

Ethics

Intermediate article

John M Doris, University of California, Santa Cruz, California, USA
 Stephen Stich, Rutgers University, New Brunswick, New Jersey, USA

CONTENTS

Introduction: what is ethics?

The relation of cognitive science to ethics

Applied ethics

Normative ethics

Meta-ethics

The methods of ethics

Ethics is the philosophical study of standards of conduct and value. Such inquiry often encounters questions regarding human motivation, affect, and cognition that can be addressed by empirical and theoretical work in the cognitive sciences.

INTRODUCTION: WHAT IS ETHICS?

Contemporary philosophical ethics is customarily divided into three subdisciplines, here ordered by increasing theoretical abstraction: applied ethics, normative ethics, and meta-ethics (see Darwall, 1998 for a useful survey).

Applied ethics addresses concrete ethical problems in areas such as law, medicine, business, and agriculture. While work in applied ethics may involve the advocacy of particular positions, philosophers often aim only to clarify, through philosophical methods of argumentation and analysis, the issues surrounding such controversial topics as abortion, euthanasia, capital punishment, and the development of genetically modified organisms.

Normative ethics investigates obligation, virtue, and value; very often, work in this area consists in attempts to adduce general claims about the sorts of things that are morally valuable and general principles for the regulation of conduct and choice. In contemporary normative ethics, three approaches are especially prominent. Proponents of virtue ethics, inspired especially by Aristotle, maintain that human excellence is the fundamental evaluative consideration; conduct is to be judged by reference to such virtues as courage and temperance. Consequentialists and Utilitarians, working in the tradition of Bentham and Mill, argue that right actions are those that, of the options available, maximize the total quantity of a valuable end such as pleasure or happiness. Followers of Kant, or Kantians, argue that an act is morally permissible only in so far as it proceeds from a principle that

could coherently serve as a law governing the choice of all rational agents.

Meta-ethics concerns the structure of moral reasoning and justification, the sources of moral motivation, and the nature of moral concepts. Perhaps the most persistent issue in meta-ethics concerns the objectivity of ethical inquiry. Is ethical discourse merely the statement or expression of preference or emotion, as various brands of skepticism about moral objectivity maintain? Or is there some method for discovering objectively true answers to ethical questions, or conferring rational justification on ethical judgments, as moral realists have argued?

THE RELATION OF COGNITIVE SCIENCE TO ETHICS

Historically, many moral philosophers have been skeptical about the relevance of the sciences to ethics. In the spirit of Hume (1778/1740, p. 469), it is argued that because scientific inquiry is primarily descriptive or factual and ethical inquiry is primarily prescriptive or normative, there is an unbridgeable 'logical gap' between the *is* of the sciences and the *ought* of ethics. However much the sciences may tell us about human beings, how the beings in question should comport themselves remains – to borrow an influential formulation of Moore's – an open question. For example, contemporary life science has discovered much about the development and functioning of the fetus, yet this has, in the eyes of many, done little to resolve the abortion controversy.

But there have always been those who viewed this gap between *is* and *ought* as illusory, and much recent philosophical work suggests that the barrier between description and prescription may not be as impermeable as followers of Hume and Moore suppose (Railton, 1995). Certainly moral philosophers often offer arguments – such as those

regarding the nature of moral motivation and cognition – whose presuppositions appear amenable to empirical investigation. In so far as this is the case, cognitive science bears a strong *prima facie* relevance to ethics. Accordingly, we will not attempt to adjudicate the continuing philosophical controversy regarding the relation of scientific to ethical inquiry, but will instead identify interfaces between cognitive science and ethics that merit further empirical study and philosophical discussion.

APPLIED ETHICS

Many issues in applied ethics raise important epistemological questions – questions about what people can and cannot be reasonably expected to know. Other issues raise closely related questions about what people do and do not understand. The cognitive sciences have produced empirical findings relevant to such questions, and these findings often play a central role in debates in applied ethics. In this section we will give two brief illustrations.

One of the first domains in which experimental psychology was applied to an important ethical issue was the debate over reliance on eyewitness testimony in legal proceedings. Traditionally, the identification of a defendant by an eyewitness to a crime has played a central role in Anglo-Saxon law. However, almost 100 years ago, Munsterberg conducted a number of dramatic, though poorly controlled, experiments in which many eyewitnesses to a staged ‘crime’ gave seriously inaccurate reports of what they had seen. In a book that generated considerable controversy in the legal profession and beyond, Munsterberg (1908, p. 194) argued that it is ‘astonishing that the work of justice is ever carried out in the courts without ever consulting the psychologist and asking him for all the aid which the modern study of suggestion can offer’.

More recent work, using much improved methods, has demonstrated that people’s memories of events can be affected in dramatic and disquieting ways by information (or *misinformation*) that they encounter after the event. In one experiment, conducted in train stations and other naturalistic settings, Loftus (1979) and her students staged a ‘robbery’ in which a male confederate pulled an object from a bag that two female students had temporarily left unattended and stuffed it under his coat. A moment later, one of the women noticed that her bag had been tampered with and shouted: ‘Oh my God, my tape recorder is missing.’ She went on to lament that her boss had loaned it to her and that it was very expensive.

Bystanders, most of whom were quite cooperative, were asked for their phone numbers in case an account of the incident was needed for insurance purposes. A week later, an ‘insurance agent’ called the eyewitnesses and asked about details of the theft. Among the questions asked was ‘Did you see the tape recorder?’ More than half of the eyewitnesses remembered having seen it, and nearly all of these could describe it in detail – this despite the fact that *there was no tape recorder!* On the basis of this and other experiments, Loftus concludes that even casual mention of objects that were not present, or of events that did not take place (for example, in the course of police questioning), can significantly increase the likelihood that the objects or events will be incorporated into people’s memories. Other studies have shown that people are often mistaken when they identify someone as a participant in an event that they witnessed, and that both white and black subjects are considerably less accurate at identifying people of the other race (Loftus, 1979; Lloyd-Bostock and Clifford, 1983; Wells and Loftus, 1984).

Findings like these, along with a number of highly visible cases in which people have been convicted of crimes they did not commit on the basis of eyewitness testimony, have led one leading legal scholar to recommend that judges should be required ‘to direct the jury that it is not safe to convict upon eyewitness evidence unless the circumstances of the identification are exceptional or the eyewitness evidence is supported by substantial evidence of another sort’ (Devlin, 1976). Though this recommendation has not been widely accepted, the recent use of DNA testing to reverse wrongful convictions that were often at least partially based on unreliable eyewitness testimony has once again focused attention on the need for reform. Although some philosophers have questioned the wisdom of relying on empirical studies in this context (Coady, 1992), the research raises an important and troubling question: when, if ever, is it ethically defensible to allow legal decisions to be influenced by such an empirically suspect source of evidence?

Another area of research in cognitive science that has important implications for the legal system, and also for issues about medical and academic decision-making, focuses on the reliability of predictions about an individual’s future behavior based on interviews and ‘clinical’ judgments. In some American states, interviews with parole board members play a central role in deciding whether an offender will be granted a parole. In one state system, for example, interviewers are

asked to make a prediction of the risk that the offender will commit another crime. A study by Carroll *et al.* (1988) showed that psychologists who had never spoken with the prisoners could make much more accurate predictions by using purely statistical 'actuarial' methods and basing the prediction entirely on a few background factors, such as number of previous convictions. This finding is one of well over 100 studies comparing clinical to actuarial methods for predicting a wide range of behavioral phenomena, from academic success to the prognosis for patients admitted to a mental hospital. In just about all of these studies, actuarial methods are at least as good as clinical prediction, and often they are much better. These findings have led to considerable discussion about *why* the human cognitive system does such a bad job on these tasks, despite the often powerful conviction on the part of 'expert' judges that their professional experience enables them to make better predictions (Nisbett and Ross, 1980; Dawes, 1994). Though we do not currently have a good understanding of the mechanisms that give rise to clinical predictions, their unreliability raises serious questions about the moral justification for using such predictions in making important decisions.

NORMATIVE ETHICS

Making defensible moral judgments requires a sensitivity to ethically salient features of the environment: one is not likely to act compassionately if one is oblivious to suffering, nor is one likely to act justly in complex circumstances without some awareness of what justice requires. It appears as though individuals vary on this dimension; just as some are especially attuned to the dictates of etiquette or fashion, others seem especially attuned to the demands of morality.

This sort of observation has been prominent in philosophical thought regarding the virtues, or excellences of character. According to Aristotelians such as McDowell (1979), virtue involves a 'reliable sensitivity' to ethically salient considerations. The requisite sensitivity must be both *reliable* and *robust*; while even morally mediocre people may manifest ethical sensitivity in particular instances, the virtuous person manifests appropriate ethical sensitivity in all circumstances where she can reasonably be expected to do so, regardless of whether this exercise is difficult or easy.

However appealing the foregoing view may be, experimental social psychology suggests that ethical sensitivity is typically far from robust.

A large body of research indicates that people's moral judgments and behavior are extraordinarily sensitive to features of the situation in which they are embedded, including many features that most people would agree are morally irrelevant. (See Doris (2002) for details and references.) Experimental studies of helping behavior provide some particularly disconcerting examples. Finding a bit of change or hearing the noise from a lawn mower can make the difference between helping and not – results that likely would have astounded Aristotle. In a classic demonstration by Darley and Batson, subjects performed tasks at two separate sites. The behavior of interest occurred when subjects walked from one site to the other, passing the experimenters' assistant slumped in a doorway, apparently in some distress. Before leaving the first site, subjects were told either that they were running late ('high hurry'), were right on time ('medium hurry'), or were a little early ('low hurry'); thus members of each group experienced a different degree of time pressure while travelling from one site to the next. Helping behavior varied markedly with degree of hurry: 63 per cent offered help in 'low hurry', 45 per cent offered help in 'medium hurry', and only 10 per cent offered help in 'high hurry'. Evidently some subjects wanted to help but reluctantly ceded to the demands of punctuality, while for others (more germane here), time pressures muted their sensitivity to ethically salient stimuli – in some instances subjects stepped unconcernedly over the stricken 'victim', apparently without registering the scene. There is no reason to think that the subjects in these studies were aberrant in their sensitivity to situational factors; it is quite reasonable to suppose that most people exhibit very substantial situational variability in their moral judgments and behavior, whatever we are antecedently inclined to say about their character.

These empirical issues do not by themselves show that a conception of virtue as involving a robust and reliable ethical sensitivity is of no use to ethics. Perhaps virtue is to be understood, as some philosophers have suggested, in terms of a rarely attainable ideal that serves to focus people's ethical aspirations (Blum, 1994, pp. 94–6), or perhaps the empirical findings do not unsettle a suitably sophisticated moral psychology of virtue. Or perhaps, as others have argued, results such as those just recounted suggest major revisions for philosophical thinking on virtue and character (Doris, 1998, 2002; Harman, 1999). It remains to be seen how compelling these alternatives are, but it is obvious that virtue ethics is in a fertile tension with empirical work.

META-ETHICS

Meta-ethics is concerned, in part, with the nature of moral concepts, the structure of moral reasoning, and the nature of moral motivation, and there are many ways in which the findings and theories in cognitive science are relevant to these concerns. In this article we have space for only a single example.

As we noted earlier, the Kantian tradition in moral philosophy emphasizes the role of reason in ethics, and many Kantians have invoked reason to explain why people should be motivated to act morally. The explanation they propose is that it would be irrational not to be motivated to do what one morally ought to do (or what one believes one morally ought to do), in much the same way that it would be irrational not to believe the conclusion of a sound argument (Nagel, 1970, p. 3). One famous obstacle to proposals of this sort is suggested by Hume's (1777/1966, pp. 282ff) hypothetical example of the 'sensible knave' – a person who recognizes that it is morally wrong to be dishonest in a situation where he would benefit from dishonesty, but who is not moved at all by this judgment. More recent writers have urged that Hume's sensible knave is more than a philosophical fiction since there actually are psychopaths who know the difference between right and wrong but simply have no motivation to *do* what is right.

Those who wish to defend the link between reason and moral motivation have adopted two quite different strategies, both of which appear to make substantive empirical assumptions of the sort that cognitive science might test (Nichols, 2002). The first strategy relies on a claim about our ordinary concept of moral judgment: the concept of moral judgment entails that 'agents who make moral judgments are motivated accordingly' (Smith, 1994, p. 66). Philosophers who adopt this strategy recognize that psychopaths may *say* that something is 'morally required' or 'morally wrong' and that they may be quite sincere. But if psychopaths are not motivated in the appropriate way, their words do not mean what non-psychopaths mean by these words and the concepts they express with these words are not the ordinary moral concepts that non-psychopaths use. This strategy only works, of course, if it is true that our ordinary moral concepts *require* that people to whom the concepts apply have the appropriate sort of motivation. And, for two quite different reasons, it is far from clear that this is the case. First, there is considerable disagreement in cognitive science about whether and how concepts are structured, and about how we are to determine when something is built into

or entailed by a concept (Margolis and Laurence, 1999). Second, there has been almost no serious empirical work aimed at uncovering the structure of ordinary people's moral concepts. There has, however, been work on ordinary people's epistemic concepts, and that work suggests that people in different ethnic and socioeconomic groups use significantly different epistemic concepts (Weinberg *et al.*, forthcoming). If the same is true for moral concepts, it will pose a serious challenge for the conceptual approach to the problem posed by sensible knaves and psychopaths.

A second strategy for defending the connection between reason and moral motivation takes the link to be empirical rather than conceptual. Those who adopt this approach maintain that it is an empirical fact that the faculty of reason, when functioning normally, generates motivation to do what one judges one ought to do, just as it is an empirical fact that the faculty of reason generates an inclination to believe the conclusions of arguments one judges to be sound. If psychopaths fail to have the appropriate moral motivation, it is argued, then there must be something wrong with their reasoning faculty. But is this true? The answer is far from clear.

Blair (1995) has shown that psychopaths *do* exhibit surprising deficits on a cluster of cognitive tasks that have been used frequently by psychologists who study moral development. In these tasks subjects are presented with descriptions of various transgressions such as a child hitting another child, or a child leaving the classroom without the teacher's permission. In the moral development literature, transgressions of the first sort are labeled 'moral' while those of the second sort are labeled 'conventional'. From quite early on in childhood, normal children distinguish moral from conventional transgressions on a number of dimensions: they view moral transgressions as more serious, they explain why the acts are wrong by appeal to different factors (harm and fairness for moral transgressions, social acceptability for conventional transgressions), and they understand conventional transgressions, unlike moral transgressions, to be dependent on authority. If the teacher says there is no rule about leaving without permission, children think it is OK to leave without permission. But if the teacher says there is no rule against hitting other children, they do not think that hitting is acceptable. What Blair found was that incarcerated psychopaths do not draw the moral/conventional distinction. Though normal children and normal adults (including a control group of non-psychopath prisoners) have no

trouble classifying new cases in one category or the other, psychopaths fail to do so.

These results might well give some support to the hypothesis that psychopaths have a reasoning deficit, and thus that they do not pose a problem for those who maintain that a properly functioning reasoning faculty always generates some motivation to do what one believes one ought to do. But, as Nichols (2002) has pointed out, the issue cannot be so easily resolved, because psychopaths have also been shown to have affective responses that are quite different from those of normal subjects. When shown distressing stimuli (e.g. slides of people with dreadful injuries) and threatening stimuli (e.g. slides of an angry man wielding a weapon), normal subjects exhibit much the same suite of physiological responses. Psychopaths, by contrast, exhibit normal physiological responses to threatening stimuli, but abnormally low physiological responses to distressing stimuli (Blair *et al.*, 1997). Thus, Nichols argues, it may well be that what is wrong with psychopaths is not that their reasoning system is abnormal, but that their affect system is abnormal, and that it is their affective abnormalities that are responsible for their inability to draw the moral/conventional distinction. If that is the case, then Hume's challenge continues to pose a major problem for those who think there is a link between reason and moral motivation. Resolving this issue will require conceptual clarification on how to draw the boundary between reason and affect, and on what counts as an *abnormality* in each of these domains. It will also require much more empirical work aimed at understanding exactly how psychopaths and non-psychopaths differ. Clearly, this important issue in meta-ethics cannot be addressed responsibly without taking account of the increasingly rich body of empirical findings generated by the cognitive sciences.

THE METHODS OF ETHICS

'Intuition pumps' or 'thought experiments' are among the most prominent elements of philosophical method. The technique is to present a hypothetical example and attempt to elicit some philosophically telling intuition; if the 'experiment' is successful, it may be concluded that competing theories must account for the resulting intuition. In such cases, those forwarding the thought experiment insist that for a theory to be viable, it must be consistent with the intuition elicited by the hypothetical example; if the theory is inconsistent with the intuition, the hypothetical is supposed to be a

counterexample to the theory. This is a central rhetorical strategy in ethics, as evidenced by debates over Utilitarianism. Consider, for example, an influential intuition pump proposed by Williams (1973, pp. 97–100), who asks us to imagine the following circumstances: George, a new PhD in chemistry in dire need of a job to support his family, is offered a chance at a job researching chemical and biological weapons (CBW). George is strongly opposed to CBW research, but if he does not take the job, it is likely to go to a man with rather alarming enthusiasm for CBW research. Williams maintains that Utilitarians are bound to say that George should take the job, presumably because doing so will both help his family and hinder CBW research, thus maximizing total available welfare. But Williams argues that 'many of us would certainly wonder whether ... that could possibly be the right answer at all'. Williams apparently thinks that in this case people will conclude that the Utilitarian prescription does not properly account for the value of George's 'integrity', a commitment to principle that would be undermined if he took the job.

In response, the Utilitarian may deny that ethical theory should be constrained by intuitions. But this is a rather unappealing expedient; a theory that makes no reference to intuitions risks becoming divorced from the experience of ethical life. On the other hand, it is also perilous to deploy intuitions as a constraint on ethical reflection. Given the ubiquity of moral controversy, cultures and individuals may vary dramatically in their moral intuitions; how is it to be decided which intuitions, and whose should serve as constraints? For example, Haidt *et al.* (1993) found evidence suggesting that Americans' ethical intuitions varied with their socioeconomic status. Returning to George the chemist, is Williams right to suppose that 'many of us' value integrity over food on the table, or is this intuition limited to the 'many of us with sufficiently high socioeconomic status'? A further look at cognitive science presses this question still harder: experimental work suggests that even within a single individual, moral intuitions may vary in quite arbitrary ways.

In an important study, Tversky and Kahneman (1981) presented a group of subjects with the following problem:

Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences of the programs are as follows:

If Program A is adopted, 200 people will be saved.

If Program B is adopted, there is a 1/3 probability that 600 people will be saved, and a 2/3 probability that no people will be saved.

A second group of subjects was given an identical problem, except that the programs were described as follows:

If Program C is adopted, 400 people will die.

If Program D is adopted, there is a 1/3 probability that nobody will die and a 2/3 probability that 600 people will die.

On the first version of the problem most subjects thought that Program A should be adopted. But on the second version most chose Program D, despite the fact that the outcome described in A is identical to the one described in C. The disconcerting implication of this study is that the moral decisions we make are strongly influenced by the manner in which the options are described or *framed*. The general worry is that ethical intuitions may be determined by features of thought experiments that are quite ethically irrelevant. To return once more to George the chemist, are responses to the thought experiment determined by reflection on the ethical substance of the case, or by ethically irrelevant features of the richly textured example? Until we can say with confidence how a given ethical intuition is generated, we should be hesitant to rely on it in argument.

This suggests a dilemma for philosophical ethics: it must either eschew appeal to thought experiments that have not been evaluated empirically, or eschew appeal to intuitions altogether. As we've already said, the abolition of intuitions is a problematic solution: it threatens to distance philosophical ethics from the experience of ethical life, and thereby alienate the study of ethics from the concerns of actual ethical agents (see Williams, 1985, pp. 93–119, esp. 116–119). But continued reliance on intuitions in the face of the difficulties just surveyed presents its own difficulty: intuitions, and the thought experiments generating them, must be subjected to systematic empirical scrutiny before they can be appealed to in adjudicating theoretical disputes. In other words, if practitioners of philosophical ethics wish to rely on intuitions, they must work to develop a cognitive science of ethics. If we are right, the implications of this are radical: philosophers must either depart from the traditional methods of ethics by abandoning the use of intuitions, or depart from the traditional methods of ethics by pursuing experimental work in cognitive science.

References

- Blair R (1995) A cognitive developmental approach to morality: investigating the psychopath. *Cognition* **57**: 1–29.
- Blair R, Jones L, Clark F and Smith M (1997) The psychopathic individual: a lack of responsiveness to distress cues? *Psychophysiology* **34**: 192–198.
- Blum LA (1994) *Moral Perception and Particularity*. Cambridge, UK: Cambridge University Press.
- Carroll J, Winer R, Coates D, Galegher J and Alibrio J (1988) Evaluation, diagnosis, and prediction in parole decision making. *Law and Society Review* **17**: 199–228.
- Coady C (1992) *Testimony: A Philosophical Study*. Oxford, UK: Clarendon Press.
- Darwall S (1998) *Philosophical Ethics*. Boulder, CO: Westview.
- Dawes R (1994) *House of Cards*. New York: The Free Press.
- Devlin P (Chair) (1976) Report to the Secretary of State for the Home Department of the Departmental Committee on Evidence and Identification in Criminal Cases. London, UK: Her Majesty's Stationery Office.
- Doris JM (1998) Persons, situations, and virtue ethics. *Noûs* **32**: 504–530.
- Doris JM (2002) *Lack of Character: Personality and Moral Behavior*. Cambridge, UK and New York, NY: Cambridge University Press.
- Haidt J, Koller S and Dias M (1993) Affect, culture, and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology* **65**: 613–628.
- Harman G (1999) Moral philosophy meets social psychology: virtue ethics and the fundamental attribution error. *Proceedings of the Aristotelian Society* **99**: 315–331.
- Hume D (1777/1966) *Enquiry Concerning the Principles of Morals*. New York, NY: Open Court (1966).
- Hume D (1740/1978) *A Treatise of Human Nature, 2nd edn*. Oxford, UK: Oxford University Press.
- Lloyd-Bostock S and Clifford B (1983) *Evaluating Witness Evidence*. New York, NY: Wiley.
- Loftus E (1979) *Eyewitness Testimony*. Cambridge, MA: Harvard University Press.
- Margolis E and Laurence S (1999) *Concepts*. Cambridge, MA: MIT Press.
- McDowell J (1979). Virtue and reason. *Monist* **62**: 330–350.
- Munsterberg H (1908) *On the Witness Stand*. New York, NY: Doubleday, Page.
- Nagel T (1970) *The Possibility of Altruism*. Princeton, NJ: Princeton University Press.
- Nichols S (2002) Is it irrational to be amoral? How psychopaths threaten moral rationalism. *Monist* **85**: 285–304.
- Nisbett R and Ross L (1980) *Human Inference*. Englewood Cliffs, NJ: Prentice-Hall.
- Railton P (1995) Made in the shade: moral compatibilism and the aims of moral theory. *Canadian Journal of Philosophy Supplementary Volume* **21**: 79–106.
- Smith M (1994) *The Moral Problem*. Oxford, UK: Blackwell.

- Tversky A and Kahneman D (1981) The framing of decisions and the psychology of choice. *Science* **211**: 453–463.
- Weinberg J, Nichols S and Stich S (forthcoming) Normativity and epistemic intuitions. *Philosophical Topics*.
- Wells G and Loftus E (eds) (1984) *Eyewitness Testimony: Psychological Perspectives*. Cambridge, UK: Cambridge University Press.
- Williams BAO (1973) A critique of utilitarianism In: Smart JJC and Williams BAO, *Utilitarianism: For and Against*. Cambridge, UK: Cambridge University Press.
- Williams BAO (1985) *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press.
- Gilligan C (1982) *In a Different Voice: Psychological Theory and Women's Development*. Cambridge, MA: Harvard University Press.
- Goldman A (1993) Ethics and cognitive science. *Ethics* **103**: 337–360.
- Harman G (2000) *Explaining Value and Other Essays in Moral Philosophy*. Oxford, UK: Oxford University Press.
- Kagan J and Lamb S (eds) (1987) *The Emergence of Morality in Young Children*. Chicago, IL: University of Chicago Press.
- Kohlberg L (1981) *Essays in Moral Development*, vol. 1: *The Philosophy of Moral Development: Moral Stages and the Idea of Justice*. San Francisco, CA: Harper & Row.
- Sober E and Wilson DS (1998) *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.
- Stich S (1993) Moral philosophy and mental representation. In: Hechter M, Nadel L and Michod RE (eds) *The Origin of Values*, pp. 215–228. New York, NY: Aldine de Gruyter.
- Wright R (1994) *The Moral Animal: The New Science of Evolutionary Psychology*. New York, NY: Pantheon Books.

Further Reading

- Alexander R (1987) *The Biology of Moral Systems*. New York, NY: Aldine de Gruyter.
- Flanagan O (1991) *Varieties of Moral Personality: Ethics and Psychological Realism*. Cambridge, MA: Harvard University Press.

Event-related Potentials

See **Auditory Event-related Potentials**
