

A photograph of a man from behind, wearing a dark suit jacket and white shirt, looking out over a vast, green, rolling landscape. In the distance, there are snow-capped mountains under a dramatic, cloudy sky.

What *is* Morality?

META-ETHICS IN PLAIN TALK

What is Morality?

meta-ethics in plain talk

Version 1.0, March 2009



This book is licensed under the Creative Commons Attribution-Noncommercial 3.0 license. You may copy and modify this work freely, as long as you (1) attribute this original work to Luke Muehlhauser, and (2) don't charge money for any derivative works.

<http://commonsenseatheism.com>

<http://lukeprog.com>

Cover photo by Flickr user samsnet, used under Creative Commons license:

<http://flickr.com/photos/samsnet/132997008/>

Many thanks to Alonzo Fyfe, who spent decades developing the theory defended in this book, and who proofed early drafts of this book.

Free audiobook version also available, at <http://commonsenseatheism.com>.

Contents

Chapter 1	Chapter 6
Good and Bad.....3	Morality as a Myth.....18
Chapter 2	Chapter 7
Gods.....8	Moral Value that Really Exists.....21
Chapter 3	Chapter 8
Virtue.....10	Objections.....30
Chapter 4	Chapter 9
Duty.....12	How to Make the World a Better Place.....37
Chapter 5	
The Greatest Good.....15	

CHAPTER 1

GOOD AND BAD

§

Some people don't care what is good or bad. They don't care if they make the world a better place, a worse place, or if they leave no mark on it at all.

This book is for people who *care*, and who want to make the world *better*.

If you're one of those people, you have chosen a hazardous course. Billions of people before you wanted to make the world a better place, but they actually made it *worse*.

During World War II, German and Japanese soldiers thought they were fighting to make the world better. If you had been born in 1910 in Japan, you would have thought the same thing. All those people died making the world *worse*, despite their best efforts to make it better.

Billions of well-meaning people defended and spread tribalism, slavery, sexism, and racism,

because they honestly thought those things were moral.

Muslims and Christians have oppressed women and persecuted gays, because they thought that was what God wanted - that doing so would make the world a better place.

All these people made the world worse when they were trying to make it better. If you're going to make the world a better place, you're going to have to be smarter or more careful than them, somehow. That's not going to be easy.

If you want to make the world better, you have to know what "better" is. Is capitalism better? Is abortion better? Do some prejudices - like prejudices against pedophiles - make the world better? Can preemptive war make the world better? Does spreading the morals of Christianity, Islam, Confucianism, Buddhism, or Jainism make the world better? Are closed romances better than open ones? Is lying better? In which circumstances?

Most people think they know the answers to these questions. You probably think so, too. But think of all those people who made the world worse. They also thought they knew the answers.

And here's the thing. If you're like most people, you make your moral decisions the same way they did. You close your eyes, shut out distractions, and ask your conscience. And then your conscience - your *moral feeling* - delivers you the answer.

That is exactly how racists, sexists, bigots, fascists, religious zealots, and violent people throughout history answered their moral questions.

We're going to have to find a more accurate way to answer moral questions.

That's where meta-ethics comes in.

Meta-Ethics

Meta-Ethics is the practice of asking the big questions about morality. Does morality exist? How so? What does "good" mean? How can we know what is good?

It is no use asking "Is rape wrong?" if we have no idea if "wrongness" actually exists or what makes something wrong.

When people say "theft is wrong," they seem to mean that an act like theft emits a kind moral radiation - tiny particles we might call "goodons" and "badons" - and that we have evolved a sense to detect those particles. Somehow the act of typing commands into a financial computer to embezzle money emits the same kind of moral radiation as driving a truck through the wall of a department store and running off with designer jewelry. Moreover, the universe somehow "knows" the intent of your "heart," and that also influences

which moral radiation is given off by each of your actions.

Is *this* what we mean when we talk about morality? Does this accurately describe *what really exists* in the universe? If not, what kind of moral values *do* exist?

Those are the questions of meta-ethics. If we can find solid answers to them, we may be able to do better than most people who have come before us, and really *know* - not just *feel* - how to make the world a better place.

The First Question

Do moral values exist? That is the First Question of meta-ethics. If moral values do not exist, then morality is a myth. It is a fantasy story we tell ourselves, like the tales of Zeus or Luke Skywalker. If moral values do not actually exist, then the other questions of meta-ethics do not matter much.

In this short book, we'll look at several attempts to answer the First Question - several theories of how moral values might really exist.

How are we going to pick the winner? How do we know which theory, if any, will truly help us make the world better?

We can't choose the winner based on how well it conforms to what our conscience tells us. Your conscience is just a product of evolution and culture. Two centuries ago, your conscience would have told you that racism and sexism were good. Your conscience cannot be trusted.

We are not looking for the theory that we *like*. We're looking for a theory that - just like any good scientific theory - accurately describes our universe. You "become" a utilitarian or a Kantian the same way you "become" an atomist. It doesn't matter whether you *like* the idea that matter is made of tiny atoms; it only matters whether or not that is *true*.

So, do moral values exist? If so, what are they? Let's see if we can find out.

CHAPTER 2

GODS

§

The oldest moral systems to address the First Question - “Do moral values exist?” - are religious ones. According to most religions, moral values exist because they come from the gods. The gods tell us what is right and wrong.

Suspiciously, the gods of every tribe and nation always agreed with those in power, and also with their ignorance of the world (a flat earth, an earth-centric universe, magical explanations for everything instead of a germ theory for disease or a neurological theory for mental illness). Did these morals come from the gods, or from the ancient people who first told us about their gods?

Besides, we can’t tell which gods are real and which ones aren’t. And even if we could, how would we know they are telling us what is really moral? What if God himself is evil? God is so much smarter and more powerful than we are; how would we know? Or what if God tells us a bunch of random nonsense just to see if we’ll play along? That doesn’t seem too far-

fetched with some gods - like Yahweh, who supposedly commanded us not to shave, gather sticks on Saturdays, or wear mixed fabrics. If we are told, "Do this because God says so," we can always ask, "and why should I do what God says?"

We must also ask, "Can moral values really be nothing more than a person's whim?" If God decided to command rape, would that make rape good? Or perhaps God merely passes on to us moral values that exist beyond him. But then the foundation of moral values is just as mysterious as ever.

Because of these problems, most philosophers - including religious ones - do *not* think we can get moral values from God.

There is also the slight problem that gods do not exist.

But religion was a primitive and ignorant attempt to answer moral questions. It was a nice try for its time. Perhaps we have found better answers in the age of science and reason.

CHAPTER 3

VIRTUE

§

Some philosophers think that we should not measure *actions* as good or bad, but *character*. Goodness is that which is done out of a virtuous character.

But then, what is virtuous? Virtue is any habit or quality that allows an agent to achieve its purpose. The virtues for an axe are sharpness and durability. The virtues for a hunting dog are a sensitive nose, stealth, obedience, and more.

So, to know what is virtuous for a human, one must know what the purpose of a human is. Perhaps our purpose is to pass on our genes, in which case ambition, sex, and child-rearing are virtues. Maybe it is to live in a stable society, in which case cooperation, fairness, and compliance are virtues.

Philosophers have had different ideas of what human purpose is, and have therefore proposed very different lists of virtues. How could we choose which list is more accurate?

No philosopher has been able to show where human “purpose” comes from, or why one purpose or set of virtues is a better description of what really exists than another.

Thus, virtue theory gives no account of how moral values might objectively exist in the real world. We can't make the world a better place by promoting certain virtues, because we have no reason to think that one list of virtues is more “real” than another. Let us move on.

CHAPTER 4

DUTY

§

If you want to get from San Diego to Los Angeles, you *should* travel north. If you want to survive, you *should* eat food. If you want to get elected to public office, you *should* tell people what they want to hear.

Those are practical imperatives. Some philosophers think that *moral imperatives* are those things that should or should not be done *regardless* of any “if” condition. For example, you *should not murder*, regardless of whether you want to get to Los Angeles or survive or get elected or anything else.

That sounds nice, but what *is* this moral duty? Where does it come from? Does it really exist, or is it another fantasy like the gods?

Some philosophers talk as though right actions are those that bring the universe back into balance. But nobody has explained what, exactly, is out of balance, how to measure it, or

how we know that certain actions restore that balance.

Immanuel Kant said it is our duty to act in ways that we could make a universal maxim. We should not lie because we wouldn't want everyone to lie all the time.

There are at least three problems here. First, how do we know what the maxim of our action is? If I were to lie to the Nazis and tell them I am *not* hiding any Jews, what is the maxim we should evaluate? "Always lie?" Perhaps it is "Always lie to Nazis." Or maybe "Never admit to hiding Jews." It could just as well be, "Always answer a knock at the door by denying that you are hiding Jews."

Second, Kant's theory does not seem to fit everyday life. Right now I'm writing a book on a Dell laptop. Does that mean I think everyone should be writing a book on a Dell laptop?

Third, *why* is it our duty to act in ways we could make a universal maxim? Does universality have intrinsic value? How?

Intrinsic duty

Philosophers of duty-based ethics say we all have certain rights and duties. We have a right to life. A right to property. A right to free speech. A duty not to murder. A duty not to rape.

Where are they getting these lists of rights and duties? Are they observing something that exists in the universe, or are they just listing things that would seem to create the kind of society *they* want to live in? Once again, it seems we are depending on moral feelings, not on a careful study of what really exists.

Does “duty” really exist, or is it something we made up to influence others? Do “rights” exist? Do we have any way to tell that a “right to life” exists for humans but not apes, or fish but not sponges? How are philosophers measuring these rights and duties?

Until someone can show that “duties” and “rights” exist as intrinsic properties of our universe, they will remain as mysterious as gods and fairies.

CHAPTER 5

THE GREATEST GOOD

§

Other philosophers think that “good” is not based on rights and duties, but on the consequences of an action. A right action is one that produces “the greatest good for the greatest number.”

What is the greatest good? Maybe it is *happiness*. In that case, an action is good if it causes more happiness than unhappiness. Or maybe “good” means *pleasure*. If so, an action is good if it causes more pleasure than pain. Some philosophers think an action is good if it satisfies preferences. Others think that because the universe is doomed to chaos, the greatest good is that which increases the organization and complexity of the universe.

Intrinsic value

There are many problems with “greatest good” theories. But the most basic one is this: on

what grounds should we try to maximize happiness instead of pleasure, or preference satisfaction instead of complexity? It is as if these things have some kind of intrinsic moral value. It is as if they radiate goodness.

In fact, we have no reason to think any of these have intrinsic value. What is intrinsic value? How do we detect it? How is “intrinsic value” any less mythical than the gods? If we could crack certain actions open, would we find something called “intrinsic value” hiding inside? I think not.

The is-ought gap

Another problem is something called the “is-ought gap.” It’s also a problem with the other moral theories we’ve considered. In the 18th century, David Hume pointed out that anyone who argues for a particular moral theory starts out with a bunch of “is” statements, and then at the very end suddenly makes an “ought” statement, without explaining how he got from “is” to “ought.”

For example, someone might argue:

1. Scientists agree that humanity’s carbon pollution *is* causing global warming.

2. Computer models show that global warming, if it continues, *is* going to cause widespread devastation to the planet.
3. Therefore, humans *ought* to stop polluting the air with carbon.

But how did we jump from “is” statements to “ought” statements? We were talking about one thing and then jumped to a totally different topic, without any bridge between the two.

Philosophers encounter this problem whenever they try to come up with moral theories that are grounded in reality. You can describe the world with “is” statements as much as you want, but how can you make the magical leap to “ought”?

So, it seems that “greatest good” theories also fail to give an account of moral values that actually exist in the world.

CHAPTER 6

MORALITY AS A MYTH

§

There are many nice theories about what is good and bad, but none of them have been able to show us that moral values *actually exist* according to good evidence. This has led many philosophers to conclude that morality is nothing but a fairy tale.

Perhaps morality is a myth, like the gods. It was primitive man's attempt to explain his primal *feelings* about good and bad, and to say they were something more.

Perhaps morality is a tool used by those in power to maintain a peaceful society. They can say the laws exist not just to keep order, but because they correspond to some moral laws that are woven into the fabric of the universe.

Or perhaps morality is nothing more than the moral feelings of a particular culture, group, or person. If so, witch-burning and wife-beating are moral if they are accepted in that culture. Racism is moral for the Ku Klux Klan because that is the moral code for their group.

Child-rape is moral for the pedophile because that's what *he* thinks is good for *him*.

Perhaps morality is based on goals. If you want to get from San Diego to Los Angeles, you should go north. If you want to get a job for which you are underqualified, you should lie on your resume. If you want to take the resources of a weaker country, you should invade it.

Or perhaps morality refers to nothing at all. “Good” and “bad” are meaningless words. We cannot judge or praise anyone for anything.

These ideas are pretty unsatisfying because they don't fit our moral feelings. We *want* to condemn the pedophile for child-rape, and we want to praise Gandhi for his non-violent liberation of India.

But if moral values do not actually exist in the universe, what else can we say? We may not *like* moral relativism, but if it is *true*, well then it is *true*. At first we did not like to hear that we were not at the center of the universe, but it happens to be true that we live on a tiny planet on the outskirts of a minor galaxy among billions of others. Many people do not *like* the fact that life ends at death, but it happens to be true. We have accepted other hard truths before, and found a way to move on. Perhaps we will have to do that with the non-existence of moral values; accept it, and move on.

Knowing that moral values do not exist would not stop people from *acting* morally, anyway.

Many relativists are kind, honest, and generous people. We still like the feeling of helping other people, even if that action does not really radiate goodons.

So, is this the situation we must admit? Is morality a fairy tale we have to leave behind when we “grow up” as a species, just like we must leave behind our primitive ideas about gods and souls and the afterlife? Is there really no way to make the world “better,” because “better” is an illusion?

Maybe. But maybe not.

CHAPTER 7

MORAL VALUE THAT REALLY EXISTS

§

We have been searching the universe for moral value, and come up empty-handed. But don't give up yet! Let's keep looking.

First, we should know what "moral value" might look like. Does it have legs? What color is it? We have to know what we are looking for if we are going to find it. What *is* moral value?

Morality is concerned with *reasons for action*. Reasons for action to feed the poor. Reasons for action to be kind to others. Reasons for action to not torture children.

But we've had a tough time finding reasons for action that really exist. Intrinsic value doesn't seem to exist. Neither do duties or gods.

So are there any reasons for action that *actually exist*?

Yes.

Desires exist.

In fact, as far as we know, desires are the only reasons for action that *do* exist.

If the poor did not desire to be fed, there would be no reason for action to feed them. If children did not desire to avoid torture, there would be no reason for action to avoid torturing children. If we all had desires not to eat food but to soak up sunlight, then we would have no reason for action to feed the poor. Instead, we'd have reasons for action to give them access to sunlight. Desire is the source of all moral value.

Direct and indirect value

There are two ways something can have value.

A thing has *direct value* when it is desired. Sunsets, relationships, peace, candy, adventure, drugs - all these things acquire direct value when they are desired.

A thing can also have *indirect value* when it *tends to bring about* something that is desired. An attitude of good humor tends to bring about pleasant feelings and healthy relationships that are desired. Democracy tends to bring about personal freedoms that are desired.

Good and bad

Since desire is the only reason for action that exists, something is “good” if it fulfills the desires in question. Something is “bad” if it *thwarts* the desires in question. It’s that simple.

Isn’t this relativism?

So far, it sounds like we’re just talking about relativism. Under this framework, it is “good” for Johnny to bring a gun if “the desires in question” are Johnny’s desires to rob a bank. A rape would be “bad” for the victim, who desires to not be raped, but perhaps equally “good” for the rapist, who desires to rape.

But so far, we are only talking about *generic* goodness, not *moral* goodness. We’ll get to moral goodness in a moment. For now, just notice that the above statements about Johnny and the rapist are *objectively true*. If you say that it is “good” for the victim to be raped, relative to her desires to not be raped, you are *objectively wrong*. Her desires to not be raped are reasons for action that *exist* in the real world, and it is objectively *true* that her being raped thwarts her desires to not be raped. This is not a matter of opinion. As long as we are clear on what our words mean, claims about “good” and “bad” have objective truth value, because there are some reasons for action that really exist - namely, desires - and certain states of affairs in the real world really do fulfill or thwart those desires.

Moral value

When we talk about *moral* value, though, we are talking about something universal. My desire for Angelina Jolie to sleep with me does not mean she is morally obligated to sleep with me. Universal moral claims require a consideration of *all* desires.

So when talking about universal morality, “the desires in question” are *all* desires. So, “morally good” means “such as to fulfill more and greater desires than are thwarted, among *all* desires.” And “morally bad” means “such as to thwart more and greater desires than are fulfilled, among *all* desires.”

For there is no reason to exclude certain desires from the evaluation. We cannot even exclude the rapist’s desires to rape. No desire is intrinsically better or worse than any other desire, because intrinsic value does not exist.

Instead, we must evaluate the moral value of desires in the exact same way we evaluate the moral value of everything else! We ask, “How well does this desire fulfill or tend to fulfill other desires?”

Evaluating desires

In fact, the evaluation of moral claims always starts with the evaluation of desires. We do not

start by evaluating actions or laws or ideas, but *desires* - for desires are the source of all moral value, the only reasons for action that exist.

It's not so strange to evaluate the moral value of desires. Actions, laws, policies, tools, and movies are all good or bad according to their tendency to fulfill or thwart desires. Desires, too, are good or bad according to their tendency to fulfill or thwart desires.

A desire to rape is bad because it tends to thwart more and greater desires than it fulfills. A desire to show kindness is good because it tends to fulfill more and greater desires than it thwarts.

Malleable desires

We can be more specific, because *ought* implies *can*. Nobody can say that I “ought” to stop the oncoming tsunami if there is no way I *can*. So when we evaluate desires, we can only evaluate *malleable* desires - desires that can be changed.

There is no point in saying the human desire for water is good or bad, since that desire cannot be changed. To say that I “ought to not desire water,” implies that I *can* “not desire water,” which is false.

So, morality is concerned with the evaluation of *malleable desires*.

How to evaluate actions

Still, you might be disappointed that we have to evaluate *desires* instead of *actions* (because desires are the source of all moral value, not actions). Not to worry. You can judge actions as “right” or “wrong,” but not directly.

A right action is one that a person with good desires would perform. A wrong action is one that a person with good desires would *not perform*.

We can judge the morality of other things in the same way. A good law is one that a person with good desires would enact. A bad law is one that a person with good desires would reject. A good movie is one that a person with good desires would watch. A bad movie is one that a person with good desires would *not watch*. (Keep in mind we are talking about *moral* values, now. Aesthetic values are usually not defined the same way.)

Objective vs. subjective

You might think I’m saying that morality is subjective, since “good” and “bad” depend on whatever people happen to desire. This comes from a confusion about what the words

“subjective” and “objective” mean.

Morality *does* depend on desires. If there were no desires, there would be no moral value in the universe. And if everybody desired to be surrounded by deafening noise, then it would be morally right to carry a blasting boombox everywhere you went. In this sense, morality is subjective.

But that’s not what most people mean by “subjective.” Subjective morality usually means that each of us gets to choose for ourselves what is good and bad, and nobody can be wrong. Morality is *not* subjective in this sense. As we saw before, you are *objectively wrong* if you claim that “rape is good.” Why? Because rape is an action that a person with good desires would *not* perform. Instead, rape comes from a bad desire; a desire that tends to thwart more and greater desires than it tends to fulfill. Since desires are the only reasons for action that exist, we have real and universal *reason for actions* to diminish or eliminate the desire to rape in others.

Also, consider the word “objective.” Some people use the phrase “objective morality” to refer to some kind of intrinsic value written into the fabric of the universe. But intrinsic value doesn’t exist. In this sense, objective morality doesn’t exist.

But that’s not what most people mean by “objective.” Objective morality usually means that

moral statements can be *true* or *false* in the same way that scientific statements can be true or false. In this sense, morality *is* objective, as we saw above.

Finally, remember that even though morality depends on individual desires that evolution and culture happened to produce, nevertheless morality is *universal* because moral judgments refer to *all* the reasons for action that exist: *all* desires.

There is another sense in which morality is universal. Morality is concerned with desires that *everybody* should have, or that *nobody* should have. The question of whether rape is a good or bad desire is answered by asking, "What if everybody had a desire to rape?" Likewise, we can ask, "What would happen if everybody had a desire to see through their own biases?"

So, morality is both *objective* and *universal*. And not because that's a comforting thought, but because that is what we find when we look at what *actually exists* in the *real universe*.

Summary

Now we have a theory of how objective and universal moral value really exists in the universe. This moral theory has a name, by the way. It's called "desire utilitarianism."

There are many competing moral theories, but I think desire utilitarianism is the only one

that gives an *accurate* account of objective moral value that really exists. You'll have to decide for yourself. In the next chapter, we'll look at some common objections to desire utilitarianism.

CHAPTER 8

OBJECTIONS

§

People have raised many objections to desire utilitarianism, but so far they have all turned out to be misunderstandings about what desire utilitarianism actually says. Let's consider some of these misunderstandings.

Aren't you just saying that desire fulfillment has intrinsic value?

No. That would be "desire fulfillment act utilitarianism," not desire utilitarianism.

Desire fulfillment does not have intrinsic value. Moral value comes from reasons for action, and it so happens that desires are the only reasons for action that exist.

Each desire is its own reason for action, and we must consider *all* of them, for no desire is intrinsically better or worse than any other.

Let's use an example to make clear the differences between "desire utilitarianism" and "desire fulfillment act utilitarianism."

Say there are 20 sadists who desire to torture a child, and one child who desires to *not* be tortured. Desire fulfillment act utilitarianism would say it is good for the sadists to torture the child, because that is the act that would maximize desire fulfillment.

In contrast, desire utilitarianism evaluates desires themselves, based on their tendency to fulfill or thwart other desires. The desire to torture children tends to thwart more and greater desires than it fulfills, so it is a bad desire. A person with *good* desires would *not* torture children, so torturing children is a wrong action to perform.

That's the difference between desire utilitarianism and desire fulfillment act utilitarianism.

But remember, we should not accept desire utilitarianism because it gives us answers that feel good, like "torturing children is bad." History and biology show us that we have not evolved a "sixth sense" to detect moral value directly through our feelings. So, we can't accept or reject *any* moral theory based on how well it fits our moral feelings. *Our moral feelings could be totally wrong.* For example, they seem to have been wrong in the past, when everyone accepted tribalism, sexism, racism, and homophobia as morally good.

The only way to choose a moral theory is to ask if it *accurately describes what really exists in*

the universe, the same way we would accept a scientific theory. In the earlier chapters I showed why only desire utilitarianism accurately describes how moral value exists in the universe, and *that* is the only good reason to accept desire utilitarianism; not because it is nice, but because it is true.

Can desire utilitarianism bridge the is-ought gap?

Earlier we talked about the is-ought gap. How do we make the leap from talking about how something *is*, to talking about how something *ought* to be? It seems like an unbridgeable gap.

But as we've seen, *ought* is not some new kind of affirmation. It is merely another type of *is* statement. Specifically, *ought* means "There exist reasons for action such that..." Notice, that's just a specific type of *is* statement, and it refers to something that exists in the real world: reasons for action. Desires are the reasons for action that exist.

Why should I accept your definitions of value and morality?

One might object that it all depends how you define morality. If I define morality as "reasons for action" and somebody else defines morality as "whatever God wills," we could argue over definitions forever, and neither of could ever *prove* that our definition is correct. So, isn't morality subjective? How can I say that desire utilitarianism is correct, and the other

moral theories are wrong?

Recently, astronomers have debated the meaning of the word “planet.” One definition would include Pluto as a planet. The other definition would say Pluto is not a planet.

But now, imagine going to an astronomy convention to claim that we could argue over the definition of “planet” forever and we cannot prove which definition is correct, therefore people can adopt whatever theory of planet formation they like. You would be laughed off the stage.

Words do not change reality.

Earlier, I took note of what people usually mean when they use moral terms. Moral terms are used to recommend certain actions. In other words, morality refers to reasons for action.

It so happens that desires are the only reasons for action that exist, but that is not part of the definition of morality. That’s just an empirical fact. When we look at the universe for reasons for action, desires are the only ones that exist. Intrinsic values, gods, categorical imperatives, intrinsic duties - those things don’t exist. But desires do.

If society decides to use the word “morality” to mean something else - or if the Spanish take over the world and we use the word *moralidad* instead - that doesn’t change the fact that

reasons for action exist, and there are *universal* and *objective* reasons for action to, for example, promote kindness and discourage rape.

Likewise, someone could define morality to refer not to reasons for action but to intrinsic values or God's commands or something, but then he will be referring to things that, as it happens, *do not exist*. So, he will not be very effective at making moral claims about the real world using such a definition.

The definition of the word “morality” is not important. What matters is whether or not the things identified by desire utilitarianism exist and have relationships with each other as described.

What about animals, plants, and bacteria? Don't they have desires for life and growth, too?

If animals, plants, and bacteria have desires, they are just as morally relevant as human desires.

Desire, as far as we can tell, is a brain state. Plants and bacteria do not even have nervous systems. So, they do not have desires.

But *animals* have brains. So, do animals have desires?

Yes, probably. The problem is, we don't know which ones. Do worms have desires? Do insects? Clams? Jellyfish? Mice? Dogs? Where do we draw the line?

This is a scientific question, and it needs a scientific answer. What do desires look like in the brain? How do they function? Which creatures have them? Only patient research in neuroscience and related fields will give us the answers.

We have a similar but easier scientific problem with regard to human newborns. Does a recently fertilized egg have desires? No; it doesn't have a brain. A recently fertilized egg doesn't have any desires to thwart.

The human brain develops gradually from about 20 days after conception, and is quite mature by day 90. At what point does this human brain develop desires, and how many, how quickly? Again, we must await the results of science.

Back to animals. We can be sure that at least *some* animals have desires, and their desires require moral consideration. There is no reason to discount the desires of other species. We might as well discount the desires of other races, which is silly. All desires create reasons for action.

So, does this burden humanity with taking care of every last worm? Must we protect every prey from every predator? Should we set up a welfare system for cattle? Surely this is too

burdensome, too inconvenient to even think about!

But this is not a good objection. Desire utilitarianism is a theory about what exists; what is true about reasons for action in the universe; what is true about morality. A complaint about its unpleasantness is no objection to its truth. We cannot say that the sun will not explode in 5 billion years because the thought of our destruction is *unpleasant*.

Second, it may not be true that desire utilitarianism prescribes that we protect every little worm. “Animals rights” is a complex topic. You can read more about it at the links provided at the end of the book.

CHAPTER 9

HOW TO MAKE THE WORLD A BETTER PLACE

§

So now that we have a sketch of what moral value there is in our universe, how do we use this knowledge to make the world a better place? What *ought* we to do?

As we've seen, "ought" questions are about reasons for actions that exist, and it turns out that desires are the only reasons for action that exist. So, in any given situation we *ought* to do what a person with good desires would do. A good desire is one that tends to fulfill more and greater desires than it thwarts, considering *all* desires.

In contrast, a bad desire tends to *thwart* more and greater desires than it fulfills, such as the desire to rape. Alonzo Fyfe explains this best. He says:

We can see the problem with the desire to rape by imagining that we have control over a knob that will generally increase or decrease the intensity and spread of a desire to rape throughout a community. To the degree that we

increase this desire to rape, to that degree we increase the desires that will be thwarted. Either the desires of the rapist will have to be thwarted, or the desires of the victims will have to be thwarted. The more and stronger the desire to rape, the more and stronger the desires that will be thwarted.

The best place to turn this knob is down to zero – so that there is no desire to rape. If this were the case, then no victims will have their desires thwarted through rape, and there would be no rapists who would have to go through the frustration of having a desire to rape go unfulfilled. This is a desire that people generally have reason to weaken or to eliminate.

So, we have reasons for action to turn down the knob on the desire to rape. That is, we have reasons for action to condemn the desire to rape, to teach our children respect for others so they'll be less likely to develop a desire to rape, etc.

This is how we make the world a better place. We turn *down* the knob on *bad* desires (desires that tend to thwart other desires), and turn *up* the knob on *good* desires (desires that tend to fulfill other desires).

It might be useful to always think of desire utilitarianism in terms of these knobs. Otherwise, it's *very easy* to slip into thinking that “the right act is the one that fulfills the most desires,” which is false. The right act is *not* the one that fulfills the most desires. The right act is the

one that a person with good desires would perform, and a good desire is one that tends to fulfill more and greater desires than it thwarts (think of the knobs).

Obviously, if we instill in others desires that tend to fulfill other desires, the world will become a better place. If more people desired to earn an honest living, that would tend to fulfill more desires. So, we have reasons for action to promote the desire to earn an honest living by praising and rewarding that desire, and perhaps by condemning or punishing the desire to earn a dishonest living.

But turning knobs is not the only good we can do. We could, for example, earn an honest living ourselves. That's the whole point of promoting the desire to earn an honest living in the first place - so that people like us will be more likely to earn an honest living, which fulfills more desires than thievery.

What should you do? Do what a person with good desires would do. So the question becomes: "What should I desire?"

A good desire is one that fulfills or tends to fulfill more and greater desires than it thwarts. A desire to earn an honest living fulfills more desires than it thwarts. A desire to promote the desire to earn an honest living also fulfills more desires than it thwarts. These are both good desires.

But the question “What should I do?” is a complex question, often requiring research on human and animal desires that has not been done yet. I wish there were easy answers, like “Do whatever the Bible says” or “Act in accordance with this short list of virtues.” But those moral systems happen to be objectively incorrect about this universe.

Below, I’ll give some resources for how to figure out what you should do if desire utilitarianism is correct about reasons for action that exist. But here are some quick ideas:

Seek truth, promote truth-seeking, and condemn whatever inhibits truth-seeking. Since we always act to fulfill the greatest of our desires *given our beliefs*, it matters hugely that our beliefs are true. If we desire to swim and falsely believe the pool is full of water, we’ll dive straight into a big hole of cement.

More practically, if we desire a stable economy and falsely believe that Libertarian strategies will deliver a stable economy, we may end up thwarting millions of desires instead of fulfilling them. If we desire good health and believe that New Age superstitions or religious prayers are more effective than scientific medicine, we may end up thwarting more desires than we fulfill.

Truth is a big deal, and we have reasons for action to promote desires that bring us to truth, such as a love of critical thinking, a love of logic, a love of science, etc. We also have reasons

for action to condemn deceit, dishonest science, dishonest philosophy, intellectual recklessness, etc.

Another good desire is the aversion to violence as a response to words. We should avoid responding to words with violence, and we should condemn that response in others, so that fewer people will desire to respond to words with violence. Gandhi was pretty effective at changing this desire in others. Acting violently is more distasteful to more people after Gandhi.

A third good desire would be to understand forces of nature that thwart desires on a massive scale, like hurricanes, climate change, and viruses. If we can understand these things, we may be able to prevent them from thwarting so many desires. This requires the promotion and practice of well-funded, rigorous science.

Where to go from here

This has been a plain talk introduction to meta-ethics, and desire utilitarianism in particular. I encourage you to be as critical as possible of this book and the moral theory it presents. If you find a flaw in it that I have not - if you discover a way in which it is *not true* about the universe - then you will have done myself and the world a service.

But remember, this book is only the most basic introduction to desire utilitarianism. You probably have many objections that are actually misunderstandings, and have already been answered. To understand the theory better, and what its consequences are to all areas of life, you should read the work of the person who first formulated the theory: Alonzo Fyfe.

Alonzo Fyfe has written some books about desire utilitarianism. Unlike most philosophers, Alonzo writes in plain talk as much as possible, but he also provides precise statements when necessary, like “In order for something (S) to have direct value it must be the case that there is a desire that P, and P is true of S. S has indirect value if S has a tendency to bring about T, there is a desire that P, and P is true of T.”

You can find links to these books at Alonzo’s blog, AtheistEthicist.blogspot.com, where he also writes about the latest developments in his theory.

Also, I’ve written a list of common questions about and objections to desire utilitarianism, and I’ve linked to answers for all of them. If you have questions about desire utilitarianism, you should *definitely* check that page first: <http://commonsenseatheism.com/?p=776>.

Remember, no matter which moral theory is correct, one way you can make the world a better place is by spending some time learning about moral theories and figuring out which one is correct. If you hitch your wagon to the wrong theory, you could easily make the world

a worse place when you think you’re making it better, just like most people who have come before you.

Conclusion

We have surveyed the evidence about what really exists in the universe, and discovered that objective and universal moral value *does* exist. We *can* make the world a genuinely better place.

Our ancestors tried and usually *failed* to make the world a better place, because they didn’t have a correct notion of what “better” was. We need not repeat the mistakes of every generation that has come before us, because we are now equipped with knowledge of what “better” really is.

So, let’s go make the world a better place!