Facts as Determinants for Morality: How Human Action, Reason, Cognition, and Social Activity Determine What We Ought To Do

Michael S. Dauber
New York University

Introduction

It is a central axiom of moral philosophy that "ought implies can." We think that if someone can't do something, they ought not to be held morally responsible if they don't do it, and if someone *can* do something, they've done something morally wrong if they have not. Determining what we ought to do, then, is a question about facts: what is the world like? What is it possible for us to do? In terms of human action, this involves how the mind works and how people behave. Studies in neuroscience and social dynamics have shown and suggested a number of things about human activity.[1] Jonathan Haidt's moral reasoning experiments have strongly suggested that human beings don't make moral judgments on the basis of reason: instead, they have an emotional intuition of some sort, and then rationalize those feelings in some sort of reason-based explanation. Joshua Greene's work with neuroimaging builds on this idea and suggests that people make primarily emotional judgments when assessing personal moral violations, and make reasoned judgments about more collective dilemmas. Other studies, like those of Benjamin Libet, have provided tentative evidence for determinism, the idea that history and physics are ultimately responsible for the choices we make, and that personal volition plays little to no role in human action beyond a phenomenological sense of control.

This paper will explore some of the ways in which studies in neuroscience determine what we ought to do and inform both the role of the humanities and how we ought to teach the humanities going forward. I will not provide an exhaustive analysis of all of the relevant neuroscientific developments. Rather, I will provide a brief sampler of some of the most notable studies from recent years, with an eye toward painting a picture of how human nature (interpreted through psychology and neuroscience) informs what we can and cannot do, and therefore what we ought and ought not to do. The first section will explore the work of Jonathan Haidt and his Social Intuitionist Model. Through various thought experiments and interviews with test subjects, his work suggests that many of us rely on intuitions to make moral judgments, not reasons. In the second section, I discuss Joshua Greene's work on the mechanisms in play when we make deontological and consequentialist judgments. As with Haidt's work, understanding the manner in which we make certain types of judgments can help explain some of the ways we might purposely motivate people to behave in certain ways. In section III, I analyze some of the studies conducted on free will, specifically

---

[1] Guy Kahane has argued that the argument that neuroscience and psychology influence what we ought to do faces problems because there might still be a wide array of things we are capable of doing, and an even wider array of things we might be able to convince ourselves to do. I will not fully address this argument here. However, it seems that narrowing down the field of possible actions using the insights of these studies would still be helpful, if only by process of elimination. See Guy Kahane, "Is, Ought, and the Brain" in *Moral Brains: The Neuroscience of Morality*, ed. S. Matthew Liao (Oxford University Press, 2016).

Benjamin Libet's seminal study that suggests we do not have the power to decide to act, but merely to refrain from acting.

Once our brief survey of moral cognitive science has concluded, we can proceed to how we might use this information. In section IV, I suggest that these developments determine what we ought to do, both in a positive and negative sense: if we know that people cannot be motivated to act in certain ways or care about certain things, we have a more specific picture of what we ought to encourage people to do and care about. In the final section, I suggest the ways in which these findings can be applied to teaching and studying the humanities. Assuming, for the sake of argument, that these studies are correct, the proper place of the humanities lies not in descriptive accounts of human action, but in prescriptive accounts: if science can tell us how we behave, the role of the humanities is to tell us how we ought to behave.

Before moving forward, I should say a few words about my own views regarding each of these studies, and how we ought to view them for the purposes of this paper. I do not claim to believe any of the studies in this paper are true, although I am sympathetic to many elements of Jonathan Haidt's work. The papers presented here are merely used as examples of prominent studies in neuroscience, psychology, and sociology that, if they do indeed yield true conclusions, would have a great impact on how we view ourselves, our behavior, and the world around us. The central claims of this paper, namely that we ought to alter our reading and instruction of the humanities if these studies give accurate results, are contingent on the truth of the studies themselves. If we have strong reasons to doubt that one or more of these studies are correct, and that alternative accounts are viable, we would not teach some neuroscientific and psychological theories as if they were definitively true. The general claim that neuroscience, psychology, and sociology influence the humanities and normativity if these fields can help us understand ourselves and our behavior, however, remains untouched.

I. Jonathan Haidt, Moral Dumbfounding, The Social Intuitionist Model, and Moral Foundations Theory

Jonathan Haidt's book *The Righteous Mind: Why Good People Are Divided By Politics And Religion* is a synthesis of several decades of research in moral psychology. Widely construed, moral psychology is the study of how we make moral judgments and think about moral dilemmas. Haidt's work has focused specifically on exploring the emotional side of moral thinking: whereas traditional discourse on morality from Kant to Kohlberg has focused largely on the role of reasoning, Haidt believes the action takes place in the emotional centers of the brain. His work is significant for two major concepts: "Moral Dumbfounding" and the Social Intuitionist Model.

In order to test what happens when people make certain kinds of moral judgments, Haidt designed a series of interview questions designed to test our moral intuitions. Interview subjects are asked to assess whether or not a specific action is morally right or morally wrong. Many people cite harm as the chief justification for moral impermissibility: if someone is harmed, then the action is morally wrong. Haidt became particularly interested in testing why people sometimes think possible actions are wrong even if nobody is harmed. One of the most famous examples they presented subjects with runs as follows:

Julie and Mark, who are sister and brother, are traveling together in France. They are both on summer vacation from college. One night they are staying alone in a cabin near the beach. They decide that it would be interesting and fun if they tried making love. At the very least it would be a new experience for each of them. Julie is already taking birth control pills, but Mark uses a condom too, just to be safe. They both enjoy it, but they decide not to do it again. They keep that night as a special secret between them, which makes them feel even closer to each other. So what do you think about this? Was it wrong for them to have sex?[2]

Haidt found that even though nobody was harmed in the above example (no baby born of incest was created and no psychological trauma was forced on either partner), 80% of respondents still condemned the two siblings' actions. Further, respondents "seemed to be flailing around, throwing out reason after reason, and rarely changing their minds when [the interviewer] proved that their latest reason was not relevant."[3] One quote from an interview stands out in particular. After all of the respondent's justifications have been refuted or challenged, the respondent says the following: "Um…well…oh, gosh. This is hard. I really—um, I mean, there's just no way I could change my mind but I just don't know how to—how to show what I'm feeling, what I feel about it. It's crazy!"[4]
Subjects in these cases are *morally dumbfounded*, according to Haidt: they feel that something is morally wrong or objectionable, but they are unable to explain or fully understand why beyond an innate feeling of disgust, anger, sadness, etc. Haid writes that "[subjects] made moral judgments quickly and emotionally. Moral reasoning was mostly just a post hoc search for reasons to justify the judgments people had already made."[5] Haid interprets these findings as indicating a much greater role for emotions and intuitions in (moral) judgments than the rationalist tradition cared to acknowledge.[6]

To explain and map the empirical roles of emotions, intuitions, and reasoning, Haidt proposed the social intuitionist model. The main thrust of the theory is to reverse the traditional direction of rationalist models. While rationalist models generally assert that some sort of stimulus triggers us to contemplate the rightness or wrongness of an action, with emotions produced as a product of the deliberative process, the social intuitionist model argues the opposite: "Intuitions come first and reasoning is usually produced after a judgment is made, in order to influence other people." The upshot of the argument is that it explains how the process of forming judgments occurs, while also explaining why we seem to have strong reasons to hold that judgment. The role of reasoning, according to Haidt, is to influence the actions of others (Haidt calls thinks of

---

[2] Jonathan Haidt, *The Righteous Mind*, p. 45.
[3] Haidt, pp. 45-46.
[4] Haidt, p. 47.
[5] Ibid.
[6] As Haidt notes in his book, other researchers had begun to chip away at the rationalist paradigm as well. In particular, Antonio Damasio's experiments on patients with injuries to the ventromedial prefrontal cortex (vmPFC) prove striking. Patients with vmPFC injuries are incapable of feeling emotions. While they were able to recognize what we normally consider morally right and wrong actions, this knowledge did nothing to motivate them to act, suggesting that emotion is both necessary for motivating us to act and in helping us make judgments. Without the information provided by emotions, such patients were incapable of both complex moral evaluations and mundane judgments about what to do in everyday situations. See Antonio Damasio. *Descartes' Error: Emotion, Reason, and the Human Brain* (New York: Putnam, 1994).

reasoning as a mental lawyer or press secretary, coming up with justifications for our intuitions to defend us, the clients).

Haidt's findings and the findings of his contemporaries put emotional theories of motivation and moral cognition back on the table as strong alternatives to rationalist theories. As Haidt gathered more and more data about how we make moral judgments, he noticed that our reactions to moral dilemmas focused on a specific set of values. Specifically, he initially noticed that people's intuitions rested on five foundations[7]: Care/Harm (we should "protect and care for" people), Fairness/Cheating (benefits should be distributed in a just[8] manner), Loyalty/Betrayal (fidelity to a group, following up on one's promises, etc.), Authority/Subversion (accordance with beneficial hierarchies, respecting members with higher places in the hierarchy), and Sanctity/Degradation (originally about avoiding contamination, but later expanded to include taboos and a respect for objects, people, places and practices considered holy or proper).

Further, Haidt noticed that people of different political persuasions made judgments that emphasized different moral foundations. Liberals tended to place greater emphasis on the Care/Harm and Fairness/Cheating foundations, while Conservatives also factored in Loyalty/Betrayal, Sanctity/Degradation, and Authority/Subversion (and, adjusting for his later model, Proportionality and Liberty). While it is certainly realistic for members of different parties to emphasize foundations characteristically associated with the opposite party, many partisan positions track this division. Consider, for example, conservatives' opposition to flag burning: many conservatives believe burning the flag is morally wrong because it is disrespectful (Authority/Subversion), because it fails to display solidarity with service men and women who defend our country (Loyalty/Betrayal), and because the flag, the quintessential symbol of our country and values, deserves respect (Sanctity/Degradation). Next consider taxation: Liberals tend to think taxes should be higher on richer citizens to support higher government spending on welfare programs (Care/Harm in providing services for people and Fairness as Equality in attempting to elevate people in lower economic positions), while conservatives tend to people shouldn't be taxed more heavily just because they make more money (Fairness as proportionality in that the people in question should keep what they earned, and Liberty in that people should be able to keep and spend what they earn as they please).

If Haidt's theories are correct, people are driven more by emotions than by rational engagement, at least when they initially make moral judgments. These judgments tend to rest on a specific few kinds of foundations, and liberals and conservatives tend to draw on different foundations when they react to moral dilemmas and propositions.

II. Joshua Greene and Neuroimaging in Moral Judgments

---

[7] He later expanded the model to include six foundations to accommodate different conceptions of justice/fairness: liberals interpret fairness as equality, conservatives as proportionality. He later added "Liberty" to better capture this difference.

[8] As mentioned in footnote six, this is complicated by how liberals and conservatives think of justice. I here use just in a vague sense to mean "Justice, however one thinks of it." The problem will later be solved when discussing the addition of Liberty in Haidt's reformulated model.

One of the most significant technological advancements in recent decades is functional magnetic resonance imaging, commonly referred to as fMRI, a brain scanning technique that enables scientists to track the activity of different locations in the brain. In "The Secret Joke Of Kant's Soul," Joshua Greene argues that characteristically deontological judgments rely on emotional intuitions, whereas consequentialist judgments are more rational, further suggesting that many of our judgments are less rational than traditionally thought.

Before digging into the studies themselves, it is prudent to define some of the concepts and terms Greene employs. The first is the standard philosophical dichotomy between consequentialism and deontology. Consequentualism, widely construed, is the view that what makes an act morally right or wrong is its consequences, often expressed in terms of pleasure, pain, and welfare. Deontology, on the other hand, does not concern itself with the consequences of potential actions: rather, deontology focuses on the duties we might have and whether or not our actions follow moral rules. Greene's central claim is thus that when we focus on consequences we make rational judgments, and when we focus on duties and rules we have emotional reactions.

Considering "Trolley Problems," thought-experiments designed to test the intuitions of particular subjects, frequently brings out these intuitions. The standard problem runs something like this: a trolley is speeding out of control down a track. There is a switch-track up ahead. If the train continues on its current path, it will strike and kill five people. Before this occurs, however, you can pull a lever and divert the train onto a sidetrack, where it will strike and kill one person. The subject is then asked whether or not it would be moral to throw the switch. Consequentialists would say that one was required to throw the switch because one death is preferable to five. Deontologists, on the other hand, might think it would be wrong to throw the switch because one would be actively killing someone. The true force of the example lies in its ability to draw out contradictory intuitions in subjects. While many people think it would be morally permissible to throw the switch in this example, many of these people find it impermissible with one variation: instead of throwing a switch, one must push a fat man off a bridge to block the train's path. While the result is still ultimately the same (one death rather than five), many people believe there is a morally significant difference. Greene believes this is because of the difference between *personal* and *impersonal* harms: when pushing the fat man off the bridge, one feels a moral personal connection, a more direct role in the death of the person one chooses to kill, and this seems to make a morally significant difference. Greene takes this to show that Deontological judgments are more concerned with personal harms, and Consequentialist judgments are associated with impersonal harms. He writes that our aversion to personal harm seems to fall along emotional lines: "Given that personal violence is evolutionarily ancient, predating our recently evolved human capacities for complex abstract reasoning, it should come as no surprise if we have innate responses to personal violence that are powerful but rather primitive. That is, we might expect humans to have negative emotional responses to certain basic forms of interpersonal violence, where these responses evolved as a means of regulating the behavior of creatures who are capable of intentionally harming one another, but whose survival depends on cooperation and individual restraint."[9]

---

[9] Greene, p. 43, citing Sober & Wilson, 1998, and Trivers, 1971.

With this picture in mind, Greene used fMRI scans of subjects' brains when making moral judgments about trolley problems to assess what kind of judgments the thought experiment elicits. Just as predicted, examples designed to force speakers to contemplate personal harms (i.e. deontological judgments) produced greater activity in regions of the brain associated with emotional intuitions, specifically "the posterior cingulate cortex, the medial prefrontal cortex, and the amygdala,"[10] as well as the superior temporal sulcus. [11] Examples that focused on impersonal consequences produced more activity in regions of the brain associated with rational calculation, most notably "the dorsolateral prefrontal cortex and inferior parietal lobe."[12] Greene also found evidence to support the theory that making rational judgments in cases designed to produce emotional judgments takes more time and effort. If one were to make a consequentialist judgment and ignore the personal disgust one might feel at pushing the fat man off the bridge, for example, it would take more time to produce that judgment. Sure enough, Greene's research confirmed the theory: "Trials in which the subject judged in favor of personal moral violations took significantly longer than trials in which the subject judged against them, but there was no comparable reaction time effect observed in response to impersonal moral violations."[13]

The upshot of Greene's work is a deeper understanding of the ways in which we make judgments when responding to different kinds of moral dilemmas. When dealing with personal, visceral harms (or, to use Haidt's terminology, instances of disgust and perhaps even sacrilege), we make intuitive, emotional judgments. On the other hand, when we make impersonal, consequentialist judgments, we are more rational. I will not take a definitive stance on which kinds of judgments are superior, or which we ought to make. However, I will note a connection between Greene's and Haidt's work. Haidt argues that the role of reasoning is to convince us of what we ought to do when our emotional reactions lead us astray. If Haidt is right, Greene's work suggests that consequentialist judgments should sometimes be favored over deontological ones. In one sense, this conclusion is trivial: trolley problems themselves are designed to show that we sometimes ought to prefer consequentialism to deontology, and vice versa. In this sense, nothing new has been shown. However, in another sense of the conclusion, Greene's and Haidt's work shows that there might be more substance behind consequentialist judgments because they are produced through reasoning, not rapid emotional responses.

III. Studies in (the Lack of) Free Will

Whether or not humans have free will is one of the oldest and most fascinating debates in philosophy, both for its role in understanding human nature and in determining how to construct a just society: if one could not have acted otherwise when one murdered one's neighbor, is it truly fair to punish that person? Determinism is the view that all of our actions and all of the events in life are necessarily causally determined by everything that came before them. This is true both of scientific, physical causes (chemical reactions, particle collisions, etc.) and of historical events (i.e.

---

[10] Greene, p. 44.
[11] See Allison, Puce, & McCarthy, 2000 and Saxe, Carey, & Kanwisher, 2004a.
[12] Greene, p. 44.
[13] (Greene et al., 2004; Greene et al., 2001).

Columbus landing in the Caribbean determined my eventual birth). If determinism is true, although it might seem that I have the ability to freely choose to do anything whatsoever, this is merely an illusion: every thought, feeling, and action in my life is a necessary outcome of all of the events that preceded it.

In a ground-breaking study published in the *Journal of Consciousness Studies*, Benjamin Libet compared when subjects became consciously aware of their intention to act with when their brain's sent the signal to initiate action. The central idea the study relies on is that if free will is true, subjects consciously decide to act. If that's the case, the argument goes, they would become aware of their desire and decision to act before they begin the action at hand. If determinism is true, this would not be the case: subjects would not "have a say" in the matter, and would thus begin to act before consciously deciding to act. To see whether or not this is true, Libet instructed his test subjects to make some voluntary movement (a wrist flick, a hand twitch, etc.) at any time they wished, and then to report when they felt the desire to act. During this process, Libet measured when the subjects actually began to act using an electromyogram (EMG). According to Libet, "Freely voluntary acts are preceded by a specific electrical change in the brain (the 'readiness potential', RP) that begins 550 ms before the act. Human subjects became aware of intention to act 350–400 ms **after** RP starts, but 200 ms. before the motor act. The volitional process is therefore *initiated* unconsciously."[14] This suggests that we are not consciously aware of the decision to act, and thus that free will is an illusion if we must be consciously aware of our actions before executing them.

The other implication of the study, according to Libet, is that we can decide *not* to act even once the intention has been formed. As mentioned above, the process of beginning to act (i.e. intention -> awareness of intention -> action) begins before we are aware of it: we already intend to act and our body becomes ready to act before we are consciously aware of these processes. However, Libet's test subjects still became aware of their readiness to act before the action actually began. Libet suggests the following role for traditional notions of free will: "Potentially available to the conscious function is the possibility of stopping or vetoing the final progress of the volitional process, so that no actual muscle action ensues. Conscious-will could thus affect the outcome of the volitional process even though the latter was initiated by unconscious cerebral processes. Conscious-will might block or veto the process, so that no act occurs."[15] On such a conception, although we might not become aware of our initial intention to act, we can still ultimately restrain ourselves from committing the act itself. For a colloquial example, consider the response one has when attempting to anticipate a red light turning to green when driving: one watches for the light to turn green so one can drive straight, but then the turn light turns green instead. One feels one's leg muscle twitch, but one is able to restrain oneself from easing off the gas at the last moment. Libet observed something similar in his test subjects:

> The subjects in our experiments at times reported that a conscious wish or urge to act appeared but that they suppressed or vetoed that. In the absence of the muscle's electrical signal when being activated, there was no trigger to initiate the computer's recording of any RP that may have preceded the veto; thus, there were no recorded RPs with a vetoed intention to act. We were, however, able to

---

[14] Benjamin Libet, "Do We Have Free Will?" in *Journal of Consciousness Studies*, 6, No. 8–9, 1999, pp. 47–57, p. 47.
[15] Libet, pp. 51-52.

show that subjects could veto an act planned for performance at a pre-arranged time. They were able to exert the veto within the interval of 100 to 200 msec. before the pre-set time to act (Libet et al., 1983b). A large RP preceded the veto, signifying that the subject was indeed preparing to act, even though the action was aborted by the subject.[16]

If these findings are correct, the proper way to view free will and human volition is as a restraining force, not as a positive drive. Now, it is important to maintain a healthy skepticism about Libet's results: there are serious methodological questions to be answered before accepting his conclusions completely. There are also questions about the degree to which the results actually disprove the existence of free will (it seems more like Libet has merely showed that free will does not manifest itself in the same way as we initially believed) I will not fully engage those questions here. However, I will present one alternative explanation for his findings that might put our traditional sense of free will back on the table. It seems plausible that it could be evolutionarily advantageous for an organism to be able to begin to act as quickly as possible. One needn't think of a complicated example to see the merits of this ability. Suppose you are walking through the jungle, and happen upon a bloodthirsty leopard. Upon seeing it, you might reasonably conclude that it would be wise to run away as quickly as possible. Assuming that the leopard is quite fast, and is likely to be much faster than you, any delay in the escape process would be costly. It therefore seems it might be advantageous for the body to ready itself to run away as quickly as possible, especially if waiting for conscious confirmation might cost you your life. This distinction calls to mind Daniel Kahneman's "two systems approach to judgment and choice,"[17] in which we all have a fast, intuitive cognitive system and a slow, more deliberative cognitive system. Each has its own role and serves a different evolutionarily advantageous function, but both might be thought of as different forms of cognition, and the fact that one occurs very rapidly does not disprove the causal connection between the agent and the judgment or action. The mere appearance of the intention to act before one is consciously aware of that intention does not show that one did not initiate that process voluntarily, just that one was not consciously aware of it, and that might actually be very helpful to us. If this is true, or even if this is merely plausible, we have reason to be skeptical of Libet's results: one can show that the intention to act is present before one consciously decides to act without proving that free will is a myth.

Before proceeding, it might be helpful to say a few words about some more popular attempts to challenge the existence of free will. Malcolm Gladwell's bestseller *Blink: The Power of Thinking Without Thinking*[18] discusses what Gladwell calls "snap decision-making": we sometimes make judgments without being able to explain how or why. Perhaps the most interesting example in the book is that of a few archaeologists who knew an allegedly ancient artifact was a fake. Just by looking at the statue, some archaeologists were able to guess that the statue was a forgery even though they were unable to explain why, and despite testimony from other experts who believed the statue to be authentic. Gladwell calls this lack of awareness "the locked door," which we might consider a synonym for the myriad unconscious processes the brain carries

---

[16] Libet, 52.
[17] See Daniel Kahneman. *Thinking, Fast and Slow* (MacMillan, 2011), p. 13.
[18] See Malcolm Gladwell. *Blink: The Power of Thinking Without Thinking.* (Back Bay Books, 2007).

out. Gladwell's central thesis is that these unconscious processes can be very helpful, but that they can also be quite dangerous: in one of his stronger examples, he discusses the case of Amadou Diallo, who was shot by police when reaching for the identification they asked him to produce. We sometimes act automatically based on bad information or because we are under stress, often with disastrous results. Gladwell ultimately suggests that we need to learn when to use snap judgments and, if possible, to be able to consciously override them.

Like Libet's study, Gladwell's work challenges the degree to which we freely act, and whether or not we are fully aware of the processes that cause our actions. Both writers suggest that emphasis should be placed on developing restraint in order to avoid actions we ought not to commit. We will return to this conclusion in section V.

IV. Getting People To Be Moral

I believe there are two major impacts of neuroscience and moral psychology on morality and the humanities. The first is that, assuming the results of recent neuroscientific advancements and experiments are correct, we now have a better understanding of how people behave and how to get people to behave in certain ways. This is particularly important for the is/ought distinction mentioned at the beginning of this paper. If we have a better understanding of the facts about human nature and how we behave, we can make more realistic attempts to influence human behavior for the good, and we may even come to a better understanding of the most moral action in any given situation.

Take, for example, Haidt's social intuitionist model. Suppose you are attempting to raise money for refugees from the Syrian civil war. Next suppose you, like many philosophers and psychologists, believed that the best way to convince people to do things was to make a rational appeal: you try to convince people that the refugees are really suffering and that helping people in need is the right thing to do[19], and you show them some charts and graphs that would show them how their contributions would help improve people's lives. Assuming that the social intuitionist model and other studies indicating the prominence of emotional intuitions are correct, this would be a suboptimal way to get people to donate. Instead, one would be better served by making emotional appeals: naming a victim, telling his or her story, and showing a picture of a person potential donors might help. In fact, this example has actually played out in the media in recent months: a photo of a Syrian boy covered in blood and dust sitting in an ambulance after a drone strike in Aleppo influenced public discourse so much that it was even mentioned by Hillary Clinton in one of the presidential debates.

It might also be the case that knowledge of how we are motivated to act can influence what we try to convince people to do. There are many vulnerable people that we ought to help, from wounded veterans and children with cancer to victims of environmental disasters and religious persecution. If Haidt's and Greene's work is correct, people are motivated to act on different moral foundations and by different

---

[19] We can even take this a step further and provide a financial argument: Slavoj Zizek has argued that helping refugees from the Middle East may be required as "the cost of business" for global capitalism. See Slavoj Zizek. *Refugees, Terror and Other Troubles with the Neighbors: Against the Double Blackmail* (Melville House, 2016).

stimuli. Liberals, for example, generally emphasize the Care and Fairness foundations, while Conservatives also emphasize Loyalty, Authority, and Sanctity. If we know that Conservatives will be more motivated to donate to causes that involve Loyalty or Sanctity, it might be wiser to focus our efforts on convincing them to donate to charities or efforts that focus on those foundations rather than those that rely on more liberal conceptions of Fairness as equality (i.e. we might be more successful in getting Conservatives to donate to help wounded veterans than to homeless people). Building on Greene's work, if we know that a person thinks more like a consequentialist than a deontologist, we ought to make more rational appeals than emotional appeals, and to encourage such people to donate to solve more impersonal, distributive social problems than problems at the individual, personal level. If the consequences matter in at least some situations, then helping a greater number of people will be desirable. It seems likely that understanding how people think about moral issues might enable us to better influence their behavior.

Turning to the more philosophical assessment of assessing right and wrong, if we know that a possible action will be more effective (i.e. will produce a better outcome with respect to the relevant moral criteria), we have a better picture of what we ought to do. So, for example, suppose I am speaking to a large crowd of people in order to persuade them to donate more to different charities. If I know that moral judgments are based more on emotional responses than on reasoning, I might appeal to the audience's emotional side, rather than giving them the cold, hard facts and reasons at hand. To use Haidt's terminology, I would address myself to elephants, not riders. Or consider another example, in which I am trying to convince a crowd of donors to contribute more money toward one of a few different charities. Suppose I also know that most people in the crowd lean to the left politically, and that they are likely to place more value on the harm foundation than the loyalty foundation, for example. Knowing this, it would be more effective for me to persuade them to contribute their money toward a charity that will have the best possible returns on investment, rather than toward a charity that will help fewer people, but will fulfill some other sort of moral duty. The point in these examples is that if one knows a certain tactic will work based on the psychological make-up of the audience, I have a moral obligation to use that tactic, or to solicit donations toward that charity, etc.

V. Prescriptive Humanities

The second impact of these studies is on education. How ought we to teach the humanities in light of these developments? The humanities present a rich tradition of thought about a wide range of issues. Much of the literature in the humanities falls into two categories: descriptive and prescriptive literature. Descriptive literature focuses on describing and understanding the nature of the world and how we behave. Prescriptive literature attempts to influence behavior by making suggestions about what we ought to do or how we ought to solve a problem.

The humanities have rich deposits of both descriptive and prescriptive literature. Often, classical texts in the humanities contain a mix of both accounts. Perhaps the most notable example in philosophy is Aristotle's *Nichomachean Ethics*, a book dedicated to describing moral (virtuous) character, and thus to influence how we ought to behave. I will not engage in deep exegesis here, but will provide a brief summary of the book for

the purposes of illustration. Aristotle says that we have theoretical, scientific, and practical knowledge, and that morality gives us practical knowledge about how we ought to behave in the world. He describes morally good actions as virtuous, and that virtuous actions are actions that a virtuous person would take in the situation at hand. Virtuous actions fall between two extremes. For example, if one is witnessing a bank robbery, one can either 1) run away screaming, 2) stay put but do nothing, or 3) fight the robber. If you are a trained Navy SEAL that has a machine gun and the robber only has a knife, you would be a coward if you ran away. If one stayed to fight, one would be courageous (here, virtuous). If the roles were reversed, however, the calculus would be different. If one charged at the machine-gun bearing robber, one would be quite reckless. The point is that the virtuous action is the action that is appropriate in a given situation after accounting for the facts about the people involved and the circumstances at hand. The implicit normative suggestion in the text is that we all ought to act in such a manner.

The above text gives both a descriptive and a normative account of different kinds of knowledge and of morality. Other thinkers like Immanuel Kant and David Hume also attempt to describe human nature, the way we make moral judgments, and how we ought to behave and what constitutes moral behavior. The former influenced the later descriptive accounts of psychologists like Lawrence Kohlberg, whose work was the bedrock of the strong rationalist tradition of moral judgment, while the latter influenced thinkers like Jonathan Haidt, who have argued for emotion-driven accounts. If the studies discussed above are correct, we have reason to think some descriptive accounts of human nature and cognition are wrong, and that certain accounts (moral judgments as emotions, the illusion of free will, etc.) are correct. If this is true, and we can show that one descriptive account (or even a small number of descriptive accounts) is correct, then this descriptive account should be taught in humanities classes as a true account of human behavior. We might still teach traditional and alternative accounts, but these would be taught with a greater emphasis on the history of the humanities, rather than as alternative, plausible theories. This is analogous to the study of medicine, in which students are taught that that previous generations believed medicine to involve balancing the four humors: we teach the theory in the context of history, not as scientific fact.

Recent developments in neuroscience and psychology should also inform our approach to normativity in the humanities. If we know that people behave a certain way (i.e. we know the way the world *is*), then the questions that remain to be settled are how we *ought* to be. It is this area, above all others, in which many disciplines in the humanities excel. Through art, literature, history, and philosophy, students gain strong insights into what they ought to care about and what they ought to do. Accordingly, the humanities should be taught in a manner that best creates an environment for discussion of normative ideas. As with the medicinal example above, more time and energy should be focused on the normative aspects of the humanities, rather than the descriptive aspects. We should read *The Merchant of Venice* not merely as a means of understanding Shylock's society and experiencing more of Shakespeare's work, but to sample the arguments against anti-Semitism. We should read *Brave New World* not merely as enjoyable dystopian fiction, but as a chilling commentary on the dark side of commercialized, uniformed scientific progress without control. We should read the history of the world not just to learn what has already happened, but to learn what may happen in the future of similar patterns are maintained. Neuroscience cannot tell us the

answers to questions about what we ought to do. The humanities are uniquely equipped to complete the epistemological circle.

Conclusions

I have argued that recent developments in neuroscience and psychology have given us a better picture of human nature, giving us a more complete understanding of the way the world is and of what we are ultimately capable of. If the experiments of Jonathan Haidt, Joshua Greene, Antonio Damasio, Benjamin Libet, and other contemporary thinkers are correct, humans are largely emotion-driven individuals that often lack full control over their actions. If this is true, the proper role for reason and control is in training ourselves to avoid or ignore many of our initial impulses, and instead to take a more reasoned, deliberate approach to our lives. I have argued that if these findings are correct, we should focus teaching the humanities as a means of solving these normative questions, rather than as providing accurate descriptive accounts of human nature.

Bibliography

Aristotle. *Nichomachean Ethics*, 2ⁿᵈ edition, trans. Terence Irwin (Hackett Publishing Company, 1999).

Antonio Damasio. *Descartes' Error: Emotion, Reason, and the Human Brain* (New York: Putnam, 1994).

Malcolm Gladwell. *Blink: The Power of Thinking Without Thinking.* (Back Bay Books, 2007).

Joshua Greene. *Moral Tribes: Emotion, Reason, and the Gap Between Us and Them*. (Penguin Books, 2013).

Joshua Grene, "The Secret Joke of Kant's Soul" in *Moral Psychology*, vol. 3: *The Neuroscience of Morality*, ed. W. Sinnott-Armstrong, 35-79 (Cabridge, MA: MIT Press, 2008).

Jonathan Haidt, "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment"(2001) in *Psychological Review*, 108: 814-34.

Jonathan Haidt. *The Righteous Mind: Why Good People Are Divided By Politics And Religion* (New York: Vintage Books, 2012).

David Hume. *A Treatise of Human Nature*. (London: Penguin, 1969/1739-40).

Aldous Huxley. *Brave New World*, Reprint Edition (Harper Perennial, 2006).

Guy Kahane, "Is, Ought, and the Brain" in *Moral Brains: The Neuroscience of Morality*, ed. S. Matthew Liao (Oxford University Press, 2016).

Daniel Kahneman. *Thinking, Fast and Slow* (MacMillan, 2011).

Immanuel Kant. *Grounding for the Metaphysics of Morals*, 3ʳᵈ ed. Trans. J. W. Ellington. (Indianapolis: Hackett, 1993/1785).

Lawrence Kohlberg, "Stage and Sequence: The Cognitive-Developmental Approach to Socialization" in *Handbook of Socialization Theory and Research*, ed. D.A. Goslin, 347-480 (Chicago: Rand McNally, 1969).

Benjamin Libet, "Do We Have Free Will?" in *Journal of Consciousness Studies*, 6, No. 8–9, 1999, pp. 47–57.

William Shakespeare. *The Merchant of Venice*, Folger Shakespeare Library (Simon & Schuster, 2009).

Slavoj Zizek. *Refugees, Terror and Other Troubles with the Neighbors: Against the Double Blackmail* (Melville House, 2016).