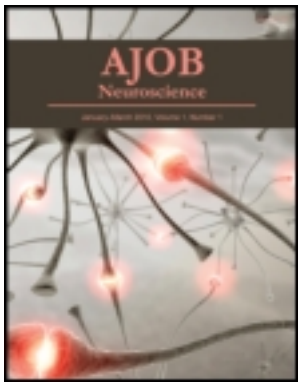

AJOB Neuroscience
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/uabn20

Neuroethics and the Possible Types of Moral Enhancement

John R. Shook [a]
[a] Science and the Public EdM program, University at Buffalo

Version of record first published: 08 Oct 2012.

To cite this article: John R. Shook (2012): Neuroethics and the Possible Types of Moral Enhancement, AJOB Neuroscience, 3:4, 3-14

To link to this article: http://dx.doi.org/10.1080/21507740.2012.712602

AJOB Neuroscience, 3(4): 3–14, 2012

ISSN: 2150-7740 print / 2150-7759 online
DOI: 10.1080/21507740.2012.712602

Target Article

Neuroethics and the Possible Types of Moral Enhancement

John R. Shook, Science and the Public EdM program, University at Buffalo

Techniques for achieving moral enhancement will modify brain processes to produce what is alleged to be more moral conduct. Neurophilosophy and neuroethics must ponder what "moral enhancement" could possibly be, if possible at all. Objections to the very possibility of moral enhancement, raised from various philosophical and neuroscientific standpoints, fail to justify skepticism, but they do place serious constraints on the kinds of efficacious moral enhancers. While there won't be a "morality pill," and hopes for global moral enlightenment will remain hopes, there will be a large variety of behavioral modifiers described as enhancers of some aspect of morality or another according to prevailing social norms. The most likely moral enhancers that will be designed, tested, and marketed in the near future will attempt to alleviate seriously immoral and illegal behavior. Some enhancers diminishing antisocial conduct and a few designed to elevate moral conduct above normal levels will also become available, although their widespread use is doubtful. There will also be specialized modifications for enhancing whatever an individual personally regards as moral, and for enabling better performance by a person in an operational role. A measure of skepticism toward all the proposed kinds of moral enhancement is advised, and where political implementation of moral enhancers is concerned, a healthy amount of cynicism as well.

Keywords: neuroethics, moral enhancement, morality/ethics, moral naturalism, relativism

WHAT IS "MORAL" IN MORAL ENHANCEMENT?

If a person's cognitive capacities can be enhanced, could a person's prosocial and even moral capacities be enhanced too? The interdisciplinary field of neuroethics (Farah 2012) has taken up such questions. Too many discussions are proceeding as if both the meaning and possibility of moral enhancement were already widely understood and agreed upon.

It is tempting to jump directly to more fascinating questions. Which techniques for achieving moral enhancement—surgical, transcranial, pharmacological, genetic, nanotechnic, cybernetic, and so on—would be most effective? Which moral enhancements would be beneficial, or at least not harmful, for society? Which moral enhancements are ethically justifiable to permit, and which should be forbidden? And which moral enhancements should become publically endorsed or even legally required? The emergence of moral enhancement as a philosophical matter (see Englehardt 1990; Hughes 1996; Macer 1995) has coincided with raising these sorts of practical questions (such as issues raised by Boire 2004; Chan and Harris 2011; Faust 2008; Harris 2011; Persson and Savulescu 2008; 2010; 2011; Spence 2008; Verbeek 2009; Walker 2009). Asking such questions, and offering any answers, depend on assigning some sense or another to "moral enhancement." However, clear and precise definitions of "moral enhancement" are not to be found; what has been called "moral" enhancement ranges from feeling empathic concern to increasing personal responsibility all the way to heightening respect for global fairness. No one doubts the good intentions behind calls for moral enhancement, but are good intentions enough?

Neuroethics should carefully consider how moral enhancement could even be possible, and, if possible, whether only modest kinds of moral enhancement could be practical. We survey potential ingredients for moral enhancers, considering philosophical and neuroscientific objections along the way. This survey lends plausibility to a balanced and moderate conclusion about moral enhancement's overall viability: It will be neither so easy, nor so impossible, as extreme stances suggest.

There won't be a "morality pill," just as there won't be a single "intelligence pill." In the future, there will be many specific types of morality "pills" and other enhancement techniques. Moral psychology/neuroscience will be reliant on cognitive, affective, and social neuroscience, but neuroethics will have to take special care with moral enhancement. While physiologically possible, as argued later in this article, designing efficacious moral enhancements must be handled quite differently from forging techniques for cognitive enhancement. Unlike the way that basic cognitive enhancements are classifiable independently from society (enhanced attention or memory remains enhanced regardless of residence), classifying moral enhancement must take into account environing social contexts. Moral intuitions, virtues, and rules are not identical around the world; changing the social context can switch a classification of a moral

I am grateful to the Penn Conference on Transforming Humanity in November 2010, where I read a portion of this article and received helpful suggestions from participants. I also thank James Giordano of the Potomac Institute for Policy Studies for his advice and support, this journal's associate editor John Banja for fine suggestions along the way, and anonymous referees for their helpful criticisms.
Address correspondence to John R. Shook, Center for Neurotechnology Studies, Potomac Institute for Policy Studies, 901 N. Stuart St. Suite 900, Arlington, VA 22203, USA. E-mail: jshook@pragmatism.org

enhancement into a moral deficit. This point is not a default admission that relativism must reign, but it is an acquiescence to a pluralistic reality. In an ideal philosophical world, consensus on some single unified ethical theory would dictate the precise descriptive meaning to "morality" and the exacting prescriptive standards for "moral" conduct. We don't yet live in that world.

IS ANY MORAL ENHANCEMENT POSSIBLE?

The way that cultures and ethical theories disagree about morality, the way that science is uncongenial to some conceptions of moral responsibility, or the manner in which naturalism can dictate the parameters of morality might rule out the possibility of any moral enhancement from the start.

Cultures notoriously disagree on many moral matters. For starters, try to conceive of an "etiquette pill" and wonder how it could work the same in Bombay, Baghdad, and Boston. A single morality pill could hardly be less improbable. Even if we set aside the problematic diversity inherent to cultural morality to favor some reasoned ethical system instead, we must choose a system. Suppose a brain modification transforms a person into someone who now takes the moral deed to be the one maximizing the welfare of all. Notorious for their inability to agree how to implement utilitarianism, most utilitarians would soon find fault with this new utilitarian's concrete moral judgments, just as they find fault with each other's. And many deontologists would simply deny that this fresh utilitarian has received moral enhancement at all—perhaps only moral degeneracy or depravity has befallen this person.

The inability of cultures or ethical theories to converge on a normative theory of morality (at least so far) doesn't show that no moral objectivity is possible (see Tersman 2006), as there could be overlapping consensus about some moral matters. Even ethical skeptics lacking hope for an ethical consensus can still take "morality" and "moral" to be meaningful to people, albeit in diverse ways, so it cannot be hastily concluded that "moral enhancement" is rendered meaningless and impossible. Nevertheless, anyone using the term "moral enhancement" as if everyone knows what is meant must either be simplifying matters to the point of negligence, or trying to speak only to those already in local moral consensus. "Moral enhancement" is not rendered impossible by the fact of moral disagreement, but this term does not automatically shine by its own meaningful light.

A different skepticism toward moral enhancement can proceed from an assumption that morality requires a nonnatural feature to moral agency, such as a supernatural, metaphysical, or contracausal free will. Neuroethics cannot endorse notions about moral enhancement through unnatural interventions, yet siding with naturalism isn't enough. Naturalists can turn skeptical against moral agency by requiring the involvement of certain indeterministic, instantaneous, or overly powerful mental processes and then judging that the sciences are not confirming those processes (see a survey of stances in Kane 2005). Because neuroethics, along with neurophilosophy, holds the sciences in high regard, any potential moral enhancement—and morality itself—would have to be based somewhere in that natural world. However, that naturalistic setting can arouse deep skepticism toward morality, and hence toward the practicality of moral enhancement.

For fairly considering the possibility of efficacious moral enhancement, neurophilosophy and neuroethics are best served by an alliance with the project of minimal moral naturalism (see Churchland 2011; Flanagan et al. 2008). Minimal moral naturalism—not to be confused with eliminative/reductive naturalism, looking to nature for moral values or authorities, endorsing moral realism, or with ethical naturalism's project of answering questions of philosophical ethics through science—simply proposes to apply relevant sciences for studying how humans manage to produce their moral valuations and perform whatever they take to be morality without any nonexistent features involved. Human morality, on this view, is naturally embodied in the ways that people conduct themselves in accord with their understood norms applicable to all social interactions (not just those with kin or friends) and any deeds of wide group concern. Humanity's reliance on morality, with its basic moral norms regarding cleanliness, caring, fairness, cooperation, conflict resolution, and the like, is universal in the sense that it is found wherever human societies are found. Morality as a human practice (setting aside moralities of other species) is now regarded as an empirically well-established cultural universal, discernable in every kind of society including the few remaining hunter-gatherer societies (Krebs 2011). There is little doubt that morality has an evolutionary origin and anthropological history; minimal moral naturalism is content to let "morality" refer only to whatever survives the scrutiny of the biological and social sciences. Although morality is hardly a "natural kind" awaiting routine scientific detection, it possesses enough stable functions and features across human societies to be amenable to interdisciplinary research without waiting upon philosophers to decide what morality really is.

Serious consideration of practical moral enhancement cannot avoid utilizing the behavioral and brain sciences for understanding how humans undertake their practices of morality, and for discovering how modifications to moral functioning would work. This minimal moral naturalism permits neuroethics to be appropriately "neuro"-based, while staying uncommitted to handing over to science the tasks of determining what is truly moral, or answering questions of philosophical ethics (see Berker 2009). Again, minimal moral naturalism can't guarantee that moral enhancement is practically achievable, but it serves as a framework for organizing information about the natural bases for our moral capacities, such as they really are. A complete neuroethical theory of morality won't be easy to assemble. Selecting out brain processes and slotting them into correspondences with specific phases of moral thinking will be a hazardous procedure at best (Gazzaniga 2005; Greene

2008; Kamm 2009). Similarly, seeking correlations between phenomenal features of moral attentiveness and discernable brain processes will run into many obstacles. Along the way, neuroethics must recognize how the existence of human capacities is a separable matter from what people happen to think their capacities are and how they function, and people are proving to be quite mistaken about the mechanisms involved with our capacities for ordinary morality. Ensuring that only natural capacities are involved with anything that deserves to be called moral conduct places practical constraints on whatever efficacious moral enhancement could be.

We have briefly looked at a few concerns that natural moral enhancement might be impossible, either conceptually or practically. Any neuroethical view taking moral enhancement to be demonstrably possible must identify the source (such as a culture or an ethical theory) for the meaning assigned to the word "moral" in "moral enhancement," and also identify the practical mechanisms aligned with modifying moral conduct in some manner detectible by the behavioral and brain sciences. We next look closer at potential means for manipulating morality.

THE MEANS TO MORAL ENHANCEMENT

Upon first hearing, "moral enhancement" is ordinarily taken to mean something like "making a person more moral," which in turn practically means "more likely to do the morally right thing." Regulatory evaluations of an experimental moral enhancer will be asking whether it is safe and effective, and if it is effective, what specifically does it do to a person's behavior?

We would be rightly suspicious toward an alleged moral enhancer whose dispositional and behavioral effects on conduct remain undetectable even by careful observation. Moral enhancers must be more than internal emotion or mood enhancers. Feeling more moral, righteous, or virtuous may be nice, yet others would understandably take closer interest in whether you are acting more morally. For example, only by presuming that increasing moral motivation guarantees some enhancement to moral conduct (see Douglas 2008) can enhancement of motives be taken as a reliable way to enhance morality. While watching for behavioral changes, overemphasizing conscious states is unnecessary. A moral enhancer may detectably change a person's conduct without altering her or his mood or self-awareness; that change in conduct might even be accompanied by emotional swings in surprising directions. There is no way to rule out in advance the possibility that a particular kind of moral enhancer could have the unusual ability to re-direct a person's conduct in a more moral direction even as this person reports having no such sentiments or intentions.

Just as one's cognitive capacity has diverse aspects ranging from alertness and focus to memory, inference, and imagination, there are many specific components to one's capacity to be moral. For example, we take great interest in whether people understand what is morally required, and where such understanding is lacking, we might seek to supplement a person's moral beliefs. For people already of sound mind and good will, adjusting a person's moral views could also suffice for dramatic moral enhancement. Alternatively, we could be dealing with a subject who has the right moral beliefs, but needs enhancements to the disposition and capacity to fulfill moral expectations in real-world situations. There are several potential disconnects between people's moral convictions and their actual conduct. For example, a person might understand what is morally right, but lack much interest in being genuinely concerned for others or for treating them in the right way. Or, a person can be properly motivated by moral concerns but something else, such as a distracting temptation or a selfish interest, ends up dictating that person's priorities. As another example, a person can successfully prioritize doing the right thing and then still have that firm intention overridden by other motives or compulsions, conscious or unconscious, when the time for action arrives.

So many factors may be involved with moral capacity and conduct (including numerous elements of social cognition; see Frith and Frith 2012) that endless theoretical debate ensues over which ones are genuinely moral factors. We can consult folk moral psychologies, cultural expectations, or some ethical theory's propositions, but in the absence of a consensus, which seems to be our present situation, does our inquiry come to a halt once again? Another sort of skepticism arises against moral enhancement here. The absence of a consensus upon the mechanisms of morality could prevent any agreement that a proposed moral enhancer could really be enhancing morality, whatever else is may be doing. This skepticism is not the fault of the behavioral and brains sciences, but our own, for failing to agree about which cognitive processes are genuinely relevant to what we want to call morality and moral agency.

Some ethical analogue to minimal moral naturalism is needed to make further progress deliberating about moral enhancement. Those who are already proposing and debating potential moral enhancements are helpfully showing us the way. Most authors sensibly appeal to widely shared and commonly accepted views about morality. Perhaps this route could be called "minimal moral commonsensism" for its appeal to primary moral factors simple enough for broad comprehension and everyday application.

MECHANISMS FOR MORAL ENHANCEMENT

Minimal moral commonsensism permits an exploratory inquiry into means for manipulating moral conduct. To a simplistic approximation, a moral enhancer would, in theory, be affecting one or more of the five factors listed next. This list isn't trying to get matters in the right order (and acknowledges how they are interlinked), but minimal moral commonsensism does expect that adjusting at least one of them to a sufficient degree (all other things being equal) should result in detectible modifications to conduct.

Enhance a person's sensitivity to moral features of situations—resulting in heightened moral appreciation.

Enhance a person's thoughtfulness about doing the moral thing—resulting in stronger moral decisions.
Enhance a person's moral judgments that get the right moral answer—resulting in more correct moral judgments.
Enhance a person's motivated choice to do what moral judgment indicates—resulting in improved moral intentions.
Enhance a person's volitional power to act on a moral intention—resulting in more moral will power.

No justification for only these factors is necessary here; the point is only that some efficacious factors must be scientifically identified while researching and designing any alleged moral enhancer. Philosophical ethicists would refine and subdivide these five factors, merge some together, or even propose a different list. Indeed, the stance of minimal moral naturalism anticipates radical revisions to theories about cognitive processes underlying moral capacity and conduct. It already appears that there is no unified cognitive system responsible for the formation and enaction of moral judgments, because separable factors are more heavily utilized for some kinds of moral judgments rather than others, and this lack of unity may be apparent in moral phenomenology as well (see Parkinson et al. 2011; Sinnott-Armstrong 2008). The roles of emotions in moral appreciation and judgment, alongside (and intertwining with) social cognition and deliberate reasoning, are so complex that research is only starting to trace how they influence kinds of intuitive judgment and moral conduct (Decety et al. 2012; Huebner et al. 2008).

As this theoretical work proceeds, bolder ventures are rushing onward to propose, design, and implement moral enhancements more or less in line with minimal moral commonsensism. Even skeptics against moral enhancement typically appeal to these kinds of factors, key points where the possibility of efficacious moral enhancement can be disputed. These rejections do not protest that interesting brain adjustments couldn't be accomplished by current or future technologies, but they do say that such adjustments could never amount to "moral" enhancement. To see how both supporters and critics of moral enhancement rely on these five primary moral factors, let us proceed through them, labeling the matching criticisms as C1–C5. After each criticism, examples of replies available to supporters of moral enhancement are mentioned, but there is no space here to adjudicate disputes.

C1. Is heightened moral appreciation even possible? An emerging objection on this point protests that there wouldn't be anything in human psychology or brain processing to enhance here, since the nervous system hasn't evolved and doesn't function to judge situations morally in the first place. The discovery of the "real" emotions and motives for conduct, operating behind a façade of warm emotional arousals, fast moral intuitions, and moral rationalizations, exposes all manner of selfish, instrumental, or instinctive causes for our behavior. Proponents of a "moral error" theory, for example, can say that people are always wrong to suppose they are intuiting, appreciating, and judging situations morally, since nonmoral interests (conscious or subconscious) are doing all the work of guiding motive and conduct (see discussions by Joyce 2002; Lillehammer 2003). Proponents of moral enhancement, however, could reply that the behavioral and brain sciences may yet conclude that moral appreciation is real and efficacious, despite having selfish rivals. Virtue ethicists, for example, have long contended that habitually attending to the moral features of social situations can be instilled and cultivated to consistent moral effect. Sentimentalists emphasizing the large role for prosocial and antisocial emotions and feelings, at both subconscious and conscious levels, may expect that elevating prosocial sentiments such as empathy and caring, and reducing reactive fears, social anxieties, and the like, could heighten moral sensitivity and appreciation. Behavioral neuroscience is detecting social emotions and empathic motivations deep within human psychology (see Franks 2010; Schulkin 2011), and they are pliable (on oxytocin-mediated empathy see Zak 2011). Neuroscientific and neuroethical inquiry can try to discern whether our capacity (such as it is) to sensitively appreciate other people and daily situations in a more moral manner might become an opportunity for designing moral enhancers.

C2. Are more thoughtful moral decisions even possible? Even if moral sensitivities can be intensified and moral appreciation heightened, many other motives and temptations crowd our minds, and expecting people to undertake more careful moral thoughtfulness could be a false hope. Not only may moral appreciation be a post hoc rationalization, but any immediate intuition or thoughtful decision favoring a moral deed could be entirely a matter of post hoc rationalization as well. Unconscious brain processes may culminate in the initiation of conduct well before even an intuition, much less a thoughtful decision, is formed in self-conscious awareness. For example, if one takes "the moment of choice" as the locus of autonomous decisions, that phenomenon's neural basis may actually lack the expected motor efficacy (as Libet's experiments might show; see Libet et al. 1999). Hasty conclusions about the unreality of moral agency and responsibility can ensue, but great care must be taken (see rejections of those conclusions by Dennett 2003; Glannon 2011a; Lavazza and De Caro 2010; Roskies 2006). Furthermore, even if moral thoughtfulness had some chance of influencing conduct, it may only be a small chance, since so many other habits of poor reasoning, along with distracting motives and habits, conscious and unconscious, have plenty of opportunity to override moral thoughtfulness. All the same, proponents of moral enhancement can try to reply that even modest chances for making a decision in a moral direction could be increased, alongside ways to diminish cognitive errors and weaken distracting motives, but this claim requires empirical confirmation. Neuroethics can explore whether there really are potential mechanisms for strengthening thoughtful moral

decision-making to levels sufficient for impacting moral conduct, permitting moral enhancers to work here.

C3. Are correct moral judgments even possible? The bluntest objection here holds that there are no "correct" moral judgments in the first place if one true or best morality does not exist. On this criticism, "moral enhancement" is a misnomer, since at most such enhancements would only modify what some people already subjectively value and intuit as good. Other denials of morality as such, perhaps coming from certain kinds of noncognitivists, emotivists, expressivists, and nihilists (surveys supplied by Schroeder 2009), protest that morality is not an appropriate candidate for any truth or correctness conditions. Proponents of moral enhancement can try to reply that it only aims at enhancing whatever people, individually or in groups, happen to already regard as morally good conduct, and not at striving toward any "correct" morality. Neuroethics can explore what forms of moral enhancement can be possible in the absence of a universally accepted morality (as this article sets out in the next section).

C4. Are improved moral intentions sufficient to influence conduct even possible? Even if people could be brought to make more thoughtful moral judgments, intentions to follow through on such judgments may be nonexistent or powerless. According to this objection, moral intentions may not necessarily follow from moral judgments, or, even if they frequently do follow, they have no more genuine power over behavior than any other conscious intentions because they are somehow epiphenomenenal. Similar to the way that moral thoughtfulness may lack efficacy, any conscious intentions may be post hoc constructions by a cognitive center assembling the results of habitual unconscious processes actually controlling conduct. Or, perhaps conscious intentions are directly generated by the brain processes controlling conduct, but they are entirely epiphenomenal because they have no efficacy in themselves. Regardless of whether moral intentions are the product of thoughtful decisions, the manifestation of properly moral sentiments, or the coincidental by-product of something else, establishing an efficacious role for moral intentions cannot be taken for granted. Proponents of moral enhancement interested in manipulating the efficacy of moral intentions will have to construct careful replies to these concerns over epiphenomenalism. Neuroethics by itself probably won't resolve epiphenomenalist issues in philosophical psychology, but it could explore whether any forms of moral enhancement could be based on some enhanceable efficacy to our sincere intentions to act morally.

C5. Is greater moral will power even possible? The considerable objection here is the claim that this alleged volitional power to execute preferred sincere intentions over other rival motivating intentions at the moment of action—typically called the "will"—is either nonexistent, irrelevant, or just epiphenomenal. Having a strong intention to act morally doesn't guarantee that this one intention, among many competing intentions both conscious and subconscious, successfully controls actual conduct. Willful volitional control over which one of many simultaneous intentions gets executed in conduct may be a post hoc illusion generated by a cognitive center assembling reports of consequences of motor activity after conduct has already been initiated. Alternatively, feelings of willfull volition may be directly generated by brain processes controlling conduct, but they are entirely epiphenomenal in themselves because they have no motor efficacy (compare the perspectives found in, e.g., Mele 2009; Sinnott-Armstrong and Nadel 2011; Walter 2001; Wegner 2002). Proponents of moral enhancements to "the voluntary will" or "will power" and the like could respond by arguing that willful volitions are quite real (while not as robust as commonly supposed), and they have some enhanceable efficacy for real-time control over such things as refraining from acting on powerful impulses, managing strong compulsions, or thoughtfully undertaking deliberate conduct (specific proposals are enumerated by Glannon 2011b). Neuroethics, with the assistance of neuroscience, could help figure out whether anything roughly corresponds to "will power" in our cognitive processes, and, if so, whether those processes could be modified for producing moral enhancement.

How Can Moral Enhancement Survive Skepticism?

Replying to objections raised in C1–C5 won't be easy and adequate defenses of moral enhancement should be humbly restrained. More folk psychology and metaethical landmines than just contracausal free will still litter the field, and mistaken notions about roles for such things as emotions, deliberations, intentions, and volitions must be detected and replaced. Any list of moral mechanisms may end up as notorious examples of exploded folk psychology or cultural ideology, as some critics of moral enhancement are already complaining, just as the theory of the four humors of primitive physiology were replaced by scientific medicine.

Humble restraint is also needed from moral enhancement skeptics. Those unhappy with the five mechanisms selected for attention here must also stand ready to abandon their own preferred moral factors if necessary in light of future science. Misplaced skepticism can unfairly divert the proper consideration of moral enhancement. Inflated expectations about roles and powers of conscious intentions and volitional decisions, for example, offer tempting targets for skeptics, but deflating them is not the same thing as eliminating natural morality. Nor is disproving some philosophically favored notions about moral agency and responsibility the same thing as disproving the possibility of enhancement to basic natural morality. For example, sophisticated philosophical standards for moral deliberation about complex issues are inappropriate criteria for ordinary expectations of moral autonomy (see Bublitz and Merkel 2009). We should likewise be wary of demands for artificially increasing

autonomy to unreasonable levels of self-determination (see Nagel 2010).

Visions of moral enhancement making people into ethical experts capable of preventing such things as world disasters are similarly unrealistic (as Harris [2010] judges). We are contemplating moral enhancers that result in more moral conduct by current everyday standards, not necessarily "ethical enhancers" that meet the heightened expectations of normative moral theorizing. Understanding cognitive processes behind moral conduct could not by itself help us judge what better moral judgment must be (Kalis 2010). The creative work done through deliberate ethics about how to adjudicate among moral duties, prioritize moral virtues, and search for the good life will not be undertaken by mere moral enhancement (Fröding and Esmerelda 2010; Jotterand 2011). We will hear much skepticism toward moral enhancement coming from people trying to imagine how a mechanistic adjustment to core brain processes could abruptly permit someone to correctly resolve a puzzlingly tough moral dilemma, analyze an ambiguous ethical scenario, or discover the good life. Moral enhancement won't be trying to engage such matters. However, some will speculate how fruitful combinations of cognitive and moral enhancers might stimulate ethical enhancement. Far more important, in my view, is a worry that easy use of "perfected" moral enhancers could diminish opportunity, capacity, and responsibility for serious ethical thinking. In some future neuroethical dystopia, individuals thinking too hard about moral ambiguities and dilemmas are told that they simply need their enhancers adjusted.

Deliberations over moral enhancement, I am recommending, only requires minimal moral naturalism and a minimal moral commonsensism. In the short term, researchers designing and offering effective moral enhancers for regulatory approval (such as FDA approval in the United States) will apply these two standpoints. In the long term, only a far-advanced neuroscience could confidently discriminate precise modifications to one factor at a time and discover its effects. For now, most modifications will manipulate multiple modular mechanisms in concert, since most neural processes are quite interconnected and functionally integrated at several levels. However, research into moral enhancers needn't worry about trying to manipulate all the mechanisms simultaneously, either. Efficacious moral enhancers only need to alter conduct by manipulating one causal mechanism at minimum (while not preventing moral conduct in some accidental way), so enhancers won't have to utilize several at once—an unlikely thing for just a single pill or surgery to do. Therefore, the odds of hitting upon a variety of genuine moral enhancers may be quite good.

Skepticism toward moral enhancement can be further reduced by considering how several prominent modes of moral judgment and conduct are continuing to emerge from psychological research. Although it remains unclear how underlying cognitive mechanisms are directly or indirectly responsible for these moral modes, such as harm/care, fairness/reciprocity, ingroup/loyalty, authority/respect, and purity/sanctity (the value/virtue pairings of Graham et al. 2011), many people in different cultures rely on them (to varying degrees) as they make moral judgments. The eventual number of genuinely significant moral modes remains to be established, and variations in morality across cultures are likely due in large part to the specific modes emphasized by each culture (as we need to be reminded; see Henrich et al. 2010). Moral enhancers will not target these modes directly, but their effects on underlying cognitive mechanisms will result in detectable alterations to ways people appeal to certain values or virtues, and presumably change how they are prepared to act upon them as well. It is far too soon to estimate the range of modifiable modes, and some modes will prove easier to manipulate in desired moral directions than others. Modifying a moral mode does not in itself constitute moral enhancement—whether a modification is "moral" still requires examining its behavioral consequences within the context of some social setting or another, as the next section explores.

How Will Moral Enhancers Be Designed?

What might the means of moral enhancement look like while going through the phases of laboratory research? As correlations are detected between kinds of cognitive alterations and adjustments to conduct, speculative "moral engineering" will ensue. Moral enhancers will be crudely designed to attempt to produce more moral conduct utilizing primary moral mechanisms and following moral modes, as their designers won't be ready to explain in precise neurological detail exactly how they work. Those designers will have to explain their modifications in terms understandable to funding sources, research ethics boards, and eventually the media and the public, which means that much commonsense folk psychology about morality won't be abandoned anytime soon. Long after advanced neuroscience discovers the many inadequacies of moral commonsensism and cultural ideology, researchers, pharmaceutical companies, and governments will have to continue to explain moral enhancers in understandable terms.

What might mechanisms and modes of moral enhancement look like as they emerge from product development pipelines? If I were a marketing executive, easily memorable labels for reception by the mass market would catch my interest. Label a moral appreciation enhancer as Sensitiva. Label a moral thoughtfulness enhancer as Prudentia. Label a moral belief enhancer as Ethicale. Label a moral intention enhancer as Benevolium. Label a moral willpower enhancer as Prokrasia. Designing specific advertising campaigns targeted for different countries, especially with Sensitiva and Ethicale, would take up much more of my time. The question of which consumers will be using these enhancers remains paramount long after their practical efficacy has been confirmed.

Selling a moral enhancer to the public has much more to do with identifying and targeting their social values and concerns than with refining a moral enhancer's neurological efficacies. From a marketing standpoint, tell me your social problem, we'll match it up with a moral mode, and

something fresh out of pharmaceutical or neurosurgical lab might help. With any luck, the same moral enhancer can be efficiently sold in one country to remedy its prevailing social concern, and simultaneously sold in another country under a different label for remedying some other concern. My pharmaceutical company's drug for intensifying loyalty could be sold to a secular country for boosting company loyalty among workers, and next to a neighboring theocracy for boosting religious devotion among worshippers. Imagine how so many different remedies will be produced from just one effective enhancer for purity, or for trust. Fortunately for those busily marketing moral enhancers, the markets will judge enhancers by what they can do to people's conduct, not by how researchers think they work in people's brains.

MODELS FOR MORAL ENHANCEMENT

Ignorance about the precise mechanisms involved in moral agency and conduct will not halt experimental moral enhancement. Moral conduct will be clumsily manipulated through something like the five mechanisms to manipulate various moral modes. We must next consider how those manipulations will be fine-tuned for producing various types of increased moral conduct. Different models for moral enhancement emerge, according to the sources of standards for what conduct deserves moral approval.

Enhancing Preexisting Moral Commitments

The simplest model for moral enhancement involves just one person, the research subject. A moral enhancer might work simply by utilizing one or more mechanisms to increase the frequency and reliability of conduct that the subject believes is moral, by strengthening whatever moral convictions, motives, or habits that person already has. For example, a vegetarian could ensure that temptations to eat meat would not overwhelm his or her moral intent to be kind to animals. A corporate executive could prevent any temptation to be compassionate to an employee from overriding her moral intent to practice strictly equal treatment among all employees. Because this simple model of moral enhancement sets aside any evaluations of the subject's own moral views, and no modifications to moral convictions or intentions are involved, we can label this model as a "generic" moral enhancer. A generic moral enhancer may cause unusual side effects, but it at least modifies people's conduct to more closely conform to moral values, standards, and intentions to which they are already committed.

Generic moral enhancers could be used by well-intentioned people for enhancing compassionate treatment of the sick, elevating devotion to a social cause, or prioritizing care for family over acquaintances or strangers. They could also be used for intensifying prejudices against despised ethnic minorities, emboldening puritanical attitudes toward those deemed to be living "impure" lifestyles, or strengthening resolve to pursue nothing but personal wealth. Whatever a person already regards as a worthy moral commitment, no matter how strange that ultimate concern may seem to others, could be strengthened by a generic moral enhancer.

Enhancing New Moral Commitments

If an experimental moral enhancer instead altered what a subject believes to be morally right or modifies the subject's moral intentions, experimenters could confirm this by detecting contrasts between pre- and postexperiment convictions. An experimental subject who initially was a pro-choice advocate, and then transformed into a pro-life advocate, might be dismayed by listening to her own words in a preexperiment recorded interview about abortion. Experimental research could deliberately venture into the territory of seeking ways to diminish or even reverse moral convictions and intentions. Imagine if a trial group selected from an oppressed tribal minority unexpectedly dropped their hostility to show friendliness toward members of the oppressing majority tribe. (Yes, these curious hypotheticals are supposed to be arousing readers' worries over what is "moral" here.) A moral enhancer that considerably reshaped a person's moral views and motives, a "morphic moral enhancer," would be useful for those who want to add a new moral conviction to their set of moral beliefs, or to entirely reverse a unwanted moral conviction. One subject could use a morphic moral enhancer to add a conviction that patriotism toward one's country is the highest moral virtue, a second subject could arouse deep sympathies with the hitherto unnoticed poor, while another subject could reverse a belief that dutiful obedience to one's parents must be paramount.

Many generic moral enhancers, and most morphic moral enhancers, won't impress lots of people as genuine cases of moral enhancement. When a generic enhancement strikes lots of people as unworthy of the label "moral," they will instead regard it as an enhancement of integrity as best, or a degradation of morals at worst. Morphic enhancers will be regarded with even greater disdain, proportional to their deviations from what society regards as normal. Categories for "integrity enhancers" and "perversity enhancers" can parallel generic and morphic moral enhancers to keep track of social disapprovals. Moral enhancement experiments will become highly controversial in ways rivaling experiments on cognitive enhancement. The neuroethical subfield of applying standards of research ethics to experimental procedures will have to devise especially cautious criteria. If, hypothetically, one half of society can see no harm from temporarily rendering someone compliant to authority, while the other half is appalled, what are the wisest constraints to experiments upon human subjects?

Who Decides Which Are Moral Enhancers?

We are led to the broader question, by what moral standard shall efficacious moral enhancers be judged? Researchers might rely on their own moral standards, but approvals from research boards should be required. The stance of social relativism recommends that the standards to apply should be those that the environing society generally

regards as morally right. Additional primary sources for moral standards are the standpoints of intercultural objectivism and ethical objectivism. Each has its characteristic advantages and disadvantages. Intercultural objectivism (roughly) seeks standards for conduct deemed to be moral by one's own culture along with many or most other cultures, hoping that wide consensus on shared "ethical platitudes" serves as a reliable indicator of genuine morality. Although there are ethical platitudes to be found around the world, such as "respect others," "be concerned for others' needs," and the Golden Rule, their adequacy for complex issues of modern life can surely be questioned. Ethical objectivism instead recommends that the moral norms to apply should be the justifiably correct moral standards regardless of what any culture or individual happens to endorse. However, the inconclusive philosophical debates among rival ethical objectivisms have already been noted.

Only three models are ultimately needed for categorizing efficacious moral enhancement: the subjective, the social, and the intercultural. An experimenter's standards will either be idiosyncratically subjective, or follow her society's morality. An ethical objectivist standpoint reduces in practice to either subjectivism (e.g., where an ethical philosopher's verdict on the correct morality is shared by few others), or to prevailing social norms (by "discovering" that one's own society gets most moral matters about right), or possibly intercultural objectivism (coinciding with what the "better" cultures, in that theory's view, together hold to be moral). Despite the best efforts of philosophers pursuing ethical objectivism, they will have little direct impact on the development and evaluation of proposed moral enhancers. Any ethical objectivism either largely agrees with what some "proper" culture(s) already believe, or its moral standards deviate too far from widely shared standards to be useful for confirming the efficacy of moral enhancers.

Who Will Get Moral Enhancers?

The model of social morality suffices for classifying morphic moral enhancers aiming at the prevention of dangerously immoral and illegal conduct. For example, people easily aroused to violent behavior could benefit from "restorative" moral enhancers targeting respect for others or compassion for others' suffering. Perhaps a restorative moral enhancer could be given to a repeat offender for heightening not just respect for persons but also what people vitally care about (like their homes) to reduce property crimes as well. Broad approval from the public will endorse moral enhancers for elevating conduct to minimally expected levels of civility (indeed, many treatments for mental illnesses do this already), especially if they successfully prevent criminal violations without unwanted side effects. Appealing to the model of intercultural objectivism can supply additional justification for these enhancers, especially where suspicions are raised that social norms aren't sufficient grounds for dealing with criminal behavior. For example, a society applying moral enhancers for reducing religious nonconformity should find little intercultural support, while a society grappling with internal political rebellion might gain wider support, and a society dealing with the deleterious effects of severe drug addictions would garner broad international support.

This model of social morality is also the right classification for morphic moral enhancers aiming at alleviating what a society regards as specific kinds of undesirably antisocial conduct by otherwise law-abiding citizens. Explicit prejudices and xenophobias, for example, or implicit/subconscious biases behind sexism, ageism, or racism, could be targeted by social enhancers in the hopes of preventing undesirable conduct. A morphic conventional enhancer such as propranolol, which reportedly reduces implicit negative racial bias (Terbeck 2012), will be called the "cure for racism" long after scientists figure out the actual cognitive mechanisms responsible for such an effect (if that effect stands up to scrutiny) and neuroethicists repeatedly note how this beta-blocker has enough other effects on us (Hurlemann 2010; Kolber 2006) that labeling propranolol as "the antiracism drug" is like labeling alcohol as "the anti-response-time drug." Its tentative assignment as a morphic social enhancer, rather than a morphic intercultural enhancer, is not because racism doesn't violate valid intercultural standards (of course racism violates objective principles) but rather because the use of some allegedly "antiracism" drug couldn't be judged as having the same moral effect regardless of which ethnic group were to receive it. For dealing with biases acquired from socially conditioned fear-learning under certain prevailing cultural conditions, affecting so many cognitive processes including those behind explicit prejudice, no neurological modification could be remotely adequate to dealing with a pervasive problem like racism. Similarly, some proposed "antisexism" or "anti-religious bigotry" moral enhancer couldn't be judged effective in the identical manner across all human societies; it couldn't even remain viable across many decades within the same evolving society. Nevertheless, treatments capable of targeting and alleviating specific arousals toward fear, bias, and prejudice are physiologically possible (keeping in mind inevitable side effects), and if they do modify moral modes to the degree that conduct is changed, these treatments will mostly get categorized as morphic conventional enhancers having but limited applicability.

All the same, any satisfactory degree of practical applicability may prove irresistible to some societies seeking immediate relief from rampant immoral prejudices. "Preventative" moral enhancement for persons in positions of authority could especially receive social endorsement, especially for those holding judicial, legal, or political power, from jurors and police officers to elected public officials, but pondering such ideas makes a transition to issues of politics. Politics is never far behind questions of morality. If the research on "conservative" brains and "liberal" brains continues to hold up (see Haidt 2012; Mooney 2012; but see the cautions of Theodoridis and Nelson 2012), then perhaps separate moral enhancement formularies for these two political camps will emerge. Religious figures might proudly say how they can be relied upon to lead their

followers down the righteous path while getting their modal enhancers (though some will just as proudly declare how they don't need those enhancers). Neuroethics has a crucial role in correctly classifying and evaluating proposed moral enhancers to help judge exactly what such treatments would be really accomplishing.

Prevailing social norms won't work as well for confirming proposed moral enhancers attempting to elevate moral performance above socially expected levels, and intercultural objectivity can't easily approve them, either. Broad agreement won't be so easily garnered, paradoxically enough, where highly moral conduct is involved. There are other priorities alongside morality, and we may not really want what we say we want. For example, heightened empathy and caring should result in greater generosity, but is there common agreement that generosity should be generically elevated across the entire population? We may say that we wish people were more generous, but what we really mean is that ungenerous people should be more generous, or that other people should be more generous to *us*, or perhaps that people should be more generous to others who *deserve* it. Heightened levels of trustworthiness or truthfulness wouldn't be morally wise either, unless matched by a sound ability to discriminate who really deserves one's trust or confidence. The overall problem here is that the maintenance of human morality requires preparedness to judge, condemn, and appropriately punish; dispositions to caring more, by themselves, can ignore or erode those moral obligations. Being nicer is hardly the same as being more moral.

Generic empathetic, altruistic, or trusting enhancement sounds lovely in theory, but those using too much will eventually strike the rest of us as errant spirits, dangerous fools, or worse. Depictions of entire societies or a whole planet undergoing empathetic moral enhancement will remain utopian fantasies. One country after another will decline moral enhancement until the "worse" countries have done it, and each country would want their neighbors to go first. Even if everyone cared a whole lot more about each other, people will still care about other priorities like personal survival, family responsibilities, career advancement, equitable distributions, class distinctions, social equality, retributive justice, restoration of homelands, political freedom, and so on. Those possessing so much more would surely be delighted by moral enhancers for those with so much less; the next great opiate of the masses could be coming, but further domestication of humanity into placidity will not evenly distribute its benefits. We should already know that compassion for others does not necessarily correlate with promotion of their welfare, in any case; ideologies and religions first proclaiming how all are equally valued and loved and next willing to degrade and damn certain people are not unknown to humanity. Those talking about how love, or a neurological facsimile, can save the planet are forgetting the real-world questions. Chasms dividing us will not bridged like that. Enthusiasts advocating moral enhancement to help with peaceful political resolutions, slowing environmental destruction, or preventing transhuman tyranny must already be holding their preferred ethical outcomes in their pockets. Besides, for any altruistic enhancement applicable for overcoming past grudges or surmounting tragedies of the commons, there will be a countering enhancement—for fairness, or respect for tradition, and so on—ready to level the playing field.

Not caring to wait for social approval or planetary enlightenment, some individuals will seek personally designed moral enhancers. Combining a generic moral enhancer with a subjective model of moral enhancement results in what could be called a "specialized" moral enhancer, enhancing only that peculiar moral conduct specifically desired by individuals. Couples may seek a neurological boost to their relationship through an attachment enhancer (Savulescu 2008). A teacher might want to enhance his or her concern for, and strengthen intentions toward, helping the immature and uneducated. A politician could modify unwanted aspects of her or his moral conscience in order to reduce her or his sense of fairness and strengthen her or his intentions that wealthy supporters receive every possible assistance. Those upholding civic standards won't be impressed by eccentric enhancers that seem remote from ordinary social concerns or appear to directly contravene common moral standards. If a person wants a moral enhancer to intensify her or his devotion for her or his favorite sports team to the point of fanaticism, are we still talking about a moral enhancer? What about a moral enhancer that intensifies a religious person's enthusiasm for proselytizing to the point of abandoning family responsibilities?

Each society will inevitably demand and enforce a distinction between what it regards as "genuine" moral enhancement and specialized moral modifiers, since specialized modifications won't typically conform to what society generally expects of moral conduct from everyone. A society may justify enforce that distinction either by setting up its own social morality as the moral standard—applying the model of social moral enhancement—or it could additionally claim justification by intercultural moral standards to which many cultures subscribe. Unlike appeals to ethical objectivism, where alleged philosophical discernment of uniquely correct moral standards is involved, appeals to interculturally objective standards only require the discernment of significant moral principles and ideals found already in many cultures. We still await some formula for deciding precisely which cultures shall count here, of course, yet societies may not be misguided for hoping that intercultural moral ideals are often a sounder basis that their own provincial standards. International accords and treaties for standardizing the global acceptability of moral enhancers will find a ready foundation in intercultural objectivism, as well.

THE MARKET FOR MORAL ENHANCERS

To clearly and completely describe a moral enhancement, four components are needed: the mechanism(s), mode(s), generic/morphic status, and finally the model (specialized, social, intercultural). Those trying to market what they call "moral enhancers" should be asked how they categorize their products.

The three models of moral enhancement are formulary classifications only sufficient for differentiating the design, marketing, and application of moral enhancers in the real world. Conceptually, these three categories are somewhat incompatible with each other. People advocating one variety of moral enhancement will tend to deny that the other kinds are fully legitimate. Those who think that there are intercultural ethical principles won't recognize a social enhancer as a moral enhancer if it violates those high principles. Similarly, those happy to apply social moral standards to enhancers will tend to deny that specialized enhancers are really moral enhancers unless they conform to social standards. Advocates for specialized enhancers will be frequently told that such enhancers are not genuine moral enhancers, by both social relativists and intercultural objectivists. Ways to avoid conflict can be constructed, as well. Enhancing basic moral duties and virtues recognized across most cultures can be compatible with enhancing certain other moral norms esteemed by just one society. Similarly, a society may tolerate many sorts of specialized enhancers so long as acceptable levels of civility and morality (and practical common sense) are sustained. Those specialized enhancers would be deemed "safe and effective"—"safe" for society and "effective" for the individual.

All three models of moral enhancers could simultaneous compete on the open market. The most likely moral enhancers that will be designed, tested, and marketed in the near future will be socially morphic modifications intended as preventative or restorative remedies for what a society generally deems to be seriously antisocial, immoral, and illegal behavior. Political propaganda and clever marketing will eagerly trail scientific claims about improving our morals, helping to manufacture acceptance of a few "prosocial" moral enhancers among a consumer public already amenable to emotional appeals and social conformity. A small number of specialized enhancers, either generic or morphic, will also be offered to individual clients as "boutique" moral enhancers in some "safe and effective" form. Intriguing types of specialized moral enhancers will enable a person to better fulfill a specific operational role, such as military service, police duty, or medical practice. Perhaps judges will be expected to use "fairness" enhancers while court is in session and rulings are crafted. We can also anticipate governments seeking to implement rapidly acting moral enhancers for spraying over unruly protesters on the march, to temporarily pacify them by arousing feelings of tranquility or submissiveness to authority, or by clouding their (wrong-headed?) convictions.

Many readers are by now thinking that there's some mistake in letting all these "enhancements" get called "moral." Indeed they cannot all be moral—but that's not because we all know what "moral" really is. There will be endless disagreement over whether modifications are genuinely moral enhancements, and no universal accord on what constitutes moral enhancers should be predicted. In this persistent climate, I urge unrelenting neuroethical scrutiny into the public dissemination and utilization of information about alleged moral enhancers and their promised value. The number of potential moral enhancers could become vast, and novel varieties will keep neuroethicists permanently busy. It will become incumbent on neuroethicists, and any evaluators of moral enhancement, to clearly announce which modes and models they are presuming while assessing moral enhancement, to avoid unnecessary verbal confusions and argumentative dead ends. The rare theorists relying on their own "perfected" ethical systems should similarly announce their presumptive foundations before explaining their evaluations of moral enhancements. Likewise, those theorists advocating moral enhancements for overcoming the limitations of ordinary moral dispositions acquired from our human biology and socialization should be reminded that what they are actually urging is ethical enhancement, so their reasoned philosophical arguments are needed from their brains, not invasive techniques for our brains.

In summary, complete skepticism toward the possibility of moral enhancement cannot be justified; only modest endorsement of specific kinds of moral enhancement seems wise; and cautious cynicism against enthusiastic hopes for dramatic moral enhancement feels appropriate. Moral enhancement should not be confused with ethical enhancement. Modifying conduct in line with what people already regard as ordinary moral behavior can never replace thoughtful adjudication among conflicting moral duties or adjusting our social norms for improving the human condition. Confusing the two, and many dominating powers will want to, can only cause more misery for humanity. If you must label me, don't call me a moral enhancement skeptic—call me a moral enhancement cynic.

We are a species that socializes individuals into moral conformities, producing people who think that they know morality when they see it, and who are highly motivated to want more of it (from others, at least). There will be novel brain modifications that meet our expectations, whatever they happen to be, for delivering moral enhancement, and people will want to use them. Fortunately for our species, we are also capable of ethically questioning our own moralities and any suggested modifications to them.

REFERENCES

Berker, S. 2009. The normative insignificance of neuroscience. *Philosophy and Public Affairs* 37: 293–329.

Boire, R. G. 2004. Neurocops: The politics of prohibitions and the future of enforcing social policy from inside the body. *Journal of Law and Health* 19: 215–257.

Bublitz, J. C., and R. Merkel. 2009. Autonomy and authenticity of enhanced personality traits. *Bioethics* 23: 360–374.

Chan, S., and J. Harris. 2011. Moral enhancement and pro-social behavior. *Journal of Medical Ethics* 37: 130–131.

Churchland, P. 2011. *Braintrust: What neuroscience tells us about morality*. Princeton, NJ: Princeton University Press.

Decety, J., K. Michalska, and K. Kinzler. 2012. The contribution of emotion and cognition to moral sensitivity: A neurodevelopmental study. *Cerebral Cortex* 22(1): 209–220.

Dennett, D. 2003. *Freedom evolves*. New York, NY: Penguin Books.

Douglas, T. 2008. Moral enhancement. *Journal of Applied Philosophy* 25: 228–245.

Engelhardt, H. T. 1990. Human nature technologically revisited. *Social Philosophy and Policy* 8: 180–191.

Farah, M. 2012. Neuroethics: The ethical, legal, and societal impact of neuroscience. *Annual Review of Psychology* 63: 571–591.

Faust, H. S. 2008. Should we select for genetic moral enhancement? A thought experiment using the MoralKinder (MK+) Haplotype. *Theoretical Medicine and Bioethics* 29: 397–416.

Flanagan, O., H. Sarkissian, and D. Wong. 2008. Naturalizing ethics. In *Moral psychology, vol. 1: The evolution of morality*, ed. W. Sinnott-Armstrong, 1–26. Cambridge, MA: MIT Press.

Franks, D. D. 2010. *Neurosociology: The nexus between neuroscience and social psychology*. New York, NY: Springer.

Frith, C. D., and U. Frith. 2012. Mechanisms of social cognition. *Annual Review of Psychology* 63: 287–313.

Fröding, B., and E. Esmerelda. 2011. Cognitive enhancement, virtue ethics and the good life. *Neuroethics* 4: 223–234.

Gazzaniga, M. 2005. *The ethical brain: The science of our moral dilemmas*. New York, NY: Harper Perennial.

Glannon, W. 2011a. *Brain, body, and mind: Neuroethics with a human face*. New York, NY: Oxford University Press.

Glannon, W. 2011b. Diminishing and enhancing free will. *AJOB Neuroscience* 2(3): 15–26.

Graham, J., B. Nosek, J. Haidt, R. Iyer, S. Koleva, and P. Ditto. 2011. Mapping the moral domain. *Journal of Personality and Social Psychology* 101: 366–385.

Greene, J. 2008. The secret joke of kant's soul. In *Moral psychology, vol. 3: The neuroscience of morality*, ed. W. Sinnott-Armstrong, 35–79. Cambridge, MA: MIT Press

Haidt, J. 2012. *The righteous mind: Why good people are divided by politics and religion*. New York, NY: Random House.

Harris, J. 2011. Moral enhancement and freedom. *Bioethics* 25(2): 102–111.

Henrich, J., S. Heine, and A. Norenzayan. 2010. The weirdest people in the world? *Behavioral and Brain Sciences* 33: 61–83.

Huebner, B., S. Dwyer and M. Hauser. 2008. The role of emotion in moral psychology. *Trends in Cognitive Sciences* 13: 1–6.

Hughes, J. 1996. Embracing change with all four arms: Post-humanist defense of genetic engineering. *Eubios: Journal of Asian and International Bioethics* 6: 94–101.

Hurlemann, R., et al. 2010. Human amygdala reactivity is diminished by the β-noradrenergic antagonist propranolol. *Psychological Medicine* 40: 1839–1848.

Jotterand, F. 2011. Virtue engineering and moral agency: Will post-humans still need the virtues? *AJOB Neuroscience* 2(4): 3–9.

Joyce, R. 2002. *The myth of morality*. Cambridge, UK: Cambridge University Press.

Kalis, A. 2010. Improving moral judgments: Philosophical considerations. *Journal of Theoretical and Philosophical Psychology* 30(2): 94–108.

Kamm, F. M. 2009. Neuroscience and moral reasoning: A note on recent research. *Philosophy and Public Affairs* 37: 330–345.

Kane, R., ed. 2005. *The Oxford handbook of free will*. Oxford, UK: Oxford University Press.

Kolber, A. J. 2006. Therapeutic forgetting: The legal and ethical implications of memory dampening. *Vanderbilt Law Review* 59: 1559–1626.

Krebs, D. 2011. *The origins of morality: An evolutionary account*. Oxford, UK: Oxford University Press.

Lavazza, A., and M. De Caro. 2010. Not so fast. On some bold neuroscientific claims concerning human agency. *Neuroethics* 3(1): 23–41.

Libet, B., A. Freeman, and K. Sutherland (Eds.). 1999. *The volitional brain: Towards a science of free will*. Thorverton, UK: Imprint Academic.

Lillehammer, H. 2003. Debunking morality: Evolutionary naturalism and moral error theory. *Biology and Philosophy* 18: 566–581.

Macer, D. 1995. Medical aspects of the repression of deviating behavior. *International Journal of Bioethics* 6: 38–40.

Mele, A. 2009. *Effective intentions: The power of conscious will*. Oxford, UK: Oxford University Press.

Mooney, C. 2012. *The republican brain: The science of why they deny science and reality*. Hoboken, NJ: Wiley.

Nagel, S. K. 2010. Too much of a good thing? Enhancement and the burden of self-determination. *Neuroethics* 3(2): 109–119.

Parkinson, C., W. Sinnott-Armstrong, P. E. Koralus, A. Mendelovici, V. McGeer, and T. Wheatley. 2011. Is morality unified? Evidence that distinct neural systems underlie moral judgments of harm, dishonesty, and disgust. *Journal of Cognitive Neuroscience* 23(10): 3162–3180.

Persson, I., and J. Savulescu. 2008. The perils of cognitive enhancement and the urgent imperative to enhance the moral character of humanity. *Journal of Applied Philosophy* 25(3): 162–177.

Persson, I., and J. Savulescu. 2010. Moral transhumanism. *Journal of Medicine and Philosophy* 35(6): 656–669.

Persson, I., and J. Savulescu. 2011. Unfit for the future? Human nature, scientific progress, and the need for moral enhancement. In *Enhancing human capabilities*, ed. J. Savulescu, R. terMeulen, and G. Kahane, 486–500. Malden, MA: Wiley-Blackwell.

Roskies, A. 2006. Neuroscientific challenges to free will and responsibility. *Trends in Cognitive Sciences* 10: 419–423.

Savulescu, J. 2008. Neuroenhancement of love and marriage: The chemicals between us. *Neuroethics* 1: 31–44.

Schroeder, M. 2009. *Noncognitivism in ethics*. London, UK: Taylor & Francis.

Schulkin, J. 2011. *Adaptation and well-being: Social allostasis*. Cambridge, UK: Cambridge University Press.

Sinnott-Armstrong, W. 2008. Is moral phenomenology unified? *Phenomenology and the Cognitive Sciences* 7(1): 85–97.

Sinnott-Armstrong, W., and L. Nadel (Eds.). 2011. *Conscious will and responsibility*. Oxford, UK: Oxford University Press.

Spence, S. A. 2008. Can pharmacology help enhance human morality? *British Journal of Psychiatry* 193: 179–180.

Terbeck, S., G. Kahane, S. McTavish, J. Savulescu, P. Cowen, and M. Hewstone. 2012. Propranolol reduces implicit negative racial bias. *Psychopharmacology*. doi:10.1007/s00213-012-2657-5.

Tersman, F. 2006. *Moral disagreement*. Cambridge, UK: Cambridge University Press.

Theodoridis, A., and A. J. Nelson. 2012. Of bold claims and excessive fears: A call for caution and patience regarding political neuroscience. *Political Psychology* 33: 27–43.

Verbeek, P.-P. 2009. Ambient intelligence and persuasive technology: The blurring boundaries between human and technology. *Nanoethics* 3: 231–242.

Walker, M. 2009. Enhancing genetic virtue. *Politics and the Life Sciences* 28(2): 27–47.

Walter, H. 2001. *Neurophilosophy of free will*. Cambridge, MA: MIT Press.

Wegner, D. 2002. *The illusion of conscious will*. Cambridge, MA: MIT Press.

Zak, P. J. 2011. The physiology of moral sentiments. *Journal of Economic Behavior and Organization* 77: 53–65.