34. Churchland, P. S. *Brain-wise: Studies in Neurophilosophy* (MIT Press, Cambridge, Massachusetts, 2002).
35. Churchland, P. M. Towards a cognitive neurobiology of the moral virtues. *Topoi* **17**, 83–96 (1998).
36. Wallis, J. D., Anderson, K. C. & Miller, E. K. Single neurons in prefrontal cortex encode abstract rules. *Nature* **411**, 953–956 (2001).
37. Casebeer, W. D. & Churchland, P. S. The neural mechanisms of moral cognition: a multiple-aspect approach to moral judgment and decision-making. *Biol. Philos.* **18**, 169–194 (2003).
38. Moreno, J. D. Neuroethics: an agenda for neuroscience and society. *Nature Rev. Neurosci.* **4**, 149–153 (2003).
39. Cacioppo, J. T. *et al*. (eds) *Foundations in Social Neuroscience* (MIT Press, Cambridge, Massachusetts, 2002).
40. Rachels, J. *The Elements of Moral Philosophy 4th Edn* (McGraw-Hill, Columbus, 2002).
41. Lapsley, D. *Moral Psychology* (West-view, Boulder, 1996).

Acknowledgements
I thank P. M. Churchland and P. S. Churchland for their close reading of the manuscript and invaluable advice about its structure and content. In addition, J. Greene's sceptical remarks were extremely helpful. J. Moll also provided useful preprints of his team's work in this area.

Online links

FURTHER INFORMATION
MIT Encyclopedia of Cognitive Sciences:
http://cognet.mit.edu/MITECS/
limbic system | moral psychology | social cognition | theory of mind
Online Bibliography of Cognitive Science and Ethics:
http://hem.passagen.se/ollekillen/cogethics.html
Access to this interactive links box is free online.

OPINION

From neural 'is' to moral 'ought': what are the moral implications of neuroscientific moral psychology?

Joshua Greene

Many moral philosophers regard scientific research as irrelevant to their work because science deals with what is the case, whereas ethics deals with what ought to be. Some ethicists question this is/ought distinction, arguing that science and normative ethics are continuous and that ethics might someday be regarded as a natural social science. I agree with traditional ethicists that there is a sharp and crucial distinction between the 'is' of science and the 'ought' of ethics, but maintain nonetheless that science, and neuroscience in particular, can have profound ethical implications by providing us with information that will prompt us to re-evaluate our moral values and our conceptions of morality.

Many moral philosophers boast a well-cultivated indifference to research in moral psychology. This is regrettable, but not entirely groundless[1]. Philosophers have long recognized that facts concerning how people actually think or act do not imply facts about how people ought to think or act, at least not in any straightforward way. This principle is summarized by the Humean[2] dictum that one can't derive an 'ought' from an 'is'. In a similar vein, moral philosophers since Moore[3] have taken pains to avoid the 'naturalistic fallacy', the mistake of identifying that which is natural with that which is right or good (or, more broadly, the mistake of identifying moral properties with natural properties). Prominent among those accused by Moore of committing this fallacy was Herbert Spencer, the father of 'social Darwinism', who aimed to ground moral and political philosophy in evolutionary principles[4]. Spencer coined the phrase 'survival of the fittest', giving Darwin's purely biological notion of fitness a socio-moral twist: for the good of the species, the government ought not to interfere with nature's tendency to let the strong dominate the weak.

Spencerian social Darwinism is long gone, but the idea that principles of natural science might provide a foundation for normative ethics has won renewed favour in recent years. Some friends of 'naturalized ethics' argue, contra Hume and Moore, that the doctrine of the naturalistic fallacy is itself a fallacy, and that facts about right and wrong are, in principle at least, as amenable to scientific discovery as any others. Most of the arguments in favour of ethics as continuous with natural science have been rather abstract, with no attempt to support particular moral theories on the basis of particular scientific research[5,6]. Casebeer's neuroscientific defense of Aristotelian virtue theory (this issue) is a notable exception in this regard[7].

A critical survey of recent attempts to naturalize ethics is beyond the scope of this article. Instead I will simply state that I am sceptical of naturalized ethics for the usual Humean and Moorean reasons. Contemporary proponents of naturalized ethics are aware of these objections, but in my opinion their theories do not adequately meet them. Casebeer, for example, examines recent work in neuroscientific moral psychology and finds that actual moral decision-making looks more like what Aristotle recommends[8] and less like what Kant[9] and Mill[10] recommend. From this he concludes that the available neuroscientific evidence counts against the moral theories of Kant and Mill, and in favour of Aristotle's. This strikes me as a *non sequitur*. How do we go from 'This is how we think' to 'This is how we ought to think'? Kant argued that our actions should exhibit a kind of universalizability that is grounded in respect for other people as autonomous rational agents[9]. Mill argued that we should act so as to produce the greatest sum of happiness[10]. So long as people are capable of taking Kant's or Mill's advice, how does it follow from neuroscientific data — indeed, how could it follow from such data — that people ought to ignore Kant's and Mill's recommendations in favour of Aristotle's? In other words, how does it follow from the proposition that Aristotelian moral thought is more natural than Kant's or Mill's that Aristotle's is better?

Whereas I am sceptical of attempts to derive moral principles from scientific facts, I agree with the proponents of naturalized ethics that scientific facts can have profound moral implications, and that moral philosophers have paid too little attention to relevant work in the natural sciences. My understanding of the relationship between science and normative ethics is, however, different from that of naturalized ethicists. Casebeer and others view science and normative ethics as continuous and are therefore interested in normative moral theories that resemble or are 'consilient' with theories of moral psychology. Their aim is to find theories of right and wrong that in some sense match natural human practice. By contrast, I view science as offering a 'behind the scenes' look at human morality. Just as a well-researched biography can, depending on what it reveals, boost or deflate one's esteem for its subject, the scientific investigation of human morality can help us to understand human moral nature, and in so doing change our opinion of it.

Neuroscience and normative ethics

There is a growing consensus that moral judgements are based largely on intuition — 'gut feelings' about what is right or wrong in particular cases[11]. Sometimes these intuitions conflict, both within and between individuals. Are all moral intuitions equally worthy of our

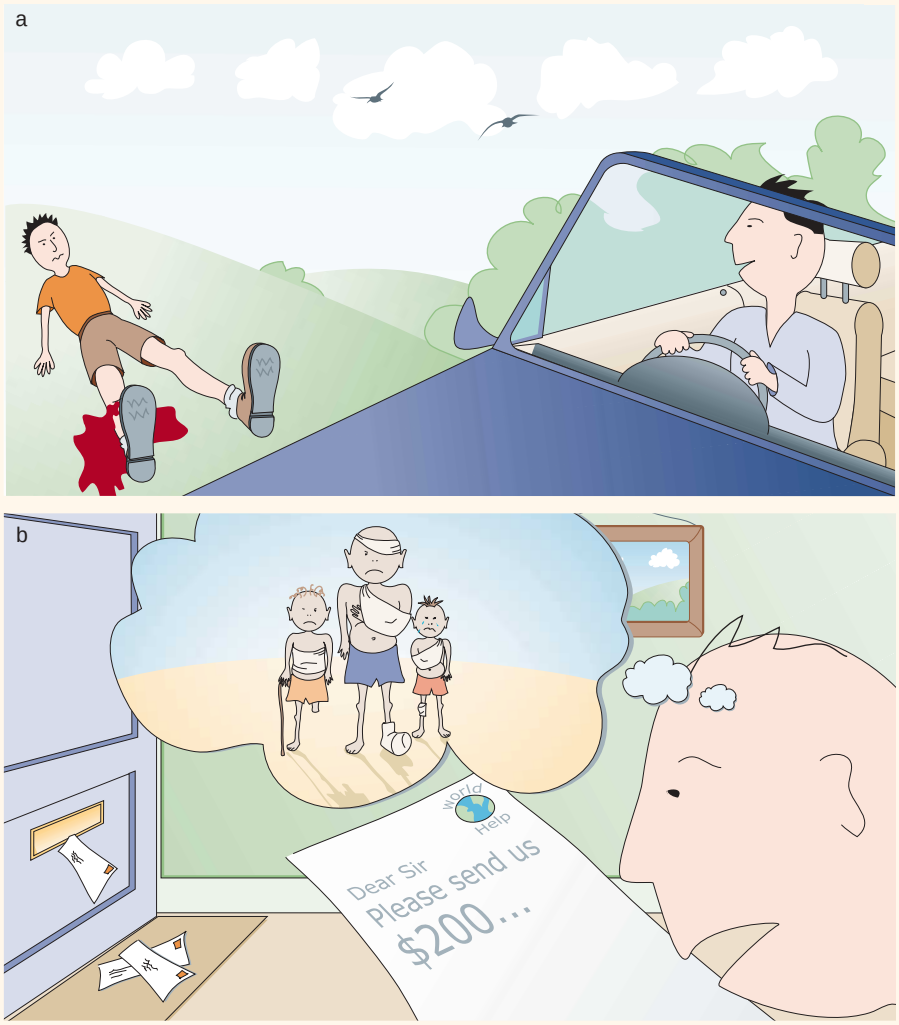

Figure 1 | **Moral dilemmas and contradictions.** People believe (**a**) that it would be deeply wrong to abandon a bleeding stranger by the side of the road in order to preserve one's leather car seats, but (**b**) that it is morally acceptable to spend money on luxuries when that money could be used to save the lives of impoverished people. Some philosophers have questioned these beliefs, arguing that our obligation to help the world's poor is as strong as our obligation to take a bleeding stranger to the hospital[12,13]. The author argues that neuroscience can help to explain why we have this pair of putatively inconsistent attitudes and that an improved understanding of these attitudes might lead us to change them.

allegiance, or are some more reliable than others? Our answers to this question will probably be affected by an improved understanding of where our intuitions come from, both in terms of their proximate psychological/neural bases and their evolutionary histories.

Consider the following moral dilemma (adapted from Unger[12]). You are driving along a country road when you hear a plea for help coming from some roadside bushes. You pull over and encounter a man whose legs are covered with blood. The man explains that he has had an accident while hiking and asks you to take him to a nearby hospital. Your initial inclination is to help this man, who will probably lose his leg if he does not get to the hospital soon. However, if you give this man a lift, his blood will ruin the leather upholstery of your car. Is it appropriate for you to leave this man by the side of the road in order to preserve your leather upholstery (FIG. 1a)?

Most people say that it would be seriously wrong to abandon this man out of concern for one's car seats. Now consider a different case (also adapted from Unger[12]), which nearly all of us have faced. You are at home one day when the mail arrives. You receive a letter from a reputable international aid organization. The letter asks you to make a donation of two hundred dollars to their organization. The letter explains that a two-hundred-dollar donation will allow this organization to provide needed medical attention to some poor people in another part of the world. Is it appropriate for you to not make a donation to this organization in order to save money (FIG. 1b)?

Most people say that it would not be wrong to refrain from making a donation in this case. And yet this case and the previous one are similar. In both cases, one has the option to give someone much needed medical attention at a relatively modest financial cost. And yet, the person who fails to help in the first case is a moral monster, whereas the person who fails to help in the second case is morally unexceptional. Why is there this difference?

About thirty years ago, the utilitarian philosopher Singer argued that there is no real moral difference between cases such as these two, and that we in the affluent world ought to be giving far more than we do to help the world's most unfortunate people[13]. (Singer currently gives about 20% of his annual income to charity.) Many people, when confronted with this issue, assume or insist that there must be 'some good reason' for why it is alright to ignore the severe needs of unfortunate people in far off countries, but deeply wrong to ignore the needs of someone like the unfortunate hiker in the first story. (Indeed, you might be coming up with reasons of your own right now.)

Maybe there is 'some good reason' for why it is okay to spend money on sushi and power windows while millions who could be saved die of hunger and treatable illnesses. But maybe this pair of moral intuitions has nothing to do with 'some good reason' and everything to do with the way our brains happen to be built.

To explore this and related issues, my colleagues and I conducted a brain imaging study in which participants responded to the above moral dilemmas as well as many others[14]. The dilemma with the bleeding hiker is a 'personal' moral dilemma, in which the moral violation in question occurs in an 'up-close-and-personal' manner. The donation dilemma is an 'impersonal' moral dilemma, in which the moral violation in question does not have this feature. To make a long story short, we found that judgements in response to 'personal' moral dilemmas, compared with 'impersonal' ones, involved greater activity in brain areas that are associated with emotion and social cognition. Why should this be?

An evolutionary perspective is useful here. Over the last four decades, it has become clear that natural selection can favour altruistic instincts under the right conditions, and many believe that this is how human altruism came to be[15]. If that is right, then our altruistic instincts will reflect the environment in which

they evolved rather than our present environment. With this in mind, consider that our ancestors did not evolve in an environment in which total strangers on opposite sides of the world could save each others' lives by making relatively modest material sacrifices. Consider also that our ancestors did evolve in an environment in which individuals standing face-to-face could save each others' lives, sometimes only through considerable personal sacrifice. Given all of this, it makes sense that we would have evolved altruistic instincts that direct us to help others in dire need, but mostly when the ones in need are presented in an 'up-close-and-personal' way.

What does this mean for ethics? Again, we are tempted to assume that there must be 'some good reason' why it is monstrous to ignore the needs of someone like the bleeding hiker, but perfectly acceptable to spend our money on unnecessary luxuries while millions starve and die of preventable diseases. Maybe there is 'some good reason' for this pair of attitudes, but the evolutionary account given above suggests otherwise: we ignore the plight of the world's poorest people not because we implicitly appreciate the nuanced structure of moral obligation, but because, the way our brains are wired up, needy people who are 'up close and personal' push our emotional buttons, whereas those who are out of sight languish out of mind.

This is just a hypothesis. I do not wish to pretend that this case is closed or, more generally, that science has all the moral answers. Nor do I believe that normative ethics is on its way to becoming a branch of the natural sciences, with the 'is' of science and the 'ought' of morality gradually melding together. Instead, I think that we can respect the distinction between how things are and how things ought to be while acknowledging, as the preceding discussion illustrates, that scientific facts have the potential to influence our moral thinking in a deep way.

Neuroscience and meta-ethics

Philosophers routinely distinguish between ethics and 'meta-ethics'. Ethics concerns particular moral issues (such as our obligations to the poor) and theories that attempt to resolve such issues (such as utilitarianism or Aristotelian virtue ethics). Meta-ethics, by contrast, is concerned with more foundational issues, with the status of ethics as a whole. What do we mean when we say something like "Capital punishment is wrong"? Are we stating a putative fact, or merely expressing an opinion? According to 'moral realism' there are genuine moral facts, whereas moral anti-realists or moral subjectivists maintain that there are no such facts. Although this debate is unlikely to be resolved any time soon, I believe that neuroscience and related disciplines have the potential to shed light on these matters by helping us to understand our common-sense conceptions of morality.

I begin with the assumption (lamentably, not well tested) that many people, probably most people, are moral realists. That is, they believe that some things really are right or wrong, independent of what any particular person or group thinks about it. For example, if you were to turn the corner and find a group of wayward youths torturing a stray cat[16], you might say to yourself something like, "That's wrong!", and in saying this you would mean not merely that you are opposed to such behaviour, or that some group to which you belong is opposed to it, but rather that such behaviour is wrong in and of itself, regardless of what anyone happens to think about it. In other words, you take it that there is a wrongness inherent in such acts that you can perceive, but that exists independently of your moral beliefs and values or those of any particular culture.

This realist conception of morality contrasts with familiar anti-realist conceptions of beauty and other experiential qualities. When gazing upon a dazzling sunset, we might feel as if we are experiencing a beauty that is inherent in the evening sky, but many people acknowledge that such beauty, rather than being in the sky, is ultimately 'in the eye of the beholder'. Likewise for matters of sexual attraction. You find your favourite movie star sexy, but take no such interest in baboons. Baboons, on the other hand, probably find each other very sexy and take very little interest in the likes of Tom Cruise and Nicole Kidman. Who is right, us or the baboons? Many of us would plausibly insist that there is simply no fact of the matter. Although sexiness might seem to be a mind-independent property of certain individuals, it is ultimately in the eye (that is, the mind) of the beholder.

The big meta-ethical question, then, might be posed as follows: are the moral truths to which we subscribe really full-blown truths, mind-independent facts about the nature of moral reality, or are they, like sexiness, in the mind of the beholder? One way to try to answer this question is to examine what is in the minds of the relevant beholders. Understanding how we make moral judgements might help us to determine whether our judgements are perceptions of external truths or projections of internal attitudes. More specifically, we might ask whether the appearance of moral truth can be explained in a way that does not require the reality of moral truth.

As noted above, recent evidence from neuroscience and neighbouring disciplines indicates that moral judgement is often an intuitive, emotional matter. Although many moral judgements are difficult, much moral judgement is accomplished in an intuitive, effortless way. An interesting feature of many intuitive, effortless cognitive processes is that they are accompanied by a perceptual phenomenology. For example, humans can effortlessly determine whether a given face is male or female without any knowledge of how such judgements are made. When you look at someone, you have no experience of working out whether that person is male or female. You just see that person's maleness or femaleness. By contrast, you do not look at a star in the sky and see that it is receding. One can imagine creatures that automatically process spectroscopic redshifts, but as humans we do not. All of this makes sense from an evolutionary point of view. We have evolved mechanisms for making quick, emotion-based social judgements, for 'seeing' rightness and wrongness, because our intensely social lives favour such capacities, but there was little selective pressure on our ancestors to know about the movements of distant stars.

We have here the beginnings of a debunking explanation of moral realism: we believe in moral realism because moral experience has a perceptual phenomenology, and moral experience has a perceptual phenomenology because natural selection has outfitted us with mechanisms for making intuitive, emotion-based moral judgements, much as it has outfitted us with mechanisms for making intuitive, emotion-based judgements about who among us are the most suitable mates. Therefore, we can understand our inclination towards moral realism not as an insight into the nature of moral truth, but as a by-product of the efficient cognitive processes we use to make moral decisions. According to this view, moral realism is akin to naive realism about sexiness, like making the understandable mistake of thinking that Tom Cruise is objectively sexier than his baboon counterparts. (Note that according to this view moral judgement is importantly different from gender perception. Both involve efficient cognitive processes that give rise to a perceptual phenomenology, but in the case of gender perception the phenomenology is veridical: there really are mind-independent facts about who is male or female.)

Admittedly, this argument requires more elaboration and support, and some philosophers might object to the way I have framed the issue surrounding moral realism. Others

might wonder how one can speak on behalf of moral anti-realism after sketching an argument in favour of increasing aid to the poor. (Brief reply: giving up on moral realism does not mean giving up on moral values. It is one thing to care about the plight of the poor, and another to think that one's caring is objectively correct.) However, the point of this brief sketch is not to make a conclusive scientific case against moral realism, but simply to explain how neuroscientific evidence, and scientific evidence more broadly, have the potential to influence the way we understand morality. (Elsewhere I attempt to make this case more thoroughly[17].)

Understanding where our moral instincts come from and how they work can, I argue, lead us to doubt that our moral convictions stem from perceptions of moral truth rather than projections of moral attitudes. Some might worry that this conclusion, if true, would be very unfortunate. First, it is important to bear in mind that a conclusion's being unfortunate does not make it false. Second, this conclusion might not be unfortunate at all. A world full of people who regard their moral convictions as reflections of personal values rather than reflections of 'the objective moral truth' might be a happier and more peaceful place than the world we currently inhabit[17].

The maturation of human morality will, in many ways, resemble the maturation of an individual person. As we come to understand ourselves better — who we are, and why we are the way we are — we will inevitably change ourselves in the process. Some of our beliefs and values will survive this process of self-discovery and reflection, whereas others will not. The course of our moral maturation will not be entirely predictable, but I am confident that the scientific study of human nature will have an increasingly important role in nature's grand experiment with moral animals.

Joshua Greene is at the Department of Psychology and the Center for the Study of Brain, Mind, and Behavior, Princeton University, Princeton, New Jersey 08544, USA.
e-mail: jdgreene@princeton.edu

doi: 10.1038/nrn1224

1. Doris, J. M. & Stich, S. P. in *The Oxford Handbook of Contemporary Analytic Philosophy* (eds Jackson, F. & Smith, M.) (Oxford Univ. Press, New York, 2003).
2. Hume, D. *A Treatise of Human Nature* (eds Selby-Bigge, L. A. & Nidditch, P. H.) (Clarendon, Oxford, 1739/1978).
3. Moore, G. E. *Principia Ethica* (Cambridge University Press, Cambridge, 1903/1959).
4. Spencer, H. *Data of Ethics* (Kessinger, Belle Fourche, 1883/1998).
5. Sturgeon, N. L. in *Essays on Moral Realism* (ed. Sayre-McCord, G.) 229–255 (Ithaca, Cornell University Press, 1988).
6. Darwall, S., Gibbard, A. & Railton, P. Toward a fin de siecle ethics: some trends. *Philos. Rev.* **101**, 115–189 (1992).
7. Casebeer, W. D. Moral cognition and its neural constituents. *Nature Rev. Neurosci.* **4**, 841–847 (2003).
8. Aristotle. *Nicomachean Ethics* (translated by Irwin, T.) (Hackett, Indianapolis, 1985).
9. Kant, I. *Groundwork of the Metaphysic of Morals* (translated by Patton, H. J.) (Harper and Row, New York, 1785/1964).
10. Mill, J. S. *Utilitarianism* (ed. Crisp, R.) (Oxford Univ. Press, New York, 1861/1998).
11. Haidt, J. The emotional dog and its rational tail: a social intuitionist approach to moral judgment. *Psychol. Rev.* **108**, 814–834 (2001).
12. Unger, P. *Living High and Letting Die: Our Illusion of Innocence* (Oxford Univ. Press, New York, 1996).
13. Singer, P. Famine, affluence, and morality. *Philos. Public Affairs* **1**, 229–243 (1972).
14. Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M. & Cohen, J. D. An fMRI investigation of emotional engagement in moral judgment. *Science* **293**, 2105–2108 (2001).
15. Sober, E. & Wilson, D. S. *Unto Others: The Evolution and Psychology of Unselfish Behavior* (Harvard Univ. Press, Cambridge, Massachusetts, 1998).
16. Harman, G. *The Nature of Morality* (Oxford Univ. Press, New York, 1977).
17. Greene, J. D. *The Terrible, Horrible, No Good, Very Bad Truth About Morality and What to Do About It.* Ph.D. Thesis, Princeton University (2002).

Acknowledgements:
Many thanks to W. Casebeer, A. Herberlein and L. Nystrom for their valuable comments.

Online links

MIT Encyclopedia of Cognitive Sciences:
http://cognet.mit.edu/MITECS/
limbic system: http://cognet.mit.edu/MITECS/Entry/markowitsch
moral psychology: http://cognet.mit.edu/MITECS/Entry/deigh
social cognition: http://cognet.mit.edu/MITECS/Entry/gilbertd
theory of mind: http://cognet.mit.edu/MITECS/Entry/gopnik

Online Bibliography of Cognitive Science and Ethics:
http://hem.passagen.se/ollekillen/cogethics.html

Biography

Bill Casebeer is an assistant professor of philosophy at the US Air Force Academy. He took his interdisciplinary Ph.D. in cognitive science and philosophy at the University of California at San Diego in 2001. He is a Major in the US Air Force, where he also serves as an intelligence analyst. His research interests include the neural mechanisms of moral judgement, naturalized ethical theory, the philosophy of biology and psychology, and the dynamics of terrorist recruitment. He is co-authoring a book on violent non-state actors (such as terrorist organizations), and is accomplishing fMRI work in the cognitive neurobiology of moral decision-making.

ONLINE GREENE

Biography

Joshua Greene is a cognitive neuroscientist and philosopher whose experimental research uses functional magnetic resonance imaging to study the neural bases of moral judgement. He received his Ph.D. in philosophy from Princeton University in 2002 for work examining the foundations of ethics in light of recent advances in psychology, neuroscience and evolutionary theory. He is currently a postdoctoral fellow in the Department of Psychology and Center for the Study of Brain, Mind and Behavior at Princeton University.