

Can Psychologists Tell Us Anything About Morality?

John M. Doris, Edouard Machery and Stephen Stich on why ethicists ought to listen to them

Lately, it's (again) become fashionable to raise questions about the relevance of human psychology to human morality. On its face, the notion seems singularly unpromising—isn't what we're like relevant to what we ought do? But in fact, this idea is an old one, which has been discussed by moral philosophers for centuries.

We therefore enter discussion with an unwelcome sense of déjà vu. But the recent exchange between Tamsin Shaw and Steven Pinker and Jonathan Haidt in the *New York Review of Books* compels us to consider the old idea in its most recent incarnation, which holds not that human psychology doesn't matter for morality, but that what *professional psychologists* tell us about morality doesn't matter for morality.

Often, this view is supported more by insinuation than argument. Allegedly, academic psychology is in crisis: *Psychologists are complicit in human rights abuses! Psychologists fabricate data! Psychology experiments don't replicate!*

Yes, yes, and yes. Some psychologists accept morally dubious employment. Some psychologists cheat. Some psychology experiments don't replicate. *Some*. But the inference from *some* to *all* is at best invalid, and at worst, invective. There's good psy-

chology and bad psychology, just like there's good and bad everything else, and tarring the entire discipline with the broadest of brushes won't help us sort that out. It is no more illuminating to disregard the work of psychologists *en masse* on the grounds that a tiny minority of the American Psychological Association, a very large and diverse professional association, were involved with the Bush administration's program of torture than it would to disregard the writings of all Nietzsche scholars because some Nazis were Nietzsche enthusiasts! To be sure, there are serious questions about which intellectual disciplines, and which intellectuals, are accorded cultural capital, and why. But we are unlikely to find serious answers by means of innuendo and polemic.

Could there be more substantive reasons to exclude scientific psychology from the study of ethics? The most serious—if ultimately unsuccessful—objection proceeds in the language of “normativity”. For philosophers, normative statements are prescriptive, or “oughty”: in contrast to descriptive statements, which aspire only to say how the world is, normative statements say what ought be done about it. And, some have argued, never the twain shall meet. While philosophers haven't enjoyed enviable suc-

cess in adducing lawlike generalisations, one such achievement is Hume's Law (we told you the issues are old ones), which prohibits deriving normative statements from descriptive statements. As the slogan goes, "*is* doesn't imply *ought*."

Many philosophers, ourselves included, suppose that Hume is on to something. There probably exists some sort of "inferential barrier" between the *is* and the *ought*, such that there are no *strict logical entailments* from the descriptive to the normative. At the same time, moral philosophy – the philosophical subfield most commonly engaged in the business of normativity – is shot through and through with descriptive claims; on the pages churned out by the journals and university presses, recognisably moral argument is everywhere buttressed by recognisably empirical observations about humans and their world.

While it may seem surprising that philosophers so often paddle along in apparently blissful ignorance of a widely acknowledged law of thought, we rather doubt it is surprising to the casual observer uncorrupted by Hume. Before entering the seminar room (or even after) does anyone think that what ought be done about Big Tobacco is altogether *unrelated* to facts about the health and mortality implications of smoking? Moral philosophy is a messy business, and is seldom, if ever, deductively airtight: *is* does not entail *ought*, but answers to questions about how it is most reasonable for humans to think and act are everywhere structured by facts about the circumstances and psychologies of those thinking and acting.

It's possible that some moral philosophers will wish to take issue with this (perhaps because they think moral philosophy

should trade only in "ideal" theory), but it is important to notice that they are taking issue with how moral philosophy is usually done – and has been done, from its beginnings in antiquity (we doubt, for example, that we've so far said anything that Aristotle would disapprove of). More importantly, that's how moral philosophy *ought* be done: the best theories in ethics and moral psychology are likely to contain various admixtures of fact and value. Indeed, this supposition enjoys the support of its own venerable dictum—commonly attributed to Kant—"ought implies can": it is either nonsensical or unfair to say your infant ought to refrain from crying when she's hungry, if she can't reasonably be expected to do so.

The inescapable presence of facts in moral theorising does not get us the conclusion that any facts of interest to moral philosophy are the facts uncovered by scientific psychology. That can only be established by figuring out what the scientific facts are, and seeing how they might be of interest for theorising in ethics. Doing so requires detailed discussion of both the philosophy and psychology, an interdisciplinary endeavour which by now boasts a large literature. While exponents of this methodology insist that scientific psychology can *inform* ethical thought, they do not (as perhaps should go without saying) contend that scientific psychology can *replace* ethical thought. Still less, do they suggest that psychologists be enshrined as "moral experts" to whom the rest of us owe deference (an aspiration, by the way, we've never heard psychologists of our acquaintance expressing). Their conviction is simply that the scientific investigation of human psychology can enrich discussion of human morality.

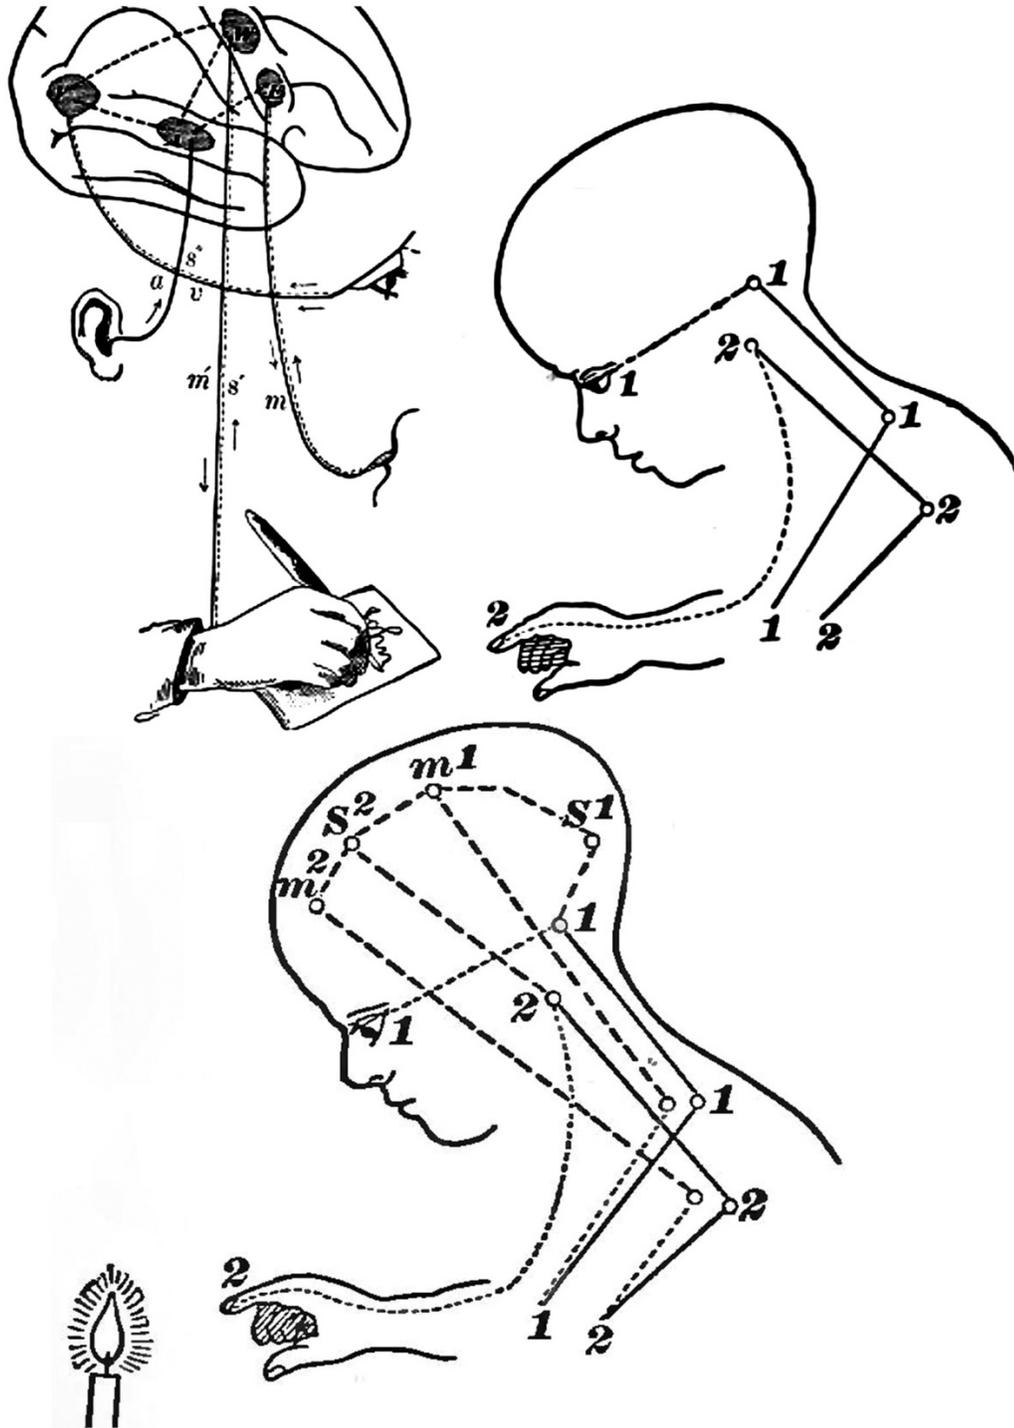
Consider psychological egoism, the view that the ultimate goal of all human behaviour is our own self-interest. Psychological altruism, by contrast, maintains that the ultimate motivation of *some* human behaviour is the well-being of others. In the seventeenth century, Thomas Hobbes offered what many took to be a powerful case for egoism, and in the following two centuries the great Utilitarian philosophers and social reformers, Jeremy Bentham and John Stuart Mill, were convinced that hedonism, a version of egoism maintaining that humans are capable of only two ultimate goals, experiencing pleasure and avoiding pain, was the correct account of human motivation. But, of course, the Utilitarians also thought that what we *ought* to do is whatever will produce the greatest happiness for the greatest number of people. So their normative theory often required that people behave in ways that, according to the Utilitarians' psychological theory, might be impossible. What to do? Part of Mill's answer was that we should engage in manipulative and draconian social engineering designed to instil fear of displeasure from the "Ruler of the Universe" (though Mill himself was an atheist).

The normative landscape for Utilitarians (as well as for Kantians and others) looks very different if egoism is false. But is it? In the four centuries since Hobbes, philosophers contested the question using anecdotes, intuitions, and *a priori* arguments that convinced almost no one. Then, about 40 years ago, experimental social psychologists turned their attention to the debate between psychological egoists and psychological altruists. It's been a long, hard slog. Designing experiments that provide per-

suasive evidence for psychological altruism or egoism is a challenging project. By now, however, there is an impressive body of findings – beautifully discussed in Daniel Batson's *Altruism in Humans* – suggesting that Hobbes and his fellow psychological egoists are wrong. Humans *are* capable of purely altruistic motivation, and there are social interventions that can encourage or discourage altruistic behaviour. Does anyone really think that this impressive body of empirical work does not have an important role to play in normative theorising?

A more recent example can be found in John Mikhail's important book, *Elements of Moral Cognition*. Mikhail, who is a philosopher, a cognitive scientist and a law professor who teaches human rights law, has spent over a decade studying the sorts of moral dilemmas featured in Joshua Greene's path-breaking early work using neural imaging to study moral reasoning. From this, Mikhail makes an impressive, albeit controversial, case that all normal humans share an important set of innate moral principles. Then, building on John Rawls' influential account of when moral principles are justified, Mikhail argues that if there are pan-cultural innate moral principles, then on a Rawlsian account of justification, those principles are justified.

Mikhail concludes that his account provides the much-needed intellectual underpinning for the doctrine of universal human rights that has played a central role in international law since the end of the Second World War. Is he right? Not surprisingly, opinions differ. But it would be outrageous to argue that Mikhail's work should not be considered in normative theorising about human rights simply because he relies on



Illustrations from the first edition of *The Principles of Psychology* by William James

psychological experiments.

Another area where scientific work may inform public policy concerns the psychology of disgust. In his influential article, “The Wisdom of Repugnance” (1997), Leon Kass, Chairman of the President’s Council on Bioethics until 2005, claimed that “in crucial cases ... repugnance is the emotional expression of deep wisdom, beyond reason’s power fully to articulate it.” To simplify, disgust, which he calls “repugnance”, is for Kass a sort of moral sense that can identify morally impermissible actions even when arguments founder: there are wrongs that people can *feel* even though they can’t rationally articulate the basis of the wrongness. While we have no interest in attributing such views to Kass, it bears noting that others have deployed similar arguments to justify their condemnation of homosexuality and opposition to same-sex marriage.

Recent empirical research raises serious questions about the wisdom of disgust. According to Daniel Kelly’s sophisticated synthesis of 30 years of psychological research on disgust in his book *Yuck!*, the contemporary human disgust reaction results from the fusion of two distinct mechanisms: one dedicated to identifying poisonous foods, and one dedicated to identifying sources of parasite transmission, including microbes and disease carrying agents. Kelly contends that this protective response to cues indicating a risk of poisoning or contamination has been co-opted in many other psychological domains, including the psychology of moral judgement. Yet disgust is still triggered by socially learned poisoning and contamination cues: the influence of these visceral cues makes disgust a distorting influence on moral judgement, causing us to con-

flate what our culture regards as “icky” with what is “morally repugnant”. Here again, we’re unable to take seriously the idea that psychologically informed theorising should not play an important role in normative thinking: whatever conclusions one reaches, is there a credible argument for excluding considerations like those Kelly raises in debating “the wisdom of repugnance”?

In addition to illuminating extant moral questions, research in experimental psychology enables us to identify new ones. Recent work has revealed that all of us are infested with a motley of surprising “implicit,” or unconscious, biases. Many people, including people who are committed to racial equality, may nonetheless exhibit prejudicial thought, such as associating black faces with negative words and white faces with positive words. There is a growing body of evidence suggesting that these implicit biases also affect our behaviour, though we are very often completely unaware that this is happening. Moral philosophers have long been concerned to characterise the circumstances under which people are reasonably held to be morally responsible for their actions, and a common theme in this discussion is that ignorance may serve as an excuse. Are we morally responsible for behaviour that is influenced by implicit biases, if these tendencies are ones of which we are not aware? That’s a question that has sparked heated debate, and it is a question that could not have been responsibly asked without reference to the empirical findings reported by psychologists.

Perhaps the first place where substantial numbers of contemporary moral philosophers did what we’re suggesting they ought to be doing, and took an in-detail look at

scientific psychology, was in discussions of moral character and virtue. A familiar idea in philosophical ethics (and, we dare venture, in life) is that good character is an enduring bulwark against doing wrong, and a reliable source of doing right. The compassionate person, it's tempting to think, won't be cruel, even when distracted or provoked, and will behave with appropriate kindness even when doing so comes at personal cost. Trouble is, there's lots of psychology – most famously, the oft-replicated Milgram experiments – indicating that moral failure is rather *easily* induced: distressingly slight situational pressure may result in ordinarily decent people doing less than decent-- even appalling -- things. The psychology has led “character skeptics” like us to question the privileged role reserved for character in much ethical thought, whether it be academic philosophy's “virtue ethics” tradition, or the many discussions of character in popular political writing (where we're told people “vote character”).

We readily admit our skepticism isn't mandatory. On the contrary. It is extremely controversial: there is now a large and lively philosophical literature debating the issues. Our point here is that if moral philosophers had taken seriously blanket prohibitions against engaging scientific psychology, this literature—a literature, we hasten to add, situated in and around *normative* ethics—*would not have been possible*. Very arguably, philosophical discussion of moral character is thriving like never before. And very arguably, this is in substantial measure due to a widespread willingness to take scientific psychology seriously.

Hopes of descriptive – or normative – purity are doomed; the descriptive and

normative are inextricably intermingled wherever moral psychologists and moral philosophers practice their crafts. As shown in the few examples we have adduced (as well as the many others we might have), this is a very good thing: ignoring the *is* – emphatically including the *is* as revealed by the empirical psychology – is frequently a recipe for disastrous theorising about the *ought*.

Psychology, we're sure, can, and ought, do better. And the same, we're just as sure, is true of philosophy. But doing better won't be effected by rigidifying disciplinary boundaries that are of little more than administrative interest. It will be effected by thinking closely and charitably about the best the disciplines have to offer.

John M. Doris is professor in the philosophy–neuroscience–psychology program and philosophy department, Washington University in St. Louis. His books include Lack of Character: Personality and Moral Behavior, and Talking to Our Selves: Reflection, Ignorance, and Agency.

Edouard Machery is distinguished professor in the department of history and philosophy of science and director of the Centre for Philosophy of Science at the University of Pittsburgh. Among his many books are Doing without Concepts and Philosophy Within Its Proper Bounds (forthcoming) from Oxford University Press.

Stephen Stich is board of governors distinguished professor of philosophy and cognitive science at Rutgers University, and honorary professor of philosophy at the University of Sheffield. His books include From Folk Psychology to Cognitive Science, Deconstructing the Mind, and two volumes of Collected Papers.