



DOING ETHICS WITH OUR BRAINS

John
Hooker

We are told that we live in a post-truth age.

When the facts get in the way, we turn to “alternative facts” that serve our purposes. Rather than listen to another point of view, we focus only on arguments and talking points that support our ideology. Not everyone is like this, of course, but it seems to capture the tenor of the times. Worst of all, it exacerbates the polarization that so many worry about, because we can’t find common ground.

The root problem, in my view, is a gradual abandonment of rationality. We can’t reach consensus because we no longer acknowledge a rational basis for resolving disputes.

Ethics was an early casualty of this retreat from reason. Over a period of decades we have gradually surrendered to the notion that ethics is merely a matter of personal opinion or personal values. I have my view and you have yours, and that’s the end of it. This assumption permeates not only popular culture, but the academic world in which I live. My colleagues will tell you, almost to a person, that there is no objective basis for resolving issues in ethics as there is in other fields. We can do physics or mathematics with our brains, but not ethics.

It is hard to overstate how dangerous this idea is. For centuries, ethics was our primary intellectual tool for reaching consensus. Some of the smartest people who ever walked the earth realized this and invested much of their energy in ethical thought. They include Confucius, Socrates, Plato, Aristotle, Adi Shankara (the leading exponent of Hindu philosophy), Siddhartha Gautama (the Buddha), Thomas Aquinas, Immanuel Kant, and a host of more recent thinkers.

We have forgotten this tradition of ethical reasoning just when we need it most. Our survival in a crowded and fast-moving world depends on a vast web of complex interlocking systems: production, supply chains, commerce, energy, transportation, communication, and legal regulation. These systems would collapse in minutes without the cooperation of countless individuals, and we can cooperate only if we agree on rules of conduct that we find reasonable and just. The task of ethics is precisely to help us reach this kind of agreement. Just as engineering provides the intellectual basis for the physical components of our world, ethics must provide the intellectual basis for social cohesion. Simplistic platitudes and gut feeling won't do the job. We need a subtle and sophisticated theory that can deal with the complexities of modern life.

We may doubt that ethics can meet this challenge, but this is to be expected, because few of us have ever seen a rigorous ethical argument. I didn't see one until I was in graduate school. So we naturally conclude that rigor is not possible in ethics. Granted, we can't find all the tools we need by reading Plato or Kant, which is where ethics courses too often stop. But we don't expect to find the physical formulas we need by reading Copernicus, either. We must plug into more recent thought.

We have forgotten the tradition of ethical reasoning just when we need it most.

Ethical Reasoning

We humans are very talented at rationalizing our behavior; even terrorists justify what they do. If we are to reach consensus on ethics, it is vital to distinguish mere rationalization from correct analysis. The only way to accomplish this is to agree on a few bedrock principles of ethical reasoning before we consider any specific issues. Then we must stick with these principles when we analyze a dilemma, even when we don't like the outcome. We do this in other fields and can do it in ethics.

Let me illustrate one of these principles. I sometimes ask my students why cheating on an exam is wrong. Some say it's wrong because you might get in trouble. But suppose people are cheating and get away with it. Does this make it OK? Most students would say no, particularly when others are cheating and getting the high grades. Some say that cheating is wrong because you might be unqualified for your career. But suppose the exam has nothing to do with your career. Does this make cheating OK? Again, most students would say no. The depressing fact is that, in my many years of teaching, I have never encountered a student who can explain to me why cheating is wrong. This is not the fault of the students, but of the surrounding culture in which they live. Our ethical discourse has become so primitive that we can't justify a simple rule of conduct we learn from childhood.

Here is a quick explanation of why cheating is wrong. I begin with two basic premises: reason is universal, and we have reasons for acting the way we do. Now if I cheat on an exam, I do it for certain reasons. Let's say I cheat because I can get away with it, I want a good job, and the resulting good grade will help me land one. But since reason is universal, I must grant that these same reasons justify cheating for anyone to whom they apply. I know perfectly well that if everyone with these reasons acted on them, one of two things would happen. Either cheating would become general practice, so that grades would be meaningless, or the school would crack down on cheating, and I wouldn't be able to get away with it.

So I am caught in a contradiction. One part of my brain says I should cheat for these reasons, but another part of my brain says that others should not cheat for these same reasons, because then I would no longer be able to get away with cheating and benefit from it. I can't have it both ways. If the reasons justify cheating for me, then I must grant that they justify cheating for all the others, again due to the universality of reason. This, in a nutshell, is why cheating in this situation is wrong. In the ethics business, we say that it violates the generalization principle: I must be able to believe rationally that the reasons for an action are consistent with the assumption that everyone with the same reasons performs the same action.

Obviously I must say more to justify this principle fully. There are other principles as well, based on maximizing utility and respecting autonomy. Applying these principles can be tricky in hard cases, and it takes years of experience to build the necessary skills. Yet no one expects physics and chemistry to be easy, and we shouldn't expect ethics to be easy, either. The world is complicated, and ethics must be complicated enough to deal with it.

Skeptical Reactions

Some will insist that silly little rules like the generalization principle can't deal with the messiness of the real world. For example, the principle presumably tells us not to lie, because if everyone lied, no one would believe the lies. A standard counterexample is the tragic situation of Anne Frank and her family. They were holed up in an Amsterdam office building in the 1940s to hide from Nazis. When the police came to ask about their whereabouts, their collaborators in the building lied and said they didn't know. Ethics supposedly tells us this is wrong. But think about it. The reason for lying to the police is that it will help conceal the whereabouts of the Franks. If everyone with these reasons lied, the police probably wouldn't believe any of the lies.

Yet the lies would serve their purpose, which is to conceal the Franks' whereabouts. If everyone pleads ignorance, the police are as clueless as ever, which means that lying conforms to the generalization principle in this case. The lesson here is that applying ethical principles is a significant intellectual task, no less than applying Maxwell's equations in physics.

I also hear the opposite response. Ethical analysis may be possible, but we don't need it. We know in our gut that cheating is wrong and that Anne Frank's protectors should lie. But do we know in our gut which videos YouTube should take down? Do we know how much personal data Facebook should collect or sell? Do we know which asylum seekers the government should allow into the country? Do we know whether it is okay for a bakery to refuse to cater a gay wedding? We may think we know, but others have directly opposite opinions, and we fall back into polarization.

No one expects physics and chemistry to be easy, and we shouldn't expect ethics to be easy, either. The world is complicated, and ethics must be complicated enough to deal with it.

Even dilemmas that appear clear in retrospect may be foggy at the time. A famous example is the case of the Ford Pinto exploding gas tank, which dates back to the 1970s. Ford decided not to fix the defective gas tank at a cost of \$11 per car, even though repair would have avoided a projected 180 deaths in fiery explosions. Ford's reasoning was that the cost of fixing 12.5 million cars substantially exceeds the expected benefit, if a generally accepted value is placed on human life. This decision was made by well-meaning and conscientious managers on the basis of the utilitarian principle I mentioned earlier. We don't like their conclusion, but what exactly is wrong with it?

The problem is that Ford violated autonomy. Its managers were rationally constrained to believe (and acknowledged) that deaths and serious injury would result from a failure to recall the cars. This is a violation of autonomy unless drivers and passengers knowingly and rationally consent to the risk of driving these cars. While selling nondefective cars is also certain to result in death and injury, customers assume a known risk when they ride in one. They do not, however, assume the risk of riding in a car with a defect they don't know about. I am glossing over some details, but we can see Ford's basic mistake: it applied only one of the tests for ethical conduct.

The Premises

There may also be skepticism about the two premises I began with. One of them is that reason is universal. How do I know this? Maybe I don't, but we don't mind accepting it as a working hypothesis for all of science. I assume nothing more for ethics.

The second premise is that we have reasons for acting the way we do. I make this assumption because it allows us to distinguish free human action from the instinctive behavior of, say, a bumblebee. A bee's behavior is totally determined by chemical and biological causes, so we don't view it as freely chosen action. Yet human behavior is no less determined by physical and biological causes. Our experience of deciding how to act is itself the product of neural mechanisms. If I decide to move my finger while I lie in an MRI machine, the machine operators can observe my brain making this decision before I am aware of making it. My brain has already made up my mind when I make up my mind. Then how can we distinguish my seemingly free action from the buzzing of a bee? Maybe there is no distinction, and therefore no role for ethics?

There is nothing new about this conundrum, of course. The philosopher Epicurus described the conflict between freedom and determinism over 2000 years ago. Fortunately, a resolution suitable for ethics has evolved in the last couple of centuries, namely a "dual standpoint" theory originally suggested by Kant. It characterizes free action as behavior that can simultaneously be explained in two ways: (a) as the result of physical causes, and (b) as the conclusion of a reasoning process. A bee's behavior can be explained only in the first way, while we humans admit the second kind of explanation as well. Human action can be explained by identifying the reasons that led to the agent's choice of action. These reasons are not psychological causes or motivations, but reasons that the human agent consciously takes to justify the action.

Ethics follows immediately. Because the agent's reasons must explain why the action was taken, they must be coherent. Self-contradictory reasons don't explain anything. If the reasons behind an action are inconsistent, as when they violate the generalization principle, the agent is not acting at all. Its behavior is ethically equivalent to a twitch, or buzzing of a bee. On the other hand, if the agent has coherent reasons, it is acting autonomously and therefore ethically. The ethical imperative is really a call to freedom, a call to exercise autonomy. When we tell lies unethically, break promises, or harm others, we are reducing ourselves to animals that merely behave and do not act.

Autonomous Machines

The theory I have sketched is a deep and powerful one that lays a foundation for ethical reasoning in real life. It is also remarkably futuristic, because it provides ethical tools for the coming age of autonomous machines. The analysis nowhere presupposes that an agent must be human. It only presupposes that autonomous action has two kinds of explanation. If a machine's actions can be explained by its programming and as based on reasons that it adduces for the actions, the machine can be an autonomous agent.

No machine today is autonomous in this sense; self-driving cars and the like are programmed but not autonomous. Yet truly autonomous machines may be just over the horizon. Suppose, for example, I someday purchase an intelligent household robot that can explain its conduct. If the robot neglects to do the dishes, for example, I might ask why. The robot responds that it is beginning to develop rust in its joints and believes that washing dishes will make the problem worse. When I ask how the robot knows about the rust, it explains that its mechanic discovered the problem during a regular checkup and advised staying away from water until a rustproof coating can be applied. If I routinely carry on with the robot in this fashion, I am rationally regarding it as an autonomous agent—not as human or as a person, but as an agent.

This has some interesting implications. One is that I have to treat my robot decently. We owe robots duties that follow from their agency, namely generalizability and respect for autonomy

The ethical imperative is really a call to freedom, a call to exercise autonomy.

(utilitarian duties may not follow solely from their agency). We normally must not lie to them or break our promises to them. We must not throw them into the trash while still functional, because this violates autonomy. On the other hand, machines owe similar duties to us, and they will observe these duties to the extent they are truly autonomous. They will be honest with us, and they will not “take over” and oppress us, because this violates our autonomy. Maybe living alongside totally ethical beings is not such a bad prospect.

Toward Ethical Consensus


The retreat from rationality in ethics may have been based on good intentions. It may have been originally motivated by a desire to deal with the very diversity of viewpoints I am trying to address. If we view ethics as merely a matter of personal values, and if we avoid affirming ethical absolutes, then we don’t offend anyone by “imposing our values on others.” We live and let live, and we can all get along.

Yet the purpose of ethics is not to impose our values on others. Just the opposite. The purpose is to arrive at behavior norms on which we can all agree. Ethics is a negotiation tool, not a judgmental tool. Nor does ethics affirm “absolutes,” if this means rules of conduct that apply in all situations. It only imposes formal consistency requirements on the rationales behind our actions. It doesn’t insist, for example, that lying is always wrong, but recognizes that the ethics of lying depends on the context in subtle ways. The only absolutes in ethics are the two premises I began with: the universality of reason, and the necessity of reasons for autonomous action.

I think we have learned that we cannot *just* live and let live. Our lives are too tightly interconnected. Rather, we must use our brains to identify ground rules that we can all accept because they are reasonable. **Maybe if we work toward this goal, it will begin a reversal of our retreat from rationality in general.** 📖



Info



Ready to dig deeper into this idea? Buy a copy of [Taking Ethics Seriously.](#)

Want copies for your organization or for an event?

We can help: customerservice@800ceoread.com 800-236-7323

About the author

John Hooker is a T. Jerome Holleran Professor of Business Ethics and Social Responsibility, and Professor of Operations Research, at Carnegie Mellon University. He has also held visiting posts at several universities, most recently the London School of Economics and the State University of Campinas, Brazil. He brings his extensive background in philosophy and logic to the rigorous analysis of ethical dilemmas, and his background in management science to making sure the dilemmas are realistic. In addition to his blog, [Ethical Decisions](#), he has published over 170 research articles, eight books, and five edited volumes on ethics, philosophy, operations research, and cross-cultural issues, including [Business Ethics as Rational Choice](#) and [Working across Cultures](#). He is the founding editor-in-chief of the world's only academic journal dedicated to teaching business ethics, and he developed the ethics program in the Tepper School of Business at Carnegie Mellon University.



Powered by the love and tender care of 800-CEO-READ, ChangeThis is a

vehicle for big ideas to spread. Keep up with the latest developments in business books and ideas at 800ceoread.com

This document was created on June 13, 2018 and is based on the best information available at that time.

The copyright of this work belongs to the author, who is solely responsible for the content. This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs License. To view a copy of this license, visit Creative Commons. Cover image from Adobe Stock.

Share this

Pass along a copy of this manifesto to others.



Subscribe

Sign up for e-news to learn when our latest manifestos are available.

