Press.

Blair, R.J.R. (1995): 'A Cognitive Developmental Approach to Morality: Investigating the Psychopath', *Cognition* 57, 1–29.

Blair, R.J.R. (1996): 'Brief Report: Morality in the Autistic Child', *Journal of Autism and Developmental Disorders* 26, 571–579.
Blair, R.J.R. (1999): 'Psychophysiological Responsiveness to the Distress of Others in Children in Autism', *Personality and Individual Differences* 26, 477–485.
Blair, R.J.R., Sellars, C., Strickland, I., Clark, F., Williams, A., Smith, M. and Jones, L. (1996): 'Theory of Mind in the Psychopath', *Journal of Forensic Psychiatry* 7, 15–25.
Bloom, P. (2004): *Descartes' Baby: How the Science of Child Development Explains What Makes Us Human*, New York: Basic Books.
Boyer, P. and Barrett, H.C. (2005): 'Domain Specificity and Intuitive Ontology', in D. Buss (ed.), *Handbook of Evolutionary Psychology* (pp. 96–118), Hoboken: Wiley.
Chalmers, D. (1996): *The Conscious Mind*, Oxford: Oxford University Press.
Churchland, P. S. (1996): 'The Hornswoggle Problem', *Journal of Consciousness Studies* 3, 402–408.
Damasio, A. (1994): *Descartes' Error: Emotion, Reason, and the Human Brain*, New York: Putnam.
Dennett, D. (1981a): 'Intentional Systems', in *Brainstorms: Philosophical Essays on Mind and Psychology* (pp. 3–22), Cambridge: MIT Press.
Dennett, D. (1981b): 'Mechanism and Responsibility', in *Brainstorms: Philosophical Essays on Mind and Psychology* (pp. 233–255), Cambridge: MIT Press.
Dennett, D. (1987): 'True Believers: The Intentional Strategy and Why It Works', *The Intentional Stance* (pp. 13–36), Cambridge: MIT Press.
Dennett, D. (1991): *Consciousness Explained*, New York: Little, Brown.
Dolan, D. (2002): 'Emotion, Cognition, and Behavior', *Science* 298, 1191–1194.
Fisher, H. (2004): *Why We Love: The Nature and Chemistry of Romantic Love*, New York: Holt.
Frith, U. (2003): *Autism: Explaining the Enigma*, 2nd edition. Oxford: Blackwell.
Gallagher, H., Jack, A.I., Roepstorff, A. and Frith, C. (2002): 'Imaging the Intentional Stance', *NeuroImage* 16, 814–821.
Gergeley, G. and Csibra, C. (2003): 'Teleological Reasoning in Infancy: The Naïve Theory of Rational Action', *Trends in Cognitive Sciences* 7, 287–292.
Greene, J. and Haidt, J. (2002): 'How (and Where) Does Moral Judgment Work?', *Trends in Cognitive Sciences* 6, 517–523.
Happé, F. (1994): 'An Advanced Test of Theory of Mind: Understanding of Story Characters' Thoughts and Feelings by Able Autistic, Mentally Handicapped, and Normal Children and Adults', *Journal of Autism and Developmental Disorders* 24, 129–154.
Hare, R. (1991): *The Hare Psychopathy Checklist – Revised*, Toronto: Multi-Health Systems.

Hare, R. (1993): *Without Conscience: The Disturbing World of the Psychopaths Among Us*, New York: Guilford Press.
Hegarty, M. (2004): 'Mechanical Reasoning by Mental Simulation', *Trends in Cognitive Sciences* 8, 280–285.
Jack, A.I. and Robbins, P. (2004): 'The Illusory Triumph of Machine over Mind: Wegner's Eliminativism and the Real Promise of Psychology', *Behavioral and Brain Sciences* 27, 665–666.
Jack, A.I. and Roepstorff, A. (2002): 'Introspection and Cognitive Brain-mapping: From Stimulus-response to Script-report', *Trends in Cognitive Sciences* 6, 333–339.
Jack, A.I. and Shallice, T. (2001): 'Introspective Physicalism as an Approach to the Science of Consciousness', *Cognition* 79, 161–196.
Kuhlmeier, V., Bloom, P. and Wynn, K. (2004): 'Do 5-Month-old Infants See Humans as Material Objects?', *Cognition* 94, 95–103.
Ledoux, J. (1998): *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*, New York: Simon and Schuster.
Leslie, A.M. (1994): 'ToMM, ToBy, and Agency: Core Architecture and Domain Specificity', in L. Hirschfeld and S. Gelman (eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture* (pp. 119–148), Cambridge: Cambridge University Press.
Leslie, A.M. and Thaiss, L. (1992): 'Domain Specificity in Conceptual Development: Neuropsychological Evidence from Autism', *Cognition* 43, 225–251.
Levine, J. (1983): 'Materialism and Qualia: The Explanatory Gap', *Pacific Philosophical Quarterly* 64, 354–361.
Martin, A. and Weisberg, J. (2003): 'Neural Foundations for Understanding Social and Mechanical Concepts', *Cognitive Neuropsychology* 20, 575–587.
Michotte, A. (1963): *The Perception of Causality*, London: Methuen.
Moll, J., de Oliveira-Souza, R. and Eslinger, P.J. (2003): 'Morals and the Human Brain: A Working Model', *NeuroReport* 14, 299–305.
Moll, J., de Oliveira-Souza, R., Eslinger, P.J., Bramati, I.E., Mourão-Miranda, J., Andreiuolo, P.A. and Pessoa, L. (2002): 'The Neural Correlates of Moral Sensitivity: A Functional Magnetic Resonance Imaging Investigation of Basic and Moral Emotions', *Journal of Neuroscience* 22, 2730–2736.
Moll, J., Eslinger, P.J. and de Oliveira-Souza, R. (2001): 'Frontopolar and Anterior Temporal Cortex Activation in a Moral Judgment Task', *Arquivos de Neuro-Psiquiatria* 59, 657–664.
Nagel, T. (1974): 'What Is It Like to Be a Bat?', *Philosophical Review* 82, 435–450.
Premack, D. and Woodruff, G. (1978): 'Does the Chimpanzee Have a Theory of Mind?', *Behavioral and Brain Sciences* 4, 515–526.
Preston, S.D. and de Waal, F.B.M. (2002): 'Empathy: Its Ultimate and Proximate Bases', *Behavioral and Brain Sciences* 25, 1–20.

Richell, R.A., Mitchell, D.G.V., Newman, C., Leonard, A., Baron-Cohen, S. and Blair, R.J.R. (2003): 'Theory of Mind and Psychopathy: Can Psychopathic Individuals Read the ''Language of The Eyes''?', *Neuropsychologica* 41, 523–526.

Salmon, N. (1986): *Frege's Puzzle*, Cambridge: MIT Press.

Singer, P. (1976): 'All Animals Are Equal', in T. Regan and P. Singer (eds.), *Animal Rights and Human Obligations* (pp. 148–162), New York: Prentice-Hall.

Singer, T., Wolpert, D. and Frith, C. (2003): 'Introduction: The Study of Social Interactions', in C. Frith and D. Wolpert (eds.), *The Neuroscience of Social Interaction: Decoding, Imitating, and Influencing the Actions of Others* (pp. xiii–xxvii), Oxford: Oxford University Press.

Spelke, E.S. (1988): 'The Origins of Physical Knowledge', in L. Weiskrantz (ed.), *Thought Without Language* (pp. 168–184), Oxford: Oxford University Press.

Spelke, E.S., Phillips, A. and Woodward, A.L. (1995): 'Infants' Knowledge of Object Motion and Human Action', in D. Sperber, D. Premack and A.J. Premack (eds.), *Causal Cognition: A Multidisciplinary Debate* (pp. 44–78), Oxford: Oxford University Press.

Turiel, E. (1983): *The Development of Social Knowledge: Morality and Convention*, Cambridge: Cambridge University Press.

Tye, M. (1999): 'Phenomenal Consciousness: The Explanatory Gap as a Cognitive Illusion', *Mind* 108, 705–725.

Wall, P.D. (2002): *Pain: The Science of Suffering*, New York: Columbia University Press.

Wolfe, T. (2000): 'Sorry, but Your Soul Just Died', in *Hooking Up* (pp. 89–109), New York: Farrar, Straus and Giroux.

Philip Robbins
Department of Philosophy
Washington University in St. Louis
MO, 63130-4899,
USA
E-mail: probbins@wustl.edu

Anthony I. Jack
Department of Neurology
Washington University in St. Louis
MO, 63130-4899,
USA
E-mail: ajack@npg.wustl.edu