

Kwame Anthony Appiah

Experimental moral psychology

How new is experimental philosophy?¹ Wasn't Descartes, whose "mechanical philosophy" aimed to overturn Aristotelianism, really an experimental philosopher? After all, much of his attention was devoted to geometry and optics, and for a period he was revered among scholars as, principally, a sort of mathematical physicist. (That's why the one reference to him educated people mostly know is in talk of the "Cartesian" coordinates he helped invent.) He also spent much time and energy dissecting cows and other animals. Only later was he repositioned as, centrally, a theorist of mind and knowledge, whose primary concern had to do with the justification of belief. In *The Passions of the Soul* (1649), Descartes aimed to solve what we now think of as the canonically philosophical puzzle about the relation between the soul and the body by way of an empirical hypothesis about the role of the pineal gland. Without the pineal – as Nicolaus Steno pointed out in 1669 – Descartes has no story of how mind and body are functionally integrated.²

I don't want to overstate the case: before the disciplinary rise of modern philosophy, one can readily trace dis-

tinctions – between, say, reason and experience, speculation and experiment – that seem cognate to our way of organizing knowledge. Descartes gives us hope when he refers to "first philosophy," and he famously maintained that "all philosophy is like a tree, of which the roots are metaphysics, the trunk is physics, and the branches, which grow from this trunk, are all of the other sciences, which is to say medicine, mechanics, and morals."³ Yet even here we can see that his taxonomy isn't quite ours: morals, to us a division of philosophy, is to Descartes a practical endeavor on a par with medicine.

By the next century, the growing prestige of experimentation was apparent everywhere. The encyclopedist D'Alembert praised Locke for reducing metaphysics to what it should be: *la physique expérimentale de l'âme* – the experimental science of the spirit. And Hume subtitled his great *Treatise of Human Nature*, as we don't always remind ourselves, *Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects*. The point is not just that the canonical philosophers belong as much to the history of what we now call psychology as to the genealogy of philosophy. It is that the "metaphysical" and the psychological claims are, insofar as we insist on dis-

© 2009 by Kwame Anthony Appiah

tinguishing them, profoundly interdependent. Their proper place as ancestors of both modern disciplines is reflected in the fact that many of the claims they make about the mind – including those claims that *are* thought to be of current philosophical relevance – are founded in empirical observation, even if they are not often founded in experiment. They depend on stories about the actual doings of actual people, on claims about how humanity actually is. Hume’s *History of England* – five volumes of empirical information, elegantly organized – has rightly been seen as expressing philosophical ideas about morality and politics and the human mind.

Intellectual historian Knud Haakonssen has argued that our modern conception of the discipline is presaged in the epistemological preoccupations of Thomas Reid and Immanuel Kant.⁴ But Reid himself was emphatic in his suspicion of mere conjecture. Every real discovery, he says, is arrived at by “patient observation, by accurate experiments, or by conclusions drawn by strict reasoning from observation and experiments, and such discoveries have always tended to refute, but not to confirm, the theories and hypotheses which ingenious men had invented.”⁵

As for Kant, it is anything but historically anomalous that the professor from Königsberg, to whom we owe the analytic-synthetic distinction, worked avidly on both sides of the putative divide. Herder admired Kant first of all for his lectures on geography. The founder of “critical philosophy” elaborated theories of the winds and of the earth’s rotation, and dispensed advice about the training of the young. “Games with balls are among the best for children,” he wrote. (“How did he know?” you might wonder.)⁶ Kant possessed the resources for a conceptual partition between what

we think of as philosophy and psychology, but not a vocational one. He was preoccupied with the professionalization of philosophy, but not with avoiding the empirical world. To historians of psychology, Hume and Kant both figure large; and some hold that contrasting Humean and Kantian traditions in scientific psychology continue to this day: as a measurement-driven experimental psychology on the one hand, and the schemas of cognitive science on the other.

One thing we humanists must do is remind ourselves always of the continuities of our intellectual traditions, even when we are trying to produce something new. When talking about recent experimental philosophy, it is important, I think, to start by remembering that we are continuing a long tradition.

There are at least two rather different ways in which experiment is currently being brought to bear in moral philosophy. One way – which I associate with the work of Joshua Knobe – continues the project of conceptual analysis, as in Knobe’s well-known work on intention. Here experiment figures in exactly the way that Austin and Wittgenstein would have thought mad or impertinent. If conceptual analysis is the analysis of “our” concepts, then shouldn’t one see how “we” (or representative samples of us) actually mobilize concepts in our talk? One of the most interesting things that disciplined and responsible experiments on usage reveal is that the real world contains scads of folk (like the students we meet in our introductory classes) who simply aren’t inclined to say the things that Austin or Wittgenstein would have supposed all competent speakers should say.

In Knobe’s most famous experiment, subjects were asked to consider two sce-

narios. In the first, the chairman of a company is asked to approve a new program that will increase profits and also help the environment. “I don’t care at all about helping the environment,” the chairman replies. “I just want to make as much profit as I can. Let’s start the new program.” So the program is launched and the environment is helped. The second story is identical – except that the program will *hurt* the environment. Once again, the chairman is indifferent to the environment, and the program is launched in order to increase profits, with the expected results. When Knobe presented these scenarios to subjects in a controlled experiment, he found that when the program helped the environment, only 23 percent agreed that the chairman had “helped the environment intentionally.” When the program harmed the environment, though, 82 percent agreed that the chairman had “harmed the environment intentionally.” This pattern recurred when various other scenarios were tested.⁷

One striking feature of this result – irrelevant to Knobe’s main point, but important nevertheless – is that roughly one in five people in each case had apparently managed to grow up with very different intuitions from the rest of us. Both intuitions have their advocates. And it seems to me that the right answer isn’t a matter of a head count. This work would be valuable and suggestive even if it skipped the actual experiments. (“It would be natural to say,” Knobe might have written, “that the chairman in one situation had harmed the environment intentionally, whereas . . .”) The experimental evidence enforces a useful modesty about how much weight to give one’s personal hunches, even when they’re shared by the guy in the next office.

But, as I say, there’s a second sort of experimental philosophy. The paradigm here might be Josh Greene’s work in moral psychology, in which he and his colleagues have studied the fMRI images of the brains of people thinking through so-called trolley cases, in which people are asked what they think the right thing is to do when a trolley, whose driver is unconscious, is bearing down on a group of people on a track. In a typical trolley case, you are offered the option of diverting the trolley from a track where it will kill six people to another track where it will kill only one. In the past, philosophers have used these scenarios to ask questions about how we make moral choices. But Greene’s experiments look at more than patterns of response to scenarios, more even than statistics about concept use. Greene is not just carrying out with a sample of speakers the thought-experiments that might once have been done by the philosopher in her armchair. Rather, he is looking, too, at evidence about the neural or psychological processes that underlie our responses.⁸

One thing we can say about this work is that it is really only an extension of something that we have been doing for a long time anyway; it’s just the further pursuit of the philosophy of – in this case, moral – psychology. Philosophers in cognitive science have been doing this sort of thing for a very long while. They’ve been looking at work in experimental psychology to see what it suggests about how our minds work, and they’ve sometimes suggested experiments and hung around while they were being done. And one reason that many of the canonical philosophers figure also in the histories of psychology is that, though they may have had fewer instruments and less rigorous methods, they were often relying on

empirical propositions about the mind and its brain all along.

It seems to me, however, that we risk losses along with the gains to be had from the work of experimental philosophy, unless we hold onto what is best in the traditions of conceptual analysis. One problem with methods, like all tools, is that they can lead us to focus too little attention on the problems with which they cannot help. If you have a hammer, as they say, everything can begin to look like a nail. I think that happened with the development of the trolley paradigm in moral philosophy – long before it was connected with experimental moral psychology, in the days when trolley intuitions were pursued exclusively in the armchair. An experimental moral philosophy attentive solely to the neurophysiology of the ways we think through trolley problems would give us a distorted picture of our ethical lives.

A few decades ago, Edmund Pincoffs pointed out the historical novelty of what he called “quandary ethics,” asking hard questions about the disciplinary distortions it had engendered.⁹ Quandary ethics, as the name suggests, took the central problem of moral life to be the resolution of quandaries about what to do. One of its favorite methods was to examine stylized scenarios, like the trolley problem, and figure out what we should do and why.

I am not always a foe of quandaries; but, especially when we’re trying to come to grips with the larger subject of *eudaimonia* – the question of what it is for a human life to go well – students of the moral sciences, including the experimental philosophers, should recognize how stark the limitations of quandaries are. To turn to them for guidance in the arena of ethics, conceived at its broadest, is like trying to find your way around at night with a laser pointer.

Consider this one difficulty: in all of those trolley cases, the options are given in the description of the situation. But in the real world, situations are not bundled together with options. Instead, the act of framing – describing a situation, and thus determining that there’s a decision to be made – is itself a moral task. It’s often *the* moral task. Learning how to recognize what is and isn’t an option is part of our ethical development. For example, part of the point of the stringency of the prohibition against murder, Anscombe once observed, was “*that you are not to be tempted by fear or hope of consequences.*”¹⁰ So a proper response to situations like these would be to look, first, for other options. To understand what’s wrong with murder is, in part, to be disinclined to take killing people as an option. If we want to learn about normative life from stories, I suspect that the most helpful ones are going to come from movies, novels, and the like, in which characters have to understand and respond to complex situations, not just pick options in an SAT-style multiple-choice problem. In life, the challenge is not so much to figure out how best to play the game; the challenge is to figure out what game you’re playing.

I offer this brief objection not because this argument is inaccessible to the experimental moral philosopher. Indeed – as I began by insisting – the difficulties with quandary ethics long antedate recent experimental moral psychology. My point is a different one: that in thinking about what we learn from the experiments about ethical life, un-experimental philosophy is still very much in order. Nothing I have said about quandaries depends on any experimental evidence.

So far, then, just a caution. The armchair remains an important research

tool, and work done in armchairs remains central to our subject – as, of course, it does to most: mathematics is not the only way to turn coffee into theories. And so, let me end with a brief sketch of just one case where it seems to me experimental psychology has been extremely helpful in thinking about our moral lives. It has to do with some lessons of social psychology; and my account here will draw heavily on John Doris’s excellent book *Lack of Character*.

Social psychologists are mostly “situationists”: they claim (this is a first stab at a definition) that a lot of what people do is best explained not by traits of character, but by systematic human tendencies to respond to features of their situations that nobody previously thought to be crucial at all.¹¹ They think that someone who is, say, reliably honest in one kind of situation will often be reliably dishonest in another. They’d be unsurprised, for example, that Oskar Schindler was mercenary, arrogant, hypocritical, calculating, and vain sometimes . . . but not always; and that his courage and compassion could be elicited in some contexts but not in others.

Now, to ascribe a virtue to someone is, among other things, to say that she tends to do what the virtue requires in contexts where it is appropriate.¹² An honest person will resist the temptations to dishonesty posed by situations where, say, a lie will bring advantage, or failing to return a lost wallet will allow one to buy something one needs. Indeed, our natural inclination, faced with someone who does something helpful or kind – or, for that matter, something hostile or thoughtless – is to suppose that these acts flow from their character, where character is understood as a trait that is consistent across situations and, therefore, insensitive to differences in the

agent’s environment, especially small ones. But situationists cite experiments suggesting that small – and morally irrelevant – changes in the situation will lead a person who acted honestly in one context to do what is dishonest in another.

In the past thirty years or so, psychological evidence for situationism has been accumulating. Back in 1972, Alice M. Isen and Paula Levin found that, if you dropped your papers outside a phone booth in a shopping mall, you were far more likely to be helped by someone who had just had the good fortune of finding a dime waiting for him in the return slot. A year later, John Darley and Daniel Batson discovered (in perhaps the most famous of these experiments) that Princeton seminary students, even those who had just been reflecting on the Gospel account of the Good Samaritan, were much less likely to stop to help someone “slumped in a doorway, apparently in some sort of distress,” if they’d been told that they were late for an appointment. More recently, Robert Baron and Jill Thomley showed that you were more likely to get change for a dollar outside a fragrant bakery shop than standing near a “neutral-smelling dry-goods store.”¹³

Many of these effects are extremely powerful: huge differences in behavior flow from differences in circumstances that seem of little or no normative consequence. Putting the dime in the slot in that shopping mall raised the proportion of those who helped pick up the papers from one out of twenty-five to six out of seven – that is, from almost no one to almost everyone. Seminarians in a hurry are six times less likely to stop like a Good Samaritan.¹⁴ Mindful of these examples, you should surely be a little less confident that “she’s helpful” is a good explanation next time someone stops to assist you in picking

up your papers, especially if you're outside a bakery!

I am not going to worry about whether we can give an account of the results that better comports with our common sense about why people, ourselves included, do what they do. The question I want to ask is, why should ethical theory care about these claims at all?

Suppose I give you change because (in part) I just got a whiff of my favorite pastry. Of course, if I had a settled policy of never giving change, even that pleasant aroma wouldn't help. So there are other things about me – the sorts of things we would normally assess morally – that are relevant to what I have done. But let's suppose that, other things being equal, if I hadn't had the whiff, I'd have ignored your plaintive plea to stop and change your dollar for the parking meter. Pleased by the ambient aroma, I was inclined to do what, according to the virtue theorist, a kind or helpful or thoughtful person – a virtuous person – would do; and I acted on that inclination. A typical virtue theorist will think I have done the right thing because it is the kind thing (and there are no countervailing moral demands on me). But, on the situationist account, I don't act out of the virtue of kindness. Does this act accrue to my ethical credit? Do I deserve praise in this circumstance or not? Have I or haven't I made my life better by doing a good thing?

A situationist might encourage us to praise someone who does what is right or good – what a virtuous person would do – whether or not they did it out of a virtuous disposition, but only for instrumental reasons. After all, psychological theory also suggests that praise, which is a form of reward, is likely to reinforce the behavior. (What behavior? Presumably not helpfulness, but being helpful

when you're in a good mood.) For we tend to think that helping people in these circumstances, whatever the reason, is a good thing.¹⁵

But the virtue ethicist cannot be content that one acts *as if* virtue ethics is true. And we can all agree that the more evidence there is that a person's conduct is responsive to a morally irrelevant feature of the situation, the less praiseworthy it is. If these psychological claims are right, very often when we credit people with compassion, as a character trait, we're wrong: they're just in a good mood. And if hardly anyone is virtuous in the way that virtue ethics conceives of it, isn't the doctrine's appeal eroded? Given that we are so sensitive to circumstances and so unaware of that fact, isn't it going to be wondrously difficult to develop compassion, say, as a character trait? We can't keep track of all the cues and variables that may prove critical to our compassionate responses: presumably the presence or absence of the smell of baking is just one among thousands of contextual factors that will have their way with us. How, if this is so, can I make myself disposed to do or to feel the right thing? I have no voluntary control on how aromas affect me. I cannot be sure that I will have a free dime show up whenever it would be a fine thing to be helpful.

There are some philosophers, among them the aforementioned John Doris, who take the social-science literature about character and conduct to pose a serious and perhaps lethal challenge to the virtue ethicist's worldview. For one thing, our virtue theorist faces an epistemological difficulty if there are no actually virtuous people. As in all spheres of thought, so in moral deliberation: we sometimes need to think not only

about what the right answer is, but also about how we discover what the right answers are. Rosalind Hursthouse, in her book *On Virtue Ethics*, has argued:

- 1) The right thing to do is what a virtuous agent would do in the circumstances.
- 2) A virtuous person is one who has and exercises the virtues.
- 3) A virtue is a character trait that a person needs in order to have *eudaimonia*, in order to live a good life.¹⁶

No interesting version of virtue ethics holds that doing the right thing is all that matters; we should want to be the kind of person who does the right thing for the right reasons.

Still, Hursthouse and others insist that virtue ethics isn't *entirely* "agent-centered," rather than "act-centered"; it can also specify what the right thing to do is – namely, what a virtuous person would do. How are we to follow that advice? If we were fully virtuous, we would find ourselves disposed to think, act, and feel the right things. But we are not. If we knew someone who was virtuous, we could see what she would do, I suppose. Given the depressing situationist reality, however, maybe no actual human being really is (fully) virtuous. And even if a few people did get to be virtuous against all the odds, we would have to have some way of identifying them, before we could see what they would do. So we would need, first, to know what a good life looks like, and then we would need to be able to tell, presumably by reflecting on actual and imaginary cases, whether having a certain disposition is required for a life to be good – and required not in some instrumental way, as nourishment is required for any life at all, but intrinsically.

If experimental psychology shows that people cannot have the sorts of character traits that the virtue theorist has identified as required for *eudaimonia*, there are only two possibilities: she has identified the wrong character traits, or we cannot have worthwhile lives. Virtue theory now faces a dilemma. The problem for the idea that we have gotten the wrong virtues is a problem of method. For virtue theory of the sort inspired by Anscombe, we must discover what the virtues are by reflection on concepts. We can, in principle, reflect on which of the stable dispositions that psychology suggests might be possible – being helpful when we are in a good mood, say – are constitutive of a worthwhile life; or which – being unhelpful when we aren't buoyed up by pleasant aromas – detract from a life's value. But to concede *that* is to accept that we'll need to do the experimental moral psychology before we can ask the right normative questions. On this horn of the dilemma, virtue theory will find itself required to take up with the very empirical psychology it so often disdains.

On the other horn of the dilemma, the prospect that we cannot have worthwhile lives makes normative ethics motivationally irrelevant. What is the point of *doing* what a virtuous person would do if I can't *be* virtuous? Once more, whether I can be virtuous is obviously an empirical question. Once more, then, psychology seems clearly apropos.

Still, we should not overstate the threat that situationism poses. The situationist account doesn't, for example, undermine the claim that it would be better if we *were* compassionate people, with a persistent, multitrack disposition to acts of kindness. Philosophical accounts of the character ideal of compassion, the conception of it as a virtue, need make

no special assumptions about how easy or widespread this deep disposition is. Acquiring virtue, Aristotle already knew, is hard; it is something that takes many years and most people don't make it. These experiments might confirm the suspicion that compassionate men and women are rare, in part because becoming compassionate is difficult. But difficult is not the same as impossible; and perhaps we can ascend the gradient of these virtues only through aspiring to the full-fledged ideal. Nor would the ideal be defeated by a situationist who busily set about showing that people whom we take to exemplify compassion – the Buddha, Christ, Mother Teresa – were creatures of environments that were particularly rich in the conditions that (according to situationists) elicit kindly acts.

Finally, we could easily imagine a person who, on the virtue ethicists' view, was in some measure compassionate, and who actually welcomed the psychologists' research. Reading about these experiments will only remind her that she will often be tempted to avoid doing what she ought to do. So these results may help her realize the virtue of compassion. Each time she sees someone who needs help when she's hurrying to a meeting, she'll remember those Princeton seminarians and tell herself that, after all, she's not in *that* much of a hurry; that the others can wait. The research, for her, provides a sort of perceptual correction akin to the legend you see burned onto your car's rearview mirror: *objects may be closer than they appear*. Thanks for the tip, she says. To think that these psychological claims by themselves undermine the normative idea that compassion is a virtue is just a mistake.

We might also notice what the situationist research *doesn't* show. It doesn't

tell us anything about those seminarians (a healthy 10 percent) who were helpful even when rushing to an appointment; perhaps that subpopulation really did have a stable tendency to be helpful – or, for all we know, to be heedless of the time and careless about appointments. (Nor can we yet say how the seminarians would have compared with, say, members of the local Ayn Rand society.) There could, consistent with the evidence, be a sprinkling of saints among us. Some will dispute whether the dispositions interrogated by social psychology can be identified with the normative conception of character traits elaborated by the classical virtue theorists.¹⁷ And, of course, the situationist hypothesis is only that, in explaining behavior, we're inclined to overestimate disposition and underestimate situation. It doesn't claim that dispositions don't exist.

None of these caveats wholly blunts the situationist point that the virtues, as virtue ethicists conceive them, seem exceedingly hard to develop, which must leave most of us bereft of *eudaimonia*. But virtue ethics is hardly alone in assigning a role to elusive ideals. Our models of rationality are also shot through with such norms. Recall the nineteenth-century hope that, in the formula, logic might be reduced to a "physics of thought." What succeeded that project was an approach captured in another formula, according to which logic is, in effect, an "ethics of thought."¹⁸ It tells us not how we do reason, but how we ought to reason. And it points toward one way of responding to the question we have posed to the virtue ethicist: how might we human beings take seriously an ideal that human beings must fall so far short of?

If you have been following debates about the role of ideals in cognitive psychology, you might think the answer is

to treat claims about virtues as moral heuristics. But there are many difficulties, I think, for this view.¹⁹ Here is one: for faithful Aristotelians, this whole approach, in which we seek moral heuristics that will guide us imperfect creatures to do what a virtuous person would do, is bound to look very peculiar. Virtue ethics wants us to aim at *becoming* a good person, not just at maximizing the chance that we will do what a good person would do. The contrast with familiar cognitive heuristics is striking.

For cognitive heuristics are, so to speak, twice dipped in means-end rationality. First, the right outcome is defined by what someone equipped with ideal means-end rationality, someone possessed of infinite cognitive resources, would do. Second, we then apply means-end rationality to determine how people with limited cognitive resources can maximize their chances of doing what's right according to the first test. When we try to concoct a heuristic of virtue, we must start, analogously, by defining the right outcome as what someone ideally virtuous would do. Since we're not ideally virtuous, the heuristics model now introduces means-end rationality to maximize your chance of doing what's right by the first test. The trouble is, of course, that virtue ethics requires that we aim at the good for reasons that aren't reducible to means-end rationality. With the cognitive heuristic, what matters is the outcome. But if virtue ethics tells you that outcomes aren't the only thing that matters, then you cannot assess heuristics by means-end rationality – that is, by looking at the probability that they will produce certain outcomes.

To be sure, the fact that virtues are meant to be constitutive of a life of *eudaimonia* – so that they are traits necessary to make our lives worthwhile –

is consistent with the view that a virtuous person's life will have good effects as well. Perhaps a life of virtue will be an enjoyable life, too: Aristotle certainly thought that a fully virtuous person would take pleasure in the exercise of virtue. But the value of the virtues does not come just from the good results of virtuous acts or from the enjoyment that virtue produces; it is intrinsic, not instrumental. A virtuous life is good because of what a virtuous person *is*, not just because of what she *does*.

We can distinguish, then, between *having a virtue* and *being disposed to do the virtuous act over a wide range of circumstances*. We can distinguish, in particular, *being an honest person*, someone who has that virtue as the virtue ethicist conceives of it, and *being someone who, across a wide range of circumstances, behaves as an honest person would*. Suppose honesty matters in my life because it promotes reliability and thereby helps me support the flourishing of others. If that were so, I might explore some alternative possibilities by which I might refrain from deceiving others. Perhaps someone has developed a Bad Liar pill, which will impair my capacity for successful deception; or perhaps our town has collectively decided to add the drug to the water supply, as a moral counterpart to fluoridation. Equivalently, we could try to heighten our ability to detect deception. Either strategy amounts to a similar trade-in: a scenario in which I strive to be honest in all situations is exchanged for a scenario in which I can usually be relied upon not to deceive others.

It would be a mistake to deny the instrumental significance of honesty; but doesn't our moral common sense recoil at the idea that honesty matters *only* because of this instrumental significance? (There's a question for Knobe to pur-

sue.) Denying that significance courts moral narcissism; but reducing honesty's importance to that instrumental significance threatens to replace the ethical subject with the object of social engineering. In all events, my aim is not to fine-tune the dictum that an action is right if it's what a virtuous person would do. I have only tried to illustrate how alien that dictum is to what made the eudaemonist tradition appealing in the first place.

I have no doubt, then, that we are learning things worth learning from all sorts of experimental philosophy. I have no doubt, too, that it is a bad idea, if you are interested in the sorts of questions these philosophers (and their friends in psychology and economics) are addressing, to ignore their work. There is even good reason, as I have argued, to think of what they are doing as much more continuous with the past practices of major philosophers than the paradigm of con-

ceptual (or, I might add, phenomenological) analysis would suggest. But it remains the case that responding to the experiments requires the sort of careful examination of arguments, the making of distinctions, the reflection on unactualized possibilities that are also a part of the tradition and can be found in – to construct a deliberately eclectic list – Aristotle, Aquinas, Descartes, Leibniz, Hume, Kant, Hegel, Frege, Husserl, James, Russell, and Sartre. The enthusiasts for the experiments should insist that the armchair would be a much less interesting and productive place if they were not going about their business, too. I agree; I have been agreeing all along. Indeed, insofar as method is concerned, I have only this modest pluralist suggestion: that we would do well to sustain a variety of traditions of reflection on questions that matter to us. Unless you already know all of the answers, you don't know for sure which questions are worth asking.

ENDNOTES

¹ This essay is based on material from my book *Experiments in Ethics* (Cambridge, Mass.: Harvard University Press, 2008).

² See Nicolaus Steno, *Lecture on the Anatomy of the Brain*, introduction by Gustav Scherz (Copenhagen: Nyt Nordisk Forlag, Arnold Busck, 1965), 12 et seq.

³ René Descartes, (1647) "Lettre-Préface de l'édition française des Principes"; available at <http://www.ac-nice.fr/philo/textes/Descartes-LettrePreface.htm>.

⁴ Knud Haakonssen, "The Idea of Early Modern Philosophy," in *Teaching New Histories of Philosophy*, ed. Jerry Schneewind (Princeton: University Center for Human Values, 2004), 108.

⁵ Thomas Reid, *Essays on the Intellectual Powers of Man*, Essay I., iii, "Of hypothesis," in *The Works of Thomas Reid*, vol. I (New York: Published by N. Bangs and T. Mason, for the Methodist Episcopal Church, J. and J. Harper Printers, 1822), 367–368.

⁶ *Kant on Education*, trans. Annette Churton (Ann Arbor: University of Michigan Press, 1960), 61.

⁷ Joshua Knobe, "The Concept of Intentional Action: A Case Study in the Uses of Folk Psychology," *Philosophical Studies* 130 (2006): 203–231.

- ⁸ Joshua D. Greene, R. Brian Sommerville, Leigh E. Nystrom, John M. Darley, Jonathan D. Cohen, "An fMRI Investigation of Emotional Engagement in Moral Judgment," *Science* 293 (5537) (2001): 2105 – 2108.
- ⁹ Edmund Pincoffs, "Quandary Ethics," *Mind* 80 (1971): 552 – 571; see also his *Quandaries and Virtues: Against Reductivism in Ethics* (Lawrence: University of Kansas Press, 1986).
- ¹⁰ G. E. M. Anscombe, "Modern Moral Philosophy," *Philosophy* 33 (1958); reprinted in G. E. M. Anscombe, *Ethics, Religion, and Politics* (Minneapolis: University of Minnesota Press, 1981), 34.
- ¹¹ See Lee Ross and Richard E. Nisbett, *The Person and the Situation* (Philadelphia: Temple University Press, 1991).
- ¹² John Doris, *Lack of Character: Personality and Moral Behavior* (Cambridge: Cambridge University Press, 2002), 16 – 19.
- ¹³ Alice M. Isen and Paula F. Levin, "The Effect of Feeling Good on Helping: Cookies and Kindness," *Journal of Personality and Social Psychology* 21 (1972): 384 – 388; John M. Darley and C. Daniel Batson, "'From Jerusalem to Jericho': A Study of Situational and Dispositional Variables in Helping Behavior," *Journal of Personality and Social Psychology* 27 (1973): 100 – 108; Kenneth E. Matthews and Lance K. Cannon, "Environmental Noise Level as a Determinant of Helping Behavior," *Journal of Personality and Social Psychology* 32 (1975): 571 – 577; Robert A. Baron and Jill Thomley, "A Whiff of Reality: Positive Affect as a Potential Mediator of the Effects of Pleasant Fragrances on Task Performance and Helping," *Environment and Behavior* 26 (1994): 766 – 784. All are cited in Doris, *Lack of Character*, 30 – 34, 181.
- ¹⁴ And people are about one tenth as likely to help someone behind a curtain who has had what sounds like an accident if there's someone else standing by who does nothing; Bibb Latane and Judith Rodin, "A Lady in Distress: Inhibiting Effects of Friends and Strangers on Bystander Intervention," *Journal of Experimental Social Psychology* 5 (1969): 189 – 202, as cited in Doris, *Lack of Character*.
- ¹⁵ Of course, its being good because it helps doesn't mean it isn't bad overall: suppose you're a nasty person who offers change only because you know you're being watched by someone who has promised to give you fifty bucks if you ever do anything generous. That's blameworthy: you're trying to fake generosity.
- ¹⁶ Rosalind Hursthouse, *On Virtue Ethics* (Oxford: Oxford University Press, 1999). "Virtue Ethics," *The Stanford Encyclopedia of Philosophy*, fall 2003 edition, ed. Edward N. Zalta; available at <http://plato.stanford.edu/archives/fall2003/entries/ethics-virtue/>.
- ¹⁷ See, for example, Julia Annas, "Virtue Ethics and Social Psychology," *A Priori* 2 (2003): 20 – 59; and Rachana Kamtekar, "Situationism and Virtue Ethics on the Content of Our Character," *Ethics* 114 (2004): 458 – 491, which, however, acknowledges, at 461, that "virtue ethics can benefit from considering the particular situational factors that social psychology suggests have a profound influence on behavior."
- ¹⁸ Theodor Lipps, (1880) *Die Aufgabe der Erkenntnistheorie und die Wundt'sche Logik I.* (The task of epistemology and Wundtian logic I.), *Philosophische Monatshefte*, 16, 529 – 539; cited in Edward S. Reed, *From Soul to Mind: The Emergence of Psychology from Erasmus Darwin to William James* (New Haven: Yale University Press, 1997), 188.
- ¹⁹ For more of them, see my *Experiments in Ethics*, chap. 4.