

Distribution Agreement

In presenting this thesis as a partial fulfillment of the requirements for a degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis in whole or in part in all forms of media, now or hereafter now, including display on the World Wide Web. I understand that I may select some access restrictions as part of the online submission of this thesis. I retain all ownership rights to the copyright of the thesis. I also retain the right to use in future works (such as articles or books) all or part of this thesis.

Yaseen A. Jamal

April 16, 2013

Beginning a Science of Morality:

How the link between well being and neurobiology can define our values

by

Yaseen A. Jamal

Dr. Richard Patterson

Advisor

Department of Philosophy

Dr. Richard Patterson

Advisor

Dr. Susan Bredlau

Committee Member

Dr. Arri Eisen

Committee Member

Dr. Melody Siegler

Committee Member

2013

Beginning a Science of Morality:

How the link between well being and neurobiology can define our values

by

Yaseen A. Jamal

Dr. Richard Patterson

Advisor

An abstract of a thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment of the requirements of the degree of
Bachelor of Sciences with Honors

Department of Philosophy

2013

Abstract

Beginning a Science of Morality:

How the link between well being and neurobiology can define our values

By Yaseen A. Jamal

Moral Science is a novel approach to normative ethics introduced by neuroscientist, Sam Harris. This consequentialist framework utilizes a scientific process to define moral values and rests on a factual understanding of human well being. Though logically sound, this system lacks empirical support. Specifically, Harris defends the potential for science to discover moral values, but he does not adequately explain how a developed science of morality might function.

This thesis aims to construct the empirical groundwork necessary to substantiate Harris' argument by beginning a new science of morality. Chapter one explores conceptions of human flourishing and concludes with two sets of values. Chapter two seeks to define the proposed values by drawing on relevant ideas from neuroscience and biology. Chapter three elucidates the hypothetical construction of the proposed moral science and examines its applications.

Beginning a Science of Morality:

How the link between well being and neurobiology can define our values

by

Yaseen A. Jamal

Dr. Richard Patterson

Advisor

A thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment of the requirements of the degree of
Bachelor of Sciences with Honors

Department of Philosophy

2013

Table of Contents

Introduction.....	1
Chapter One: Human Flourishing.....	7
Chapter Two: Defining Values.....	33
Chapter Three: Science and Morality.....	53
Works Cited.....	64

Introduction

This thesis aims to begin the construction of a science of morality, a novel approach to identifying normative values through empirical methods. This project is a scientific exploration of the “moral landscape,” a hypothetical space that represents the entirety of human conscious experience. The idea of modeling consciousness as a landscape was first proposed in 2011 by neuroscientist, Sam Harris; in his book, Harris defends the factual nature of the moral landscape, and he argues that a scientific exploration of this space can inform ethical discourse. However, many of Harris’ critics emphasize that the lack of empirical demonstration defines his philosophy as naive. In order to validate his framework, the present thesis will use the moral landscape as a platform for scientific investigation, aiming to discover moral values and beginning their empirical definition. In essence, a new science of morality will be constructed and the broader implications of the project will be discussed, thereby facilitating the advance of this nascent framework from a philosophic hypothesis to a functional scientific theory.

In order to begin this task, the philosophy supporting the moral landscape must first be made clear. Stemming from the defense presented in Harris’ book, the general argument for the moral landscape begins with a simple observation. When considering human conscious experience, it is clear that there is a range of possibility; not all conscious experiences are alike. This point can best be understood by simply imagining the best conceivable life and comparing it to its opposite. Regardless of individual opinion of what this life might entail, and before any terms are defined, the difference between these two lives seems self-evident.

The next thing to notice is that the differences between conscious states are factual differences. Assuming that the human experience is directly and consistently related to states of the brain (indeed, modern neuroscience indicates that this is a plausible assumption), then the

range of all possible conscious experiences must relate to the range of all possible states of the brain. Therefore, this range must be finite, though incomprehensibly large, and differences in the quality of life necessarily relate to differences in neural structure or function.

The above recognitions are all that are required to admit the existence of a moral landscape, a hypothetical, finite space representing the range of all possible conscious states. In this model, the peaks of the landscape represent the most positive human experiences, the heights of human flourishing and the ultimate aims of human behavior; in contrast, the valleys represent the worst possible conscious states, those indicative of human suffering. Human life involves the totality of movement along this landscape via values, the actions/thoughts that drive movement towards specific points on the moral landscape. Further, this space is an empirical one, and so it is necessarily open to scientific study. In principle, the entirety of this space can be explored by science, though the question of whether such an immense task can be fully accomplished remains unanswered. The reason this landscape is termed “moral” will become clear in the following discussion.

Upon admitting the existence of the moral landscape, the conversation formally shifts to ethics. Harris asserts that every individual strives for a good life; that is, each person wants to live a life worth living. He claims that the good life is absolutely better than the worst life, and that all individuals would prefer the good life over its opposite. According to Harris, any individual who claims otherwise is either ignorant of what the good life means or has a mental illness (psychopath/sociopath). Additionally, any possible human desire, regardless of its moral quality, can be better fulfilled in the context of the good life rather than the bad, so the assertion that all individuals strive toward the same goal in life seems plausible. That is, assuming all humans aim for a life of flourishing, rather than suffering, seems like a fair proposition.

The precise definition of the good life is intentionally left open, but it should be understood that this notion, by intuition, relates to the most positive kinds of human experience rather than the negative kinds of experience. If it did not, then it would not longer be “the good life,” because its basic definition is linked with heights of human fulfillment. To clarify, Harris does not provide reasons why we should value human flourishing over suffering. This is not a problem, however, because the motivation to increase well being appears universal. Further, the good life is a broader concept that encompasses a sum of experiences. Claiming that all humans strive to increase well being does not necessarily imply that all humans only strive for positive experiences. To reiterate, the concept of a good life does not need to be fully elucidated in order to gain a sense of what kinds of experience (suffering or flourishing) characterize it; although there may be individual differences in conceptions of a good life, positive conscious experiences are a necessary condition for all.

Harris goes on to contend that morality can only be understood in the context of human conscious experience. Any conception of ethics or the good life must relate to movement on the moral landscape. Quality of experience is represented by the moral landscape along its vertical axis (higher points on the landscape represent conscious states of increased well being, while lower points represent the opposite), thus morality generally concerns moving away from valleys and toward peaks. Accordingly, moral values are those behaviors, thoughts, or practices that promote this movement. Values that do not affect movement on this space are, for Harris, completely irrelevant to a discussion of morality. The only reasonable domain of values is the fully characterized moral landscape; other potential domains of value (i.e. religion, *a priori* philosophical systems, or culture) must overlap with this domain in order to hold ethical significance.

Given these definitions in light of the factual nature of the moral landscape, there simply must be right and wrong ways to maximize well being; that is, there are necessarily right and wrong answers to moral questions, since morality directly relates to conscious experience and its underlying biology. Because the conscious experiences represented by the peaks and valleys are factual ends, there must be correct routes toward them. Therefore, Harris argues that values, which are behaviors/thoughts that drive movement toward peaks, are empirical concepts and necessarily await scientific discovery.

Although Harris mentions the importance of modern neuroscience in understanding conscious experience, he fails to explain how science will illuminate the moral landscape. The present thesis seeks to address this question by beginning a new science of morality. The philosophic reasoning described above has set up the foundations for this thesis, but I will not be specifically defending Harris' philosophy. For the purpose of the current study, Harris' primary argument will be assumed to be valid and sound, and the remainder of this work will expand on the aforementioned considerations.

A science of morality seeks to illuminate the moral landscape: first, values which promote individual movement toward peaks on the moral landscape will be identified; second, proposed values will be further detailed through a discussion of relevant biology and neuroscience; third, the utility of this knowledge for an ethical system will be clarified, and the full relationship between science and morality will be delineated. Here, the dual function of the present thesis is clear: a scientific exploration of the moral landscape aims to empirically substantiate Sam Harris' philosophy and ignite a new science of morality that may continue into the future.

Chapter one will begin with a consideration of potential peaks on the moral landscape and continue with an identification of values. In this chapter, different theories of human flourishing will be evaluated for compatibility with the moral landscape. Consequently, two sets of moral values will be proposed.

Chapter two will continue the development of the nascent moral science through a second look at the two sets of values offered in the first chapter. First, the benefit of empirically defining values will be made clear. Next, the two sets of offered values will be clarified in terms of relevant biology and neuroscience, and the form of future work in moral science will be suggested.

Finally, chapter three elucidates the progression of the moral science from its philosophic origins to its real-world applications. The findings from chapters one and two will be summarized and their potential role in ethical conversation will be considered. If successful, the completed picture of how a science of morality might operate, along with its normative potential in academic or societal discourse, will become clear. Finally, the full extent of contact between science and morality, apart from the prior work, will be discussed.

It is worth noting that a particular feature of moral science is especially prominent in the current work. Like any other scientific project, the current thesis is an imperfect effort to better understand natural phenomena and to utilize its findings for the benefit of mankind. In the true spirit of scientific inquiry, it lends itself to critique, peer-review, and healthy skepticism.

Therefore, this discussion primarily serves to ignite the future conversation about the relationship between science and morality that will perpetuate the advance of this novel discipline. Scientists and philosophers, students and professors, governments and the public alike can draw from this potential field of study to spark global awareness about the need for civilized moral discourse. It

is my hope that a science of morality, similar to chemistry, physics, biology, or any of the other sciences, can improve our understanding of the natural world, enhance the global progress of modern society, and continue to provide the immense sense of meaning and joy that has sustained its development thus far.

Chapter One: Human Flourishing

A science of morality is charged with the task of identifying and fully understanding normative values. This task will operate via scientific exploration of the moral landscape, a factual model that represents the range of all possible conscious experiences. The moral landscape has already been shown to be empirically grounded during Sam Harris' philosophic defense, and will serve as the primary foundation for this thesis. Specifically, the illumination of this space (i.e. the identification of moral values) will involve recognizing those patterns of thought or behavior which encourage individual transition toward peaks of human flourishing and clarifying those patterns in terms of neurobiology.

To achieve this goal, the current chapter will first explore five conceptions of human flourishing and narrow this list through reconciliation with the nature of the moral landscape. Next, these chosen "peaks" of the landscape will be used to identify possible values, and observational studies will be used to support their utility in increasing well being. Finally, several important counter-arguments will be addressed. This exercise aims to serve as the first demonstration of how a science of morality might discover normative values, and carries the virtue of empirically substantiating Sam Harris' philosophic arguments.

To begin the scientific search for normative values, we must first gain a sense of what the peaks of human flourishing might look like, so that we can better recognize the patterns of thought and behavior (values) that strive toward them. Again, these peaks are representative of conscious experiences. The broader concept of human flourishing, or the good life, represents the entirety of lived conscious experience (i.e. the sum of movement along the moral landscape), and implies engagement with values that aim toward peaks on the landscape. Still, it is not necessary

to define the totality of conscious experiences in the good life in order to discover which values lead to the positive experiences that characterize it.

The purpose of a science of morality is to identify values, not to fully elucidate human well being, which is a much more complex task. To clarify, the well being of an individual can be understood as the particular movement of that individual through the slopes of the moral landscape, while values are the driving forces of this movement. Flourishing would involve a movement driven by values that aim towards peaks on the moral landscape and, to some extent, an experience of those peaks, while suffering would involve movement that is primarily driven by behaviors that progress downwards on the moral landscape.

The reason the five conceptions of peaks of human flourishing need to be considered is quite straightforward. If the moral landscape is an accurate representation of human consciousness, then it is most likely that individuals are not completely knowledgeable of its highest peaks or the values that lead toward them. It is possible for individuals to be wrong about how to direct their own well being, since movement on the moral landscape has been shown to be a matter of fact. Therefore, when illuminating the range of possible conscious experiences, a science of morality must consider all options for theories of flourishing and select the most likely to accurately reflect the actual peaks. This method is the most reliable way to begin exploring the moral landscape, which, again, is an empirical space that is best explored by an empirical method, because without this method, the identification of values will function arbitrarily.

Dr. Daniel Haybron is a professor of philosophy who offers what he claims to be a comprehensive taxonomy of theories of human well being. He lists five broad families of conceptions of human flourishing, and it appears that any individual perspectives about what constitutes a good life will likely fall into at least one of these categories. He expands on the

well-known taxonomy drafted by philosopher Derek Parfit in 1982, which included hedonistic, desire, and objective list theories, by including additional eudaimonistic and “authentic happiness” theories (Haybron, 22). For the present study, it is helpful to consider each of these in relation to the moral landscape in order to establish the potential peaks that scientific moral values will aim toward.

The hedonistic view of human flourishing recognizes pleasure as the ultimate good, meaning that pleasurable experiences define the peaks on the moral landscape. In this case, the broader conception of human flourishing suggests that an individual should live (move within the moral landscape) via values that strive for pleasant experiences. This view equates well-being with a mere feeling and, in general, emphasizes actions that promote pleasant experiences above all else (Haybron, 22). Although there are a variety of views about what constitutes pleasant experiences, all reflect our intuitive tendencies to seek out hedonic gains. This theory may seem attractive because it offers a seemingly plausible way to enjoy life, as its aim is enjoyment itself. However, it should be clear that the kind of human flourishing at the center of the present moral system, the peaks of lived experiences and highest aims for a meaningful life, will probably require more than simply feeling good. The difference between the best conceivable life and the worst, mentioned during the introduction, almost certainly involves more variables than the pleasantness of experience.

To better understand this refutation, consider Robert Nozick’s “experience-machine,” a hypothetical instrument that humans could plug into and experience a simulated reality consisting of all of the pleasurable conscious experiences an individual could want, given that there would be no way to know the difference between reality and the simulated world. The thought experiment asks, if given the choice, whether there is reason to not plug into the

machine. If hedonistic experiences are the only factors that matter for human well being, then there would be no reason to choose reality instead of the experience-machine; logically, the simulated reality is the better route to the good life, assuming the machine can deliver more pleasurable experiences than reality itself. Therefore, if we can conceive of any logical reason to not plug into the machine, then there are likely other values involved in human well being in addition to the pursuit of pleasure (Haybron, 18).

Nozick offers three of his own reasons explaining why reality might be preferred over the experience-machine. First, he asserts that our desire to engage in activity is logically prior to our desire for the experience of doing the activity. Second, he asserts that we prefer to live as a human being, rather than as a body with a stimulated mind. Third, he suggests that the experience-machine is restricted to only the kinds of pleasurable experiences humans can conceive of, whereas reality offers unrestricted potential for experience (Haybron, 18). Thus, because there are logical reasons to not want to plug in, there must be other qualities of conscious experience that distinguish degrees of well being, so the peaks on the moral landscape likely represent more than just experiences of pleasure. There are certainly other objections to the hedonist view of human flourishing, but Nozick's thought experiment appears to be the most relevant.

Similar to the hedonistic theory, the desire theory emphasizes individual autonomy in crafting a meaningful life, but includes values in addition to those that encourage pleasant experience. This approach seeks to empower individuals to achieve their own conceptions of a good life, and assumes that people are fully knowledgeable about how best to improve their own well being (Haybron, 23). The desire theory is appealing because it maintains the importance of individual freedom and allows for an infinite number of personal conceptions of well-being. It is

important to recognize that these desires originate in culture, past experiences, and personality (since moral science has not fully developed), so there is a high probability that individual desires will not reflect the correct ways to move toward peaks on the moral landscape. The development of moral science, however, can enable these individuals to form the correct desires.

The individual freedom to value just about anything, which the desire theory maintains, conflicts with the factual nature of moral truth (human well being). If individuals are ignorant about the peaks and valleys of experience, how might they know which desires will most likely lead to their flourishing? It is entirely possible that someone could want things that are irrelevant or detrimental to their well-being simply because they are ill-informed about the consequences on their conscious experience. Therefore, a theory that merely values the fulfillment of desires, rather than the effects of those desires on conscious experience, hardly seems like a genuine approach to well-being or even morality. To reiterate, the highest peaks on the moral landscape more likely reflect the instrumental goal of desire fulfillment in general rather than the experience of fulfilling desires itself. This goal is likely a kind of transient conscious experience that is more meaningful than the feeling of desire satisfaction, assuming that personal desires are inconsistent with the correct values suggested by the moral landscape. However, if personal desires happen to match factual values, or if personal desires are influenced by values developed by a science of morality, then the experience of desire fulfillment as a peak on the moral landscape seems more reasonable.

Assuming individuals are ignorant about the moral landscape, the desire theory of well-being is restricted to one kind of conscious experience, desire fulfillment. The development of moral science can help enlighten these individuals about which values lead to upward movement on the moral landscape, but because desire fulfillment theory emphasizes individual autonomy in

value generation, there is still a chance that the individual could form values that are irrelevant or detrimental to conscious experience, simply because the ultimate source of values must be the individual.

This problem is addressed by the next approach, but it is important to recognize that even if individuals produce values that have some bearing on conscious experience, the emphasis on individual autonomy could lead to all sorts of different interpretations and applications of the values generated by moral science, which would miss the point of relaying the values in the first place. Thus, desire fulfillment does not seem like the most plausible goal for a normative ethics, because individual autonomy in the identification of values simply must be limited, given that there is a factual basis for human well being and that individuals could form values that are plainly wrong. Notice that the freedom to engage in scientific values is different than the freedom to generate values; the former is an inalienable human right, while the latter is the responsibility of moral science.

The third conception of peaks of human flourishing is a more specific formulation of the second, and is termed the authentic happiness theory. This approach still empowers individual autonomy, but it recognizes that values must affect conscious experience and focuses on a broader idea of happiness, including overall life satisfaction and positive conscious experiences. In this case, individuals are informed about the primacy of conscious experience for well-being and shape their desires to align with their reflection. To reiterate the difference between this view and the second, desire formation theory emphasizes individual satisfaction of any desire, while authentic happiness encourages individual satisfaction of desires that impact their conscious experience (Haybron, 23-4). Specifically, this means that the individual must form a desire, such as becoming CEO of a company, with the assumption that fulfilling this desire will result in their

happiness, or at least in some kind of positive experience. This is not a condition assumed by the second approach.

The authentic happiness theory of human flourishing seems only partially compatible with a science of morality because, like the desire fulfillment theory, it emphasizes that values must be autonomously generated rather than supplied by an external source. Given that the moral landscape is a factual space, it is very possible that individuals could generate values which, unlike desire fulfillment theory, consequence on individual conscious experience, but still do not promote well being. The virtue of illuminating the moral landscape is that a scientific method can correctly identify values based on the factual understanding of conscious states. Even if individuals were informed about the moral landscape, it seems more likely that science can generate reliable values rather than individuals creating their own.

The authentic happiness theory, though more compatible with the moral landscape than the desire fulfillment theory due to its consequentialist criteria, still leaves open the possibility for erroneous judgments about values that improve well being. The fundamental shortcoming here lies in the individual's absolute freedom. If the peaks on the moral landscape are factual ends, and values can be correct or incorrect routes toward these ends, then it seems counter-productive to maximize human well being by encouraging individuals to think of their own routes to well being. To reiterate, associating the authentic happiness theory with the peaks on the moral landscape seems to run counter to the goal of moral science, which involves the effort to identify values that reliably indicate movement towards the genuine peaks (which will be identified later in this discussion).

Even if individuals were informed about the moral landscape, but free to design their own values, the possibility for individuals to select the wrong values is still open, because of

individual interpretation of the information relayed by the science of morality. Again, the current effort is to identify what kinds of conscious experiences represent the peaks on the moral landscape, not to fully delineate what the good life entails. Though the authentic happiness theory includes the necessary criteria for values to affect movement on the moral landscape, the emphasis on individual generation of values defeats the purpose of a science of morality; it incorrectly assumes that the greatest human happiness is found when individuals generate their own values to shape their own conscious experience. The moral landscape indicates that different values have different consequences on conscious experience; autonomously generated values cannot all be equally sufficient for a good life. The authentic happiness theory fails to recognize this fact, thus, this theory probably does not reflect the peaks on the moral landscape.

The fourth approach is different from the first three in that it asserts that individuals do not know what is best for their own well-being, and restricts their freedom to generate values which they think might improve their well being. Rather, those who are knowledgeable about well-being are the most apt to make normative prescriptions. Evidently, this approach was made popular by the ancient Greeks (Plato and Aristotle, in particular) and is termed the eudaimonistic theory. This theory operates by examining human nature and delineating human capacities, consequently seeking values that fulfill this nature; by definition, this life of satisfying human potential is the life worth living (Haybron, 24). However, this theory may inappropriately de-emphasize the importance of positive affect and subjective experience in well-being. It seems possible that a conception of eudaimonia be offered, but without any degree of individual autonomy, it is unreasonable to assume that a eudaimonistic view of well being is equally fulfilling for every individual.

Though the moral landscape would benefit from identifying eudaimonistic goals as peaks, since values would be easier to identify, at least a minimal degree of flexibility (individual autonomy) must be incorporated in order to account for individual differences in the subjective experience of eudaimonistic behaviors. For example, instructing 100 individuals to run a mile will most likely result in a variety of subjective experiences of the run, so the individual must be given at least some autonomy to customize the eudaimonistic behavior to match his or her unique capacities.

Though eudaimonistic values do not require everyone to engage in the same specific actions, they must be general enough to be applied universally. Indeed, Aristotle acknowledged that different people would have to give different amounts in order to qualify as generous individuals, so autonomy is clearly incorporated into the eudaimonistic theory of well being. Further, Aristotle recognized that exercising generosity must a pleasant experience, so the previously established criteria of values bearing on conscious experience also seem to be incorporated into the concept of eudaimonia. At this point, eudaimonistic activity that incorporates individual freedom seems like a promising way to understand the peaks on the moral landscape. Still, there is one final theory of human flourishing left to consider.

The fifth conception of the peaks of human experience involves a list of abstract goods whose possession indicates well-being. This approach is referred to as list theory and identifies the realization of notions such as friendship, knowledge, or achievement as intrinsically meaningful. This view is different from desire theory and similar to eudaimonistic theory in that it is offered by those studying human nature, rather than by individuals (Haybron, 26-7).

If full autonomy is allowed in this theory, then the possibility for individuals to be wrong is once again open. If freedom is instead limited, then scientists are responsible for determining

the most likely peaks on the moral landscape, and their determination is more reliable because of rational, observational, and verifiable methods. However, this theory of the peaks on the moral landscape leads to a problem: it fails to identify why the goods are intrinsically meaningful and seems to stop short of the sense of fulfillment that these goods reflect. While Aristotle acknowledged that the exercise of natural human capacities is fundamentally tied to experience, the list theory fails to make this recognition. Therefore, a list of goods produced by examining human capacities without identifying the consequences of this behavior in terms of conscious states has the potential to be irrelevant or contradictory to the moral landscape.

The nature of the moral landscape, and the philosophic framework supporting the science of morality, centers around the subjective conscious experience, so identifying the long-term sense of fulfillment that comes from engaging in eudaimonistic behavior is crucial for list theory to accurately reflect the peaks on the moral landscape. In addition, the population of goods on a list theory of well-being loses a necessary factual anchor by failing to acknowledge the primacy of subjective perception, and thereby loses its relevance to the present work. Like most of the other potential conceptions of peaks on the moral landscape, this approach misses the point of a science of morality, which is to lucidly define values that promote the greater sense of fulfillment only implied by list theories. Accordingly, this suggests that the peaks on the moral landscape likely reflect a sense of fulfillment that can only come from exercising natural capacities rather than possessing a list of eudaimonistic goods.

To summarize the discussion of theories of well-being so far: the hedonistic theory does not hold because peaks of human experience probably involve more than pleasant feeling; thus, hedonistic satisfaction appears instrumental to more meaningful experiences. The desire theory and list theory do not seem useful in relation to the moral landscape because these theories are

restricted to a single kind of peak experience, preference satisfaction; it ignores the range of other possible experiences that satisfying different preferences will lead to, including those which consequent in more/less positive experience than the experience of preference satisfaction. Finally, the authentic happiness approach fails to recognize the factual nature of values by granting too much autonomy to individuals. Again, these four theories are likely not accurate representations of the peaks in the moral landscape. If they were representative of peaks, then the science of morality would be undermined, since its task is to scientifically reveal correct routes toward experiences characteristic of human flourishing and to teach these values to individuals. The goal is not to merely relay the information to individuals so that they can generate their own values. The final reason these theories are not compatible with a science of morality is because they may ignore the primacy of conscious experience in well being.

Given the nature of the moral landscape, and the conclusions about possible theories of well being outlined above, it appears that hedonistic, desire, authentic happiness, and list theories have serious shortcomings, which restrict their potential definition as peaks of human experience. To reiterate, these theories mistakenly assume that individuals are wholly capable of determining the correct values to improve human well-being, and/or the theories ignore the consequences of values on conscious states. However, Aristotle's notion of eudaimonia appears most plausible; it certainly recognizes the importance of subjective experience for well being, it agrees with the factual nature of the moral landscape by limiting individual generation of values, and it is particularly relevant to scientific study because it requires the examination of human capacities.

I argue that the peaks on the moral landscape most likely echo a careful consideration of human nature and the most fundamental human capacities as well as a degree of informed,

autonomous behavior; Aristotle's conception of eudaimonia appears to be the best approximation of the factual peaks on the moral landscape, which await scientific discovery. Further, this concept seems to most closely align with the goals of a science of morality, which is to study human capacities and prescriptively define their exercise. To clarify, the most worthwhile peaks of human experience are likely eudaimonistic; this concept is not exemplified by static states of existence in temporal isolation (simple points on the moral landscape). Rather, the concept refers to dynamic lived experiences (movement on the moral landscape), which only eudaimonistic activities can deliver.

On the moral landscape, it is essential to note that the heights of human well being are not static achievements; these peaks are fluid conscious experiences that reflect a specific kind of engagement with certain thoughts or behaviors. They represent conscious experiences that are not causally related to the values that led to them; rather, these peaks exemplify the conscious experience of eudaimonistic actions/thoughts in-and-of-themselves. In this way, our definition of values is expanded to include those that maintain an individual's position on a peak in addition to promoting movement in the direction of a peak.

Certainly, the individual will need to thoughtfully approach normative eudaimonia in order to custom fit these practices to his or her own life. Accordingly, peak experiences on the moral landscape likely require establishing the conditions for an individual to be able to think about and apply notions of eudaimonia first. For example, an individual will probably not be able to engage with Aristotelian values if he is being physically or verbally abused and continuously prevented from exercising the capacities he wishes to exercise. Alternatively, even if the individual is poised to engage with values, he will most likely not be able to reach the highest peaks of experience if he or she is not knowledgeable about eudaimonia. For example, the son of

a billionaire might have plentiful opportunities to engage in any conceivable eudaimonistic action, but without knowledge of them, the individual will likely be severely limited.

At this point, a reliable conception of peaks on the moral landscape becomes most clear: the most worthwhile human experiences are thoughtful engagements in autonomously customized eudaimonistic behaviors. Thus, well-being can be understood as intelligent, informed exercise of human capacities and fulfillment of human nature; well being is that movement through the moral landscape which is driven by values that are both eudaimonistic and autonomously modified. To be clear, the peaks on the moral landscape are most likely the conscious experiences of engaging in thoughtfully customized eudaimonistic values.

Accordingly, the values we are presently seeking now fall under three categories: values that are instrumental in empowering individuals to customize scientifically produced, eudaimonistic values to their own lives; eudaimonistic values that are intrinsically meaningful and whose exercise exemplifies increasing well being; and values that reflect both of these qualities simultaneously. Engaging in any of these values will drive individual movement toward, or maintain individual position at, peaks on the moral landscape.

At this point, it is important to clarify the sense of fulfillment that I have repeatedly alluded to. This sense is a particular kind of experience that only comes through exercise of modified eudaimonistic behaviors. Similar to Nozick's contention with the experience machine, I argue that this experience alone is not representative of a peak, but rather, the activity that brings this experience is of prime importance. This might lead to the question of whether there is only the basic feeling of fulfillment that is consistent across peaks, or whether this experience is transient depending on the activity. Because engagement with eudaimonistic values delivers the experience, and passage through the moral landscape is a continual flow of experience, certainly

there will be different degrees of experience quality as the individual comes closer to a peak on the moral landscape. Alternatively, because the experience is closely tied to eudaimonistic activity, different values will deliver different kinds of positive experiences leading to different peaks. Again, one eventual goal of a science of morality is to delineate these values to become fully knowledgeable about the character of different kinds of eudaimonistic experiences (different peaks) and understand the differences between them.

Additionally, it is important to note that the idea of fulfillment seems to imply a potential to be fulfilled. Again, this is not a problematic recognition, since the observation of human nature and understanding of natural capacities to identify eudaimonistic behaviors simultaneously acknowledges different kinds of potentials in individuals. Finally, it is essential to note that this feeling of fulfillment is not a kind of pleasure, as the Epicureans might contend. One reason is because the experience of fulfillment is not the ultimate goal of values; rather, values aim to promote specific kinds of activity that are fulfilling to engage in, so moral normativity should not be seen as exclusively instrumental in achieving meaningful conscious states.

Although defining two different kinds of values may have increased the responsibilities of a science of morality, I offer an efficient solution. Because the present work is among the first empirical projects that aim to scientifically define values, it seems reasonable to start with values that are both instrumental in promoting individual ability to thoughtfully engage in moral behavior and exemplify eudaimonistic meaning in and of themselves. Once the basics are laid out, the project can continue by potentially identifying instrumental and intrinsic values separately.

In the scientific search for values, three criteria have been established thus far. The first is that values must both exemplify eudaimonistic conceptions of the good life and enable individuals' independent, rational, and informed thought about how to best to engage with nature-fulfillment theories. Second, the values must stem from a scientific analysis of human needs and motivation in order to develop novel eudaimonistic ideas. Third, the values must be understood factually; that is, the connection between values, the human brain, and positive conscious experience should be clear. Finally, the values must be verifiable; an important benefit of a science of morality is that everything it asserts about values that lead to peaks on the moral landscape can be tested, which will enable identification of correct values to become increasingly more accurate and reliable with further research.

To satisfy these criteria, I propose that Abraham Maslow's hierarchy of human needs can reliably serve as our first set of values. Maslow was a psychologist who introduced "A Theory of Human Motivation" in 1943 in order to provide a comprehensive account for human behavior. At the time it was written, there was little sound evidence to support any claims about motivation, so Maslow utilized his own clinical observations to propose a theoretical framework which claims to categorize most, if not all, human behaviors in light of fundamental human needs. These categories are arranged in a pre-potent hierarchy, meaning that higher categories cannot be fully met unless lower categories are met. However, partial fulfillment of each level is still sufficient for an individual to begin satisfying the next level (Maslow, 370-396).

The first level of this hierarchy, and our first moral value, is the satisfaction of human physiological needs. This value represents the most basic of all human motivation and relates to the pursuit of physical health. Specifically, satisfying the physiological needs would include obtaining proper nutrition, regularly exercising with appropriate frequency and intensity, and

taking steps to prevent/treat any pathological illness. It seems likely that without these basic needs being met, a greater sense of well-being will probably not be possible. To clarify, the value here is satisfying physiological needs. This value moves individuals toward a peak on the moral landscape in that it primes them for being able to customize higher eudaimonistic behaviors in the future, and also engages them in the eudaimonistic act of satisfying physiological needs. Again, peaks on the moral landscape are conscious experiences, and so satisfying the physiological needs moves individuals closer to the experience of a peak.

The next level of the hierarchy, which, according to Maslow, will tend to dominate motivation only when the previous level is at least partially satisfied, concerns safety needs. Thus, our second value for the present moral system is the satisfaction of human safety needs. These too are fundamental drives that are at least partially pre-requisite for most other human activities. The safety needs involve establishing a lasting sense of security for the individual and at the very least, freedom from harm of any kind, including conflicts and dangerous situations. For example, safety needs may include establishing a shelter/home or securing a long-term prevention of dangerous/harmful external behaviors that affect the individual. Again, participation in this value would prepare the individual to be able to think about and apply notions of eudaimonia, and would engage the individual in the basic eudaimonistic value of securing of safety.

Continuing upwards, Maslow identifies the third level with sociality, and so the third value is the fulfilling of human social needs. These needs encompass the wide variety of social behaviors that humans may engage in. This level would likely include friendship, love, trust, communication, and many other forms of social interaction. As an instrumental value, sociality can certainly better equip individuals to intelligently think about eudaimonia, since open

communication generally fosters novel ideas and exchange of thoughts. Further, the exercise of social capacities would be engaging in eudaimonistic behavior, as sociality is deeply rooted in human nature. Again, the participation in the value of satisfying social needs will move an individual even closer to a peak on the moral landscape because satisfying social needs prepare an individual to critically apply higher eudaimonistic behaviors and engaging in this eudaimonistic act relays a conscious experience that is now even closer to our target peak on the moral landscape.

At this point in the hierarchy, it seems evident that the values are becoming more vague, meaning that the value is not specific enough to restrict variability in how individuals may satisfy the value. For example, different individuals might have different conceptions of how to satisfy the social needs, and apart from the biological markers that will be discussed in chapter two, there seems to be no way to identify specific practices that will genuinely satisfy the need. To resolve this problem, I propose that support from scientific disciplines is necessary to specify values just after their first identification, but prior to the micro-analysis that will be presented in chapter two.

To illustrate this point, consider the satisfaction of human social needs. Because we established that all eudaimonistic values must stem from observational studies of human nature, and because the satisfaction of social needs is a eudaimonistic value, any values that are more specific reformulations of the satisfaction of social needs should also be eudaimonistic and should also stem from consideration of natural capacities. Fortunately, primatologist Frans de Waal has studied primate behavior and especially primate sociality extensively, and because human behavior shares at least some degree of commonality with primates in general, his suggestions seem reasonable for the present value of social needs satisfaction.

In particular, de Waal has demonstrated that fairness and empathy are natural primate capacities, and play a regular role in primate social interactions. For example, De Waal defines empathy as a subjective experience that results from an understanding of another's experience, with the condition that there is a distinction between sense of self and other. In one experiment, he observed that the screams of a punished infant rhesus monkey cause others to approach or embrace it; here, the experience of one individual was relayed to others, and because there was a distinction between self and other, these monkeys knew that had to soothe the victim rather than soothe themselves. In another experiment, de Waal observed that capuchin monkeys will reject food offerings if they perceive that their peers are not receiving fair offers. Obviously, de Waal's conclusions stem from a much larger reservoir of experiments, but the point to recognize here is that fairness and empathy seems like reasonable values involved in the satisfaction of the social needs, and these values are particularly advantageous because they stem from an analysis of primate behavior (De Waal, 44-49).

The fourth level in Maslow's hierarchy deals with individual self-esteem, and so the next value for our moral system is the development of positive self-worth. This category is evidently more vague than the others, and will likely encompass an even broader range of behaviors that might satisfy it. Some examples of behaviors that may fit in this category include achievement of goals, validation/approval from self or others, or building of confidence in social situations. Again, the effect of satisfying this value on improving an individual's ability to think more clearly about the nature of eudaimonia should be quite clear.

Similar to the social values, it is not immediately apparent which kind of behaviors will rightly satisfy the self-esteem needs. Therefore, we must again draw from a scientific discipline that would offer a reliable route to improving self-esteem; this route must be more specific than

the satisfaction of self-esteem needs in general, but not quite as particular as the biological consideration of this value that will be offered in the next chapter. Psychologist Kristen Neff offers an interesting perspective on the present value that is worth considering. In one study, Neff highlighted the strong correlation between regular self-compassion and improved psychological well being. Specifically, Neff defined self-compassion as engaging in feelings of kindness, caring, or love towards oneself, which requires mindfulness of one's own suffering, and she cited examples of improved mental health, reduced stress, and the use of self-compassion to treat clinical illnesses such as depression, anxiety, and eating disorders as examples of the improved psychological well being that is correlated with self-compassion (Neff, 31-4).

In relation to self-esteem, she acknowledges that low-self worth is correlated with negative psychological outcomes, such as low motivation, depression, and suicidal ideation, though she seems skeptical that increased self-esteem will fully solve the problem. This is because the traditional view of self-esteem requires evaluation of competence and self-worth within a social world. Because this evaluation will not always be positive, due to the variety of ways an individual could feel incompetent in a world as socially complex as the Western one, the solution would be to modify our high regard for self-esteem as prerequisite for well being to incorporate self-compassion, which appears to be a much healthier way to form positive attitudes about oneself (Neff, 31-4). Therefore, not only has drawing from psychological studies enabled the present effort to specify the value of satisfying self-esteem needs, it has enlightened us about the potential downsides to traditional views of self-esteem and expanded our thoughts about this value. In this way, a science of morality will continue to improve with increasing knowledge about values.

The fifth and final level of this hierarchy is termed self-actualization. When describing this category of behavior, Maslow seems intentionally vague and limits his discussion by stating that this category represents those behaviors that fulfill the maximum potential of human beings. Part of the reason Maslow is vague is because he understands this category as highly relative to individuals, thus identifying the behaviors involves the individual fully learning about him or herself. This is the ultimate aim toward which the instrumental aspects of the prior values strive towards, but participation in this stage is not a specific value; rather, it encompasses the variety of potential values that an individual in this stage might engage in. Here, it is helpful to revisit a basic idea to solidify our progression thus far: the preceding stages of the hierarchy partially serve to establish the platform for individuals to think about and apply higher eudaimonistic notions and partially serve as indicators of what engaging in eudaimonistic behaviors feels like.

At this point, if the individual continues to engage in the aforementioned values, he or she will become stagnant at a point on the moral landscape just prior to the peak we have been striving towards. In order to reach this peak, the individual must utilize his or her newfound autonomous ability to customize and apply higher eudaimonistic conceptions to his or her own life. Therefore, once the first four values have been satisfied, a science of morality must supply higher conceptions of eudaimonia in order to further promote the individual's upward movement on the moral landscape and increase the likelihood that peaks of conscious experience become a factor in the individuals' well being.

One way that eudaimonistic values beyond Maslow's hierarchy can be identified is through correlation studies between habits and subjective assessments of well being. To illustrate this point, consider the long-established practice of mindfulness meditation, common to many cultures and religions, and especially prominent in Buddhist traditions. Essentially, meditative

practice involves drawing attention to the breath, maintaining emotional regulation, and developing an awareness of the mental stream of consciousness, the continual flow of thoughts, images, memories, and perceptions. Specifically, mindfulness meditation emphasizes focus on the present moment and disciplined regulation of mental function (Baer, 2003).

Recently, this practice has gained popularity in clinical interventions for mental health (Baer, 2003). For example, one study at the University of California, San Diego, measured the effects of an 8-week meditation training on clinically depressed patients. Overall, the practice led to statistically significant decreases in ruminative thinking and dysfunctional life attitudes, citing increased focusing ability and emphasis on the present as potential reasons of the effectiveness of the training (Ramel, 2004). Another study at UCLA similarly enrolled clinically diagnosed ADHD patients in an 8-week meditation program, and measured the effects on ADHD symptoms through laboratory tasks measuring attention and cognitive inhibition. The study found a reduction in self-reported ADHD symptoms, improved performance on the laboratory tests, and decreased self-reported anxiety and depression symptoms (Zylowska, 2007). Further, a study at University of California, Berkeley measured stress, rumination, forgiveness, and hope, in college undergraduates before and after two 8-week meditation programs. The students were randomly selected and a pre-test was administered to establish baseline data. The study found significant reductions in stress levels as well as increased self-reported forgiving actions (Oman, 2008). Studies such as these demonstrate a strong correlation between meditation and increased well being, thus, one potential eudaimonistic value beyond Maslow's hierarchy could be mindfulness meditation.

With a set of moral values defined, it is now important to consider the reasons these particular values serve as a good starting point for a science of morality. First, the offered set of

values agrees with our intuitive sense about what might be required to promote well-being. Second, it is a foundational set, thus it is definitely possible for other conceptions of well-being that may utilize this one as a scaffold. For example, the four values that stem from Maslow's hierarchy supply the basic conditions for an individual to reach a potential peak, but whether or not they achieve that peak depends on the application of further values that build upon the hierarchy within the stage of self-actualization. Third, analyzing behavior into specific components allows for a measureable and concrete way to approach the abstract notion of value. As we will see in chapter two, each of the values just discussed can be understood in terms of factual concepts, and by recognizing this, we will be able to establish objective measures that specify what each value requires in order to be satisfied. Finally, these values stem from a scientific examination of human nature and maintain the importance of natural human capacities, so these values agree with the eudaimonistic criteria offered in the first half this chapter.

The values established thus far lend themselves to a factual understanding through neuroscience. The components suggested in the theory are all components that have already been addressed by neuroscience in quite specific ways, thus starting with such a foundation will allow for a rapid and pragmatic advance for moral science. However, in theory, any conception of well-being should eventually be able to be understood in terms of neurobiology. It goes without mentioning that the reason this particular conception was used rather than another was because this conception stems from a scientific analysis of human nature, allows for informed autonomy, and satisfies the criteria for peaks on the moral landscape (a peak is the conscious experience of fully engaging in a customized, eudaimonistic behavior). Again, the present moral science derives its values and its function from empirical observations about humans and human behavior rather than from a transcendent or abstract source.

At this point, a starting point to approaching well-being has been defined, thus the next step is to characterize our conception with neurobiology. Again, the purpose of this effort is to define moral values as grounded in scientific fact. Before this effort can continue, however, I must address several important counter-arguments.

The first comes from people who accept the offered value system, but are skeptical of its contribution to the moral system discussed by Sam Harris. Because Harris' system rests on the premise that values can be understood factually, the main skeptics at this point are those that don't already notice the factual nature of our established values. This doubt can be easily resolved by thorough consideration of the pending transition from identifying values to factually elucidating values.

It was already established that the values at the center of the present moral system must be derived from observations about human capacities. A consideration of human motivation and behavioral psychology helped to delineate Maslow's hierarchy, which led to the identification of four moral values. In chapter two, these values will be reiterated in biological terms. Thus, science begins the moral project by identifying human values through observational consideration, and continues by deconstructing the experience of these values through neuroscience. The four values from Maslow's hierarchy are the most basic of values, but it is entirely possible that even these four can be understood in different ways. This is why a specification in terms of biology is necessary; it can produce accurate values that will, in fact, move individuals toward peaks on the moral landscape as long as the values are specific enough for individuals to concretely follow.

The second major counter-argument comes from a lack of satisfaction with the proposed values. Some will argue that there is probably more to reaching the peaks of conscious

experience than what was extracted from Maslow's theory, and these people are certainly justified in their effort. However, this is less of a counter-argument than an observation about how a science of morality operates. I am not proposing that the proposed values are sufficient to reach all peaks of conscious experience on the moral landscape. In fact, I am not even proposing that they are necessary in order to move towards a peak. Instead, I am claiming that these values allow for a pragmatic starting point because of their dual instrumental/intrinsic function and because they are already understood at a neuronal level. However, there are likely many more peaks and many more values that help individuals move toward these peaks that await our discovery.

The third major counter-argument is a more extreme version of the second. Some might argue for a different conception of well-being that is entirely unrelated to the proposed values. These people might also argue that our intuitions about well-being are incorrect, and may suggest completely alternative conceptions of human flourishing that may seem to contradict common sense. Though this creative/innovative method appears to go against the current analysis, I whole-heartedly welcome such attempts. In fact, a science of morality can greatly benefit from creative suggestions about different ways to conceive the peaks of human experience.

It seems certain that Maslow's hierarchy of needs is not the only suggestion for behaviors that improve well-being. Apart from our intuitive thoughts about which values promote well-being, it is entirely possible that there are modes of behavior that we are wholly unaware of that promote well-being in unknown ways. For example, the present model of flourishing does not involve any degree of suffering. However, many people might argue that suffering, delivered in manageable doses, can increase an individual's discipline, improve

character, and prepare the individual for future hardships. Notice that even questions as complex as this must involve facts and consequences on well being.

As such, we might be able to empirically investigate whether suffering, in the proper contexts, can lead to peaks through longitudinal studies in which children living in environments conducive to the experience of pain/negative affect are observed over many years. Specific biological processes can be monitored and associated with behavioral effects, which might indicate the long-term effects of negative experiences on subjective assessment of well being. A more likely route will involve a greater understanding of brain development. For example, if a certain value is proposed to aid individuals reach a certain kind of peak, then the experience of this peak can be biologically detailed; the factors which might influence the particular development of neural structures that allow for this experience, should, in principle, be discoverable. If the experience of suffering is necessary for a particular kind of peak experience, it should be clear via neuronal development. Still, this is only speculation, but there are factual answers, in principle, to these kinds of questions that science will be able to address through a comprehensive understanding of the moral landscape.

The point here is that it is possible to think of other, non-intuitive patterns of thought/behavior that can actually increase well-being. Thus, in order to propose alternative conceptions of value that are entirely unrelated to the ones already offered, an individual must demonstrate that it improves well-being (it is instrumental in preparing the individual for self-actualization, intrinsically eudaimonistic, or both) and that it can be construed in terms of neurobiology. If a novel characterization of human flourishing meets these criteria, then it can be considered for incorporation into the present moral science.

The preceding counter-arguments all reveal the necessity for future work to be done on comprehensively understanding what it is that brings human life toward the peaks of experience. This work may or may not involve examining the components of Maslow's theory discussed previously, but will certainly have to be grounded in data in order to be seriously considered. Evidently, well-being is quite a complex concept and we are only beginning to understand it from a scientific perspective. The values discussed previously offer a way to begin the exploration of human flourishing, though there are likely other values that either build upon or fundamentally differ from the present conception. It is important to emphasize here that complete answers are not needed for a pragmatic working system.

What I mean to convey is that the present effort is one of learning. Although the goal of a science of morality is to eventually elucidate moral values factually, the task is performed through discovery. The efforts of moral science improve with time and data, and these efforts consequent in both expanding the existing conceptions of well-being and discovering any new ones. In particular, the science of morality will have to connect the scientific recognition of values, generated via observational study, to the scientific explanation of those values, developed via known biological mechanisms.

Chapter Two: Defining Values

The goal of a science of morality is to identify values that improve human well being and to understand those values at the level of the brain. Thus far, we have considered possible conceptions of well being in order to gain a sense of what morality aims toward. As a result, we were able to identify two sets of moral values at the end of chapter one.

The first set of values, drawn from Maslow's hierarchy of human needs, involves establishing the basic conditions for living well while simultaneously preparing individuals to independently think about and apply higher eudaimonistic values to their own lives. The second set of values are the higher eudaimonistic ideas that build upon the first set of values. The sole eudaimonistic value that was offered was mindfulness meditation, although there are certainly a large number of other eudaimonistic practices that await our study. Again, to clarify definitions, values are those behaviors/practices that promote well being; these behaviors exemplify a meaningful life insofar as they relate to the exercise of natural human capacities (intrinsic values), and/or empower individuals to critically think about and apply the eudaimonistic values to their own lives (instrumental values).

The task of the present chapter is to further specify the two sets of values mentioned above by exploring their biological foundations. Here, it is worth detailing the portions of the two sets of values to be considered. The values that stem from Maslow's hierarchy are: satisfying human physiological, safety, social, and self-esteem needs, and these values seek to encourage individuals toward the stage of self-actualization. The single eudaimonistic value that will be explored is meditative practice, with the ultimate goal of achieving some degree of intrinsically meaningful experience.

The factual explication of values is a necessary step in the progression of a science of morality because of three reasons. First, the philosophic foundation of the moral landscape rests on the premise that the difference between the good life and the bad life can be understood through a full understanding of the human brain. While Harris assumed this to be true in his defense of the moral landscape, it is still necessary to empirically demonstrate this assumption for the sake of bolstering the argument. Thus, moral science needs to validate the factual basis of well being in order to function as a working ethical system. Again, our present ethical framework is one of consequentialism and realism; that is, there are right and wrong answers as to whether values improve well being, and these answers are anchored in the scientific facts that characterize the moral landscape. One aim of the proceeding discussion is to demonstrate this statement.

The second benefit of a factual understanding of moral values is that reliable measures can be developed to assess well being and these can be used to inform moral judgments. For example, when considering the harmful effects of genital mutilation, a much stronger moral argument can be made for, or against, this practice by using specific factual measures rather than an intuitive grasp of well being, which will naturally be more ambiguous. For example, genital mutilation may have long-term consequences on reproductive hormone production, ability to form trust relationships with the opposite sex, or maintenance of positive mental health.

These effects certainly occur via biological mechanisms, and so the short-term and long-term consequences of genital mutilation on well being can be assessed by using empirical measures; this evaluation provides the most reliable reasoning for any sort of judgment about whether genital mutilation is a good practice. Most likely, understanding values at the neuronal level is not always necessary to determine the general effect of a practice on an individual's well

being (upward or downward movement on the moral landscape), so the purpose of defining the neuroscience is to provide specific and objective, rather than vague and intuitive, reasoning to make a more reliable evaluation. Those individuals who might argue that genital mutilation is culturally relative could not be more wrong; factually defining the consequences of this practices on well being will validate this assertion. Once we begin using facts to anchor our thoughts about how to maximize human flourishing, the reasoning behind our decisions will become most clear and most objective, and so we will be better prepared to communicate about ethical issues.

The third reason values must be understood factually is because this process will expand our knowledge about human well being and facilitate the creative generation of new values. A thorough understanding of a given value might reveal that its underlying biological mechanisms play a similar role in other kinds of behaviors that we may not have been aware of. For example, if developing trust in other individuals were shown to be a value (hypothetically speaking), then analyzing the biology of trust would reveal oxytocin to play an important role. Continuing to learn about the mechanisms of oxytocin would reveal that it is heavily secreted during childbirth and during breastfeeding (Lee, 127).

Here, learning about a single (hypothetical) value leads to determining mechanisms that may play a role in other behaviors, which expands our thoughts about trust and forces us to consider the role of child-rearing in trust development or its effect on well-being in general as an independent value. Learning about proposed values at the level of the brain can offer new insights about human behavior and potentially commonalities with other behaviors/practices that appear to be unrelated; these commonalities can either aid in our creative generation of new values to study or strengthen our understanding of the original values which revealed them.

It is worth mentioning here that this project does not rely on a reductionist view of consciousness, which asserts neurobiological processes as sufficient, rather than simply necessary, to produce the human experience. Clearly, the relationship between mind, body, and the phenomenal world is itself a highly debated philosophical issue. The premise that is important for a science of morality is that, at the very least, conscious experience is accurately reflected in the activity of the brain. In the scientific community, there is little doubt that our lives, our thoughts, emotions, and behaviors, are fundamentally tied to functional neurobiology, but acknowledging this fact still leaves open the possibility for other models of conscious experience, such as those presented by metaphysical or phenomenological schools of thought.

The argument to follow rests on the humble assumption that there is a reliable degree of correlation between experience and neurobiology; in the unlikely event that science discovers that conscious experience and neural activity are, in fact, mutually exclusive processes, my argument will not hold. Therefore, identity theory, which addresses the relationship between mind and body, is also an object of study that will need to be considered in order for a science of morality to stand independently. However, this conversation is a philosophic one that presents itself logically prior to the exercise of moral science, and so it will not be further discussed in the present work.

At this point, the empirical definition of values can begin. The first stage of Maslow's hierarchy, and the first value of the nascent moral science, is the satisfaction of human physiological needs. Specifically, I argue that meeting these needs entails consuming a nutritious diet, engaging in regular exercise, and preventing or treating any potential diseases. Notice that the nature of these components of the present value is quite visibly factual. The biochemistry of human metabolism is very well studied, and there is a large degree of consensus

about which raw materials are necessary for optimal bodily function. According to the Harvard School of Public Health, a nutritious diet consists of appropriate quantities of carbohydrates, protein, fats, fiber, fruits/vegetables, dairy products, a lower sodium and sugar intake, and a proper supply of vitamins and minerals; nutrition is even studied to the extent that the quantities of these materials in various foods (nutrition facts labels) are made clear (hsph.harvard.edu). In addition, based on study of human physiology, the Centers for Disease Control (CDC) offers guidelines on two essential components of exercise, aerobic and muscle-strengthening, and details the appropriate frequency and intensity for different kinds of individuals (cdc.gov). Finally, the field of allopathic medicine is obviously a well-established scientific discipline, and the factual nature of preventing or treating disease should be self-evident. To summarize, the value of satisfying human physiological needs can be analyzed into components, and each component can be met via specific, empirically-based, and measurable behaviors.

It is important to note that I am not trying to comprehensively delineate the biology of human physiological needs, nor will I be attempting to do so for the remaining values, although these are certainly goals for the broader science of morality. For the present argument, I am trying to convey the idea that moral values can be understood factually, where values are those thoughts and behaviors that bear on human well-being. Though defending this idea in relation to human physiology might seem trivial, the biological specificity with which we can understand this stage best exemplifies the intention for the values to follow. Since the remainder of the proposed values involve a greater degree of subjectivity/ambiguity, the well-known empirical foundation of the value of meeting physiological needs can serve as a model for the proceeding discussion.

The second value for the present moral science involves meeting the human safety needs, the second stage of Maslow's hierarchy. One approach to factually understanding this value comes from noting the primacy of the physiological stress response during the absence of a sense of safety. The biological features of the experience of imminent threat or danger are well understood, and can be split into two broad categories: acute stress reactions and states of chronic stress. These categories are representative of unsatisfied safety needs, thus, the physiology of the stress response offers at least one empirical gauge (and there are probably others) of whether this value is being satisfied.

The acute stress reaction involves short-term physiological responses to threat, and operates through the hypothalamic-pituitary-adrenal axis. Essentially, once a danger is perceived, the hypothalamus becomes active, sending corticotropin releasing hormone to the anterior pituitary gland. Consequently, adrenocorticotrophic hormone is released, activating the adrenal cortex and triggering the distribution of cortisol throughout the bloodstream. Simultaneously, the hypothalamus activates the sympathetic nervous system (the "fight or flight" response), which prepares the body for rapid muscular action. This response involves the neurotransmitters adrenaline/noradrenaline, and its effects include increased heart rate, dilated pupils, inhibition of digestive processes, constriction of blood vessels, and accelerated respiratory processes. To reiterate, the short-term response to danger primarily involves (but is not limited to) the action of cortisol and adrenaline/noradrenaline to generate an appropriate response to the danger, which usually leads to either aggressive/confrontational behavior or fleeing the situation (McLeod, 2010).

It is immediately apparent that the stress response likely results from a perceptual act that triggers this response. Though it may seem like we first have to understand what is actually

dangerous, I argue that this is not necessary because of two reasons. First, a portion of our innate response to danger stems from non-reflective, evolved intuitions. For example, we have an aversion to rotten food, and we are intuitively aware that we need to flee from a charging rhino. Second, there is certainly a portion of the danger response that interacts with learned behavior, and these mechanisms are more sophisticated. Although these kinds of values are more difficult to pin down, a developing science of morality will be able to elucidate the complexity. Still, for the present work, it is not necessary to fully delineate short-term responses, because the more relevant concern involves long-term stress responses.

Another feature of the discussed biology is to notice is that not everyone has an optimal stress response. Therefore, the consequences on the safety needs of individuals in terms of stress response may vary, depending on individual nuances in physiology. It may appear that establishing a measure of benefit and harm via hormone levels or particular stress activation pathways might be unreliable because of this variability, but there are two reasons why this is not a problem. First, the variability between individuals is not so great that hormone levels vary to an unpredictable degree. Neuro-endocrinology is an established field of study and is certainly aware that it needs to account for individual differences. Second, if the stress response is notably variable in certain individuals, there will be other consequences on the individual's well being if the danger is real. We are not restricted to measuring the stress response in order to gauge the safety needs, as there will certainly be other consequences on conscious experience that can be measured, either related to the safety needs or any other needs in the hierarchy.

The stress response can manifest in more serious conditions when chronically active. In particular, there are negative, observable physiological consequences to long-term activity of the sympathetic nervous system and to continual distribution of cortisol throughout the body. For

example, because sympathetic nervous system activation involves suppression of parasympathetic nervous system activity, effects of prolonged stress will typically result in slowing down of the digestive system, constipation, anorexia, difficulty urinating, erectile dysfunction, and increased susceptibility to infection, due to the additional long-term suppression of the immune response (Gleitman, 2004).

Further, mechanisms of physical stress may affect mental health as well, and can lead to clinical illnesses such as post-traumatic stress disorder or panic disorder. These disorders, which are highly subjective illnesses, have already been approached scientifically; the “Diagnostic and Statistical Manual of Mental Disorders” (DSM) is the reference used by psychiatrists for a reliable assessment of mental disorder, as it was developed within an empirical framework and stems from the scientific study of mental health. Finally, the long-term effects of abnormally high cortisol secretion may manifest in other physical symptoms, including decreased cardiovascular health, increased risk for obesity, or hindrance of sexual maturation. Here, it is essential to recognize that the long-term effects of unsatisfied human safety needs are also factual and measurable. While these symptoms do not necessarily indicate chronic stress, their presence is, at the very least, cause for concern, especially since they interfere with the full achievement of our first value, the satisfaction of physiological needs. To summarize, the biological symptoms and mechanisms of the human stress response can serve as one approach to better understand how human safety needs might be met, and this empirical knowledge can be used to identify even more specific values that alleviate the stress response (Gleitman, 2004).

At this point, one might consider the possibility that stress can sometimes be a beneficial experience. Stress might heighten attention, improve performance, or build mental stamina in dealing with difficult situations. To address this concern, it is helpful to revisit the difference

between identifying values and theorizing about the good life. Values are the behaviors or thoughts that drive movement towards the peaks of experience on the moral landscape, while the good life is the totality of movement on the moral landscape driven by specific values. There are numerous opinions about what kind of total movement along this moral landscape the good life might involve, but the fact that moral values must drive this movement is clear. It is not necessary that values are restricted to upward movement to reach a peak on the moral landscape; some values might require passing through a valley in order to move towards a higher peak. Since the sum of movement on the moral landscape driven by values represents the good life, particular positions on both valleys and peaks seem like sensible possibilities.

As our knowledge of the moral landscape improves, we will increasingly understand the potential paths individuals may follow on it, and we can better discuss how valleys of human experience can be instrumental in reaching certain peaks. The discussion of safety needs, then, is restricted in the sense that it asserts that satisfying safety needs is only one particular route towards our defined peaks, and does not take into account the role of unsatisfied safety needs to potentially support a different kind of upward path that we are presently unaware of. The broader science of morality seeks to uncover all various paths towards peaks by illuminating the moral landscape, which may or may not involve downward movement. This will require a more comprehensive understanding of values and their potential outcomes, and may or may not involve considering different kinds of peaks that may be on the moral landscape. However, for the present work, we will continue to focus on the particular values aiming toward the kind of peaks defined in chapter one, acknowledging that there may be others yet to be considered.

Our next value consists of satisfying human social needs. This value, which represents the third stage (following physiological and safety needs) in Maslow's hierarchy, can be quite

difficult to approach from a biological perspective, as sociality is a complex phenomenon and entails a wide variety of behaviors that might contribute to the fulfillment of this stage. However, I want to emphasize once more that the goal of our present effort is not to comprehensively delineate all of the behaviors that might satisfy human social needs, though this is still an important task for the broader moral science. Instead, the more prevalent long-term goal is to identify the neural features that social behaviors have in common. At present, it is possible only to cite some examples in which progress has been made. In this way, a factual basis for the present value can be established, and it can be used either as a reference to gauge proper fulfillment of the social stage of Maslow's hierarchy or to generate novel methods that might satisfy the human social needs.

To clarify this claim: a science of morality can be seen as a continual specification of values. Starting from the most general value, the promotion of well-being, the present value is specified as satisfaction of social needs. This can be further detailed by first understanding the neural correlates to social experiences and later identifying even more specific values (some of which may be intuitive while others may be unbeknownst to us) which epitomize, at least in some sense, the offered correlates. This again opens the possibility for further empirical analysis of the new values, which might once again lead to enhancing the set of neural correlates. These can again be used as either a tool of reference or further value identification. In this way, a science of morality can improve with time and further empirical research.

To identify those components of sociality that are particularly susceptible to scientific understanding, it is helpful to refer to the philosopher Patricia Churchland. She has constructed a scientific account of human social behavior, drawing on research from evolutionary biology, neuroscience, and genetics/epigenetic to support her argument. According to Churchland's

research, there are three major brain processes that contribute to sociality: the urge to care for self/others, the capacity to evaluate and predict consequences of social actions, and reward/punishment system linked to application of social practices.

Churchland argues that the urge to care for others in mammals is an extension of the innate care-for-self that characterizes evolutionarily successful organisms. The neural mechanisms underlying this self-care likely expanded due to the historical appearance of the neurotransmitter, oxytocin, which has been shown to play a significant role in many kinds of caring behavior. This biological modification must have conveyed some adaptive advantage to the individual, since there are a variety of costs to caring. In addition to caring behavior, arginine vasopressin probably increased the complexity of social interaction, as this molecule has been strongly associated with attachment and pair-bonding. Churchland asserts that the basis for sociality fundamentally involves the modulation and activity of these two molecules. Still, sociality consists of the capacity to evaluate and predict the effects of potential social behavior, and this is likely related to the neural circuitry of learning. Further, the reward/punishment system, which may have developed independently, became connected to the host of other social behaviors, and fundamentally operates via the neurotransmitters dopamine and serotonin.

It is important to note that Churchland does consider the large variety of ways these mechanisms can interact, and she emphasizes that sociality can take on many forms. This should be evident in considering the diversity of possible human interactions, especially in the Western world. Churchland argues that the particular form of sociality that mammals may adopt (and indeed, there are many) depends on environmental conditions, food resources, and capacities of individual, and although detailing each of these processes is beyond the scope of the present work, it is worth mentioning because these are factual concepts; the interaction of biological

mechanisms to produce particular forms of sociality is open to scientific understanding, despite its intimidating complexity.

Scientific progress in understanding our value of satisfying social needs is not undermined by the intricacy of this kind of behavior, but it may be slowed down. Though the precise details of these facts and their influence on social behavior may be difficult to attain at the moment, these facts certainly exist, in principle, and knowledge about the significance of these facts will likely begin to be developed in the near future. Most likely, our knowledge will begin by exploring the mechanisms of basic social behaviors, such as nurturing newborns or copulation in long-term marriage, and moving on towards more complex behaviors, such as extraversion/introversion, social rank in friendship groups, or biological responses to unfamiliar social situations.

Admittedly, our discussion about the biology of human sociality is at most a feeble attempt at adequately defining the value of social satisfaction, partially because such a complex concept requires a much more lengthy treatment and partially because there is still some debate about the mechanisms of social experience in the scientific community. At the very least, our conversation has indicated that even such a complex feature of conscious experience is susceptible to scientific analysis. As research continues, our understanding of this value will increase, and we will be able to elucidate the biological basis of social behavior in the near future.

Developing a high sense of self-esteem, a value that represents the fourth stage of Maslow's hierarchy, is our next focus. Although there are a large number of factors that can contribute to this development, our approach will be similar to that for the previous values. That is, we seek to elucidate the underlying experience of self-esteem, the positive/negative affect that

is associated with an individual's sense of self, and to use this knowledge either as a reference to gauge whether the self-esteem needs are being met, or as a tool to identify further behaviors that are neurologically consistent with the established basis for the experience of self-esteem. By this point, this style of value specification that characterizes our science of morality should be most evident.

The subjective experience of self-esteem can be understood as a temporally fluid evaluative sense that relates to social rank, though there may be other ways to define this experience. Serotonin is a neurotransmitter that is known to be released in response to positive social feedback and is correlated with high self-esteem, while negative social feedback and failure to achieve are associated with inhibition of serotonin receptor binding, lower self-esteem, and increased violent behavior. In particular, one study demonstrated that serotonin levels fluctuate in response to social feedback, and correlate with the subjective evaluation of self-worth. Interestingly, these effects are not permanent, and eventually return to basal levels, indicating that continuous positive social feedback, successful achievement, or increased reputation are required to maintain a high sense of self-esteem (Wright, 1995).

Because this stage of the hierarchy can be understood as an evaluative sense, its biological nature is much easier to recognize than the complex phenomenon of sociality. Still, self-esteem involves significantly more neural processes than we have described, but the point to realize here is that the value of satisfying self-esteem needs, like our other values, has the potential to be understood factually. This recognition, along with the recognition that science can find moral values, are all that are needed to make my defense of moral science successful.

The final stage of Maslow's hierarchy, self-actualization, represents the fully revealed potentiality of an individual to think critically about achieving his or her own sense of

eudaimonia. This stage is unique from the others in that it is the most subjective, and the conscious experiences which may be involved during self-actualization can significantly vary from person to person, because this stage is entirely dependent on the individuals' application of eudaimonistic knowledge to his or her own life, which will likely involve a capacity to critically think about values as well as a capacity to engage in thoughtful discussion with others about the good life.

At this point, it would seem sensible to consider satisfaction of self-actualization needs an umbrella value, since this stage was the instrumental goal of the prior values and serves as a platform for engaging in higher eudaimonistic values. As mentioned in chapter one, in order for self-actualization to become more than just a stage of preference satisfaction, science needs to supply higher, eudaimonistic conceptions of well being that the individual in the self-actualization stage can draw from to craft a meaningful life. Without these conceptions, the individual will not be able to exercise the correct eudaimonistic practices that lead to upward movement on the moral landscape; instead, he or she will form their own preferences about the best ways to live in addition to the values in Maslow's hierarchy, which may or may not be correct in guiding movement toward peaks, either for themselves or for those around them. Thus, the value, self-actualization, includes within it the more specific, eudaimonistic values that build upon those in Maslow's hierarchy, and the ones that will be thoughtfully personalized by self-actualizing individuals.

One eudaimonistic value, described in chapter one, centers around meditative practice. The value at hand, meditation, is slowly but surely beginning to be understood by modern neuroscience. Again, our goal is to define this value scientifically, meaning that the act of

meditation itself, rather than the long-term benefits of meditation, must first be empirically outlined.

Studies suggest that experienced meditators, when monitored during a meditation session, display a variety of physiological changes during the session. First, there is a notable decrease in sympathetic nervous system activity, and associated increase in parasympathetic activity, which contributes to a peaceful state of rest. Further, there is a general decrease in overall neural activity during meditative states, along with a continual maintenance of a positive state of wakefulness. Finally, brain regions involved in internalized attention and emotional processing show increased activity (Lutz, 2008). The point to recognize here is that the value of meditation promotes a conscious state that can be recognized at the physiological level. This will enable us to discover novel methods to achieve similar conscious states, thereby increasing our knowledge about the character of this particular category of experience.

Interestingly, a popular method of studying the effects of meditation on the mind during meditative practice focuses on measuring brain wave activity. Several studies show increases in low-frequency theta and alpha waves, neural patterns that are usually observed during sleep or deeply relaxed states. The subjective experience of thoughtless awareness has also been characterized. In particular, a study of long-term Sahaja Yoga meditators revealed particular patterns of electrical activity in the brain, measured via electroencephalography (EEG), that strongly correlate with the subjective experiences of self-reported positive states (Saggar, 2012)

The empirical consideration of meditative practice offers much insight for a science of morality, aside from its obvious demonstration that the value of meditation can be understood factually. Specifically, studies of meditation demonstrate that subjective experiences strongly correlate with particular patterns of brain wave activity. During the entirety of our discussion of

values, it is apparent that much of our knowledge relates to biochemical activity. Though this provides a way to understand the mechanisms underlying certain behaviors, it leaves much to be desired in terms of measurement of subjective experience. In contrast, the scientific study of meditation is heavily based on EEG data, which is quite beneficial because it reveals particular patterns of brain activity that strongly correlate with self-reports of subjective experience. To specify the benefit of this even further, our discussion of meditation reveals a plausible way to empirically approach the complexity of subjective experience, which is the focus of our science of morality.

At this point, the values proposed in Chapter One have been analyzed and their empirical basis has been established. The first principles in Harris' moral landscape have been substantiated, since the discussion supports the proposition that values can be scientifically identified and neurologically understood. This exercise defines the nascent moral science as an authentic empirical project and offers its first moral values; the preceding discussion constructed the earliest definitive foundation for the present moral system. However, before the discussion can continue into normative application, I must address several counter-arguments.

The first critique stems from a resistance to analyzing conscious experience into biological components. Some people might argue that an idea as complex as human well being cannot, in any way, be understood in terms of neurons, chemistry, and electrical activity. These people will insist that a scientific approach is not suitable for ethical discourse. However, I see no alternative approach to fulfilling the goal of factually understanding values. If the direct connection between mind and brain is true, the range of all possible human experiences must have factual correlates; that is, the moral landscape is an empirical space susceptible to scientific exploration.

Admittedly, it is improbable that this landscape will be comprehensively understood by science, at least in the foreseeable future. The important point to consider here comes from Harris' discussion. There are correct answers, in principle, to moral questions due to the factual nature of the moral landscape. Even though answers in practice may be difficult to attain, there is no reason to not pursue them. Again, a science of morality is an effort to improve global and individual well being by learning about the brain.

Our goal is to identify correlates to well-being that are measureable and observable, for the benefit of moral decision-making, rather than to reduce the complexity of conscious experience to simple biology. Though critics might accept this view for the most obvious of moral cases, many will be skeptical of my method for more serious moral issues. In response, I argue that these kinds of problems necessarily involve factual consequences on well being, so there are answers, in principle, that await scientific grasp. This suggests that the reason certain issues are unclear is because we are simply ignorant of the full short-term and long-term consequences on well being, which are facts that await our finding.

Though I am not asserting that science can resolve all of these issues, the moral landscape will become an increasingly more effective tool for moral discourse as it develops. At the very least, science will be able to anchor any moral conversation and provide a way to objectively and reliably reason about how best to improve the well being of the individuals involved. At best, it can provide all of the relevant facts and ways to measure fulfillment of values that might enlighten us about the potential outcomes of our behavior, thereby enabling rational, objective solutions. Again, this science of morality is not for moral puzzles or armchair dilemmas; it is being developed to provide a reliable route to addressing real problems with real consequences on conscious creatures.

Another counter-argument might be that the discussed biology does not adequately support the values proposed in Chapter One. In fact, this is a very plausible proposition. However, it is important to recognize that the current work is naïve. This investigation represents the first steps of a science of morality as a functional normative system and pronounces its evolution from an exclusively philosophical framework. It is likely that the preceding discussion does not fully explain the proposed values, but the purpose was to demonstrate that this kind of work is possible and indicate how it might operate; the aim was to empirically substantiate Harris' philosophic argument. Now, the task of moral science will be to expand on the first steps offered in the present work and develop into a much more detailed, critical, and comprehensive endeavor that can better identify and understand moral values.

This counter-argument reveals a particular feature of moral science that is worth noting. Like the other sciences, the present work is subject to critique, peer-review, and significant revision, so the scientific discussion of moral values can only improve from here. In this way, the newborn moral science can evolve into a proper academic discipline. Here, the purpose of my work is again made evident: I aim to begin the formal science of morality by offering a basic demonstration of how moral science might operate, and it is my hope that this exercise, a scientific search for values, will be sufficient to ignite and sustain its future development.

The third and final counter-argument asserts that defining values empirically is either unnecessary or irrelevant to a moral discussion. I admit that the connection between the neuroscience discussed above to a normative ethics is vague, so it seems fair to hold this position. However, it is important to realize the purpose of this effort. As stated in the beginning of the chapter, defining values factually has three main benefits: substantiating the philosophy of moral science, constructing an empirical standard which can be used to make moral judgments,

and opening doors to novel values or a better understanding of neural function in relation to conscious experience. Now that moral values have been identified and scientifically explored, the present work must delineate the utility of this information for ethical application. Once this has been accomplished, Sam Harris' argument for the possibility of science discovering moral values will have been substantiated. In addition, the science of morality will begin its life as a functioning ethical framework.

To conclude, I want to highlight the potential for future scientific research and offer my suggestions about how this work might relay into a science of morality. Clearly, we have much to learn about the human brain and its connection to the conscious mind. However, modern neuroscience has gained significant momentum in the second half of the twentieth century, and there is an impressive variety of work being done to approach the complexity of the brain. Much of the work that is going into this effort is the objective revelation of the moral landscape. That is, modern neuroscience is becoming increasingly more reliable in comprehending this hypothetical space, and its output is an objective reservoir lacking any moral valence.

On the other hand, there is significant effort from other fields, including psychology, sociology, or anthropology, that specifically focuses on the science of well being, as partially indicated in chapter one. This work also operates independent of the moral science, and seeks to scientifically approach the concept of well-being, rather than its factual basis. As made evident thus far, the task of the science of morality is to draw on both of these disciplines to connect knowledge about well being and knowledge about the human brain in order to produce moral values. Therefore, the science of morality will run parallel with the advancement of modern neuroscience and the scientific study of human flourishing, drawing from these disciplines for

creative generation of values, and building on itself, as any authentic science would, for autonomous improvement of understanding.

Chapter Three: Science and Morality

Until recently, the relationship between science and morality has been unnecessarily restricted, partially due to the illusory divide between facts and values and partially because of our failure to recognize that morality fundamentally concerns the well-being of biological creatures, which suggests that there are biologically-based right and wrong answers to moral questions, or at least, to general questions about what makes human life better or worse. However, many scientists and philosophers have already started the conversation between the two seemingly disparate disciplines.

In his book, Sam Harris identifies the three projects science must take on if we are to gain a better understanding of human morality and develop a functional, relevant, and readily applicable system of ethics. The first project, and the one that has gained the most momentum thus far, investigates the evolutionary origins of moral behavior and the biological mechanisms underlying moral decision-making. The second project is “moral science”; Harris and many others (including psychologist, Steven Pinker and theoretical physicist, Lawrence Krauss) have already defended the philosophy supporting this project, and this thesis represents its first steps as an empirical process. Finally, the third project addresses existing moral convictions in society and seeks to fully educate individuals about the moral landscape and the values that will continue to be identified by a science of morality. Even though Harris asserts that this third task is the most important of the 21st century, this task has not yet been approached scientifically (Harris, 2010).

The discussion thus far has provided the foundation to carry out project two; scientific conceptions of values have been proposed and reformulated in terms of neurobiology. However, this effort is fully realized as a science of morality only when the knowledge has been

implemented into ethical dialogue; to clarify, moral science must integrate its found values into global discourse for the sake of moral progress. Now that an objective basis for evaluating individual well-being has been established, moral questions can be judged against standards for human flourishing and suffering.

One possible route for integrating scientific knowledge into normative discourse is through simple reiteration. For example, consider the possible normative applications of the values derived from Maslow's hierarchy. Through neurobiology, we have gained a factual grasp of the physiological, safety, and social needs and we have considered how this knowledge might be relevant to ethics. For example, it seems fair to assert that a continually hyper-activated sympathetic nervous system or chronically impaired oxytocin regulation are likely not features of a happy individual, partially because these features indicate that the safety and social needs are not fully being met. Thus, reasons for these undesirable consequences might be that the individual is subject to intentionally immoral behavior or is mistakenly engaged with the incorrect values that should increase well being.

It is more likely than not that these features of conscious experience, if left unattended, will lead the individual toward valleys rather than peaks on the moral landscape. Accordingly, moral action might involve intelligent response to these biological markers in order to get the individual on the correct path towards peaks. Without our factual knowledge of these markers, our moral thoughts would be restricted to vague speculation about how to improve the condition of someone experiencing negative conscious experiences, since it is biology that offers the specificity that allows for precise recognition of the problem.

The above example displays how an understanding of well-being can address moral concerns. To imagine the full potential of what this process entails, consider the model of the

moral landscape hypothesized by Harris. Over time, science will increasingly illuminate this three-dimensional space, defining the peaks and valleys of human experience in terms of biology. That is, science will factually understand which values promote well-being; consequently, this knowledge can be employed in ethical philosophy as well as recycled to improve existing conceptions of well-being. Acknowledging the science of morality in this way, as a fully realized discipline rather than a nascent proposition, reveals two important goals for the future operation of moral science.

The first goal in the future of moral science involves its pragmatic engagement in society. Evidently, our world is not in a morally ideal condition. There are certainly disagreements about how one ought to behave across cultures and even within cultures. However, these disagreements manifest in dire ways, ranging from the mistreatment of women to the genocide of small children, so there is definitely a need for global conversation about human well-being, even if the terms morality or ethics are excluded. Thus, a science of well-being can inform the governments, private institutions, and global organizations about the general condition of society and assist them in making improved decisions for our future.

Most likely, the effects of moral science will relay into policy, justice, and human rights, but it is important to note that the function of importance here is information. Human history is a clear testament to the fact that knowledge is better than no knowledge; as humans increasingly discovered and recorded knowledge, society became more civilized and better equipped to secure the well-being of its populace. I am not arguing that a fully realized science of morality has the final word on ethical issues or even that it can provide all the answers to difficult questions. Rather, I am simply arguing that the reservoir of knowledge that moral science deals with can help society tackle the numerous moral issues that face us today, given that it presents

objectively right and wrong ways to improve individual well-being. Essentially, the moral science is the most directly relevant informative process to date with the potential to improve that state of society and its people. Still, this is a lofty goal in relation to the present state of moral science, but there is a more feasible aim.

The second goal involves the natural involvement of moral science in academia. Inevitably, the volume of scientific findings in relation to conscious well-being will be of at least some importance in the academic sphere. This means that, at the very least, it will engage in the conversation about morality among the academic community across various disciplines. Universities may engage students in this dialogue, meaning that the youth, who are most empowered to make important changes in society, will be trained in rational moral thought. This academic engagement, like other controversial scientific theories, will be brought into public awareness and will probably generate reactions, garner support, and stimulate debates. Certainly, a science of morality will improve the way educated society talks about such an important issue, and has the potential to develop into an established discipline in the academic community.

Having established what the transition from scientific data to ethical responsibility might look like in its most basic form, a fully realized science of morality is now conceivable. The most immediate responsibility is the illumination of the moral landscape; science is first concerned with understanding human well-being at the level of the brain. The next task involves interpretation and reformulation of the facts into morally relevant, normative literature. Finally, the third responsibility involves the participation of primarily academic and global institutions in ethical conversation, drawing on moral science to make informed decisions.

This is the completed picture of what Sam Harris hinted at when arguing for the notion of a scientifically accessible moral landscape. It is worth mentioning here that even if primacy of

human well-being in moral philosophy is not convincing (i.e. if Harris' moral landscape is not persuasive), the present work still stands; the only thing that would change is that my effort would be termed a science of well being rather than a science of morality. Regardless of terminology, however, the importance of the current project in academia and in global society remains, and the potential benefits it can offer individuals are still very real. The reservoir of knowledge/values discovered by a science of morality is independent of whether or not the knowledge is accepted for a moral philosophy; it still offers a reliable way to measure and improve well-being.

As clarified by Harris, science and morality connect through two other projects. One project involves exploring the biological basis of morality and the other focuses on persuading people to adopt moral behavior as defined by the present system. Harris strictly limits his discussion of these projects, implying that they are not interrelated and that moral science is an independent pursuit. However, I argue that all three of these projects will influence each other because they are, in fact, closely interrelated.

The scientific investigation of human moral behavior and its origins is a worthwhile effort in itself, but it is also an important tool that can improve the way moral science operates. Specifically, knowledge about the biological mechanisms underlying moral decisions can be used to correct for potential inadequacies in human judgment. For example, psychologist Jonathan Haidt proposed the social intuitionist model (SIM) in 2001, arguing that subconscious mental processes play a more significant role in moral judgments than conscious reasoning. In essence, the SIM contends that, when making addressing a particular ethical question, human arrive at a decision via quick, automatic processes rooted in the limbic system, and conscious reasoning is used post-hoc to justify the decision. Though Joshua Greene and others contend that

social reasoning does play a role in producing these intuitive judgments, distinct from post-hoc reasoning, the emotional foundations of “moral” judgments cannot be denied (Haidt, 2001). This is an important recognition for a science of morality to make because it identifies two barriers to the public relationship with scientific values. The first deals with the range of cultural understandings of “morality”, a task that is, by definition, the focus of project three. The second is our intuitive/natural understanding about how to improve well-being, which must be addressed by project two (moral science).

Acknowledging that human intuitions about morality stem from emotional capacities rather than faculties of reason reveals an additional priority for moral science that has not yet been identified. The social intuitionist model implies that, at the individual level, decisions about well-being are primarily rooted in emotion. Therefore, in order to best put moral science into practice for individuals, we should present scientific values in emotionally pertinent ways. This means that the neuroscience of well-being needs to be thoughtfully presented in order to carry ethical significance. I argue that, given our understanding of intuitive moral judgments, the most reasonable way to do this is with a nod towards our emotional character, meaning that our scientific search for values must recognize that values must have emotional significance in order to be seriously considered by the masses as having moral relevance.

I am simply proposing that this potential science, apart from the technical literature that supports it, should also provide its findings in emotionally relevant forms. This will likely allow for greater reception and application, while the original literature will still be available. For example, the published neuroscience findings that might be generated from of a highly technical study of meditation will likely be as unintelligible for non-specialists as quantum physics is to an anthropologist, probably because of scientific jargon, required background knowledge, and

complex methods. Even the necessary meta-analysis of this literature that will be applied in academic or institutional moral conversations may not be so meaningful for most people. Therefore, just as Carl Sagan painted the cosmos as aesthetically and philosophically profound through popular media, a science of morality must paint the neuroscience of well-being as emotionally significant in order to best ignite individual ethical reflection. The key idea to take away here is that presentation is an important priority for a science of morality, unlike the other natural sciences.

The scientific study of the evolutionary/biological mechanisms of moral behavior can offer a further benefit to a science of morality: an objective reason to extend the value of well-being from the individual to others. Our first recognition in the philosophic defense of moral science was that human well-being is the only possible domain of values, but this recognition left open the question of whether this value was restricted to the individual or genuinely included the well-being of others.

Basic evolutionary theory suggests that individual organisms compete with others for resources; not all individuals can survive, thus only the fittest organisms for a particular environment can pass on genes to the next generation. Evidently, this idea suggests that individuals have evolved numerous ways to promote their own well-being to increase the chances of creating successful offspring. A quick evaluation of general individual behavior confirms this idea. For example, consider your own daily life and the frequency of behaviors that tend to improve the individual's well being, however you may define it, versus behaviors that might diminish it. It seems glaringly obvious that individuals most frequently act to improve their own well-being, and even if they don't appear to, their intentions probably suggest otherwise, since we don't always know how best to increase well-being.

One common scientific explanation for altruistic behavior relates to kin selection, which argues that it is evolutionarily advantageous for an individual to promote the reproductive success of genetically similar organisms (Churchland, 2011). In a general sense, individuals that are genetically similar can increase the chances of passing on their genes either by increasing their own reproductive success or that of their close relatives'. However, this theory does not adequately explain why humans often behave altruistically regardless of genetic similarity between involved parties. Still, the theory of kin selection is a step in the right direction toward understanding apparent human moral behavior. In fact, Patricia Churchland offers a compelling argument that expands this theory into a social account for the development of moral behavior in primates.

Churchland delineates the evolution of pro-social behavior in primates and hominids from an evolutionary perspective. She approaches morality within a social context, defining it as a four-dimensional scheme for social success that is rooted in caring and attachment, prediction of others' behavior, problem-solving in social contexts, and learning key social behaviors. Accordingly, she defines values based on the human capacity to learn and predict, arguing that values reflect modes of behavior that help individuals successfully navigate the social world. From this perspective, she utilizes relevant neurobiology, evolutionary theory, and behavioral experiments to understand the evolution of human morality (Churchland, 2011).

The above findings suggest that humans likely have an innate concern for others' well-being. This is important for a science of morality because it allows us to extend our search for values that improve individual well-being to ones that improve collective well-being. If we were fundamentally selfish, moral science would merely indicate how individuals can live meaningful lives or how institutions can secure certain values to support individuals' efforts toward personal fulfillment. Because humans are likely concerned for others at a fundamental level, the central

tenet of moral science can evolve. Thus, moral behaviors are not only those that increase individual well-being, but also those that increase collective well-being. Accordingly, a science of morality can indicate how groups and societies can thrive. To reiterate, knowledge of moral origins in our evolutionary history allows the present science of morality to expand the value of well-being.

The third and final task that science must take on involves guiding existing convictions towards the right values. In particular, science must convince people whose values are drastically different from those found by science and people whose values, regardless of their similarity to scientific values, only concern themselves. Though Harris asserted that this task is the most important one for society in the 21st century, he fails to realize its immediate relevance to the other two projects. In contrast, I argue that knowledge about the way humans already behave in the name of “morality” plus knowledge derived from moral science are both necessary and sufficient to address the question of unifying global morality.

Given that individuals are raised within cultural contexts that often provide ethical mandates that likely establish a permanent ethical core, the task of realigning existing values with a factual understanding of well-being seems difficult. It goes without saying that rationality and a positive regard for science are probably prerequisite for any intelligent discussion of morality. Accordingly, religious fundamentalists, terrorists, or sociopaths are not the primary audience for the present effort. That being said, the scientific exploration of behaviors and attitudes that resemble morality (as defined in the colloquial sense) can first ground a conversation about values. Part of the reason that agreement is difficult is because morality may be seen as something transcendent, inaccessible, or not real. By first demonstrating that, when we speak of morality, we are referring to the well-being of conscious creatures, and that the idea

of morality is a social phenomenon, and finally that it is rooted in evolutionary adaptation, we can establish the groundwork for a meaningful discussion by anchoring it to a relatable/relevant domain.

Upon admitting a natural basis of moral significance, the responsibility of realigning values falls on project two, the science of morality. At this point, a philosophical argument defending the moral landscape will ensue, establishing an objective basis for moral truth. This task will necessarily involve communicating two fundamental ideas. First, that there are differences between conscious experiences (i.e. that there is a difference between the conscious experience of human flourishing and suffering) and that these differences lawfully relate to the physiology of the brain. Acceptance of these ideas admits of the existence of a moral landscape and admits that there are right and wrong answers to moral questions.

Once the existence of a moral landscape has been agreed upon, the final task in unifying global values involves conveying the scientific illumination of the moral landscape and its associated normative conclusions. Although science does not purport to “prove” assertions or conclusively report truths, there should be some degree of consensus in the community about the correct ways to improve well-being based on fact. Acknowledging this potential is the final step in improving existing ethical convictions and culturally developed senses of good and evil, and will proceed either by realigning these convictions towards correct values or by substantiating them with a rational and empirical foundation.

The above line of reasoning represents a plausible attempt at project three, a future conversation between the morally informed and the morally ignorant. It is important to recognize that I am not asserting any final truth or logical perfection in the above discussion. I am arguing that the above represents the most convincing, and most likely to be utilized, route in the

advancement of global moral discourse. This reasoning needs to be made lucid, partially because the full relationship between science and morality (i.e. the three projects) has not been comprehensively drawn out.

When Harris mentioned the three points of contact between science and ethics, not only did he imply that the projects are independent of each other (which they are not), he also suggested that there is a lack of academic conversation between them. Clearly, all three tasks are being addressed at present; however, a reflection on their relationship needs to be elucidated for the sake of moral progress. To reiterate, Harris' third project of persuading others toward scientific values is undeniably dependent on the success of the first two projects. Even if this goal is not fully realized (i.e. people still don't accept scientific values), at the very least, it will have ignited the moral discussion that the current state of global society quite clearly needs.

Works Cited

- Harris, Sam. *The Moral Landscape: How Science Can Determine Human Values*. New York: Free Press, 2010.
- Haybron, Daniel. "Philosophy and the Science of Subjective Well-Being." *The Science of Subjective Well-Being*. New York: Guildford Press, 2008
- Maslow, Abraham. "A Theory of Human Motivation." *Psychological Review* 50 (1943): 370-396
- De Waal, Frans. *Primates and Philosophers: How Morality Evolved*. New Jersey: Princeton University Press, 2006
- Neff, Kristen. "Self-Compassion and Psychological Well-Being." *Constructivism in the Human Sciences* 9 (2004): 27-37
- Baer, Ruth. "Mindfulness Training as a Clinical Intervention: A conceptual and empirical review." *Clinical Psych.* 10 (2003): 125-143
- Ramel W, Goldin P, Carmona P, McQuaid J. "The Effects of Mindfulness Meditation on Cognitive Processes and Affect in Patients with Past Depression." *Cognitive Therapy and Research* 28. (2004): 433-455
- Zylowska L, Ackerman D, Yang M, Furtrell J, Horton N. "Mindfulness Meditation Training in Adults and Adolescents with ADHD." *Journal of Attention Disorders* (2007): 1-10
- Oman D, Shapiro S, Thoresen C, Plante T, Flinders T. "Meditation Lowers Stress and Supports Forgiveness Among College Students: A randomized clinical trial." *Journal of American College Health* 56 (2008): 569-578

- Lee HJ, Macbeth AH, Pagani JH. “Oxytocin: The Great Facilitator of Life.” *Prog. Neurobiol* 88 (2009): 127-51
- “Healthy Eating Plate.” *Harvard School of Public Health: The Nutrition Source*. 2011. Accessed April 2013 www.hsph.harvard.edu/nutritionsource
- “Physical Activity for Everyone.” *Centers for Disease Control and Prevention*. 2011. Accessed April 2013. www.cdc.gov/physicalactivity/everyone/guidelines/index.html
- McLeod, S.A. “What is the Stress Response.” *Simply Psychology*. 2010. Accessed April 2013. www.simplypsychology.org/stress-biology.html
- Gleitman H, Fridlund AJ, Resiberg D. *Psychology* 6. W.W. Norton and Company. (2004)
- Churchland, Patricia. *Braintrust: What Neuroscience Tells Us About Morality*. New Jersey: Princeton University Press, 2011.
- Wright, R. “The Biology of Violence.” *The New Yorker*. (1995): 68-77
- Lutz A, Slagter H, Dunne J, Davidson R. “Attention Regulation and Monitoring in Meditation.” *Trends Cog. Science*. (2008): 163-9
- Sagar M, King BG, Zanesco AP, et. Al. “Intensive Training Induces Longitudinal Changes in Meditation State-Related EEG Oscillatory Activity.” *Front Hum. Neurosci*. (2012)
- Haidt, J. “The emotional dog and its rational tail: A social intuitionist approach to moral judgment.” *Psychological Review* 108. (2001): 814-834