

## Introducing *Neuroethics*

Neil Levy

Received: 25 January 2008 / Accepted: 29 January 2008 / Published online: 14 February 2008  
© Springer Science + Business Media B.V. 2008

Do we really need another journal? There has been a proliferation of new journals in philosophy, ethics and related fields in recent years. One might well suspect that we have been multiplying journals not merely beyond necessity, but beyond desirability. Is yet another journal, this one with a title that may seem merely a faddish buzzword, really a good idea? Moreover, does its subject matter really require a new journal? Many people think that neuroethics is merely a subdiscipline of bioethics, and that the issues it raises are best treated in the context of existing bioethical journals.

Whether or not there are too many journals published today, it seems to me that the journal *Neuroethics* is really required. The growth of bioethics was a response to the explosion in bioethical knowledge and its applications, knowledge which promised to transform our understanding of health and well-being. New medical technologies provoked new questions—for

instance, about the beginnings and the end of life, about what it means to be human, about the risks of hubris—and provoked a growing feeling of unease. This same unease, provoked by a new range of technologies raising new kinds of questions, is replicated today by the growth of neuroscientific knowledge and its applications. But the new questions require new ways of thinking and new concepts, not merely the application or even the extension of existing bioethical concepts. Hence the need for a new discipline, and for a new journal.

Neuroethics, as I use the term (following [16]), refers to two, closely interrelated, enterprises. First, it refers to ethical reflection on new technologies and techniques produced by neuroscience (and other sciences of the mind). These questions are closely analogous to the kinds of issues that are the traditional territory of bioethics; just as the latter ponders questions about the application of new biomedical techniques (is cloning permissible? When should we turn off respirators and other life-support equipment? Should genetic enhancement be permitted?), so neuroethics attempts to answer questions about the applications of neuroscientific knowledge: does the use of psychopharmaceuticals threaten our self-conception? Should evidence from brain imaging be admissible in criminal proceedings? Are psychopaths responsible agents? And so on. These questions are analogous to bioethical issues, but they are sufficiently different to warrant the birth of a new discipline.

---

N. Levy (✉)  
Centre for Applied Philosophy and Public Ethics,  
University of Melbourne,  
Parkville, VIC, Australia  
e-mail: nllevy@unimelb.edu.au

N. Levy  
The James Martin 21st Century School,  
University of Oxford,  
Oxford, England

The second branch of neuroethics is more dramatically different from the enterprise of bioethics. It refers to the ways in which the new knowledge emerging from the sciences of the mind illuminates traditional philosophical topics: What is the nature of morality? What explains losses of self-control? When are beliefs justified? How should knowledge be pursued? These questions, going to the very heart of what it means to be a human being, have no real analogue in bioethics. The two branches of neuroethics interact, producing a generally new discipline, one to which bioethicists have much to contribute, but which is equally the province of neuroscientists, philosophers, psychologists, sociologists and lawyers (to name just a few of the disciplines abutting and feeding into neuroethics).

Biomedical knowledge promised, and still promises, to transform our understanding of life. Neuroscientific knowledge promises to transform our understanding of something yet more intimate: of what it means to be a thinking being. One way to grasp the profundity of the new challenge is by thinking about the difference between bioethics and neuroethics in Cartesian terms. Descartes famously distinguished between two fundamentally different kinds of substance, *res extensa* and *res cogitans*; extended things and thinking things, or more familiarly, matter and mind. He held that human beings, unlike all other animals, were amalgams of both kinds of substance, consisting of both matter and mind. Mind is immaterial, extensionless and immortal; it can survive the destruction of the body. It is, in short, the soul. Cartesian (substance) dualism is no longer taken seriously; the relation between the brain and the mind is too intimate for it to be at all plausible. But the temptation to identify *ourselves* with our *minds* remains strong (indeed, as Paul Bloom [5] has argued, it is possible that substance dualism is the innate default view of human beings). From this perspective, we can grasp the qualitatively different, and more radical, challenge that neuroscientific advances pose to our self conception compared to medical advances. Whereas medical advances, important as they are, deal with our *bodies*, neuroscientific discoveries promise—or threaten—to reveal the structure and functioning of our *minds* and, therefore, of our souls.

Regardless of the truth or falsity of dualism, there is a lot to be said, if not for simply identifying the self with the mind, then at least for taking the mind to be

the core of the self. And there really seems to be a sense in which neuroscience (and the other sciences of the mind) is stripping back the mysteries of the mind in sometimes disturbing ways, threatening our notion of ourselves as autonomous, rational and moral beings. In what follows, I will briefly sketch ways in which the sciences of the mind seem to force us to confront the possibility of a major shift in our self-conception, focusing on each of these fundamental categories. I will also show how these challenges pose practical and ethical questions that must be confronted.

## Rationality

It is central to our self-conception that we are rational beings. According to Aristotle, “Man is a rational animal”. We, alone of all animals, are capable of guiding our actions by reasons, and distinguishing, in the light of reason, between what is and what merely seems to be. Our rationality is not only definitive of what we most essentially are, it is also what is most prized in us, providing us with a standard to live up to. For Aristotle once again, the life of reflection was the highest to which we could aspire; for Socrates the unexamined life was not worth living and for John Stuart Mill it was better to be a Socrates dissatisfied than a pig satisfied. For us, the merely animal (unreflective) life is a life that is unworthy. But the sciences of the mind threaten our image of ourselves as rational animals.

They do this in many ways. First, they apparently show that far fewer of our actions are guided by reasons than we might have thought. The evidence here comes largely from work in social psychology, on the *automaticity* of actions. Automatic actions are effortless, ballistic (uninterruptible once initiated) and typically unconsciously initiated; that is, they are not made in response to conscious reasons of ours but are instead more like reflexes, triggered by features of the situation in which we find ourselves. In the influential terminology introduced by Stanovich [18], automatic actions are system 1 processes, not slow, effortful, conscious and deliberative system 2 processes. System 1 processes are evolutionarily more ancient; they are the kind of cognitive process we share with many other animals, whereas system 2 processes are the kind distinctive of us. If we are rational animals, and that is what distinguishes us, it is only inasmuch as

we deploy system 2 processes that this is true. The threatening finding from social psychology is not that we often deploy system 1 processes; it is that these are *by far* the more common. The overwhelming majority of human actions are caused by automatic mental processes [2]. In the light of the sciences of the mind, our claim to be rational animals suddenly looks somewhat shaky.

Worse is to come. Even when we do deploy system 2 processes, the rationality of our thought is less than we might have hoped. The evidence for this claim comes largely from cognitive psychology, especially work in the heuristics and biases traditions. *Heuristics* are mental short cuts and rules of thumb that we deploy, usually without realizing we are doing so; *biases* are the ways in which we weight the significance of information in making judgments. There is a huge mass of evidence showing that when we assess arguments or make decisions, we deploy such heuristics and biases, often in ways that mislead us. I shall mention only a few of the ways in which we assess information badly.

Human beings are pervasively subject to the *confirmation bias*, a systematic tendency to search for evidence that supports a hypothesis we are entertaining, rather than evidence that refutes it, and to interpret ambiguous evidence so that it supports our hypothesis [15]. The confirmation bias (along with a substantial dose of wishful thinking) helps to explain many people's belief in supernatural events. Suppose your hypothesis is that dreams foretell the future. The confirmation bias makes it likely that you will pay attention to confirming evidence (that time you dreamt that your aunt was unwell, only to learn that around that time she had a bad fall) and disregard disconfirming evidence (all the times when you dreamt about good or bad things happening to people you know when no such event occurred). The confirmation bias works in conjunction with the *availability heuristic*, our tendency to base assessments of the probability of an event on the ease with which instances can be brought to mind [22]. Because confirming instances are more easily recalled, memory searches, carried out in good faith, lead us to conclude that our hypothesis is true.

You may think that the tendency to believe in the supernatural is harmless and trivial. This may or may not be right (think of the occasional cases of parents preferring to have their seriously ill children treated

by new-age healers rather than qualified physicians), but there is no doubt that the kind of biases at issue here do real world harm. One instance is the recent rash of claims involving 'recovered memories' of sexual assault. There is no evidence that *any* such recovered memories were true, but we do know that many of them were false. There is therefore no reason to regard such memories as reliable. Yet on the basis of this evidence, many people were imprisoned, and many more families ruptured irrevocably. Why was there this sudden rash of recovered memories? Part of the explanation lies in the techniques used by some therapists to elicit possible repressed memories. Since they believed that these memories were deeply repressed, they encouraged their patients to visualize events they could not recall, or to pretend that they happened. But these techniques are known to be effective in producing false memories, or in otherwise bringing people to mistake imaginings for reality [13]. Why did they do this? Confirmation bias helps to explain their behaviour—they noticed that patients sometimes appeared to improve when they used these techniques, and ignored alternative explanations of these improvements (was the mere fact that someone was listening to them helping their mental state? Might the passing of time by itself be playing a role?) and ignored cases in which the techniques failed to help [19]. Ignorance of our systematic biases and cognitive limitations—for instance, on the part of patients who take the vividness of a 'memory' as evidence of its veracity, of therapists who are unaware of the need to test hypotheses systematically, and courts who take sincere memory and eyewitness testimony as irrefutable evidence—can causes great harm.

The example of repressed memory has two morals for us. First, it helps to suggests how the issues dealt with by neuroethics are practically important. Applying the knowledge gained from the sciences of the mind, in court rooms and in clinical practices, would lead to less harm and more good. Second, however, we should appreciate how disturbing is the evidence of the limitations of our rationality, the fallibility of our memory and the unreliability of our experience as a guide to reality. We think we are rational beings; we think that our memories are transcriptions of past events, we think that we have a good grasp of what the world immediately around us is like, but we may be wrong.

## Autonomy

Closely connected to our sense of ourselves as rational beings is our belief that we are autonomous choosers. We are autonomous precisely because we can distance ourselves from our beliefs and attitudes and assess them, before proceeding to act on our assessment. If we are less rational than we think, then we are less autonomous as well. But in addition to the threats to our rationality, there are independent threats to our ability to act autonomously. Roughly, we are autonomous to the extent to which we can bring our behavior into line with our considered judgments. There is evidence that our power to do so is more limited than we like to think. Once again, I have space to review only some small proportion of the evidence here, which comes from every single one of the sciences of the mind.

Consider, first, the data on *hyperbolic discounting* [1]. It is rational to discount future goods; that is, to think that a good available in the future is worth less to you right now than it would be when you actually receive it. For instance, if I offer you a dollar now or two dollars in 3 months time, you might rationally prefer to take the dollar now. This might be the rational choice for several reasons: because you cannot be certain to get the money in the future (I might be untrustworthy; you might die in the interim) or because you expect to have less need of the money then than now. Discounting is not evidence of irrationality. But hyperbolic discounting is evidence of irrationality. I discount future goods hyperbolically when my discount function is not linear; in fact, if it is mapped on a graph, it produces a highly bowed curve. When discount curves are bowed like this, they can cross, and my preferences can be highly unstable. Hyperbolic discounters experience preference reversals of the following sort: asked on Monday whether I prefer one dollar on Tuesday or two on Wednesday, I might choose to wait until Wednesday and take the two dollars. But if I discount the future hyperbolically, the closer I get to rewards the more I value them; hence on Tuesday my preferences might shift, leading me to prefer the immediately available one dollar even though I know that if I take it, I cannot have two dollars tomorrow. Since I prefer the dollar, I take it. Predictably, however, I soon regret my choice, wishing I had waited.

Hyperbolic discounting helps to explain many failures of autonomy characteristic of human beings.

Most obviously, and as the example just given suggests, it helps to explain why people typically do not save as much money for the future as they think they should. They may sincerely judge that they ought to save, but immediately available rewards prove too tempting, and they spend their money on goods they quickly regret purchasing. It also explains more extremes breakdowns of autonomy, such as addictive behavior [12]. Why do apparently rational individuals, who sincerely say that they prefer to give up their drug, often go back to it, and often long after they have gone through the pain of withdrawal? Hyperbolic discounting is an important part (though only a part) of the answer: because when the drug becomes available, they experience a preference shift. The preference shift can, after all, look so enticing, so apparently rational. ‘Just this once’, one says to oneself; after all, it’s not as though taking the drug once does much damage (think of smoking: having just one cigarette doesn’t damage one’s health much, and what little damage it does might well be outweighed by the pleasure it gives). The problem, of course, is that the situation is endlessly repeated, and ‘just this once’ becomes *always*.

Hyperbolic discounting might be a large part of the explanation for one of the greatest public health challenges facing developed nations: the challenge of obesity. Why do people persistently choose to consume more calories than they need, even when they know that their over-consumption risks shortening their lives, and they sincerely assert that they prefer living longer to eating cheeseburgers? Part of the explanation may come from the fact that fast food, highly palatable food in general, is widely and easily available, and when it is immediately available we experience a reversal of preferences. We continue to think that we should eat moderately and skip dessert, but ‘just this once’ we will indulge ourselves.

Another part of the explanation of the obesity crisis, and of losses of autonomy more generally, might be the phenomenon known as *ego depletion* [3]. Roughly, the ego depletion hypothesis is the theory that self-control is effortful, and that engaging in it draws upon a special reserve of energy. When that reserve is depleted, self-control becomes progressively more difficult. The evidence for this theory comes from studies in which subjects are divided into two groups. One group performs a task that requires self-control—say watching a funny movie without smiling—while

the other performs a task that does not require self-control (say, rating various options without actually choosing between them). Then both groups are given a common self-control task: say holding one hand in icy water (the ‘cold pressor task’) or attempting to solve an anagram puzzle that is in fact insoluble. The finding is that subjects in the ego depletion group persist a significantly shorter time at the self-control task than subjects in the control group. The conclusion of the researchers is that self-control resources are depleted when they are drawn upon, and that when self-control reserves are low, engaging in tasks that require self-control becomes much more difficult.

It may be that ego depletion is also at work in instances of hyperbolic discounting; since resisting tempting rewards draws on the reserves of self-control, ego depletion might explain the preference reversals typical of the hyperbolic discounter. There is independent evidence that ego depletion causes preference reversals, rather than merely overcoming the agent: depleted individuals not only choose immediate rewards that are tempting but which they are later likely to regret, they also choose to commit themselves to future rewards that are like this. For instance, they will not only choose to eat candy now, rather than fruit, when depleted, they will also choose trashier films to watch in a few days time when they are depleted than they would otherwise [4]. Alone or together, however, both ego depletion and hyperbolic discounting seem to threaten our autonomy, in the sense that they are significant obstacles to getting ourselves to act in accordance with our considered judgments. If autonomy requires that we be able to bring our actions into line with our considered judgments made in ideal conditions—as I have argued elsewhere—then the pervasiveness of ego-depletion and hyperbolic discounting are threats to autonomy.

Once again, the research outlined above has both immediate practical implications and philosophical implications for us. First, it helps to explain failures of autonomy, and thereby suggests strategies for preventing them in the future. These strategies are both individual and social. As individuals, we can take steps to structure our environments to prevent preference reversals and to keep our self-control resources plentiful. We can ensure that we are not tempted by immediate rewards that might cause preference reversals. For instance, we can put our money in fixed term accounts, removing the tempta-

tion to spend it on luxuries, or we can put time locks on our drinks cabinet, ensuring that we have to wait until evening before we indulge. Most simply, we can buy a small candy bar, rather than a big one; that way we know that if our resolve to save it for tomorrow crumbles, the damage to our diet will be limited. We can avoid shopping after a stressful day, when our reserves of self-control are at a low ebb. Finally, there is evidence that we can practice self-control, thereby increasing our internal reserves [14].

More can be accomplished by institutions seeking to increase agents’ autonomy. Governments can require individuals to contribute to retirement plans or social security; they can ban, limit the sale of, or tax goods that are highly rewarding and which might therefore be very tempting; they can regulate the opening hours of bars; they can restrict the content of advertising. These kinds of measures are frequently dismissed as paternalistic, but it is far from obvious whether this epithet is justified. Paradigm paternalism forces agents to act against their own wishes but for their own good; the measures envisaged instead force agents to act in ways in which they themselves endorse, if not at the moment of action, then at least in a cool hour. Far from constituting paternalistic interventions, then, they might be seen as promoting personal liberty [21].

Beyond the important practical question concerning what autonomy promoting measures are justified, however, we also need to assess the philosophical challenge to our conception of ourselves as autonomous agents posed by the scientific findings mentioned. Assessing whether the suggested interventions are paternalistic might require us to take a stand on issues of personal identity, and the right of one person-stage to bind another. We might be required to reassess what it means to choose freely and responsibly. We might need to rethink our practices of criminal justice, or perhaps even to jettison the notion of accountability. Neuroethical questions quickly lead to profound philosophical issues, allowing us to see some of the oldest and deepest intellectual challenges in a new light.

## Morality

We have briefly sketched some of the ways in which neuroethics forces us to rethink two characteristics traditionally held to be distinctive or definitive of



human beings, our rationality and our autonomy. There is a third traditional answer to the question what makes us special: we are *moral* beings. Once again, however, neuroethics sheds new light upon, and may seem to threaten, this prized characteristic of ourselves.

A great deal of work, in both psychology and neuroscience, seems to demonstrate that emotion plays a much greater role in moral judgment than most people, and philosophers, have thought. Some of this work seems to put pressure on philosophical theories of moral judgment. One of the most influential theories, the theory that (arguably at least) underlies the notion of human rights, is *deontology*, the theory, most closely associated with Immanuel Kant, that morality is basically about rights and duties. One way to understand deontology and its associated rights and duties is as follows: these rights and duties place constraints on what we may do to improve general welfare. That is, we ought always to improve welfare, *except when* doing so would infringe a right; then we have a duty to refrain from acting to improve general welfare. Consider is a well-known illustration, the famous trolley problem [8]. The problem is designed to demonstrate how rights constrain welfare maximization. In the problem, we are presented with two variants of a scenario in which we might act to maximize welfare, by saving the greater number of people:

(1) Imagine you find yourselves by the tracks when you see an oncoming trolley heading for a group of five people. The people cannot escape from their predicament and will certainly be killed if you do nothing. In front of you is a lever; if you pull it, you will divert the trolley to a side-track, where it will certainly hit and kill one person. Should you pull the lever?

Most philosophers have the intuition that we ought to pull the lever; moreover, most ordinary people, tested by the growing number of psychologists interested in morality, agree [6]. But now consider this variation on the problem:

(2) Imagine you find yourself on a bridge over the tracks when you see an oncoming trolley heading for a group of five people. The people cannot escape from their predicament and will certainly be killed if you do nothing. Next to you

is a very large man. You realize that if you push the large man onto the tracks, his great bulk will stop the trolley (whereas your slight frame will not); he will certainly die, but the five people on the tracks will be safe. Should you push the large man?

Most philosophers have the intuition that you should *not* push the large man; once again, most ordinary people agree. At first glance, this is puzzling: the cases seem to be relevantly similar. In both, you are faced with the choice of acting to save five people at the cost of one. Why should it be right to save the five in case (1), but not (2)?

The standard answer is that people have rights, including a right to life, and that pushing the large man would infringe his rights. But redirecting the trolley is not infringing anyone's rights (perhaps because we use the large man as a means to an end—were it not for his bulk, we could not stop the trolley—but since the presence of the man on the side-track is not necessary for stopping the trolley, we do not use him as a means). But recent research by neuroscientists has thrown doubt on this explanation.

Greene et al. [10] scanned the brains of subjects considering the trolley problem and similarly structured dilemmas. They found that when subjects consider *impersonal* dilemmas—in which harms caused are not up close and personal—regions of the brain associated with working memory showed a significant degree of activation, while regions associated with emotion showed little activation. But when subjects considered *personal* moral dilemmas, regions associated with emotion showed a significant degree of activity, whereas regions associated with working memory showed a degree of activity *below* the resting baseline. The authors plausibly suggest that the thought of directly killing someone is much more personally engaging than is the thought of failing to help someone, or using indirect means to harm them. But the real significance of this result lies in the apparent threat it poses to some of our moral judgments. What it apparently shows is that only some of our judgments—those concerned with maximizing welfare—are the product of rational thoughts, whereas others are the product of our rational processes being swamped by raw emotion. This result has been taken as evidence for discounting deontological intuitions, in favor of a thoroughgoing consequentialism [17].

If Greene's results seem to challenge one important class of moral judgments, revealing them to be irrational, other work seems to threaten the entire edifice of morality, conceived of as a rational enterprise. In a series of studies, Jonathan Haidt has apparently shown that ordinary people's moral judgments are driven by their emotional responses, and that the theories they offer to justify their judgments are post hoc confabulations, designed to protect their judgments [11]. We take ourselves to reason our way to our moral judgments, but in fact our reasons are just rationalizations, Haidt suggests. Together with Wheatley, Haidt has shown that inducing emotional responses using post-hypnotic suggestion influences people's moral judgments [23]. These results seem to suggest that the idea, beloved of philosophers, that morality is responsive to reasons is false. They also threaten the notion that moral argument can lead to moral progress.

Once again, the implications of this work for our self-conception are potentially dramatic. When we proudly proclaim that we are moral animals, we do not mean that our behaviour is driven by affective responses, in the kinds of ways which characterize the reciprocal altruism and sense of fairness possessed by chimps, monkeys, and even much simpler animals (see [20, 7]). Instead, we pride ourselves on a rational morality, which transcends our merely animal inheritance. This flattering image of ourselves may need heavy qualification. More immediately and practically, there may be policy implications of some of these findings. If, for instance, it can be shown that some (and only some) of our moral responses are irrational, because driven by raw emotion, then we have a powerful reason for rewriting policy to discount these responses.

I shall not pursue these questions further, having done so at length elsewhere [12]. I do not aim—or feel able—to solve these problems; I aim only to demonstrate the range, practical significance, and sheer fascination of the kinds of issues with which neuroethics is concerned. Neuroethics is at the confluence of a number of the most significant currents in recent thought; it also promises to help illuminate some of our oldest and deepest puzzles. Its importance can hardly be overstated. Hence the pressing need for a new journal.

## The Papers

I am proud to present the inaugural issue of *Neuroethics*. The papers gathered here, by some of the most im-

portant contemporary neuroethicists, reflect the range and interest of the new discipline, mixing ethical reflection with a consideration of the deep philosophical issues raised by advances in the sciences of the mind.

Martha Farah, an influential cognitive neuroscientist as well as an important neuroethicist, asks what light neuroscientific evidence can shed on the traditional philosophical problem of other minds: roughly, the question of whether we can know that others have minds at all (given that we only have direct access to their behaviour, not their thoughts). As she points out, this is not merely an abstract philosophical question, but of direct ethical relevance: having a mental life is a necessary condition of having certain kinds of ethical status.

The primary methodology used to investigate mental states neuroscientifically, of course, is functional magnetic resonance imaging (fMRI). How reliable is this technology? Some people have argued that reliance on fMRI—the belief that we can deduce mental states and functions from brain images—is phrenology revived. In her paper, Adina Roskies—a holder of doctorates in both philosophy and neuroscience, and also an important neuroethicist—examines the epistemic value of brain images. She argues for a position midway between the scepticism of those who dismiss the technology as new-wave phrenology and its uncritical supporters. Neuroimaging does give us insight into mental states and functions, but we need to be aware of the inferential distance between the image and the brain.

In their contribution, Julian Savulescu and Anders Sandberg tackle another question at once philosophical and practical: the permissibility and advisability of using psychopharmaceuticals to produce or enhance love. This is a philosophical issue, inasmuch as love is often felt to be somehow transcendent, rooted in what is finest, and most spiritual, in us. Can love survive being understood in terms of neurotransmitters and hormones? Savulescu and Sandberg suggest that it can, and set out conditions under which it might be permissible to use 'love drugs'.

Walter Glannon, author of one of the first monographs on neuroethics [9], devotes his paper to reflections on the permissibility of psychopharmaceuticals to enhance mental functions. He argues that we are still desperately short of sufficient information with regard to the risks and benefits of existing psychopharmaceuticals, but that there are grounds for—

cautiously—regarding their use, by individuals prepared to take responsibility for the consequences, as permissible. The last substantive paper is by Sheri Alpert. If in this introduction I have emphasized the ways in which neuroethics differs from other cognate disciplines, Alpert emphasizes the similarities, and especially the risk of ethical myopia which comes from failing to learn from others. Using the example of the other new ethical subfield, nanoethics, Alpert argues for the importance of learning from responses to problems that are often analogous.

The last contribution to this inaugural issue is from Cordelia Fine. Her paper is the first of the regular ‘perspectives’ papers *Neuroethics* will publish, shorter papers reflecting on recent research in the sciences of mind, for instance presenting new findings that might otherwise have gone unnoticed by people outside the field, or (as in this case) opinion pieces presenting a viewpoint. In her perspective piece, Fine tackles what she calls ‘neurosexism’; the use of neuroscientific research to demonstrate that male and female brains are radically different. Findings which seem to show dramatic sex differences always make good press, but as Fine shows, the science upon which they are based is either shaky or badly misinterpreted by those who abuse it.

*Neuroethics* welcomes new submissions, of substantive research articles by people working in all the fields which deal with scientific knowledge of the mind as well as the applications of this knowledge, and of shorter perspectives pieces. We have some exciting future issues planned already. Our aim is to sustain the high quality amply demonstrated by the papers published here.

## References

1. Ainslie, G. 2001. *Breakdown of will*. Cambridge: Cambridge University Press.
2. Bargh, J.A., and T.L. Chartrand. 1999. The unbearable automaticity of being. *American Psychologist* 54: 462–479.
3. Baumeister, R.F., E. Bratslavsky, M. Muraven, and D.M. Tice. 1998. Ego-depletion: Is the active self a limited resource? *Journal of Personality and Social Psychology* 74: 1252–1265.
4. Baumeister, R.F., E.A. Sparks, T.F. Stillman, and K.D. Vohs. 2008. Free will in consumer behavior: Self-control, ego depletion, and choice. *Journal of Consumer Psychology* 18: 4–13.
5. Bloom, P. 2004. *Descartes’ baby*. New York: Basic Books.
6. Cushman, F.A., L. Young, and M.D. Hauser. 2006. The role of conscious reasoning and intuitions in moral judgment: Testing three principles of harm. *Psychological Science* 17: 1082–1089.
7. de Waal, F. 1996. *Good natured: The origins of right and wrong in humans and other animals*. Cambridge, MA: Harvard University Press.
8. Foot, P. 1978. *The problem of abortion and the doctrine of double effect*. In *Virtues and vices*, ed. P. Foot, 19–32. Oxford: Basil Blackwell.
9. Glannon, W. 2006. *Bioethics and the brain*. Oxford: Oxford University Press.
10. Greene, J., R.B. Sommerville, L.E. Nystrom, J.M. Darley, and J.D. Cohen. 2001. An fMRI investigation of emotional engagement in moral judgment. *Science* 293: 2105–2108.
11. Haidt, J. 2001. The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review* 108: 814–834.
12. Levy, N. 2006. Autonomy and addiction. *Canadian Journal of Philosophy* 36: 427–448.
13. Loftus, E.D. 1993. The reality of repressed memories. *American Psychologist* 48: 518–537.
14. Muraven, M., R.F. Baumeister and D. Tice. 1999. Longitudinal improvement of self-regulation through practice: Building self-control strength through repeated exercise. *Journal of Social Psychology* 139: 446–457.
15. Nickerson, R.S. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology* 2: 175–220.
16. Roskies, A. 2002. Neuroethics for the new millenium. *Neuron* 35: 21–23.
17. Singer, P. 2005. Ethics and intuitions. *Journal of Ethics* 9: 331–352.
18. Stanovich, K.E. 1999. *Who is rational? Studies of individual differences in reasoning*. Mahwah, NJ: Erlbaum.
19. Tavis, C., and E. Aronson. 2007. *Mistakes were made (but not by me)*. Orlando: Harcourt.
20. Trivers, R. 1985. *Social evolution*. Menlo Park, CA: Benjamin/Cummings.
21. Trout, J.D. 2005. Paternalism and cognitive bias. *Law and Philosophy* 24: 393–434.
22. Tversky, A., and D. Kahneman. 1973. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology* 5: 207–232.
23. Wheatley, T., and J. Haidt. 2005. Hypnotic disgust makes moral judgments more severe. *Psychological Science* 16: 780–784.