Morality 2.0: How manipulating our minds could save the world by Dan Jones

Our moral instincts are failing to cope with the dilemmas of the modern world. Is it time to use tricks, pills or brain zaps to hack this outdated mind set?

IN JUNE, a new voice backed up what many scientists have been saying for a while – that climate change is caused by human activity and we have a moral responsibility to tackle it. In an historic edict, Pope Francis warned that failing to act would have "grave consequences", the brunt of which would fall on the world's poorest people. His words came as a stark reminder that global climate change is among the most pressing moral dilemmas of the 21st century.

It joins a long list. He could have added spiralling inequality, persistent poverty, death from preventable diseases and nuclear proliferation to the ethical challenges that define our times. Some are newer than others, but all could plausibly be fixed. The fact we're struggling with all of them raises a troubling question: does our moral compass equip us to deal with the threats we face today?

For many of the pope's billion-strong flock and other believers, moral judgement is an operation of the mind beyond scientific explanation. In recent years, however, psychologists and neuroscientists have gone a long way towards understanding the machinery underlying our moral thinking and behaviour. In the process they are getting to grips with what really drives how we decide what is right or wrong. Their insights are not only revealing the limitations of our moral minds, but also suggesting how we might manipulate them – by employing psychological tricks, or even pills and brain zaps.

Human moral psychology evolved over tens of thousands of years as we became an ever more cooperative, social species. Early humans living in small bands were forced to hunt and forage collectively or starve. That requires cooperation, which fuelled the evolution of cognitive facilities underlying collective action – the ability to share goals, responsibilities and rewards.

"Morality is a device for solving the social challenges of everyday life, where the basic problem is to get otherwise selfish individuals to work together as a group and enjoy the benefits of cooperation," says Joshua Greene, a neuroscientist at Harvard University.

So how do we make these everyday decisions? One key insight came from New York University psychologist Jonathan Haidt, who revealed that moral judgements are frequently driven not by rational, reflective thought but by intuitions and gut feelings.

In one study he quizzed people about the morality of various acts, from cleaning a toilet with the US flag to having sex with a chicken bought from the supermarket. Participants often said these acts were immoral and clung to these judgements even when they couldn't provide good reasons for doing so. Frequently, they'd throw their hands up and say, "It's just wrong!" – a phenomenon Haidt calls "moral dumbfounding".

This moral intuition is often fuelled by emotional reactions. Most people are repulsed by the thought of a human engaging in coitus with a deceased chicken, and that alone is enough to condemn the act. When reasoning comes into play, it is frequently to rationalise these intuitive decisions after the event.

Yet we are not slaves to emotion. "We have gut reactions that guide our judgements and behaviour," says Greene, "but we can also stop and think, and reason explicitly about situations to try to make better decisions."

"We tend to make quick-fire moral judgements based on emotion and gut reaction"

Greene has shown how this plays out in the brain by getting people to mull over dilemmas as they lie in an fMRI machine. A famous example considers whether it is morally permissible to push one person off a bridge on to railway tracks to stop a runaway train from hitting five people stuck just ahead, killing the person in the process. Most people feel a strong gut reaction to say no, which shows up as increased activity in brain regions that process social emotions. The upshot is that aversion to up-close-and-personal violence trumps the greater good.

But if you instead ask whether it's OK to flick a switch so that the runaway train is diverted from the track with the five people on it and towards another where it will hit and kill just one person, most people say yes. The moral maths is the same, but hitting a switch just doesn't feel as bad. Greene's studies indicate that responses to these kinds of dilemmas are dominated by colder, calculating processes grounded in the "rational" dorsolateral prefrontal cortex.

Greene likens our moral psychology to a digital camera. Intuitive moral sentiments are akin to the automatic settings, he says, while rational deliberation is analogous to manual mode, where you adjust everything by hand. "The automatic settings are good for the situations for which they've been programmed, making them efficient but not very flexible," says Greene. "Manual mode is flexible but not so efficient, as it takes time to punch in the settings."

In the same way that many of us rely on auto mode on our cameras because it is easier, we tend to make quick-fire moral judgements based on gut reactions. So here's the question: is auto-mode moral decision-making, which evolved to navigate small-scale social worlds, suited to handling issues that impact millions of distant strangers and future generations?

Greene thinks not. "These are very good at solving the problems of everyday life, but not global moral problems like environmental destruction or poverty in faraway places".

Take empathic concern, one of the key features of auto-mode morality. Roughly speaking, this is feeling the pain of others. It functions like a spotlight, throwing into stark relief the plight of whoever falls under its beam, and moving us to action. So you might think empathic concern is an unalloyed force for good.

"You might think empathy is an unalloyed force for good. You would be wrong"

You would be wrong, says Paul Bloom, a psychologist at Yale University. "Empathy, being a spotlight, is very narrow," he says. It illuminates the suffering of a single person rather than the fate of millions, and it is more concerned with the here and now than the future. "It's because of empathy that we care more about, say, the plight of a little girl trapped in a well than we do about potentially billions of people suffering or dying from climate change," says Bloom. The visceral reactions towards recent photographs of a dead Syrian boy washed up on a Turkish beach provide another case in point.

Empathy's shortcomings are compounded by the fact that we end up pointing its beam on causes that happen to come into our field of view – typically the most newsworthy moral issues, rather than those where we can do the most good. The response to the 2004 Boxing Day tsunami, in which some charities received more donations than they could spend, is one example.

All this sounds a bit disheartening, but there is hope. "We can use manual mode to train automatic mode," says Fiery Cushman, who runs a moral psychology lab at Harvard University. "In the past

10 years, we've learned that there's an enormous role for learning in shaping our moral intuitions."

In a study out last year, for example, Cushman asked people how they would feel about performing a mercy killing on a terminally ill man by various methods, including giving him a poison pill, suffocating him and shooting him in the face. You might expect opposition to each method would be predicted solely by the amount of suffering it causes. But Cushman found that it was better predicted by participants' aversion to performing the action (Emotion, vol 14, p 573).

Would linking flying with harming our kids make us more aware of climate change? *(Image: YAY Media AS/Alamy Stock Photo)*

So it looks like we base our instinctive moral judgements not only on our emotional reaction to suffering, but also on how the physical acts that cause it make us feel. And here's where we might be able to change ourselves. We learn to assign moral value to actions through the brain's dopamine system and the basal ganglia, and Cushman suggests we might manipulate this process to shape our instinctive moral reactions.

One approach is to deliberately seek out particular experiences, he says. Imagine an aspiring vegetarian who is concerned about animal welfare. If bacon sandwiches are proving too much of a temptation, they might watch videos documenting the mistreatment of animals. "This could change their automatic attitudes, so when they see meat in front of them they find it disgusting rather than appetising," says Cushman. The same tactics might help people forge an aversion to actions that increase their carbon footprint, say, or add to the plight of the world's poorest people.

New age of reason

But the choices made by a select few are hardly going to be sufficient. When it comes to creating the large-scale moral change required to solve the planet's greatest problems, Kwame Anthony Appiah, a philosopher at New York University, argues that "the question isn't 'What should I do?', but 'What should we do?'". And although manual-mode thinking can help us set our sights on the causes "we" should pursue, reshaping moral thinking en masse takes more than deliberation and reasoning.

Appiah has studied the history of moral revolutions such as British abolitionism – the 19th-century movement to end the transatlantic slave trade. "The rational, moral and legal arguments for ending the slave trade were well known long before abolition," says Appiah. What tipped the argument over into becoming a movement was that broad swathes of society came to feel collectively ashamed of being engaged in the trade. That shift was driven by activist groups raising awareness of the dreadful human cost and making the anti-slavery cause part of British national identity.

In Birmingham, for example, local politicians and thousands of citizens signed anti-slavery petitions, to be delivered to the British government. "They wanted to be part of a city that had done something about this trade," says Appiah. "This civic pride was a big part of the abolitionist movement."

Appiah suggests campaigns that speak to our sense of collective identity – as members of a city, nation, religion or social movement – are likely to be most effective. So why not scale up and play on our shared identity as humans to tackle problems that affect us all? "The problem with the notion

of humanity is that there's no outside group to compare against," says Appiah. "It would be a more useful notion if there were some aliens around."

Still, shame at immoral actions and pride in tackling moral problems could yet prove effective. Jennifer Jacquet, an environmental scientist from New York University, argues that we may need to recruit the power of shame to modify moral behaviour on social issues like climate change. "Aimed cautiously and well", shaming can be deployed for the greater good, she thinks.

Take BankTrack, a global network of NGOs that exposes banks involved with projects that threaten the environment and human rights. BankTrack has looked at banks lending to the coal industry, a major source of global carbon dioxide emissions, and compiled a list of the top "climate killers".

Its manifesto is simple: "By naming and shaming these banks, we hope to set the stage for a race to the top, where banks compete with each other to clean up their portfolios and stop financing investments which are pushing our climate over the brink."

However, harnessing the power of rational reflection, collective identity and shame may not be the only options for would-be moral revolutionaries. In their book Unfit for the Future, philosophers Ingmar Persson of the University of Gothenburg in Sweden and Julian Savulescu of the University of Oxford argue that our moral brains are so compromised that the only way we can avoid catastrophe is to enhance them through biomedical means.

Barclays ended financing mountaintop removal mining in 2015 after protests   *(Image: Daniel Leal-Olivas/i-Images)*

In the past few years, researchers have shown it might actually be possible to alter moral thinking with drugs and brain stimulation. Molly Crockett of the University of Oxford has found that citalopram, a selective serotonin reuptake inhibitor used to treat depression, makes people more sensitive to the possibility of inflicting harm on others. Earlier this year, for instance, Crockett and colleagues found that participants who had taken citalopram were willing to pay twice as much money as controls to prevent a stranger from receiving an electric shock (Current Biology, vol 25, p 1852).

"It might actually be possible to alter moral thinking with drugs and brain stimulation"

Biomedical enhancement may even work on complex social attitudes. Roberta Sellaro at the University of Leiden in the Netherlands has shown that delivering low-current electrical signals through the scalp to stimulate the medial prefrontal cortex, a brain area implicated in regulating social emotions, can reduce stereotyped attitudes towards members of different social groups.

Don't expect to see morality pills on pharmacy shelves any time soon, though. These studies are far from conclusive and the effects demonstrated are subtle. Even so, the fact that such behavioural modification is possible raises the prospect that someone might use it.

Of course that raises moral questions in itself – who to treat, how, and at what age? But Persson and Savulescu argue that if the techniques can be shown to change our moral behaviour for the better (who or what defines "better" is another question), then there are no good ethical reasons not to use them. Take the issue of consent, which children could not provide. "The same is true of all upbringing and education, including moral instruction," says Persson.

Moral robots

But wouldn't biomedical moral enhancement undermine responsibility by turning us into moral robots? Persson and Savulescu argue that biomedical treatment poses no more threat to free will and moral responsibility than educational practices that push us towards the same behaviour.

Yet even if moral hacking with drugs and brain zaps could be deemed ethically sound, putting it into practice is another matter. "It's not as if there's a moral circuit in the brain that you simply want to ramp up," says Greene. "Moral decision-making draws on a number of major brain circuits, so you're not going to be able to enhance people's morality in a pinpointed way."

And assuming you could get people to make the "right" choices, how would you deliver artificial moral enhancement across entire populations? Would we add drugs to the water supply? Would we fortify kids' cereal with moral enhancers?

Even if we could overcome these obstacles, Bloom insists that artificial enhancement is not the way to go. "I think people approach it the wrong way, and work under the illusion that we'd be better people if we messed with our emotions in one way or another."

Bloom argues that when it comes to tackling moral problems such as climate change, which we have no reliable instinctive way of dealing with, the best way forward is to try to spend more time in manual mode.

"Moral issues are complicated and hard, and they involve serious trade-offs and deliberation. It would be better if people thought more about them."