

What is the Function of Morality?

M.J. Cain

Published in:

***Teleology and modernity* [ISBN: 9780815351030] / edited by William Gibson, Dan O'Brien, Marius Turda (Routledge, 2019).**

Introduction

Does morality have a function in the teleological sense of that term and, if so, what is that function? When I ask the question 'what is the function of morality?' I mean the term 'function' to be understood in teleological terms. Thus, the function of morality is a matter of what morality is 'for', what its purpose is, or what goal it is designed to achieve. According to the dominant answer to that question within both philosophy and cognitive science morality does indeed have a function and that function is to facilitate cooperation. My goal in this chapter is to raise some sceptical doubts about this answer through the examination of an important version of it that has recently been developed by Michael Tomasello in his book *A Natural History of Human Morality*. But before launching directly into Tomasello's work it is important to have a solid understanding of the nature of morality, on the one hand, and of function, on the other.

Morality and moral judgment

Moral evaluation is a prominent feature of human life; we frequently morally evaluate our own actual and potential behaviour and that of our fellows. When we do this we make a moral judgement. Such judgements are often given a linguistic expression that utilises such normative terms as 'ought', 'should', 'right', 'wrong', 'good', 'bad', 'fair', 'unfair', and so on. For example, one might say 'you shouldn't have stolen that book from the library' or 'it is a good thing to donate to charity.' Moral judgment is a human universal; it takes place in all cultures and is done by all normal people¹ and we begin doing it very early in our lives.² It is also probably unique to humans as not even higher primates such as chimps make moral judgements.³ When I talk about 'morality' in this

¹ R. Joyce, *The Evolution of Morality*, Cambridge, MA, MIT Press, 2006.

² P. Bloom, *Just Babies: The Origins of Good and Evil*, London, Bodley Head, 2013.

³ M. Tomasello, *A Natural History of Human Morality*, Cambridge, MA, Harvard University Press, 2016.

chapter I am referring to the human capacity to make moral judgements and the mental machinery that directly supports that capacity.

As the capacity to make moral judgements is central to morality then the question is raised as to the nature of moral judgement. In the case of some types of mental states there is a distinction between occurrent and standing states. For example, I believe that Henry James wrote *The Portrait of a Lady*, a belief that I have held for some time and continue to hold even when I am not actively considering the proposition that Henry James wrote *The Portrait of a Lady*. As Prinz⁴ would put it, this belief is stored in my long-term memory and is a dispositional or standing belief. In the situation where I actively consider the proposition that Henry James wrote *The Portrait of a Lady* and commit myself to its truth (for example, in response to someone asking me who wrote *The Portrait of a Lady*) then my belief is an occurrent belief and this state involves the activation of the corresponding standing belief stored in long term memory. It is plausible to say that the occurrent/standing distinction applies to moral judgement in that the active moral judgements that we make when confronted by certain behaviour or when planning our actions often reflect deeper seated commitments that persist over a period of time whilst we are not actively making such judgments. Thus, for example, the vegetarian who actively judges that it would be wrong to eat the ham sandwich just offered to them is drawing upon a long-term conviction that it is wrong to eat meat.

The most important feature of moral judgements relates to their normative character. Suppose that someone makes the moral judgement that eating meat is wrong. In doing this they are not making a straightforward factual claim as to how people generally behave. Few people who judge that eating meat is wrong are under the impression that people do not generally eat meat. Indeed, many of our moral judgements are negative: they involve evaluating someone for not behaving in a particular way. Thus, moral judgments are inherently normative: one is not judging how things are but how things ought to be or how people ought to behave.

But what is involved in judging that one ought to behave in a particular way even if people do not generally behave in that way? The immediate difficulty in answering this

⁴ J. Prinz, Furnishing the Mind, Cambridge, MA, MIT Press, 2002.

question is that not all apparently normative judgements are on the same footing. Sometimes judgments that we express using words such as 'should' and 'ought' are prudential as when I tell a student who has asked me how to get an A grade in their next essay that they ought to read more widely and engage in deeper analysis of the texts they discuss. When I do this I am saying that in order to achieve their goal it is necessary to do what I advise, that it would be prudent to do what I advise if they have the goal in question. The key point is that I'm not implying that they are under any obligation to do what I advise independently of their goal to get an A; if they drop that goal then my advice is no longer relevant to them. But, as Kant⁵ emphasised, moral judgments are not conditional in that way; when we make a moral judgment we are judging how people ought to behave whatever their goals.

This might appear to imply that moral judgments are connected to legal laws or social conventions so that making a moral judgement is akin to making a judgement about what the law or social convention requires. For, people sometimes violate the demands of the law or social convention but if a person doesn't care about the law or social convention then that does not imply that they do not apply to them.⁶ Nevertheless, it would be a mistake to equate moral judgments with judgments about the law or judgments about the operative social conventions. We do typically demand of the law that it engages with morality; for example, the moral wrongness of rape is a very good reason for making rape a crime. However, morality is independent of the law in the respect that much of the law does not relate to moral matters. Moreover, it is commonplace and perfectly coherent to judge that a particular type of behaviour has a particular moral status (negative or positive) while the law does not engage with it (for example, having an affair when one is in a long-term relationship). Similarly, it is commonplace and perfectly coherent to judge as morally legitimate a type of behaviour that you know to be illegal (as, for example, many liberal people did in the United Kingdom prior to the decriminalisation of homosexual acts between men in 1967).

⁵ I. Kant, Groundwork for the Metaphysics of Morals, A. Zweig (trans), Oxford, Oxford University Press, 1785/1998.

⁶ P. Foot, 'Morality as a System of Hypothetical Imperatives' in Philosophical Review, 81, 3, 1972, 305-316.

With respect to social conventions similar points can be made. Although one might judge that we have a moral duty to follow social conventions (perhaps, because to fail to do so would upset our fellows or undermine social cohesion) it is commonplace and perfectly coherent to judge a type of behaviour as morally legitimate even though it violates social convention (for example, having children out of wedlock, or men wearing dresses) or regard it as morally problematic even though it is demanded by social convention (for example, addressing a woman by means of a title that reflects her marital status). Indeed, classic research in developmental psychology has established that early in their development children draw a distinction between the moral and the conventional domains.⁷ In short then, moral judgements are inherently normative but the normatively involved is distinct from that connected with prudence, social convention and legality.

A third feature of moral judgements relates to their impartial character; when we make a moral judgement we commit to holding everyone in the relevant group, including ourselves, to the same standards. Thus, one cannot coherently judge that everyone apart from oneself ought to pay tax or that it is wrong for anyone apart from oneself to eat meat. That is not to say that one cannot morally excuse an individual on a particular occasion but any such exceptions need to be principled as when one excuses someone who steals when they had no alternative way of providing for their hungry children.

A fourth feature of moral judgements is that they generally have a strong motivational force; as Richard Joyce⁸ puts it, they have 'practical clout'. Thus, for example, if someone judges that eating meat is wrong then they are generally motivated to avoid eating meat and to respond to others who eat meat by criticising, shunning or punishing them. For those committed to moral judgement internalism (such as Michael Smith⁹) this relationship between moral judgment and motivation is

⁷ E. Turiel, The Development of Social Knowledge: Morality and Convention, Cambridge, Cambridge University Press, 1983.

⁸ Joyce, The Evolution of Morality, 57.

⁹ M. Smith, The Moral Problem, Oxford, Basil Blackwell, 1994.

necessary whereas for moral judgment externalists such as Shafer-Landau¹⁰ the link is contingent. I will not take a stand on this issue in the context of this chapter.

A fifth feature of moral judgements is that they are strongly connected to our emotions. Put bluntly, we generally care about moral matters so that we often feel guilt or shame when we judge we have done something wrong and feel angry or indignant towards those those we judge to have done wrong. Once again, there is a debate as to whether or not this link is necessary or contingent with philosophers influenced by Hume on the one side¹¹ regarding it as necessary and their opponents¹² regarding it as contingent. I won't take a stand on this issue here.

A sixth feature of moral judgement relates to their interpersonal nature. We do sometimes make moral judgements that relate to the conduct of the individual in isolation from other people, as when someone judges that a person has a moral obligation not to waste their talents. However, moral judgement generally relates to how individuals ought to engage with their fellows. Indeed, Joyce¹³ goes so far as to say that a Robinson Crusoe would have no need for morality.¹⁴

I do not pretend to have given an exhaustive account of the nature of moral judgement in this section but I take myself to have made a number of points that are relatively uncontroversial and which collectively provide sufficient understanding of

¹⁰ R. Shafer-Landau, Moral Realism: A Defence, Oxford, Oxford University Press, 2003.

¹¹ For example, S. Nichols, Sentimental Rules: The Natural Foundations of Moral Judgment, Oxford, Oxford University Press, 2004 and J. Prinz, The Emotional Construction of Morals, Cambridge, MA, MIT Press , 2007.

¹² For Example, J. Mikhail, Elements of Moral Cognition, Cambridge, Cambridge University Press, 2011.

¹³ Joyce, The Evolution of Morality.

¹⁴ See B. Williams, Ethics and the Limits of Philosophy, Cambridge, MA, Harvard University Press, 1985 and T. Scanlon, What We Owe to Each Other, Cambridge, MA, Harvard University Press, 1998 for important expressions of the view that moral judgement inherently concerns our interactions with our fellows.

what I mean by 'morality' when I raise the question as to the function of morality. It is now time to turn to the question of how we are to understand the term 'function.'

Functions

The function of morality is a matter of what morality is 'for', what its purpose is, or what goal it is designed to achieve. This notion of function could do with some clarification. In recent philosophy, especially in the philosophy of mind and in discussions of the explanation of complex capacities, the term 'function' is often not understood in a teleological manner. For example, when philosophers of mind attempt to characterise types of mental state in functional terms what they mean by 'functional' has to do with causation.¹⁵ Thus, a functionalist about pain is claiming that what makes a mental state pain has to do with its causes and effects rather than what it is for in any teleological sense. And when Fodor,¹⁶ Dennett,¹⁷ and Cummins¹⁸ champion a form of explanation of the complex capacities of a system that involves decomposing the system into components that perform less complex functions than the system to which they belong, they understand the function of a component in terms of that aspect of what it does that contributes to the performance of the target capacity rather than what it is for. For example, explaining how an internal combustion engine works would involve identifying such parts as the cylinder, crankshaft, carburettor, and so on, and identifying what each of these parts does and how they interact so as to engender the overall behaviour of the engine.

An obvious example of a function in an alternative teleological sense comes from the domain of artefacts. For example, the function of a kitchen knife is to cut food. The function of artefacts relates to what they are designed to do, made to do or, perhaps, used to do. This design, manufacture or use depends upon the mental states of some

¹⁵ N. Block, 'Advertisement for a Semantics for Psychology' in Midwest Studies in Philosophy, 10, 1986, 615-678.

¹⁶ J. Fodor, 'The Appeal to Tacit Knowledge in Psychological Explanation,' in Journal of Philosophy, 65, 1968, 627-640.

¹⁷ D. Dennett, Brainstorms, Cambridge, MA, MIT Press, 1978.

¹⁸ R. Cummins, 'Functional Analysis', Journal of Philosophy, 72, 1975, 741-765.

intelligent designer, manufacturer or subsequent user. For example, a kitchen knife has the function of cutting food rather than some other function or no function at all because it was designed, manufactured or used with the intention of cutting food by some intelligent agent.

However, it is possible to have a function in the teleological sense without being the product of intelligent design, manufacture or use. Ever since Darwin evolutionary biologists have attributed functions to the traits of organisms. For example, the function of the heart is to pump blood around the body and the function of the zebra's stripes is to camouflage it from potential predators.¹⁹ At first blush, such talk of function might seem odd given that one of Darwin's²⁰ central aims was to explain the complexity of organisms and why they are often so well-suited to the environmental niche that they inhabit without appeal to an intelligent designer.²¹ What solves the puzzle here is the so-called etiological account of function championed by Wright.^{22 23} The basic idea is that the function of a trait is a matter of the effects in virtue of which it was selected and continues to exist within the population. Consider an example. Suppose that as a result of genetic mutation an organism had a trait not possessed by its parents or most of the other members of its population. For example, suppose the organism was an insect that had a distinctive dark colouration.²⁴ As a result of this colouration the insect

¹⁹ Actually, there is some controversy as to the function of the zebra's stripes and recent research puts pressure on this familiar claim. For example, T. Caro, Zebra Stripes, Chicago, University of Chicago Press, 2016 argues that the function of the zebra's stripes is to deter parasitic flies.

²⁰ C. Darwin, On the Origin of Species, London, John Murray, 1859.

²¹ R. Dawkins, The Blind Watchmaker, New York, Norton, 1986; D. Dennett, Darwin's Dangerous Idea, London: Penguin, 1995.

²² L. Wright, 'Functions' in Philosophical Review, 82, 1973, 139-168.

²³ The etiological account has subsequently been developed by P. Godfrey-Smith, 'Functions: Consensus Without Unity', Pacific Philosophical Quarterly, 74, 1993, 196-208, P. Godfrey-Smith, 'A Modern History Theory of Functions' in Nous, 28, 1994, 344-362, K. Neander, 'The Teleological Notion of "Function"', in Australasian Journal of Philosophy, 69, 1991, 454-468, and R. Millikan, 'In Defense of Proper Functions' in Philosophy of Science 56, 1989, 288-302. For a helpful overview see P. Godfrey-Smith, Philosophy of Biology, Princeton, NJ, Princeton University Press, 2014 and K. Neander, 'Does Biology Need Teleology' in R. Joyce, (ed) The Routledge Handbook of Evolution and Philosophy, Abingdon, Routledge, 2018, 64-76.

²⁴ This example is modelled on that of the famous peppered moth.

was much harder to detect by predators and so lived for a longer period and had more offspring than it otherwise would have done. Moreover, it had a longer life and produced more offspring than its fellows with the alternative traditional light coloration. The offspring of the insect inherited the genetic basis of the dark colouration and so had the same colouration as their parent. This bestowed upon them a similar reproductive advantage so that the new dark colouration gradually became commonplace within the population. Thus, what explains why the dark colouration became widespread in the population and why it continues to persist many generations down the line has to do with its effects with respect to hiding the insect from predators. Thus, the function of this specific colouration is that of camouflaging the insect from predators and the trait has this function without the involvement of any intelligent designer. Of course evolution by natural selection is typically an incremental process with new mutations modifying traits that are based upon earlier mutations. This is somewhat obscured in the example of the insect as described thus far. But the darkening process could occur gradually as a first mutation led to an insect being slightly darker than its fellows, a subsequent mutation in a later generation causing a slightly darker colouration, and so on, with each step providing a defence against predation. In such a case, the colouration at each stage in the incremental process has the function of camouflaging the insect from predators.

In the context of this chapter I'm going to accept the etiological account as a viable account of how traits get their function in the teleological sense of that term. My justification for this is that such an account strikes me as being tacitly endorsed by many of those prominent figures who discuss the function of morality. Moreover, the prominent objections expressed in the philosophical literature are not particularly germane. One such objection is that the teleological notion of function is explanatorily inert.²⁵ For example, if we want to explain how the insect in the above example avoids being eaten by predators or how the heart pumps blood around the body, appealing to the selectional history of the insect or the heart is going to get us nowhere. Rather, what we need to appeal to are the intrinsic properties that ground the relevant causal powers. For example, suppose that an insect that is a molecule-for-molecule duplicate of the evolved dark insect spontaneously comes into existence in the manner of

²⁵ R. Cummins, 'Functional Analysis', 1975.

Davidson's swampman.²⁶ The two insects will be just as effective at avoiding being eaten by predators and what explains this are factors such as their colouration, that of the surfaces on which they alight, the visual capacities of the potential predators, and so on. In the case of each insect, the explanation of their predation-avoiding powers will be just the same, regardless of their having quite different selectional histories. Therefore, selectional history and the functions that it grounds isn't doing any explanatory work.

The standard response to this objection²⁷ is to follow Mayr²⁸ in distinguishing between how-questions and why-questions in biology. How-questions ask how a biological system or structure does what it does (for example, how does the insect avoid being eaten?) and why-questions ask why the system or structure is the way that it is or why it continues to exist in a particular form (for example, why does the insect have a dark colouration?). Now, the response continues, appeals to selectional history and teleological function are not capable of answering how-questions but they do answer why-questions, hence they are of explanatory value. In short, the objection misses an important distinction and condemns appeals to teleological function for failing to do something that they need not do.²⁹ My goal in this chapter is to discuss a particular account of the function of morality. The advocates of that account are clearly concerned with why-questions: why is it that we humans, in contrast to all other extant species, engage in moral evaluation, what aspects of our history led to the emergence of our capacity for moral judgement? Hence, I will assume that the appeal to the distinction between why-questions and how questions defeats any challenge to the point and legitimacy of asking the question as to the function of morality.

I should also point out that to regard organisms as having traits that have functions in the teleological sense is not thereby to endorse the adaptationist view³⁰ that most

²⁶ D. Davidson, 'Knowing One's Own Mind' in Proceedings and Addresses of the American Philosophical Association, 60, 1987, 441-458

²⁷ P. Godfrey-Smith, 'Functions: Consensus Without Unity', 1993.

²⁸ E. Mayr, 'Cause and effect in biology', Science, 134, 1961, 1501-1506.

²⁹ Not all philosophers who defend the appeal to teleological function in biological explanation accept that such functions are not relevant to answering how-questions. A recent example is K. Neander, 'Functional Analysis and Species Design' in Synthese, 194, 2017, 1147-1168.

³⁰ G. Parker and J. Maynard Smith, 'Optimality in Evolutionary Biology', in Nature, 348, 1990, 27-33.

of an organism's traits are adaptations or products of evolution by natural selection or to deny that some traits are spandrels³¹ or that such factors as genetic drift³² and constraints play an important role in evolution.³³ Thus, different traits can have different origins and it is an open question as to whether any given trait, including the human capacity to moralise, has an origin such that it has a particular function.

The etiological account of function doesn't just apply to traits that emerge through a process of biological evolution for it also applies in the case of cultural evolution. Cultural evolution is a perfectly real phenomenon that involves the development of a cultural product such as a type of behaviour or an idea over time. As in the case of biological evolution something must be transmitted from one generation to the next. However, the manner of transmission will not be genetic but, rather, will involve social learning where one individual picks up the behaviour or idea in question from another individual by copying.³⁴ Such learning is to be contrasted with individual learning where an individual learns something on their own by means of, for example, a process of trial and error. Social learning often involves teaching where, for example, a parent actively attempts to help their offspring to acquire a particular behaviour or idea, but it need not involve such teaching.³⁵

³¹ A spandrel is a trait that is not an adaptation itself but is a byproduct of some other trait that is an adaptation. See Gould and Lewontin, 'The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme' in Proceedings of the Royal Society of London, Series B, 205, 1979, 581-598.

³² Genetic drift is a change in the frequency of the variants of an existing gene in a population due to random factors. See A. Ariew and R. Lewontin, 'Confusions of Fitness' in British Journal for the Philosophy of Science, 55, 2004, 347-363.

³³ A constraint is a factor that restrains or limits evolutionary change within a population. See S. Gould, The Structure of Evolutionary Theory, Cambridge, MA, Harvard University Press, 2002.

³⁴ K. Laland, Darwin's Unfinished Symphony: How Culture Made the Human Mind, Princeton, NJ: Princeton University Press, 2017.

³⁵ K. Sterelny, The Evolved Apprentice: How Evolution Made Humans Unique, Cambridge, MA, MIT Press, 2012.

As with biological evolution, cultural evolution relies upon a copying process that is largely reliable but sometimes gives rise to 'error' where the product of copying is slightly different from the source that is copied. Thus, variants are introduced into the world. Sometimes these variant behaviours or ideas will be less effective than their forebears and will be abandoned or not copied by the next generation. But sometimes they will be more effective than their competitors and so will come to dominate the population as a result of being widely copied.

An example will help make this concrete. Suppose that an individual hunter-gatherer has mastery of a widely used technique for skinning small game. Another member of the group watches them apply this technique and attempts to copy it. However, the copying process isn't perfect as the individual learns a technique that is slightly different from that which they attempt to copy. In particular, it is quicker to perform, demands less energy, wastes less meat and damages the animal skin less than the original copied technique. This benefits the individual; they and their family get to eat more, have better quality animal skin for making clothes and have more time and energy for other pursuits. The advantages of this new technique is not lost on the other members of the group who themselves begin to copy it (often accurately) so that it comes to be the dominant technique for skinning small game in the group usurping the old technique. Several generations down the line there is another inaccuracy in copying the technique during the learning process which, once again, has benefits so that a new modified version of the technique becomes the most popular within the group. This process of copying with the occasional beneficial modification continues over many generations. Thus, we have a case of the cultural evolution of a technique for skinning small game, a phenomenon that bears sufficient similarities to biological evolution to count as a case of evolution in a non-metaphorical respect. It is possible that processes of thinking and reflection were involved in the evolution of the skinning technique. For example, the modifications might not always have come about through unintentional mistakes; rather, they could sometimes be the products of thoughtful attempts to improve the old technique. Hence, cultural evolution need not be as blind, and dependent upon fortuitous accidents as its biological relative and this explains why it proceeds at a considerably faster pace.

The existence of cultural evolution means that the etiological account of function can be applied to cultural products as well as to biological traits. Hence, to claim that morality has a specific function is not to thereby commit oneself to the view that

morality is an evolved biological trait. Accordingly, amongst philosophers and cognitive scientists who argue that morality has a cooperative function one can distinguish between those who think of morality as a biological adaptation³⁶ and those who think of morality more in cultural terms.³⁷

Cooperation

Now that I have explained what I mean by the terms 'morality' and 'function' I can turn to address the question as to the function of morality head on. Humans are social beings who generally live together in groups. But we are not merely social beings; in addition we are co-operators. Many of our endeavours involve working together with our fellows to achieve a common goal that we would not be able to achieve on our own. Of course, we are not the only animals that cooperate but the range and flexibility of our cooperative behaviour is unparalleled. Cooperation has enabled us to spread across the globe and build societies based upon rich cultures and sophisticated technologies that rarely stand still.

One view which dominates the literature is that the function of morality is to facilitate and support cooperation. Here is a clear expression of that view from Jonathan Haidt.³⁸

Moral systems are interlocking sets of values, virtues, norms, practices, identities, institutions, technologies, and evolved psychological mechanisms that work together to suppress or regulate self-interest and make cooperative societies possible.

³⁶ For example, R. Joyce, [The Evolution of Morality](#), 2006 and J. Greene, [Moral Tribes: Emotion, Reason and the Gap Between Us and Them](#), London Atlantic, 2013.

³⁷ For example, M. Tomasello, [A Natural History of Human Morality](#), 2016 and K. Sterelny, 'Moral Nativism: A Skeptical Response', in [Mind and Language](#), 25 2010, 279-297.

³⁸ J. Haidt, [The Righteous Mind: Why Good People are Divided by Politics and Religion](#), London, Penguin, 2013, 314.

In a similar vein Joshua Greene writes ‘morality is a set of psychological adaptations that allow otherwise selfish individuals to reap the benefits of cooperation.’³⁹

One of the most interesting, well-developed and important versions of the view that morality has a cooperative function has been developed by Michael Tomasello⁴⁰ (2016) as part of a broader project concerned with the origins of human cognition.⁴¹ It will be instructive to consider Tomasello’s perspective on morality in some detail.

Tomasello’s goal is to provide what he calls a ‘natural history of morality’, an account of how morality developed in the human species. Morality is something that emerged in the species – and also emerges in the development of the individual – as a result of a process of intellectual reflection and subsequent insight into our situation when we attempt to cooperate. Once it has developed, morality serves to further support such cooperative endeavours making them more likely to succeed. The underlying cooperative behaviour and the processes of reflecting upon it themselves depend upon evolved psychological mechanisms without which there would be no human morality. However, these mechanisms did not specifically evolve for morality.

For Tomasello morality has two components, the morality of sympathy and the morality of fairness. The morality of sympathy is at work when we help one of our fellows out of a concern for their well-being. The morality of fairness is at work when we make judgements concerning what is fair or just or what a person deserves based upon a sense of obligation that we ought to treat our fellows fairly. The morality of sympathy has a precursor in our great ape relatives who display sympathy towards their kin and friends. However, the morality of fairness is unique to humans.

According to Tomasello, human morality emerged in two stages, both of which involved attempts to find new ways of cooperating in response to ecological changes. The first stage occurred in the region of 400,000-500,000 years ago before the arrival of modern humans (*homo sapiens*) some 150,000 years ago. Due to deforestation

³⁹ J. Greene, Moral Tribes: Emotion, Reason and the Gap Between Us and Them, London Atlantic, 2013, 23.

⁴⁰ M. Tomasello, A Natural History of Human Morality, 2016.

⁴¹ Seminal works in that project include M. Tomasello, Constructing a Language, Cambridge, MA, Harvard University Press, 2002, M. Tomasello, Origins of Human Communication, Cambridge, MA, MIT Press, 2008 and M. Tomasello, A Natural History of Human Thinking, Cambridge, MA, Harvard University Press, 2015.

humans came under competition for food from ground-dwelling apes and to survive needed to work together as pairs in hunter-gatherer ventures. Such cooperation was necessary for survival and in this context co-operators had a stake in one another as the welfare of each depended on the welfare of the other. Accordingly, any motivation for cheating was undercut and in helping another one did not run a major risk that one's generosity would not be reciprocated. Thus, Tomasello's portrayal of the plight of early human co-operators is at odds with the putatively more gloomy perspective of those who emphasize the problematic nature of reciprocation and the strength and rationality of the motivation to cheat.⁴² One way of putting this is to say that the plight of our ancestors is better represented by the Stag Hunt⁴³ – where two individuals have to choose between hunting for small game on their own or working together to capture larger prey – than the Prisoner's Dilemma⁴⁴ – where two prisoners have to decide whether or not to confess to a crime and/or identify their partner in crime to the authorities. In the Prisoner's Dilemma, the rational thing for the individual to do is to identify their partner as having committed the crime but deny their own involvement. In contrast, in the Stag Hunt the rational thing to do is to cooperate.

Within the context of pairwise cooperation humans evolved psychological mechanisms that supported such cooperation and facilitated an appreciation of their plight such as to give rise to what Tomasello calls 'second-personal morality.' Such a morality holds only between pairs of co-operators whilst they are engaged in cooperative activities. Second-personal morality has two components, one relating to sympathy – as one extends one's concern for the well-being of one's kin and friends to one's cooperative partner – and the other relating fairness – as one feels an obligation to treat one's partner fairly and as they deserve to be treated.

The central mechanisms that underpin this second-person morality are threefold: first, joint intentionality, second, second-personal agency and third, joint commitment.

⁴² Such gloominess is alluded to in the above quotations from Haidt and Greene which refer to the challenge of selfishness and the need to suppress self-interest.

⁴³ B. Skyrms, The Stag Hunt and the Evolution of Social Structure, Cambridge, Cambridge University Press, 2004.

⁴⁴ W. Poundstone, Prisoners Dilemma, Game Theory and the Puzzle of the Bomb, Oxford, Oxford University Press, 1992.

Joint intentionality is exercised when two individuals both share certain mental states and are mutually aware of that commonality.⁴⁵ In the context of hunter-gatherer cooperation this involved mutual knowledge of a shared goal, mutual knowledge of the role of each partner in achieving the goal, and mutual knowledge of the ideal way to perform each of these roles. Such mutual knowledge is stored in what Tomasello calls ‘personal common ground.’ Reflecting upon their situation both cooperative partners come to appreciate that each constitutes an ‘I’ interacting with another – a ‘you’ – to form a joint agent, a ‘we.’ Further, they appreciate that their individual roles in the ‘we’ are interchangeable so that there is a relationship of equivalence between the partners.

The second mechanism underpinning second-personal morality, second-personal agency, involves treating partners and potential partners with mutual respect and a sense of desert, evaluating and choosing partners on the basis of their collaborative potential, and being aware that others evaluate oneself in the same way. What is crucial in the exercise of second-personal agency for the emergence of morality is the individual’s recognition that in virtue of their equivalence with their cooperative partner both deserve respect from the other should they perform their role in line with the mutually known role ideals.

The third mechanism involved in second-personal morality, joint commitment, is exercised when an individual identifies with a larger body – that is, the ‘we’ formed with their cooperative partner – and explicitly grants authority to that larger body so as to accept being judged by its standards as shared in common ground and to take responsibility for judging themselves and others by those standards. Thus, as a result of making joint commitments, individuals judge both themselves and their fellows impartially rather than in a self-serving manner, something that is crucial if their behaviour is to be based on a sense of how they ought to behave rather than being merely strategic.

Second-personal morality constitutes a very local form of morality that only holds between collaborative partners whilst they cooperate. However, with further social changes ‘a full-blown, group minded, cultural morality of “objective” right and wrong applying to everyone in all situations’ (2016: 84) began to emerge some 150,000 years

⁴⁵ An appeal to joint intentionality is a recurring theme of Tomasello’s output. For example, it is central to Tomasello’s (2002) anti-Chomskyan account of language acquisition.

ago. The crucial social changes involved the formation of cultural groups that brought together small groups of foragers in order to protect them from competitor groups. Within such cultural groups each individual entered into a relationship of interdependence with many other people, most of whom they would not know directly and could only recognise as a fellow group member on the basis of their behaving in a manner characteristic of group members. With the formation of cultural groups the morality of sympathy was scaled up as individuals came to feel sympathy towards all other members of their cultural group.

To operate in a cultural context the mechanisms of joint intentionality, second-personal agency and joint commitment that underpin second-personal morality needed to be scaled up to those of collective intentionality, cultural agency and moral self-governance and it is through the exercise of these mechanisms that a fully-fledged public morality of justice and fairness emerged.

Collective intentionality involves members of a cultural group having mutual knowledge of the conventional practices of that culture and of the ideal way to perform the various roles people perform within the group. Such knowledge constitutes cultural common ground and enables individuals to interact effectively with those who they do not know intimately. For, they will know how they are generally expected to behave in any given cultural context and how the other participants in that context can be relied upon to behave, be it – to use contemporary examples – a supermarket, an academic conference, a rock concert, a job interview or a blind date.⁴⁶

Cultural agency, the second component underpinning group-minded cultural morality, is a mechanism for controlling potentially uncooperative members of the group and rests upon social norms that specify how we must behave in various situations to be cooperative. Some of these norms are merely conventional but some have a moral character that ‘comes from the underlying second-person morality of sympathy and harm, fairness and unfairness, in which the norm is grounded.’⁴⁷ A norm against stealing would be an example of a norm that has a moral character in

⁴⁶ In tying cultural conventions to mutual knowledge, Tomasello is allying himself to a view of conventions following in the tradition most associated with D. Lewis, *Convention: A Philosophical Study*, Cambridge, MA, Harvard University Press, 1969 and in stark opposition to R. Millikan, ‘Language Conventions Made Simple’ in *Journal of Philosophy*, 95, 1998, 161-180.

⁴⁷ M. Tomasello, *A Natural History of Human Morality*, 2016, 100.

virtue of its relationship to second-personal morality in contrast to norms of etiquette. An important element of cultural agency is an evolved tendency to enforce social norms on others even when we are an unaffected bystander.

Moral self-governance, the third component underpinning group-minded cultural morality, is a mechanism for self-regulation in a social context that relies upon what Tomasello (following Korsgaard,⁴⁸ calls 'reflective endorsement' along with guilt. Reflective endorsement occurs when an individual, in the course of action planning, reflects on whether a particular goal or value is a good one to pursue. Guilt arises when an individual reflects retrospectively on her actions and judges that she has done something wrong and so deserves punishment. Such moral self-governance is underpinned by an identification with the creators of the social norms in terms of which the individual judges her actual and potential behaviour so that those norms are seen as something that 'we' created for our own benefit.

With respect to the development of morality in the individual, Tomasello sketches a process that resembles that of the emergence of morality in the species. Thus, a child first develops a second-personal morality between the ages of 1 and 3 where they help spontaneously, share equally with others and honour and expect their fellows to honour joint commitments. It is only from the age of 3 that children develop a group-minded cultural morality that reflects the norms of their cultural group. But once such a morality has developed, the second-personal morality is not abandoned; rather, the group-minded cultural morality rests on top of a simpler second-personal morality and their demands can sometimes pull in opposite directions giving rise to internal moral conflict.

Several points are worth noting about Tomasello's natural history of morality. First, Tomasello does think that biological evolution is at work; the three mechanisms involved work in second-personal morality and their scaled up analogues that operate in group-minded cultural morality are either biological adaptations or are underpinned by biological adaptations. However, they are not adaptations specifically for morality but for cooperation more generally. In this respect, his view of morality echoes his view of language developed in *Constructing a Language*⁴⁹ according to which the creation

⁴⁸ C. Korsgaard, The Sources of Normativity, Cambridge, MA, Harvard University Press, 1996.

⁴⁹ M. Tomasello, Constructing a Language, 2002.

and learning of language depends upon the more general mechanisms of mind-reading and pattern recognition.

Second, Tomasello portrays the emergence of morality as a heavy duty intellectual process where individuals reflect upon their situation as they attempt to cooperate and so come to view that situation in certain moral terms, for example, as one where their fellows merit being treated in certain ways, or as one where we are all obliged to behave in certain specific ways.

Third, for Tomasello the specific content of an individual's moral outlook will reflect their experiences. This is particularly true of group-minded cultural morality as it requires the individual to come to know the cultural conventions of her community, something that in turn depends upon cultural transmission and learning. Accordingly, members of different cultures can have different moral outlooks notwithstanding the fact that their respective outlooks will always relate to sympathy and fairness. For example, as the experiments involved in playing the Ultimatum Game⁵⁰ indicate, members of different cultures hold different views as to what is fair and what is not.⁵¹ For example, some think that fairness requires an equal distribution of resources whereas others think that it is fair to distribute resources on the basis of merit.

Fourth, Tomasello thinks that group selection is at work in the moral domain. That is, the moral outlook of one cultural group might outperform that of another in virtue of facilitating more effective cooperation between its members with the upshot that the less cooperative group might die out or get assimilated into the more cooperative one. Tomasello shares the widespread scepticism as to the significance of group selection as a biological force. However, he follows Richerson and Boyd⁵² in postulating a

⁵⁰ In the Ultimatum game one player, the proposer, is given a sum of money. Their task is to offer the other player, the responder, of portion of this money. If the offer is accepted then the money is split accordingly. If the responder rejects the offer (for example, on the grounds that it is too small to count as a fair division of the money) then neither players gets to keep any of the money.

⁵¹ J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, and R. McElreath. 'In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies', American Economic Review, 91, 2001, 73-78.

⁵² P. Richerson and R. Boyd, Not By Genes Alone: How Culture Transformed Human Evolution, Chicago, Chicago University Press, 2005.

process of cultural group selection. The implications of this is that the specifics of an individual's moral outlook will typically reflect the identity of her cultural group and its historical interactions with other cultures rather than their biological nature.

Morality and Cooperation

Should we be convinced by Tomasello's perspective on the function of morality? I am sceptical for several reasons. A first objection relates to the power of moral judgement to make a relevant impact on our behaviour. In the discussion of the nature of functions above I argued that for a trait to have a particular function it must have produced relevant effects in its history, effects that benefitted organisms that had the trait. Applied to the present case, this implies that for morality to have a cooperative function, past individuals who engaged in moral judgement must have been more effective co-operators than their fellows who did not morally judge or, alternatively, more effective co-operators than they otherwise would have been had they not morally judged. Put bluntly, moral judgment must have made a difference to how our ancestors behaved and that difference must have had a positive impact on their cooperative endeavours. But how plausible is it that moral judgement had such a positive impact?

I share with Tomasello an assumption that humans are generally rational and capable of working out how best to behave in order to satisfy their interests, goals and desires. In other words, we are generally effective at strategic reasoning and this has a positive impact on how we behave. Very often morality and self interest point in the same direction; that is to say, that what an individual would do if they were acting rationally out of self-interest would be just the same as what they would do if they were acting on the basis of moral considerations. For example, morality demands that I prepare thoroughly for all the classes that I teach and am generally helpful towards my students. But so does self-interest as a poorly prepared and unhelpful teacher will receive negative evaluations from their students which will in turn lead to disapprobation from their managers and colleagues that can harm their long-term job prospects.

We saw that in order to undercut the idea that humans are prone to cheating and defection Tomasello emphasises the fact that co-operators have a stake in one another. Indeed, he sometimes characterises his position as the stakeholder model of

morality. Thus, for example, it is in the interests of a hunter-gatherer to give food to their foraging partner as they have a stake in that partner's well-being. This is important for Tomasello for it explains how the mechanisms underpinning morality can evolve be it by biological or cultural means. But if we have a stake in one another's well-being then the demands of self-interest and those of morality will typically agree just as I have claimed. For example, suppose a hunter-gatherer retrieves a rabbit from a trap they have set with a foraging partner. What the morality of fairness and justice demands is that they share the prey with their partner. It would be unfair to fail to inform the partner and consume the rabbit alone. But it would also be ill-advised from a strategic perspective. For, if one's partner goes hungry they may be too weak to be of much use on the next foraging trip. Or if the act of unfairness is discovered the partner might break off the partnership undermining the selfish individual's future food gathering prospects.

My point here is not that there is no difference between acting out of self-interest and acting on the basis of moral considerations. For, as Tomasello argues, motives matter when it comes to morality. For example, to be a moral agent or to act out of one's moral convictions one has to do what one does out of a sense of obligation and not for purely strategic reasons. It is for this reason that he denies that chimps are moral agents; even when chimps seemingly behave in line with the demands of morality (for example, when they putatively treat a fellow chimp justly) they do so for strategic reasons.

The upshot of all this is that the capacity to make moral judgements and act on them couldn't have made any difference to the behaviour of our ancestors who were already engaged in strategic reasoning. Developing a capacity to make moral judgements and act on them would have been akin to putting on braces when your trousers were firmly secured by a belt. But, given the nature of functions, for morality to have a cooperative function the capacity to make and act on moral judgements must have made a relevant difference to the behaviour of our ancestors; it must have made them behave in ways that enhanced their cooperative endeavours, behaviour that they otherwise would not have engaged in (or behaviour that their strategic but non-moralistic fellows didn't engage in). Consequently, on the stakeholder model, morality could not have secured a cooperative function.

A potential response to this objection has been developed by Joyce⁵³ in the context of defending his view that morality is a biological adaptation.⁵⁴ What Joyce claims is that we are subject to weakness of will; we often fail to do things that we want to do and know are in our best interests and we have a particular problem delaying gratification. For example, I might have an extra glass of wine when I know I will regret it the following day or give that 5-mile run a miss on a cold and rainy day despite supposedly being on a fitness drive. Thus, in cooperative situations we run the risk of failing to do what we know is in our best interests so that our cooperative endeavours break down. The advantage of morality, according to Joyce, is that moral judgements have a 'practical clout' that regular intentions or desires do not have; in other words, it is much harder to do something that one judges to be wrong or fail to do something that one judges to be right than it is in the cases where one's motivations to act are provided by mental states that do not have a moral character. This raises the question of how moral judgements have such motivational force; why is it harder to ignore a moral judgment than a regular non-moral strategic judgement? Joyce provides the familiar answer that moral judgements have an emotional component (without being purely emotional states) and it is this that provides them with their motivational force.

This raises the question of whether moral judgments are sufficiently powerful with respect to action; do they systematically outperform self-interested reasoning in terms of causing us to act. If not, then Joyce's argument collapses and we are left with the conclusion that morality doesn't have a cooperative function.

I am sceptical of the power of moral judgment to drive behaviour for the following reasons. First, it is almost a truism that humans often behave in ways that conflict with their moral judgements and convictions and this is reflected in the fact that so many of our moral judgements are negative in tone and involve judging that someone did something that they shouldn't have done and knew full well that they shouldn't have done. Second, some cognitive scientists⁵⁵ question whether reasoning plays a substantial role in arriving at moral judgements. Whether or not they are right, my

⁵³ R. Joyce, The Evolution of Morality, 2006.

⁵⁴ R. Frank, Passions with Reasons: The Strategic Role of the Emotions, New York, Norton, 1988, develops a similar line of argument.

⁵⁵ For example, J. Haidt, 'The Emotional Dog and its Rational Tails: A Social Intuitionist Approach to Moral Judgment', Psychological Review, 108, 2001, 814-834.

experiences suggest that people often engage in reasoning to get themselves out of acting on the basis of a rather inconvenient moral judgement or injunction. For example, I have heard people who are morally in favour of giving to charity justifying their not giving to charity on the grounds that the charity in question can't be relied upon to spend donations wisely. And I've heard people who condemn major corporations for exploiting loopholes in the tax system justify their own minor tax indiscretions on the grounds that the sums involved are trifling. In short, we are rather effective at employing reasoning to escape the behavioural demands of moral judgement when acting on the basis of that judgement clashes with self-interest.

Third, in case the foregoing sounds a little too anecdotal, there is important experimental data that suggests that in those cases where morality and self-interest point in the same direction, it is self-interest that is driving behaviour and when they point in different directions self-interest often wins out. Bateson et al (2006) conducted an experiment in a University common room which operated an honesty box system for paying for tea and coffee. A clearly visible notice was displayed at head height near to the tea and coffee-making equipment and supplies. As well as indicating the prices the notice had an image printed upon it which alternated from week to week between flowers and a pair of eyes. What Bateson et al. discovered were that contributions to the honesty box were considerably higher when the displayed image was of a pair of eyes than when it was of a flower. They interpreted this as indicating that the image of the eyes reminded individuals that they were potentially being watched and that their reputation would be damaged by a failure to contribute to the honesty box. Thus, a self-interested desire to protect one's reputation rather than a genuine concern for the public good was the key behavioural driver.

A second objection to Tomasello's perspective on the function of morality involves questioning the need for morality to drive our cooperative behaviour. In the previous objection I warned against the danger of overestimating the impact of moral judgement on our behaviour. What I am warning against now is the danger of underestimating the power of non-moral factors to influence our behaviour. Humans are cultural beings and an important aspect of any culture is a rich network of conventions and non-moral norms that govern behaviour. As Tomasello is keen to point out, human individuals are very effective at learning such conventions and norms (which are stored in what he calls cultural common ground) and once learned such conventions and norms have a powerful hold on our behaviour. For example, I have internalised the British norms

surrounding queuing and slavishly follow them without reflection whenever I encounter situations where they apply. Sometimes I holiday abroad in places where I know full well that the British queuing conventions do not apply. Nevertheless, I cannot help but queue like a Brit and invariably end up waiting a long time to be served. Similarly, I have internalised the convention of standing on the right of escalators when travelling on the London underground thereby leaving a clear passageway on the left for travellers in a hurry. I know full well that this convention doesn't apply everywhere. Nevertheless, when I find myself on an escalator in another context I cannot help but stand on the right and find it almost emotionally painful to stand on the left. What these examples show is the power of conventions to direct our behaviour.

If we are so prone to pick up conventions and they have such a hold on us then why would we need morality to support our cooperative endeavours? Surely, a body of conventions would be far more effective. For example, to get people to help their less fortunate fellows all you need is a convention of helping the less fortunate rather than a capacity to judge that it is morally good to help those less fortunate than oneself. It might be objected that if we did have and learn such conventions and routinely acted on the basis of them, then we would have learned a moral norm and be acting on the basis of moral considerations. However, such an objection would miss the important point that moral agency requires not merely acting in accord with moral norms but doing what one does for moral reasons (for example, out of a sense of obligation). Thus, if the mechanisms that underpin 'doing the right thing' are the same as those that underpin the following of a non-moral convention then one would not have learned a moral norm or be acting on the basis of moral considerations. In sum, there is reason to believe that we have a psychology that is such that our cooperative endeavours would be better suited by learning and following non-moral conventions than by exercising our capacity for moral judgment.

Conclusion

Michael Tomasello has developed a particularly detailed and sophisticated version of the popular view that morality has a cooperative function that promises to dominate the literature for years to come. Drawing upon an account of the nature of moral judgement and of teleological function (which explains both how a biological trait or a cultural product could have a function) developed in the first half of this chapter I raised a number of objections to Tomasello's perspective. In particular, I argued that moral

judgment does not plausibly have the required impact upon behaviour to accrue a cooperative function; typically, the demands of morality agree with those of self-interest and where they clash it is usually the demands self-interest grounded in our strategic reasoning capacities that win out. Moreover, given the human capacity to learn conventions and the power that learned conventions have, there are far more effective means of ensuring that we behave in a way that supports our cooperative endeavours.

By focussing on a particular version of the idea that morality has a cooperative function I haven't settled the question as to what the function of morality is or, indeed, whether it has any function at all. Thus, someone convinced by my line of argument is faced with a series of options. The first option is that of developing an alternative account of function that is not grounded in evolution (be it biological or cultural) and arguing that that account applies to morality thereby grounding a particular function. The second option involves arguing for an alternative non-cooperative function for morality that is grounded in evolution. The third option is that of rejecting the idea that morality has a function. I will leave the question of which of these options is the most fruitful for another occasion.

References

- A. Ariew and R. Lewontin, 'Confusions of Fitness' in British Journal for the Philosophy of Science, 55, 2004, 347-363.
- M. Bateson, D. Nettle, and G. Roberts, (2006) 'Cues of being watched enhance cooperation in a real-world setting' in Biology Letters, 2, 2006, 412-414.
- N. Block, 'Advertisement for a Semantics for Psychology' in Midwest Studies in Philosophy, 10, 1986, 615-678.
- P. Bloom, Just Babies: The Origins of Good and Evil, London, Bodley Head, 2013.
- T. Caro, Zebra Stripes, Chicago, University of Chicago Press, 2016.
- R. Cummins, 'Functional Analysis', Journal of Philosophy, 72, 1975, 741-765.
- C. Darwin, On the Origin of Species, London, John Murray, 1859.
- D. Davidson, 'Knowing One's Own Mind' in Proceedings and Addresses of the American Philosophical Association, 60, 1987, 441-458.
- 1987.

R. Dawkins, The Blind Watchmaker, New York, Norton, 1986.

D. Dennett, Brainstorms, Cambridge, MA, MIT Press, 1978.

D. Dennett, Darwin's Dangerous Idea, London: Penguin, 1995.

J. Fodor, 'The Appeal to Tacit Knowledge in Psychological Explanation,' in Journal of Philosophy, 65, 1968, 627-640.

P. Foot, 'Morality as a System of Hypothetical Imperatives' in Philosophical Review, 81, 3, 1972, 305-316.

R. Frank, Passions with Reasons: The Strategic Role of the Emotions, New York, Norton, 1988.

P. Godfrey-Smith, 'Functions: Consensus Without Unity', Pacific Philosophical Quarterly, 74, 1993, 196-208.

P. Godfrey-Smith, 'A Modern History Theory of Functions' in Nous, 28, 1994, 344-362.

P. Godfrey-Smith, Philosophy of Biology, Princeton, NJ, Princeton University Press, 2014.

S. Gould and R. Lewontin, 'The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme' in Proceedings of the Royal Society of London, Series B, 205, 1979, 581-598.

S. Gould, The Structure of Evolutionary Theory, Cambridge, MA, Harvard University Press, 2002.

J. Greene, Moral Tribes: Emotion, Reason and the Gap Between Us and Them, London Atlantic, 2013.

J. Haidt, 'The Emotional Dog and its Rational Tails: A Social Intuitionist Approach to Moral Judgment', Psychological Review, 108, 2001, 814-834.

J. Haidt, The Righteous Mind: Why Good People are Divided by Politics and Religion, London, Penguin, 2013.

J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, and R. McElreath. 'In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies', American Economic Review, 91, 2001, 73-78.

R. Joyce, The Evolution of Morality, Cambridge, MA, MIT Press, 2006.

I. Kant, Groundwork for the Metaphysics of Morals, A. Zweig (trans), Oxford, Oxford University Press, 1785/1998.

C. Korsgaard, The Sources of Normativity, Cambridge, MA, Harvard University Press, 1996.

K. Laland, Darwin's Unfinished Symphony: How Culture Made the Human Mind, Princeton, NJ: Princeton University Press, 2017.

D. Lewis, Convention: A Philosophical Study, Cambridge, MA, Harvard University Press, 1969.

E. Mayr, 'Cause and effect in biology', Science, 134, 1961, 1501–1506.

J. Mikhail, Elements of Moral Cognition, Cambridge, Cambridge University Press, 2011.

R. Millikan, 'In Defense of Proper Functions' in Philosophy of Science 56, 1989, 288-302.

R. Millikan, 'Language Conventions Made Simple' in Journal of Philosophy, 95, 1998, 161-180.

K. Neander, 'The Teleological Notion of "Function"', in Australasian Journal of Philosophy, 69, 1991, 454-468.

K. Neander, 'Functional Analysis and Species Design' in Synthese, 194, 2017, 1147-1168.

K. Neander, 'Does Biology Need Teleology' in R. Joyce, (ed) The Routledge Handbook of Evolution and Philosophy, Abingdon, Routledge, 2018, 64-76.

S. Nichols, Sentimental Rules: The Natural Foundations of Moral Judgment, Oxford, Oxford University Press, 2004.

G. Parker and J. Maynard Smith, 'Optimality in Evolutionary Biology', in Nature, 348, 1990, 27-33.

W. Poundstone, Prisoners Dilemma, Game Theory and the Puzzle of the Bomb, Oxford, Oxford University Press, 1992.

J. Prinz, Furnishing the Mind, Cambridge, MA, MIT Press, 2002.

J. Prinz, The Emotional Construction of Morals, Cambridge, MA, MIT Press, 2007.

P. Richerson and R. Boyd, Not By Genes Alone: How Culture Transformed Human Evolution, Chicago, Chicago University Press, 2005.

T. Scanlon, What We Owe to Each Other, Cambridge, MA, Harvard University Press, 1998.

R. Shafer-Landau, Moral Realism: A Defence, Oxford, Oxford University Press, 2003.

B. Skyrms, The Stag Hunt and the Evolution of Social Structure, Cambridge, Cambridge University Press, 2004.

M. Smith, The Moral Problem, Oxford, Basil Blackwell, 1994.

K. Sterelny, 'Moral Nativism: A Skeptical Response', in Mind and Language, 25 2010, 279-297.

K. Sterelny, The Evolved Apprentice: How Evolution Made Humans Unique, Cambridge, MA, MIT Press, 2012.

M. Tomasello, Constructing a Language, Cambridge, MA, Harvard University Press, 2002.

M. Tomasello, Origins of Human Communication, Cambridge, MA, MIT Press, 2008.

M. Tomasello, A Natural History of Human Thinking, Cambridge, MA, Harvard University Press, 2015.

M. Tomasello, A Natural History of Human Morality, Cambridge, MA, Harvard University Press, 2016.

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

B. Williams, Ethics and the Limits of Philosophy, Cambridge, MA, Harvard University Press, 1985.

L. Wright, 'Functions' in Philosophical Review, 82, 1973, 139-168.