

INTRODUCTION

Is cognitive penetrability the mark of the moral?

Philip Gerrans and Jeanette Kennett

Introduction

The papers collected in this special issue all address the relationship between empirical research and philosophical accounts of the nature of moral judgement. In particular, the papers all make claims about the extent to which research in cognitive neuroscience concerning the nature of cognitive processes which underlie moral judgement can support a sentimentalist or rationalist account of meta-ethics.

The papers include explicitly philosophical accounts (Joyce, Cullity, Jones) which discuss the conditions that would have to be met for one or other meta-ethical theory to be vindicated, revised or refuted. This suggests an artificially neat division of labour in which empirical research describes the cognitive processes involved in moral judgement and philosophers decide whether those processes count as moral, sentimental or rational. However, as the papers by cognitive neuroscientists (Blair, Fiddick, Stone) make clear, how one conceives of the nature of moral judgement influences both the construction of experiments designed to probe its structure (Fiddick, Blair) and the interpretation of results. For example, most cognitive neuroscientists and some philosophers (Prinz) working in the field have interpreted the results to support some version of sentimentalism in meta-ethics. Interestingly, however, two papers in the collection dispute that interpretation on a mixture of philosophical and empirical grounds (Kennett, Fine).

Our aim in the rest of this introduction is to provide a framework for evaluating those debates and to situate the papers within it so that readers can evaluate arguments without being weighed down by the sometimes technical disciplinary vocabularies or the differences of emphasis in the different papers. We first describe the essential theses of sentimentalism and rationalism. These philosophical theses might seem too abstract or gravid with normativity to be evaluated against the data of cognitive neuroscience but we do not think this is the case. In fact we can reconfigure one essential aspect of the debate between sentimentalists and rationalists as a debate about the answer to the following question, which can be directly addressed by empirical research.

Are the Processes Involved in Moral Judgement Cognitively Penetrable?

The term cognitive impenetrability was introduced by Zenon Pylyshyn to clarify a debate among cognitive scientists about the extent to which a person's knowledge,

goals and expectations could influence her perceptual experience and consequent judgements. Pylyshyn pointed out that if perception could be influenced in this way the processes involved in representing goals, beliefs and expectations could not be independent of those involved in perception. As he put it, perceptual processes would be *cognitively penetrable*.

If a system is cognitively penetrable then the function it computes is sensitive, in a semantically coherent way, to the organism's goals and beliefs, i.e. it can be altered in a way that bears some logical relation to what the person knows (Pylyshyn 1999, 343).

Cognitive processes not sensitive in this way are cognitively impenetrable. Visual illusions, for example, persist despite our beliefs that they are not veridical representations. Other phenomena are similarly cognitively impenetrable: responses to familiar faces, basic emotional responses to stereotypical stimuli, and some classes of automatic and immediate judgement. The mind has a large number of systems which perform their cognitive processing autonomously and immediately. In the cognitive science literature these impenetrable systems are known as modules: domain-specific cognitive subsystems whose computational processes are independent of goals, beliefs and expectations.

The mind also has cognitive systems which are sensitive to knowledge, expectations and goals. We can revise our beliefs, goals and plans in the face of new evidence or as a result of inferences which correct previously held assumptions. These systems are central processes and depend crucially on the cognitive functions like attention, working memory, executive functions (the ability to hold and manipulate representations in working memory) and inference (the ability to produce consistent sets of beliefs applying rules of procedural rationality).

The distinction between cognitively penetrable and impenetrable processes is very helpful in evaluating papers in this collection. We think that sentimentalism would be largely vindicated if moral judgement always depended on cognitively impenetrable processes, and in fact Blair and Prinz explicitly argue from evidence that moral judgement depends on processes, typically involving the affective systems, which are not sensitive to beliefs, expectations and goals, to sentimentalism. Fiddick's paper is not framed in these terms but he does argue that moral judgements depend on a variety of impenetrable inferential processes. On the other side, the papers by Fine, Kennett and Stone argue on a variety of grounds that there is empirical evidence to support the idea that moral judgement is sensitive to beliefs, expectations and goals. They then argue that evidence of, actual or potential, cognitive penetrability supports a version of rationalism.

Before we discuss the papers in more depth we need to briefly define what we take to be the core intuitions behind sentimentalism and rationalism and show that they can be mapped to the distinction between penetrable and impenetrable processes without doing undue violence to the project of meta-ethics.

Sentimentalism and Rationalism

For sentimentalists moral judgements are founded in the sentiments. Development and exercise of a capacity for moral judgement depends essentially on emotional responses. According to Prinz, to believe that something is morally wrong or right is just to have a sentiment of approval or disapproval towards it, where those sentiments are understood as dispositions to have a range of emotions depending on the circumstances which elicit the judgement. The point we wish to emphasise here is that the sentiments are

not understood as states of mind which are essentially subject to rational revision in the face of evidence or argument. This is equally true of naïve sentimentalist accounts which treat moral judgements as expressions of emotion and sophisticated projectivist accounts which descend from Hume. The impetus to projectivism is the need to accommodate the fact that moral judgements appear to be revisable in the face of evidence and argument without retreating to the rationalist idea that the process of revision is the same as the process by which evidential beliefs are revised.

Rationalism is the view that moral judgement is characterised by the way the reasons advanced in support of the judgement are justified. Despite their specialised subject matter, moral judgements are no different from other judgements which require the ability to acquire and evaluate evidence and test evidential propositions for consistency with a set of background beliefs. Rationalists have always argued that if putative moral judgements are based on emotions which are not sensitive to these justificatory requirements, they do not really deserve the title 'moral'.

Sentimentalism is a descriptive account and is straightforwardly testable against empirical evidence. If it turns out that the processes involved in moral judgement are not essentially sentimental, sentimentalism will have to be abandoned or modified.

In the case of rationalism the matter is not quite so straightforward since it also involves an account of what psychological facts *justify* moral judgements. Thus collecting evidence that many purportedly moral judgements are based on automatic responses to actions or situations, which bypass processes of rational justification, does not automatically falsify the claim that morality is essentially linked to a capacity for rational justification. However, if it were demonstrably the case that all (most?) purportedly moral judgements were of this kind the rationalist would have to decide whether her theory was a theory of actual human moral deliberation or an unrealisable ideal.

Thus we think that the rationalist should claim that moral judgements can be produced in humans by the type of deliberative processes required by rationalists for the justification of moral claims. This allows us to assess empirical evidence for the role in moral judgement of the type of processes implicated by sentimentalism and rationalism.

To reconnect these points with the distinction between cognitively penetrable and impenetrable processes, we can note that in order for a judgement to be actually justified in the way required by rationalism, the processes which underlie it must be *cognitively penetrable*. While it might typically be made automatically and unreflectively, if it continues to be made in that way in the face of accumulating evidence or argument to the contrary it cannot be said to be justified in the first place. If counter evidence or argument can influence the processes which underlie judgement, then the goals, beliefs and expectations of the subject, and the cognitive processes which regulate them, are playing a role in modifying a judgement. And thus moral judgement must be cognitively penetrable.

Equally, the sentimentalism of Blair and Prinz would be largely vindicated if the processes by which moral judgements were made could be shown to be impenetrable. In fact, the fact that emotional processes are implicated in moral judgement is less important to simple sentimentalism than their alleged cognitive impenetrability. If emotions were obviously cognitively penetrable then the contrast between rationalism and sentimentalism would diminish or disappear.

Nevertheless, establishing whether or not moral judgements are cognitively penetrable does not close meta-ethical debate. Empirical evidence might show that our moral concepts do not map straightforwardly to the cognitive processes underwriting moral

judgement. We then have the choice to revise our concepts to accommodate the evidence or to maintain the concepts and live with the incongruity. Thus, in the absence of *straight-forward* empirical confirmation for a particular meta-ethical theory, attenuated forms of realism, scepticism or nihilism remain viable as options for meta-ethics.

The Case for Sentimentalism

James Blair and collaborators point out that the clinical conditions of psychopathy and acquired sociopathy are cases of 'pathological conditions [which] lead to breakdowns in morality' (p. 13). People with these conditions have apparent 'moral deficits' which make them insensitive to the distress of others to the extent that they are unconcerned about inflicting harm on other people. These pathologies thus provide a chance to develop a neuropsychology of morality on the assumption that the deficits result from damage to discrete subsystems which underpin moral judgement. As with all forms of neuropsychology, the methodology requires that any such deficits leave other aspects of cognition relatively intact.

In Blair's view, what distinguishes moral rules from conventions is that violation of a moral rule produces a victim. Normal children categorise rule violations in terms of violations which involve a victim and merely conventional violations which do not. Psychopaths and acquired sociopaths fail to make this distinction. Given that a large part of morality concerns the impermissibility of harms to others, the apparent failure of psychopaths and acquired sociopaths to recognise and act on this imperative is a significant moral deficit.

A simple sentimentalist explanation might have it that the distinction reflects a difference between rules whose violation produces an emotional response and rules whose violation does not. Psychopaths would be people who lack the appropriate emotional responses and hence cannot make the distinction.

Blair et al., however, note that conventional transgressions do occasion emotions, typically anger, which is reduced when signs of contrition or appeasement are displayed. The neural system involved is the human version of phylogenetically conserved mechanisms for responding to violations of status. Psychopaths are not impaired in their ability to be angered by such transgressions although they do lack the ability to inhibit anger when the violator gives signs of distress or appeasement.

So psychopaths do have emotional, though abnormally regulated, responses to conventional transgressions. Moral transgressions, however, involve the recognition of intentional harm to a victim and not just the violation of conventions. This requires integrated functioning of systems required to detect intentions, systems which allow conditioning to aversive stimuli using the fear response and systems involved in empathetic identification. The amygdala is a structure involved in both anger and aversive response which is connected to other ventromedial structures involved in learning, memory and behavioural control (the Integrated Emotional System). When a child with an intact Integrated Emotional System sees a victim's distress caused by an intentional action, she associates her own distress and fear with that of the victim and is aversively conditioned to the harmful action. This ability to be conditioned allows her to learn to associate negative affect with the intentional infliction of harm. Psychopaths, however, lack the neurocognitive mechanisms required for this type of conditioning to occur. In particular, their responses to others' distress are abnormal (evidenced even in failure to inhibit inappropriate anger responses to status violations) and they seem to lack the ability to become aversively

conditioned. Thus Blair's hypothesis is that the circuitry linking this fundamental system for affective regulation to other cognitive systems is abnormal in psychopathy.

We leave details of the account to readers. Here we note two important features. The first is that it is essentially sentimentalist. It is the absence of suitable emotional responses which renders psychopaths unable to be conditioned to respond to others' distress and hence to make the appropriate distinction between moral and conventional judgements.

The second is that as a result of normal development, rule violations are categorised as moral or conventional *automatically*. Blair makes the point explicit when he contrasts the systems involved in moral judgement with cognitive processes which are 'effortful, controllable rational process[es]' (p. 13). In other words, moral judgements made by the Integrated Emotional System are *cognitively impenetrable*.

Jesse Prinz also appears committed to the thesis that moral judgement is cognitively impenetrable. He draws upon a range of evidence to support a series of increasingly strong theses about the role of emotion in moral judgement. He begins from the intuition that moral judgements are intrinsically action guiding and seeks to explain this by showing that they are emotional in nature. He moves from neuroimaging data which indicates that emotions *co-occur* with moral judgements, to philosophical intuitions about trolley car cases which suggest that emotions *influence* our judgements of rightness and wrongness. He provides several variants of trolley cases which suggest that the crucial factor is the proximity of a victim. Prinz's analysis of these cases is consistent with Blair's account of aversive conditioning to victim distress. Our responses to the close up and personal trolley cases would be predicted by Blair's model.

Prinz then argues that emotions do not merely influence moral judgements; they are both *necessary* and *sufficient* for moral judgement. In support of the sufficiency thesis he cites experiments by Haidt which show that subjects hypnotised to feel disgust at a neutral word judge that morally admirable characters are morally wrong in vignettes in which the word occurs. This suggests, he says, 'that a negative feeling can give rise to a negative moral appraisal without any specific belief about some property in virtue of which something is wrong' (p. 31). Other experiments suggest that the role of justification in moral judgement is after the fact. When one justification is removed we search for another, when all are removed we point to the feelings themselves. Incest is ultimately judged wrong *because* we find it disgusting. The judgement is, in our terms, cognitively impenetrable. In support of the necessity thesis he draws upon research on psychopathy which suggests that emotions are developmentally necessary for acquiring the capacity to make moral judgements and the formation of moral concepts. Finally, Prinz claims that emotions are necessary in a synchronic sense. Sincere moral judgements are expressions of underlying dispositions to have emotions. He argues against the metacognitive move made by neo-sentimentalists according to which a moral judgement is an endorsement of the agent's sentiments of approbation/disapprobation, in favour of the simpler sentimentalism we can trace back to Hume.

In these two accounts the connection between cognitive impenetrability and sentimentalism is clear. Thus one way to defend rationalism would be to provide evidence that the cognitive processes in question are not impenetrable. This, however, requires a conception of cognitive penetrability which shows how goals, beliefs and expectations can influence cognitive processes. Blair provides the starting point when he equates effortful and controllable processes with rational processes. This way of putting things avoids disputes about the nature of rationality while still providing a strong contrast

between two types of process: automatic and effortless, modularised processing and cognitively impenetrable versus effortful and controlled cognitively penetrable central processing.

The Case for Rationalism

The paper by Valerie Stone allows us to sharpen the distinction by focusing on the nature and evolutionary origin of the cognitive systems involved in effortful and controlled central processing. The crucial point of her account is that 'some very abstract cognitive abilities that are unique to our species are layered on top of phylogenetically older emotional instincts for aggression and for empathy' (p. 55). Human beings have evolved an enormous prefrontal cortex in the seven million years since we diverged from our common ancestor with chimpanzees. The main task of these frontal systems is to override automated modular functioning and allow cognitively penetrable processes to regulate social behaviour. Without these frontal systems we would be like other animals whose minds are essentially a bundle of integrated modules which perform their tasks automatically and inflexibly.

However, as the notion of layering implies, the phylogenetically older systems are not detachable from our recently acquired executive abilities. The prefrontal cortex is densely connected to, and influenced by, older systems. Our ancient behavioural propensities and emotional responses are conserved and influence most of our cognitive life. This is why reasoning is 'controlled and effortful'. It requires us to inhibit automatic processing while we search our memories and compare and evaluate information relevant to a problem, and then to control the behaviour which implements the solution arrived at by abstract forms of thought.

Sentimentalists offer an account of moral judgement in which phylogenetically ancient primate or early hominid cognitive systems do most of the work and moral theory becomes an *ex post facto* rationalisation of impenetrable responses.

Stone's account, however, suggests that the recent evolution of the human prefrontal cortex makes it unlikely that we have evolved hard-wired systems for the abstract aspects of moral cognition: the subsumption of actions, agents and effects under rules concerning appropriateness of harms. Rather, our moral cognitive repertoire is likely to contain older automated responses, sometimes inhibited or overridden by our capacities for more deliberate executive control of behaviour. Equally, our best laid moral plans arrived at by impeccable metacognitive processes might be subject to reversal or disruption when we encounter an emotionally salient stimulus which engages our automated systems.

If this picture of the mind also captures the nature of moral judgement, it suggests that moral judgements may not be monolithic and that the contributions of different types of processing might need to be teased apart. Jeanette Kennett's paper broaches this idea in her reconsideration of the extent to which psychopathy is evidence for sentimentalism. In Blair's view, psychopaths cannot construct a normal Integrated Emotional System in view of their deficits in amygdala function which prevent them from constructing circuitry required to respond appropriately to others' distress. Kennett points out, however, that this is by no means the only deficit characteristic of psychopathy. Psychopaths are also impulsive and compulsive risk takers who do not evaluate the consequences for themselves of their actions. As she points out, these failures are failures of inhibition, planning

and executive function: all aspects of effortful and controlled cognition implicated by rationalists as essential to moral cognition. '[A] person . . . who could not reflect upon whether or not his desires provided reasons for action, whose desires were unresponsive to such reflection, or who could not be guided by the results of his deliberations, through exercises of planning and self-control, would not count as a rational . . . [or] a moral agent' (p. 76).

Kennett points out that, even if the fundamental deficit is affective, its effect may be to prevent the child from acquiring the capacity for rational deliberation and self-control. Blair's argument seems to assume that the role of emotion in the normal child is to allow it to make the moral conventional distinction automatically, i.e. to make normal moral judgement cognitively impenetrable. In Kennett's view the role of empathetic response might be to sensitise the child to aspects of situations which require a controlled and effortful response, and it is this ability which psychopaths ultimately lack.

Kennett's paper, while empirically informed, is a conceptual paper. She provides an account of rationalism which makes room for considerable affective influence on decision making and shows that such influence would not undermine the standard arguments for rationalism.

Karen Jones' response to Prinz likewise tackles conceptual issues. She argues that simple sentimentalism cannot provide a satisfactory account of what it takes to be in possession of moral concepts and that this is the terrain in which meta-ethical debate is won or lost. *Some* degree of cognitive penetrability of moral judgement is required for basic moral conversability. A person who would not withdraw her moral judgement upon accepting that it is fully explained by a hypnotically induced disgust at the word 'often' would not count as being conversable with the relevant concepts. Jones concludes that 'to be viable . . . sentimentalism must make the neo-sentimental ascent' (p. 46) and incorporate this requirement for reflective endorsement. These two papers suggest the possibility of a rapprochement between conceptually defensible versions of sentimentalism and empirically adequate versions of rationalism.

Fine's paper directly confronts the empirical challenge we identified earlier. If it is the case that moral judgements are cognitively impenetrable (even if conceptual analysis shows they *ought* to be penetrable), rationalism is not an empirically adequate account.

Fine discusses a suite of recent experiments which have led many cognitive neuroscientists to endorse sentimentalism on the grounds that moral judgements are produced 'without any conscious awareness of having gone through the steps of searching, weighing evidence, or inferring a conclusion' (p. 83), i.e. are cognitively impenetrable.

Fine, however, presents evidence that these automatic judgements (typically the application of morally loaded stereotypes to people or situations) can be disrupted by attention, motivation, the demand for accuracy and whether or not the subject habitually deliberates. These cognitive factors engage controlled rather than automatic processes and allow a subject to retrieve and process non-stereotypical information and include it in the judgement process.

She also cites evidence that subjectively held values (such as a desire to be egalitarian or non-racist) influence the application of stereotypes when attentional resources are available. Recall that Pylyshyn's definition of cognitive penetrability is influenced by a person's 'goals, beliefs, or expectations'. Fine concludes that some moral judgements are penetrable under favourable conditions and thus 'there remains the possibility that sometimes the emotional dog is not wagging the rational tail, but chasing it' (p. 97).

The Case for Revision

The papers thus far considered provide a dialogue in which the contributors align themselves with sentimentalism and rationalism as we have thus far defined them. The other papers are not so readily classifiable. Larry Fiddick is concerned to show that moral cognition, while it is an inferential process, is accomplished by discrete cognitively impenetrable subsystems.

Fiddick interprets the normative aspect of morality as the application of rules prescribing what we *ought* to do in situations: deontic reasoning. He then describes experiments which show that deontic reasoning is carried out by discrete subsystems which process conditional inferences concerning social contracts and precautions. Social contracts regulate transactions in which parties act to distribute some valuable item. The conditional inferences involved are of the form: If you take a benefit you ought to perform an appropriate action in consideration (p. 112). Conditional reasoning about precautions concerns appropriate actions for avoiding hazards, yielding inferences of the form: If you want to avoid hazard you ought to take precautionary action (p. 112). Both are ecologically salient and reinforced by appropriate emotions. Violation of the hazard rule produces fear and the precaution rule produces anger, an emotion appropriate to visiting punishment on a violator.

From our point of view Fiddick's paper raises the possibility that moral judgements concerning fairness are underwritten by an affectively scaffolded process for detecting and punishing cheats which is cognitively impenetrable. How does this data bear on the debate between rationalists and sentimentalists?

On some interpretations of rationalism (Joyce forthcoming), rationalism is vindicated if the rules we apply in reaching our moral judgements are rationally justifiable. *Perhaps* the rules Fiddick says regulate judgements about social contracts and precautions are rules we would reflectively endorse. Neil Levy, in his response to Fine, suggests another way in which our cognitively impenetrable judgements might gain the backing of reason. Though he thinks that Fine does not succeed against Haidt, since it is possible that the processes of revision she outlines are themselves affectively driven, he claims that our moral judgements might still count as rational. He points out that the affective influences and cognitive biases Haidt points to in moral cognition also infect the sciences and social sciences. But judgements made by scientists are not therefore unjustified. The reason is that scientific judgements are ultimately justified, after testing and criticism, by their place in a theory developed by the scientific community as a whole. The same may be true of moral judgement. Even if individual moral judgements are made via epistemically suspect processes, 'it may nevertheless be that they are rationally held, since warrant can be transferred by testimony from a community of experts to laypeople' (p. 102).

To adopt Levy's interpretation of rationalism, however, would be to revise the account we gave earlier. On that account *ex post facto* justification, whether at the level of the individual or the community, would not be enough to satisfy the rationalist requirement that the actual processes of moral judgement by individuals be cognitively penetrable.

Garrett Cullity's paper discusses the constraints on this type of theoretical revision. He first distinguishes four aspects of meta-ethics: moral psychology, moral semantics, moral metaphysics and moral epistemology. So far we have been concentrating on questions of moral psychology but Cullity argues that empirical research is most relevant to moral

epistemology. Cullity's ultimate view is that impenetrable intuitive judgements alone do not have epistemic warrant. It follows that if the Blair/Prinz accounts exhaust the epistemology of moral judgement, our moral judgements lack warrant. However, Cullity argues for a similar conclusion to Levy, that cognitively impenetrable judgements may derive epistemic support from other warranted judgements.

Our tentative assessment of the way things stand is that neither rationalism nor sentimentalism as we initially defined them are unequivocally supported by the empirical evidence. Should we revise them or abandon them?

How might findings in moral semantics bear upon this issue? Simple sentimentalism may be right about the antecedents of our moral utterances and their connection to action, yet wrong about what we mean by them. Cullity and Joyce both draw our attention to the fact that questions in moral semantics are empirical as well conceptual. Meaning, as Cullity says 'must be ascertained by careful observation of actual usage, and convincing theorizing about that' (p. 121), though this is not to say that we can settle these questions by conducting a poll.

So it may be that when we deliberate and make moral judgements we do, as a matter of empirical fact, take ourselves to be engaging in a process which is amenable to evidence and argument. It may be that our moral semantics, and our moral concepts, are rationalist. If it turns out that this is an illusion and moral judgements are cognitively impenetrable, some rationalists will turn, not to sentimentalism but to an error theory of moral judgement. It is consistent with rationalism to maintain *both* that our moral concepts require the cognitive penetrability of moral judgement *and* that moral judgement is cognitively impenetrable. If the foundational assumption is false and there is nothing that corresponds to our concept of moral judgement then moral discourse is systematically in error. Our moral judgements cannot have rational justification. Smith (1994), for one, is explicit that if there is no possibility of rational convergence in morality we are left with moral nihilism.

When we turn to moral metaphysics and moral epistemology, other meta-ethical positions, such as a naturalism which is neither sentimental nor rationalist, may also be viable. Perhaps there are mind-independent moral properties or facts (reducible to non-moral properties) and perhaps our moral judgements are, happily, truth tracking. We could thus have epistemic warrant for our moral claims even if the truth-tracking mechanisms are impenetrable.

In his contribution, Richard Joyce argues that empirical advances can contribute to the resolution of the question of whether our moral judgements are epistemically justified. Joyce notes that moral judgements may be subject to *genealogical debunking*. If we were to discover that our moral beliefs are wholly insensitive to evidence, that would not be enough to show that they are not accurate, but it might suggest that we are not rationally justified in holding them absent a story demonstrating that their accuracy is just what would have enhanced reproductive fitness. Such a story might be available for certain innate beliefs such as our mathematical beliefs but is not available in the moral case. Joyce says that if we have an empirically confirmed theory about where our moral beliefs come from which does not imply or hold that they are true 'and their truth is not surreptitiously buried in the theory by virtue of any form of moral naturalism, then . . . we have no grounds one way or the other for maintaining these beliefs' (p. 145). Again the evidence Joyce cites may be such as to rule out certain forms of moral realism but it does not tell us that we should embrace sentimentalism or other versions of non-cognitivism rather than moral scepticism.

Now read on. . .

REFERENCES

- JOYCE, R. Forthcoming. What neuroscience can (and cannot) contribute to metaethics. In *Moral psychology: Morals in the brain*, edited by W. Sinnott-Armstrong.
- PYLYSHSYN, Z. 1999. Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioural and Brain Sciences* 22: 341–423.
- SMITH, M. 1994. *The moral problem*. Oxford: Basil Blackwell.

Jeanette Kennett, Centre for Applied Philosophy and Public Ethics, Australian National University, LPO Box 8260 ANU, Canberra, ACT 2601, Australia, and School of Philosophy and Bioethics, Monash University, Australia. E-mail: jeanette.kennett@anu.edu.au

Philip Gerrans, Department of Philosophy, University of Adelaide, Adelaide, SA 5005, Australia. E-mail: philip.gerrans@adelaide.edu.au