Morality from the Ground Up (20170315)

Alonzo Fyfe

Copyright

Table of Contents

Morality from the Ground Up
Alonzo Fyfe

Abstract: The institution of morality involves the use of the social tools of reward and punishment to mold desires and, thereby, to regulate intentional action. In this, praise serves as a type of reward, and condemnation as a type of punishment. In this article, I attempt to give some support to this idea by showing how we can derive the primary elements of what we commonly recognize as the practice of morality from some simple components. All we need are two or more intentional agents with malleable desires – desires that can be altered through interaction with the environment. The specific form of interaction we are interested in are rewards and punishments. From this, we can derive the three moral categories of obligation, prohibition, and permission; supererogatory action, the truth-bearing and the emotive components of moral statements, the dominant use of reward and punishment in morality, heroism, the concepts of "excuse" and "mens rea", and both the objective and the subjective nature of morality.

I. Introduction

I can build a complete moral system with just three ingredients. What I need are:

(1) Two or more independent agents
(2) Who engage in intentional actions (that is, act so as to fulfill their desires given their beliefs)
(3) Whose desires can be molded by rewards and punishments (including praise and condemnation).

Give me these three things, and I can build a complete moral system filled with objectively true moral propositions. It is a moral system where lying, breaking promises, theft, vandalism, assault, rape, and murder are all objectively wrong. These are no mere matters of opinion or personal taste.

Just to be clear, these initial ingredients do not include evolved moral sentiments. Nor do I need empathy or inherent dispositions to cooperate and help each other. It contains no categorical imperatives or a natural moral law. I do not need a social contract, or a god, or any giver of moral laws.

All I need is a community made up of individual intentional agents whose desires can be molded by environmental factors of a type that count as reward and punishment.

The next question I expect a reader to ask is, "What do you mean when you say you can build a complete moral system?"

What I mean is that I can account for the bulk of the components that we have come to associate with a moral system. I can explain the central role that reward and punishment – including praise and condemnation – play in the institution of morality. I can explain the fact that there are three moral categories of action (obligation, prohibition, and non-obligatory permission). I can explain the moral concepts of negligence, excuse, *mens rea*, and supererogatory (heroic) action. I can account for both the truth-bearing part of moral claims that allow us to treat them as regular propositions and argue whether they are true or false. At the same time, I can also explain their emotive component – the fact that moral claims tend to express attitudes of approval and disapproval.

Please note, I am not claiming to be able to explain our specific moral intuitions about specific actions. Those moral intuitions tend to be nothing more than the prejudices of the age. Throughout human history, we can count many times when human moral intuitions were quite comfortable with slavery, racism, and genocide. The ability to rationalize a prejudice is not a mark of success.

Instead, what I am aiming for is an understanding of how morality works, and not necessarily the conclusions that a group of people believe it supports.

To accomplish this task, I will begin by looking at an imaginary universe that contains just one person (Alph) who has just one desire (a desire to gather stones). To build a moral system, I will need to add many more people with many more desires. However, we must begin with an understanding of a lone individual with just one desire.

II. One Person, One Desire

I wish to start the task of building a complete moral system with the simplest community possible – a community of one person that has just one desire.

I invite you, the reader, to approach this simple community with the same attitude you would have to an astronomer who seeks to explain gravity first by imagining a universe with only two bodies, describing how they would move in relation to each other. We know that the real universe is much more complex, with countless bodies exerting their gravitational influence, and other forces altering the motions of things. However, this simplified view makes it easier to grasp the basic components of that more complex universe.

Similarly, physicists sometimes explain concepts such as acceleration, momentum, and inertia by imagining a system with frictionless surfaces and massless strings. These assumptions are false in the real world, but they help in explaining those concepts.

I will be using the simplified system of one person with one desire to explain some elements of intentional action.

I want to specifically warn against seeing this account as following the tradition of "man in a state of nature" that social contract theorists such as Thomas Hobbes and John Locke use to develop their moral systems. would use at the start of their theories. I am not in any way imagining that there ever was a case of a single person with a single desire as I describe here, or that we can draw direct implications about what we ought or ought not to do in the real world.

This is also not a "veil of ignorance" type of thought experiment. I am not going to argue that the conclusions that people will reach in some type of contrived situation imply that we who are living in the real world ought to adopt those same conclusions.

I simply want to take this single person with a single desire and study it for a bit, just to identify what we can say about this entity.

I will name my subject, "Alph". The single desire that Alph has is a desire to gather stones. I am also going to place Alph on a world that has only a limited number of stones to gather. This world is made up mostly of solid rock, with a few stones scattered about.

III. Desires, Beliefs, and Propositional Attitudes

We have a being (Alph) with one desire (to gather stones). What does Alph now do?

Alph does nothing.

If we give Alph a desire to gather stones, and we put Alph out in the field, Alph is simply going to sit there. He is going to have no idea what to do. He has this desire to gather stones. However, he has no idea of what a stone is, how to transport a stone, or where to take it.

If we are going to get Alph to do anything, then Alph needs to have a set of beliefs to accompany his desires – beliefs that will allow him to select the actions that will realize that which he desires. These beliefs allow him to identify the things in his environment that are stones. They allow him to know how to transport the stone. And they tell him where to place the stone - where the gathering point is at and how to navigate to that spot.

Intentional actions require both desires and beliefs.

So, what are beliefs and desires?

Beliefs and desires are propositional attitudes. That is to say, when we talk about an agent's beliefs and desires, we are talking about an agent's attitude towards a proposition.

For example, let us take the proposition, "I (Alph) am gathering stones". If Alph believes, "I am gathering stones" then Alph is in a state where he takes the proposition, "I am gathering stones" to be true. For Alph, it accurately describes the world.

If Alph is in a state of desiring "I am gathering stones", then Alph has a motivating reason to make it the case that "I am gathering stones" is true. Naturally, Alph can desire to be gathering stones while, at the same time, believing that it is not the case that he is gathering stones. Perhaps he recognizes that he is chained to the ground and cannot gather stones.

A proposition is the meaning of a sentence. Two different sentences can express the same proposition. The sentence "La casa est blanca" and "the house is white" are two separate sentences expressed in two different languages. However, so long as we assume that they are both talking about the same house, the two different sentences identify the same proposition. They both mean the same thing.

So, a propositional attitude is an attitude towards - not "a sentence", but towards whatever it is that the sentence means.

In this paper, I will frequently make use of the notation that "Alph believes that P" and "Alph desires that P". In both cases, 'P' is a proposition that identifies that which is believed or desired.

Expressing my previous point more generally, if Alph has a belief that P, then Alph has the attitude that the proposition 'P' is true – that 'P' accurately describes the universe in which Alph lives. If Alph has a desire that P, for any proposition P, then Alph has a motivating reason to make it the case that P is true or, if P is already true, then to make it the case that P continues to be true.

Beliefs merely report facts about the world. Desires motivate agents to act – either to change the world, or to prevent the world from being changed.

In our case, Alph has a desire "that I (Alph) am gathering stones". Consequently, he is motivated to make it the case that "I am gathering stones" is true – or to keep it true if he judges that it is already true. In making it the case that the proposition is true, Alph brings in all sorts of knowledge about the world – knowledge about what a rock is, now to transport a rock, and how to navigate to the gathering site. However, these beliefs do not select Alph's goal or end – they merely identify the means for achieving that end.

As long as Alph has the requisite true beliefs, Alph knows what he wants (to gather stones), he knows how to go about doing it, and he has the motivation to carry out the actions. This is enough to make it the case that Alph goes out into the field and gathers stones.

IV. Means, Ends, and Unintended Consequences

One of the things that I wish to accomplish here is to show how this makes sense of some of our commonly used terms for describing aspects of intentional action. Here, I will focus on means, ends, and unintended consequences.

We have Alph walking around his world gathering stones.

We can recognize a distinction here between Alph's end or goal and the unintended consequences of his actions.

Alph, in this example, does not have a desire to have all the stones put in one large pile. Because of his desire, his stones end up being gathered in a huge pile. However, this is an unintended side-effect of his desire to gather stones – it is not in any sense his goal or something he wants to have happen.

Recall that desires are propositional attitudes. Different propositions identify different desires. Alph's desire to gather stones is a desire that the proposition, "I am gathering stones" is made or kept true. A desire that the stones be gathered is a desire that the proposition "The stones are all collected in one location" be made or kept true. The proposition, "I am gathering stones", is not the same as the proposition, "all of the stones are gathered into one large pile". Consequently, the desire to gather stones is not the same as a desire that all of the stones end up in a big pile.

The consequences of fulfilling a desire are not always desired. In fact, it is commonly the case that some of the consequences of fulfilling a desire are unwanted.

One of the potential consequences of having sex is pregnancy. The fact that one desires sex does not imply that one desires pregnancy. In fact, an individual who desires sex might have many and strong reasons to avoid pregnancy. Alph's relationship to having a big pile of stones is the same as some lovers' attitudes towards pregnancy. It is not something the agents necessarily want – and it is something the agents can potentially have reasons to avoid. Though, in our basic example, Alph does not care about the big pile of rocks one way or another.

Here is another way to see the difference between a desire to gather stones and a desire that the stones be gathered. For the person with a desire that the stones be gathered, once they are gathered, he would sit back in contentment. He has made the proposition, "All of the stones are gathered into one big pile" true and, now, all he needs to do is to keep it true by preventing the stones from being scattered again.

In contrast, once all of the stones are gathered, the agent with a desire to gather stones is facing a problem. His desire is pushing him to return to a state in which the proposition, "I am gathering stones", is true.

One of the ways that Alph can restore a state in which, "I am gathering stones" is true is by scattering the stones he has gathered. In this case, scattering the stones is work. It is a task that Alph does not find any value in doing for its own sake, but is something he does only as a means to some other end. In other words, scattering stones has instrumental value. It has value – but as an instrument to be used in realizing something else that has value as an end.

For Alph, in this simple case, the only thing that is "good in itself" is gathering stones. Having a large pile of stones is not good in itself - it is not good in any sense at all. It is just a consequence of gathering stones. Scattering stones, also, is not good in itself. It's work. It is a chore that one must do - even though one may wish to be doing something else – so that one can return to a state in which one is realizing that which one values as an end.

We have here, then, a distinction between means, ends, and unintended consequence - all related to Alph's desire to be gathering stones. That which is valued for its own sake is any state in which the proposition P, where P is the object of a desire, is made or kept true. That which has value as a means is anything that can help to bring about a state in which such a proposition is made or kept true. Finally, in making or keeping a proposition true, a person's actions can create other truths that the agent did not desire. These are its unintended consequences.

We have, in our imaginary world, a person with the name of Alph with a single desire - a desire to gather stones. This desire, expressed in terms of a propositional attitude, is a "desire that I am gathering stones".

V. Ultimate Ends

One of the ideas that one often encounters in moral philosophy is the idea that there must be one ultimate end – one thing, and only one thing, that everybody seeks as an end in itself, and for the sake of which everything else is a means.

Aristotle told us that this sole end was eudemonia or "flourishing". According to Jeremy Bentham, it was pleasure and the absence of pain. John Stuart Mill told us that it was happiness.

In our simplified universe, Alph has only one end – to gather stones. Nothing else has value, except insofar as it provides a way to serve this one end. Eating, drinking, avoiding injury, matters to Alph only insofar as they make it possible for Alph to continue to gather stones.

However, this is true only in virtue of Alph's sole desire to gather stones. If Alph had two desires – a desire to gather stones and aversion to pain – then he would have two ends. He would then be seeking to (a) gather stones, and (b) avoid pain, and hopefully do both at the same time.

Each desire takes the form of a propositional attitude – a motivation to make or keep true a proposition P. Each desire establishes the truth of the proposition 'P' as an end. For each desire that an agent has, the agent has a different end. An agent that has three desires – a desire that P, a desire that Q, and a desire that R – has three ends; the realization of P, Q, and R.

For Alph, the desire to gather stones provides him with a motivating reason to gather stones - to make or keep true the proposition "I am gathering stones". If Alph should be asked, "Why are you gathering stones?" Alph's answer would be, "Because I want to." In this imaginary world as so far constructed there is no other reason to gather stone - no other reason exists.

Nozick's Experience Machine

Robert Nozick presents an argument that shows that it is not the case that humans have only one end such as pleasure, happiness, or some other mental state. This argument works against all brain-state theories – all theories that suggest that the only thing that matters to any person is that the molecules in their brain are placed in a particular configuration that instantiates "being happy" or "experiencing pleasure".

Nozick has asked that we consider the option of an experience machine – a machine that will artificially generate the impressions of any given state of affairs.[1] If one wants to be a successful novelist, then, by entering the experience machine, one will lay in a pod where a computer attached to sensory nerves will give one the experience of being a successful novelist. There will be nothing to inform the individual that she is not enjoying a fantastic life.

Many people presented with this option have no interest in entering the experience machine. In contemplating the option, it has little or no value. They would rather live in the real world and experience the frustrations of having a life that does not turn out entirely the way they would like it to over being fed the false impression that they were living an ideal life.

The problem comes in trying to explain this observation.

On the account being presented here, the explanation is found in the fact that people have desires that the experience machine cannot fulfill.

Recall that a desire that P is a motivating reason to act so as to realize a state in which P is true. For Alph, who has a desire to gather stones, the experience machine cannot make it the case that the proposition, "I am gathering stones" is true. Consequently, Alph has no motivating reason to enter the experience machine.

The experience machine would be a good option if the only thing that mattered to a person was that the molecules in their brain be put into a particular configuration. The experience machine, by definition, has the capacity to create and maintain that configuration. Insofar as one values happiness, or values pleasure and the absence of pain, or values satisfaction or contentment, the experience machine can provide it, and agents have reason to get hooked up to the machine.

However, if what one truly values is that the world outside of the brain have a particular configuration – that one's children are healthy and happy, that one is helping to provide medical care to impoverished children who would otherwise suffer, that one is actually entertaining an audience, and the like – the experience machine offers nothing of value. It's capacity to put matter in a particular configuration applies only to brains.

---

[1] Nozick, Robert (1974). *Anarchy, State, and Utopia*. New York: Basic Books. pp. 42–45.

I find this case to be clearer if we take the case of somebody who desires the well-being of their children. In this example, a malevolent scientist gives her two options.

Option 1: "I will hook you up to this experience machine. You will experience a life in which you are given constant assurances that your children are healthy, happy, and otherwise doing well. However, in fact, your children will be taken to a room down the hall where I will torture them mercilessly."

Option 2: "I will hook you up to this experience machine. You will appear to see your children in the next room being tortured mercilessly. You will hear their screams, see their suffering, and you will believe it. However, in reality, I will make certain that your children are provided for in every way. They will have money, a quality education, and the best medical care that will concern not only their physical health but their mental health. They will be encouraged to seek challenge and self-improvement and be given the tools to meet those challenges."

If we assume that the parent truly cares about their child's well-being, it seems that the parent would choose Option 2. Option 2 means a great deal of personal suffering. On the other hand, it is the option within which the proposition, "My children are doing well and living a good life" is true. Since desires motivate those who have them to realize the propositions that are the objects of those desires, then the parent's desire that their child have a good life gives them a motivating reason to choose the second option – not the first. The first option is only appealing to the person who desires "that I experience pleasure" – who would experience pleasure by believing that their children are well off. [2]

Objects of Desires

Another reason to doubt that we have one ultimate end that happens to be the configuration of certain molecules in the brain is found in its inability to account for the incommensurability of ends.

If I were to give you a choice between either taking $100 or $200, unless you are facing some highly unusual circumstance in which having more than $100 will thwart other desires, you would take the $200 without a second thought to the $100 option. You would have no reason to regret the option in which you had selected $100 rather than $200. This is because everything you value in having $100 is covered by the option of having $200. The person who opts for $200 has not lost anything and has no reason for regret.

However, it is often the case, when we must make a choice, that we are choosing among goods that are not substitutable. The individual who must choose whether to attend a distant university or stay in her home town with her family, friends, and significant other is not facing a choice between more or less of the same good (happiness). She is facing a choice between some of one good (advancing their education and entering into a desired profession), or some of another good (the continuing company of friends and family) that cannot substitute for it.

A couple who loses a child to illness or injury and who decides to have another child they would not have otherwise had may be able to replace a part of the value of the one child with the other. However,

---

[2] There is some indication that people prefer the "status quo" and would elect to remain in an experience machine if they were to discover they were already in one. However, even an interest in the status quo is distinct from an interest in a mental state. It would be interesting to discover if parent who are told to assume that they have always been in an experience machine, during which time their children have been made to suffer, would choose the status quo if it include the continued suffering of their children.

if they were to say, "Oh, it doesn't matter that Sally died. We have Jimmy and that is just as good," would seem odd at best.

We can account for the incommensurability of goods by the idea that each desire creates its own distinct end. A person with a desire that P, and a desire that Q, where the state "P and Q" is not possible, will regret not having P if he selects Q, and will regret not having Q if he selects P. The desire to stay in one's hometown with one's friends and family is one end. The desire to attend a distant university is another end. One must make an actual choice between them. Whatever one chooses, one must give up a set of goods – a set of ends. This brings regret.

Evolution and The Limits of Ultimate Ends

Another consideration rests in the idea that evolution will make individuals care about nothing more than the organization of molecules in the brain and not about the organization of other molecules in the world when it is the organization of molecules in the world that determine survival and reproduction. Evolution has more reason to cause us to have a concern for the organization of molecules in the world for the organization of molecules in our brains.

In fact, evolution has obviously given us several ways to determine the organization of matter in the world. The senses of sight, sound, touch, smell, and taste all give us information about the world. Our access to information about the organization of matter in our brains, however, seems limited. We have no clear organ that seems designed to give us that information, and the information we get seems unreliable.

As I will argue later, our brains also seem to be malleable. The things we can come to desire is not genetically encoded but, instead, is acquired – or at least modified – by experience. There seems to be no limit to the set of things that an individual can desire – that can become one of an agent's ends.

Evolution will likely dispose us to adopt ends that promote our evolutionary success. However, what promotes our evolutionary success depends on the environment. Consequently, there is room for a wide variety of interests consistent with evolutionary success.

In making these claims about evolutionary fitness, I am not saying that evolutionary fitness is, itself, an end. I am only saying that our brain has evolved and, as such, when it comes to describing the brain, we have reason to assume that a long history of evolution has molded its form and structure. However, for an agent with a desire that P, their end is to realize a state of affairs in which "P" is true. Evolutionary success is often an unintended consequence, not an end, and not a measure of the value of the desire that "P".

Recall that beliefs are also propositional attitudes. A person who believes, "I am happy" has the attitude that the proposition "I am happy" is true. There seems to be no limit to the propositions that can be believed. A person can believe, "I am Napoleon Bonaparte", or that water is made up of H2O, or that humans are the product of a long history of change brought about by random mutations to DNA combined with a process of natural selection.

There is no reason to hold that the ultimate objects of desire cannot be just as broad. Of course, a person can have a desire to be happy (a "desire that I am happy") - and all or almost all of us have this desire. Yet, there is no reason to rule out the possibility that an agent can also have a desire to gather stones. A person can have a desire to be having sex, or a desire to eat, or a desire to not be in pain.

Nor is it the case that a desire must necessarily take the self as an object. One can desire that their children are happy, or desire that no person go hungry. A person can have an aversion to any person having sex with somebody of the same gender (a desire that no person engage in homosexual activities) or a desire that a wilderness area be preserved even though no person will experience that wilderness. In fact, he can desire that no person experience that wilderness and that it remain in its natural state.

I will come to argue that morality depends heavily on the diversity of desires. Morality will be concerned primarily with promoting or inhibiting certain desires. It will be involved in creating in people a desire to repay debts, and aversions to telling lies. Morality will involve promoting the desire that no person go hungry and a desire to accurately distinguish fact from fiction.

In fact, I suspect that the set of propositions that can become the object of a desire is no more limited than the set of propositions that can become the object of a belief. Anything that can be believed can be desired. If a person can adopt the attitude that P is true, then a person can adopt the attitude that P be made or kept true.

In our hypothetical universe, at this point, we have one person with one desire - a desire to gather stones. In this universe, the state in which "I am gathering stones" is true is the end of all intentional actions. It is the one sole state that all beings in this universe - which, importantly, consists of Alph alone - aim. It is not happiness, or eudaimonia, or pleasure and the absence of pain, it is "I (Alph) am gathering stones" - and nothing else. And the reason it is "I (Alph) am gathering stones" is because Alph - the only intentional agent that exists - has a desire to gather stones.

VI. Reasons for Action

I am going to be writing a lot about what people have reasons to do. When I do so, I will be following Barnard Williams' model.

 I will make the connection by bringing in Bernard Williams' thesis:

> *A has a reason to φ iff A has some desire the satisfaction of which will be served by his φ-ing.*[3]

We are examining a universe that contains one person (Alph) with one desire (to gather stones). Alph has a reason to gather stones if and only if Alph has a desire the satisfaction of which will be served by gathering stones. In this case, I have stipulated that Alph has such a desire. Consequently, Alph has a reason to gather stones.

Recall that I have argued that "serving a desire" is to be understood to mean making or keeping the proposition that is the object of the desire true. If an agent has a desire that P, then any act that helps to realize a state in which "P" is made or kept true is an action that the agent has a reason to do. Alph's desire to be gathering stones is served by any action that helps to realize the proposition, "I, Alph, am gathering stones."

If Alph has gathered all of the stones together, he will then have a reason to scatter stones. Having all the stones gathered means that the proposition, "I am gathering stones" is now false. The only way to

---

[3] Williams, B., 1979. "Internal and External Reasons," reprinted in Moral Luck, Cambridge: Cambridge University Press, 1981, 101–13.

change that from "false" to "true" would be to scatter the stones. Then, he can gather them again. That is the reason for scattering stones – to restore a state in which "I am gathering stones" spoken by Alph can once again be made true.

Alph also has a reason to preserve his own life. He cannot make or keep the proposition, "I am gathering stones" true if he should die. Consequently, he has a reason to avoid walking off a cliff as he gathers stones. If he needs food and water to survive, he has reason to eat and to drink – and to grow food.

However, in our example, he has no motivating hunger or thirst. These would require the addition of new desires – hunger and thirst. In this example, he eats and drinks only as a means to the end of gathering stones and for no other reason. There is no other reason for him to eat.

Alph may or may not have reason to look in the next valley for more rocks to gather. He has a reason to look in the next valley just in case there are rocks in the next valley to gather. If there are no rocks to gather, he has no reason to look there.

We may assume that he does not know whether there are rocks in the next valley. This means that he does not know whether he has a reason to go into the next valley for rocks. He knows that there is a possibility that there are rocks in the next valley to gather. This means that he knows that there is a possibility that he has a reason to go to the next valley. The possibility of rocks generates a possibility of a reason. It does not generate an actual reason. That would require that there actually be rocks there.

A person who knows that there are no rocks in the next valley can honestly tell Alph, "You are wasting your time looking for rocks in the next valley."

We can illustrate this another way by looking at an example where a thirsty person reaches for a glass that contains what she thinks is clean water, though it is in fact poison. An observer can tell her, "You don't want to do that." When she asks why, he reports that it is poison. The point here is that his original statement, "You don't want to do that," is true. The agent may not know that she does not want to do that, and may think that she does want to do that insofar as she is thirsty and thinks that the glass contains clean water. However, in fact, she does not want to drink from the glass. She has no reason to drink from the glass.

Another way of stating this is to say that there is no such thing as an external reason – a reason that exists independent of desires and aversions. There is no reason built into the fabric of the universe, or that resides entirely within some object of evaluation. If it does not contribute to making true some P that is the object of a desire that P, then there is no reason to perform the action. Even when if does make this contribution, the only people who have a reason to perform the action are those with a desire that would be served by it.

VII. Deriving "Ought" from "Is"

At some point, many readers are going to demand that I address the challenge that one cannot derive an "ought" from an "is".

However, this challenge is only applicable to an "ought" that claims to identify some sort of intrinsic or non-material value or categorical imperative. When it comes to hypothetical imperatives, the process of deriving an "ought" from an "is" is not considered problematic. The way that this theory answers the

"ought" from "is" problem is by stating that all value claims – including, it will turn out, moral claims – are hypothetical imperatives.

An example of a hypothetical imperative is:

> *(1) If I want to get to the library via the shortest route, I should turn left on Main Street.*
>
> *(2) I want to get to the library via the shortest route.*
>
> *(3) Therefore, I should turn left on Main Street.*

This argument form is valid. Consequently, if the premises are true, then the conclusion must be true. And the premises are true - when they are true - as a matter of empirical fact. The "want to get to the library by the shortest route" is simply a description of how the person's brain is wired.

We would be guilty of an equivocation if we were then to take the conclusion and turn it into something like a categorical imperative. The fact that one ought to turn left on Main Street if one wants to take the shortest route to the library does not imply that one ought to turn left on Main Street *simpliciter* – independent of the interest to use the shortest route to get to the library. The instant each of the premises becomes false, the conclusion becomes unsupported – unless the agent has some other reason to turn left on Main Street (some other desire that could be fulfilled by such an act).

In the case of Alph, with his one desire to gather stones, he ought to gather stones. Similarly, he ought to avoid injuries, since injuries would prevent him from being able to gather stones. He ought to build a way to safely get to the bottom of the cliff – where there are a number of stones to be gathered – and haul up the rocks that he finds there. He ought to eat, when eating will serve his interests in gathering stones, and stop eating when eating no longer serves that interest. All of these "oughts" can be derived from "is" premises in a straight-forward and non-problematic way.

We will see later that moral imperatives will simply be a form of hypothetical imperative. Consequently, we will see that we can derive moral 'ought' from 'is' in a way that is just as unproblematic as we find to be the case for practical desire-based 'ought'.

VIII.   Desire Satisfaction Theories

The most common misinterpretation of this theory is to interpret it as claiming that desire satisfaction has intrinsic value and that morality consists in maximizing desire satisfaction – creating as much of it as possible. That is emphatically not what I am arguing for.

We can understand desire satisfaction in two ways.

Some desire satisfaction theories look at the felt sense of satisfaction that an agent has when he believes that a desire of his has been fulfilled. In this sense, satisfaction serves the same role that "pleasure" serves in traditional hedonism. "Frustration" serves the same role as "pain" in traditional hedonistic theories.

We may call this a subjective satisfaction theory, since the satisfaction depends on the agent's beliefs and the subjective sensation that comes from those beliefs and not on whether those beliefs are true.

That is not the type of theory I am defending here.

I am defending an objective desire satisfaction theory. This type of theory says that a "desire that P" is satisfied if and only if a state of affairs exists in which 'P' is true. A desire that one's children are healthy and happy is only satisfied, in this sense, if the proposition, "My children are healthy and happy" is true in fact. The agent may derive a strong sense of satisfaction if she believes that her children are healthy and happy. However, that satisfaction counts for nothing if the children are not healthy and happy in fact.

In talking about an objective desire satisfaction theory, there is room for a second type of mistake. The first mistake looks for the sense of satisfaction or frustration that comes from believing that a desire has been fulfilled or thwarted. The second mistake is to think that desire satisfaction itself has objective, intrinsic value. This is not true either.

Intrinsic value, in fact, does not exist.

Alph seeks one thing - to be gathering stones. He is not seeking a physical sensation of satisfaction. Nor is he seeking desire satisfaction for its own sake. Alph, with his desire to gather stones, is seeking a state of affairs in which "I am gathering stones" is true. Any state of affairs that meets this condition, "I, Alph, am gathering stones" is a state of affairs that Alph has a reason to realize. Its value is not found in the fact that Alph feels satisfaction, nor is it found in the brute fact that desire satisfaction exists. Alph finds value in it because Alph desires a state in which "I am gathering stones" is true.

To put some distance between this theory and theories that refer to the felt sense of satisfaction, I tend to use the term "desire fulfillment" to refer to a state, given a desire that P, where "P" is true. A state of affairs in which Alph is gathering stones is a state in which Alph's desire to be gathering stones is fulfilled. There is no need for a felt sensation of satisfaction. This does not imply that such a sensation cannot exist or that a person cannot desire "that I am experiencing a felt sense of satisfaction". It simply is not necessary and is not a part of our model at this point.

It may help to see the difference between desire satisfaction and desire fulfillment if we imagine Alph with a different desire. Let us give Alph a desire that the planet Pandora B come into existence. The only way for Alph to bring Pandora B into existence is to press a red button labeled "Create Pandora B". However, pressing the button will destroy Alph. He will cease to exist as Pandora B comes into existence.

For the sake of this example, Pandora B will come into existence without any form of life – without any creature capable of having desires and without the possibility that a desiring creature can evolve on Pandora B. We do not want to confuse the issue by adding additional desires.

If we look at the reasons for action that exists, Alph has a reason to press the button. Alph has a desire that P where P = "The planet Pandora B exists". Pressing the button will serve this desire. Consequently, Alph has a reason to press the button.

Alph has no reason not to press the button. Certainly, pressing the button will end his existence, but Alph has no desire to continue to exist. Nor does he have a desire that is served by his continuing to exist. This all means that he has no reason to continue to exist or to refrain from pressing the button. He

has this one desire – that Pandora B come into existence – and the fulfillment of that desire requires that he live no longer than needed to press the button.

So, he presses the button.

He ceases to exist, and Pandora B springs into existence.

The desire has been fulfilled. That is to say, the check box "Pandora B exists" can now be checked – this has been made true.

However, pressing the button creates no desire satisfaction in either sense of the term described above. Alph will experience no sense of desire satisfaction. Nor will there be any objective desire satisfaction since the instant that Pandora B comes into existence, the desire that Pandora B exists will cease to exist.

IX. Pushpin and Poetry

Jeremy Bentham wrote:

> *Prejudice apart, the game of push-pin is of equal value with the arts and sciences of music and poetry. If the game of push-pin furnish more pleasure, it is more valuable than either. (Jeremy Bentham)[4]*

In a sense, this is true.

It is not the case that pleasure is that which gives push-pin or poetry value. What matters is whether push-pin or poetry serves a desire. If push-pin serves a desire (either directly or indirectly) and poetry does not, then the agent has a reason to play push-pin and no reason to pursue poetry. The claim that agents have reasons to pursue poetry are true if and only if poetry serves an agent's desires, either directly or indirectly. The value of reading poetry and the value of push-pin both are determined by the degree to which they tend to fulfill the desires in question.

Pushpin is a game requiring two players, so it is not suitable for our imaginary world that contains just one person (Alph). However, we can make the same point by substituting the act of gathering stones for pushpin.

Prejudice apart, the act of gathering stones is of equal value with the arts and sciences of music and poetry. What gives a state of affairs in which "P" is true any end-value is that it is the object of a "desire that P". If Alph desires gathering stones, then a state of "I am gathering stones" has value for Alph as an end. If Alph has no interest in reading poetry, then a state of "I am reading poetry" has no value for Alph as an end.

In a more complex universe - one with more people and more desires, there may be elements that give more value to one over the other. Poetry promotes the ability to read, which provides instrumental value by aiding communication. Poetry may also give people a better understanding of the emotional states or of other points of view. There may be a simple second order preference that the population

---

[4] Bentham, Jeremy, *Rationale of Reward*, http://www.la.utexas.edu/research/poltheory/jsmill/diss-disc/bentham/bentham.xr18.html, retrieved 01/21/2017.

have more people who desire to read poetry and fewer who value gathering stones. These would count as additional reasons to read poetry - and to promote a desire to read poetry.

However, in our introductory world, where we have Alph, and he has one desire – a desire to gather stones – gathering stones has value (for Alph), and reading poetry does not.

In addition, an agent who has a weak desire to gather stones and a strong desire to read poetry has a stronger reason to read poetry then he does to gather stones. He would prefer reading poetry under conditions where he could do one or the other, but not both. He would, of course, prefer a way to read poetry while he gathers stones.

In the future, I will introduce a second desire and examine some of its implications - but there is still more to say about a single agent with a single desire.

It is said that we have some sort of intuition that poetry is better than push-pin or gathering stones. However, that judgment comes from our own likes and dislikes. We are the ones who prefer that Alph spends his time reading poetry rather than gathering stones. Consequently, the world in which Alph is reading poetry has more value to us, but it still has no intrinsic value and Alph still lacks a reason to realize such a state.

There is nothing in poetry that generates a reason or a natural command that it be read – by Alph, or by anybody else.

There is no good or bad, better or worse, best or worst, that does not relate an object of evaluation to a set of desires. In a world where only one desire exists – a "desire that P" - (e.g., Alph's desire to gather stones) then all goodness and badness that exists relates objects of evaluation to that desire – either to make or keep "P" true, or to make or keep it from being true.

X. G.E Moore's Beautiful World

> Let us imagine one world exceedingly beautiful.... And then imagine the ugliest world you can possibly conceive. Imagine it simply one heap of filth, containing everything that is most disgusting to us, for whatever reason, and the whole, as far as may be, without one redeeming feature. The only thing we are not entitled to imagine is that any human being ever has or ever, by any possibility, can, live in either, can ever see and enjoy the beauty of the one or hate the foulness of the other.... [S]till, is it irrational to hold that it is better that the beautiful world should exist than the one which is ugly? Would it not be well, in any case, to do what we could to produce it rather than the other? Certainly, I cannot help thinking that it would; and I hope that some may agree with me in this extreme instance. (G.E. Moore)[5]

These imaginings can serve a couple of purposes.

In one purpose, this serve as an argument against the idea that the only thing humans seek is the establishment of a particular brain state - be it happiness, satisfaction, or a surplus of pleasure over pain. It shows us that there are things that humans will choose - such as the realization of a beautiful

[5] Moore, G. E., *Principia Ethica,* http://fair-use.org/g-e-moore/principia-ethica/s.50, Retrieved 01/21/2017.

planet rather than an ugly planet - even were the options do not include manifestations of particular brain states.

We see this interpretation particularly when we focus on the phrase, "Would it not be well to do what we could to produce it rather than the other?" Moore is asking about what we have reason to do, and what we have reason to do is something other than creating or maintaining particular brain states.

On this measure, Moore's example describes something real, and the account that I am presenting here can make sense of it.

Earlier, I gave Alph a desire that the planet Pandora B exists - a planet that may well fit Moore's description of an exceedingly beautiful world. This desire is fulfilled in any universe where the proposition "Pandora B exists" is true - even in a universe where the propositions "Alph is experiencing Pandora B" becomes false.

In fact, we imagined a case in which Alph could bring the planet into existence only by destroying himself. Yet, given our initial assumptions – that he had a desire that Pandora B exist and no desire that would be thwarted by his own destruction – he had reason to bring the planet into existence, and no reason not to, even though doing so would result in his own destruction.

However, there is another interpretation of the example. It may be interpreted to mean that the value of the beautiful planet exists within the planet itself. In a universe where no being capable of experiencing it even exists, it would still be better - in some objective, intrinsic sense – that the universe with the beautiful planet exist rather than the universe with the ugly planet.

Our model of Alph and his one desire does not support this.

One point to be made against the example interpreted this way is to ask, "Which is the beautiful world, and which is the ugly world?"

In imagining the ugly world, Moore asked us to imagine a world "of filth, containing everything that is most disgusting to us".

Note the phrase, "to us".

Let us imagine another creature - one that evolved from the dung beetle. Imagine the ugly planet being a huge ball of dung. Certainly, insofar as it were within our power, we may have reason to bring about the world that is beautiful to us and not the dung world. However, the dung beetle creatures have reason to bring about world that is beautiful to them - the big ball of dung - rather than the world of our imaginings.

This creates a problem for the idea that we can completely remove the idea of some type of creature with desires and interests in determining what is beautiful. If we cannot make a judgment of which planet is beautiful without having a subject and knowing something about them, we cannot make a judgment as to which world should exist even if could never be visited or seen.

Now, let us return Alph to his original state - the one where he has one desire, a desire to gather stones.

Now, ask him which planet should exist.

He would likely answer with a shrug, though he has no reason to do even that.

To start with, the phrase "beautiful world" would confuse him. He has no preferences for anything that might go in to the construction of a beautiful world. At best, his idea world would be one in which there are a great many rocks to gather. Yet, even this would have no value if he were to also imagine that he could not gather them. Such a world could not contribute to making or keeping the proposition, "I am gathering stones" true. Because they have no relevance, they give Alph no reason for action - not even a reason to choose to answer the question.

XI. Socrates and The Pig

> *It is better to be a human being dissatisfied than a pig satisfied; better to be Socrates dissatisfied than a fool satisfied. And if the fool, or the pig, is of a different opinion, it is only because they only know their own side of the question. (John Stuart Mill)[6]*

As I continue to examine what is true in a universe with one person (Alph) with one desire (to gather stones), I want to turn my attention to another famous example from the history of philosophy. This is John Stuart Mill's quote provided above.

In the full context of this quote, Mill was arguing that some pleasures were better than others - intrinsically better.

Again, as with the discussion on Bentham, we are not concerned here with pleasure. However, we are concerned with the relative value of "a human being satisfied", "a pig satisfied", "Socrates dissatisfied", and "a fool satisfied".

Which is better?

Actually, we cannot answer this question without asking the further question, "Valuable to whom"? Once we know the answer to that question we can determine the deisres that this person or these people have. From here we can ask whether the propositions "P" that are the objects of those desires are true in states of affairs in which a human is satisfied, a pig is satisfied, Socrates is dissatisfied, and a fool is satisfied.

If there is a being that has a desire that "unsatisfied Socrates" exist and has no desire that "fool satisfied" exist, then that person would choose "unsatisfied Socrates". If, instead, that person prefers "fool satisfied", then that person has a reason to choose "fool satisfied".

We cannot tell the value of any state of affairs except as far as there is a desire that P, at which point we can then determine whether "P" is true in that state of affairs. If it is, then the agent has a reason to choose that state of affairs. If not, he does not.

Let us return to our world with its one person (Alph) and his one desire (to gather stones). Now, let us compare a case where Alph is on a world that has stones to one where Alph is on a world that has no stones. Which world has the greater value?

We still can't answer the question without first answering the questions, "to whom?" and "what does that person want?"

---

[6] Mill, John Stuart, *Utilitarianism*, Chapter 2, https://www.utilitarianism.com/mill2.htm, Retrieved 01/21/2017

The first Alph prefers the first universe. He would prefer to be in a world where he has stones to gather and has a reason to choose that over being dropped off on a world without stones. It describes a case in which "I (Alph) am gathering stones" is true.

However, if we change this and look at the situation from the point of view of a neutral observer, then the question of which that observer would prefer depends entirely on what that observer desires. If an impartial observer truly was impartial, she would have no reason to choose one world over the other. If, instead, we give this impartial observer some desires, then we look for the world within which the propositions that are the objects of those desires are true.

Now let us look at two worlds where one contains a human dissatisfied and the other contains a pig satisfied, and let our impartial observer look upon these two worlds.

Here, too, we need to identify the desires the observer has, the strengths of those desires, the propositions that are the objects of those desires, and whether those propositions are true in either world. When we do this, we will identify the universe the observer has the most and strongest reason to choose by identifying the one that fulfills the most and strongest of those desires.

The same procedure applies to determining if it is better to be Socrates dissatisfied or a fool satisfied. I think it is safe to assume that if we take the desires of Socrates as our starting point, such a person would have many and strong reasons to choose against being the lucky fool.

Once again, the decisions of an impartial observer will depend on the desires we give this imaginary person, in which case the world to choose in the world in which the propositions "P" of the most and strongest desires will be made true.

This, then, is how these comparisons among possible worlds are to be made. What matters are the desires of the person making the choice. The desires that the being the chooser becomes are relevant only to the degree that they directly or indirectly make true the propositions that are the objects of the chooser's desires.

XII. The Value of a Life

I fear that there are some potential readers – some of whom have set this paper aside long before now – because they look upon the Alph, spending his life gathering and scattering stones, as a rather pathetic life. They look at this life we are imagining in which Alph gathers stones, scatters them, and gathers them again with dread. It is absurd to say that there is any value here - to say that Alph has a life worth living.

Philosophers who write about the value of a life – a life worth living – tend to speak dismissively of a life in which an individual does nothing but count blades of grass. Even if this is what the agent desires, and the agent desires nothing else, this life is considered devoid of value.

The reader – given the reader's own interests and desires – certainly would not want to live the life that Alph lives. However, this is because a life that fulfills Alph's desire to gather stones is unlikely to fulfill many of the reader's desires. Thus, the life that Alph has reason to live – where Alph has no reason to choose any other kind of life – is not the life that the reader has a reason to live. Each of the reader's own desires has their own object – their own "proposition P" that is the object of the desire. It is unlikely that the propositions that are true of the reader's desire are made or kept true in the life that

Alph lives. However, Alph can say the same thing not only about the reader's life, but the best possible life that the reader can imagine for herself.

If the reader wants to identify "the good life", then the reader needs to take stock of her own desires and then create a life in which the propositions that are the objects of the most and strongest desires are made or kept true. If we were going to travel down that road, we would have to answer questions about such things as the timing of certain desires, and meta-desires (e.g., having an aversion at age 20 to having certain desires at age 50). Fortunately, we are not going that direction.

When I say that Alph would find value in a life of gathering stones, I am not implying that the reader either would or should find value in the same kind of life. I am saying that the life stands in a particular relationship to Alph's desires, while a different life would have that relationship relative to the reader's desires.

If I were to ask the reader to imagine Charlie, living in New York, I would say that he is near the Atlantic Ocean. In saying this, I certainly do not mean to imply that the reader is near the Atlantic Ocean, or that the reader should be near the Atlantic Ocean. Nor do I need to reader to approve of the fact that Charlie is near the Atlantic Ocean. I am only describing a relationship that exists in fact between Charley and the Atlantic Ocean. The reader is only being asked to acknowledge that this relationship exists.

Furthermore, when I say that Alph has one desire - to gather stones, then I am saying that states of affairs in which Alph is gathering stones are states of affairs that Alph has reason to bring about. I am saying that scattering stones, to Alph, would have instrumental value once all the stones are gathered. I am saying that having a large pile of stones is an unintended consequence of acting so must fulfill his desire (like getting pregnant is often an unintended consequence of having sex). I am saying that an outside observer cannot have reason to choose one world over another until that outside observer has desires, in which case the observer would choose the world where the propositions that are the objects of those desires are true.

These are all relationships that exist - no different from the proposed relationship between the location of the Atlantic Ocean to Charlie. The reader does not have to like it. The object of evaluation may stand in a relationship to the reader's desires that is significantly different from its relationship to Alph's lone desire. None of this can be used to deny the relationship between a life of gathering stones and Alph's one desire.

Everything said here would remain just as true if I were to say that Alph had one desire - to count blades of grass. Then counting blades of grass would be the only end-reason for any of Alph's actions. Alph would eat only insofar as eating was necessary to help him to continue to count blades of grass. If asked to choose between being in a world where there was no grass and a different world existing where somebody else (Betty) with a desire to count grass lived on a planet filled with grass, Alph would be indifferent between the two. This is because neither world is a world where "I am counting blades of grass" is true.

The fact that the reader has more than one desire and none of them include counting blades of grass, and has no desire with propositions that would be true in the imagined world with the imagined grass counter, and would not want to be that person is irrelevant. It is as irrelevant as her aversion to living in New York or near the ocean is to the fact that Charlie, who lives in New York, is near the ocean.

So, I ask the reader to set those feelings and those judgments aside. We are looking only at the relationships that exist in this imagined world - a world in which there is one person (Alph) with one desire (a desire to gather stones).

XIII. The Value of Truth

Truth is good. Usually.

In our simple universe, where Alph only has one desire – a desire to gather stones – Alph only has a reason to pursue something as far as it serves his desire to gather stones. On a world with a limited number of stones, scattering stones would have value since it would allow him to return to the task of gathering stones. Eating, drinking, and avoiding injury all have value since, if he did not eat or drink or became injured, he would no longer be able to gather stones.

Similarly, Alph has a reason to determine where there are stones to be gathered. He has reason to prefer true beliefs to false beliefs. If he falsely believes that a particular valley has a few hundred stones to gather, he risks wasting a great deal of time not gathering stones as he hunts for stones to gather that do not exist.

While he is out gathering stones, Alph has a reason to prefer true beliefs regarding where the gathering point is. He has reason to know, truthfully, where he was at and how to get his stone to the gathering pile. False beliefs here would doom him to wandering around aimlessly – failing to gather stones.

Jumping off a cliff as he gathers stones will put him at risk of breaking a leg, which will prevent him from gathering any more stones. Consequently, Alph's desire to gather stones gives him a reason to have true beliefs with respect to jumping off any cliffs so that he can avoid doing something that would prevent him from gathering stones. Furthermore, if there are rocks at the bottom of the cliff, he has reason to acquire knowledge about the construction of a ladder for climbing up and down the cliff, and for building a rope and bucket to haul the rocks up the cliff.

A false belief that he can fly across a chasm may have Alph laying injured at the bottom of the chasm.

Alph has reason to acquire true beliefs regarding where he can find water to drink and food to eat, as well as true beliefs that distinguish good water from bad water, good food from bad food. "Good" and "bad" in this sense is determined precisely by "that water or food that can best help me to fulfill my desire to gather stones." Food and water that makes him sick is not good food or good water in that it is not food and water that serves his desire to gather stones.

Remember, Alph only desires to gather stones. He has no taste preferences. In our model, he eats only for utilitarian reasons - to stay alive. True beliefs will help him to avoid scurvy, rickets, and other diseases that would prevent him from gathering stones. However, this requires a great many true beliefs.

In our imaginary world, believing true premises has no value for its own sake. Nothing has value – in that an agent has no reason to pursue a particular state – except insofar as it serves a desire. This applies to the pursuit of true belief. But, by and large, true beliefs are extremely useful when it comes to serving a desire – just about any desire. Alph in his world, and we in ours, all have reasons to attach a great deal of value to the acquisition of true beliefs.

Truth for its own sake has value only insofar as an agent desires truth. However, truth as a means tends to have value regardless of what else an agent may desire. Even in our own society, it is no exaggeration to say that a great many of the evils that we suffer come from agents acting on false beliefs.

XIV.   Subjectivity and Objectivity

A reader may be curious as to where I stand on the subjectivity or objectivity of value. Some may already have formed the conclusion that I am a subjectivist, since I assert that all value depends on desire. Others may assert that I am an objectivist, since the relationships that exist between Alph's desire and various states of affairs exist as a matter of objective fact.

I don't think I can answer the question of whether morality is objective or subjective. This is because the terms "objective" and "subjective" themselves have no clear definition. A great deal of ink, toner, and electronic energy gets wasted on this debate because of this.

Those who would say that this is a subjective account of value will point to the fact that, for our one agent, Alph, gathering stones has value only because he has a desire to gather stones. Furthermore, Alph's desire to gather stones gives Alph and only Alph a reason to gather stones. If Alph had a desire to scatter stones, then it would be the scattering of stones that Alph values and that which he has a reason to do. There exists no value independent of desire – no reason to act other than that which serves a desire.

All of this is true, but I am disposed to say that value is objective. After all, the statement that Alph has a reason to gather stones is objectively true. It is as objectively true as the statement that Alph has a heart rate of 72 beats per minute (assuming that Alph has a heartbeat of 72 beats per minute) or that Alph stands 1.89 meters tall (assuming that Alph stands 1.89 meters tall). If these statements are true, then they are objectively true. If the statement that gathering stones serves a desire that Alph has is true, then it is objectively true. Value, then, is objective – in this sense.

This confusion comes from the fact that "objective" and "subjective" are ambiguous. Both terms have multiple meanings, and people often get lost in these discussions thinking that they are talking about the same thing when they are not.

The two definitions of "objective" that cause the most problems are:

(1) Independent of all mental states such as belief and desire.

(2) True, regardless of whether the agent wants or believes it to be true.

Reading quickly over these two definitions, it is easy to think that they say the same thing. However, they are different. The nature of this difference makes it the case that value is subjective in the first sense, but objective in the second.

All value is subjective in the first sense. This is because all value is dependent on desire – since all true value claims describe relationships between objects of evaluation and desire. All true value claims assert that the object of evaluation either is or is not such as to fulfill the desires in question. If we eliminate all desire, then we eliminate all value.

However, none of the propositions that we have been using to talk about Alph with his one desire in his world are true because Alph believes them to be true, or because Alph wants them to be true. Alph's

attitudes towards those propositions do not determine their truth. They are not "true for Alph" and false for others. They are, instead, objectively true.

A scientist looking at this universe through his telescope would not be able to describe what he saw without postulating that Alph had a desire to gather stones. It is a statement that other scientists could then, with their own observations, support or falsify. All of the relationships that I described also exist as scientifically solid facts.

Specifically, the statement "Alph has a desire to gather stones" is objectively true (within this imaginary universe). It is as objective as the statement, "Alph has a scar on his left arm" or "Alph has a temperature of 99.1 degrees". While the first says something about the structure of Alph's wrist, and the second says something about Alph's body overall, the proposition "Alph has a desire to gather stones" reports a fact about the way his brain is structured.

According to our initial description, Alph has no desire that gathering stones have value. He has a desire to gather stones, but he has no desire that he desire to gather stones. To Alph, the fact that he desires to gather stones is unimportant, though, since he has that desire, gathering stones is important.

Similarly, I have made no claims about what Alph believes other than that he has those beliefs that allow him to make true the proposition, "I am gathering stones". There is no need for a belief that Alph values the gathering of stones. Alph need not have ever thought about the question of whether he values gathering stones. He simply goes about the chore.

Because people confuse these two concepts, they slip from the fact that values are subjective (depending on desires since they describe existing relationships between objects of evaluation and desires) to the fiction that value claims are subjective (true because the agent either believes or wants them to be true). They then use this second claim – the fiction – to argue that values are beyond debate or scientific inquiry and that an investigation into the facts of the world cannot tell us anything about what has value and what does not.

However, the former, though true, does not imply the latter. Facts about an agent's brain structure no more rules out the objectivity of value than facts about an agent's height or age rules out the objectivity of him being taller than or older than Sally.

What are we going to say, then, if we are asked the question of whether value is objective?

It's a trap.

If we answer that value is objective, we are likely to be taken as meaning that values exist independent of all beliefs and desires – as properties or entities that can be found in the world independent of all minds. People will likely object to this view – as they should. These claims are all false.

If we say that value is subjective, we are likely to be taken as meaning that value claims are true or false depending on whether agents believes them or wants them to be true or false – depending on the agent's own sentiment. These claims are also false.

In being asked the question, we are forced to choose between two answers, both of which, given the common understanding of the terms, are false.

It may be best to answer this question by saying that people are generally confused about the concepts of "objective" and "subjective" such that nobody can answer this question without spending a considerable amount of time explaining their answer. If one is pressed for an answer, the fact of the matter is that value claims can be investigated scientifically in the same way that claims about body temperature and scars can be investigated scientifically.

XV. The "Location" Analogy

I have said that the question of whether value is objective or subjective is too ambiguous to allow for a simple answer.

One way to understand these claims about subjectivism and objectivism is by analogy to statements about location.

Nothing has an absolute, intrinsic location. Absolute, intrinsic location does not exist.

In fact, you cannot give the location of anything other than by giving its position relative to something else. Denver is in Colorado. The earth orbits the sun. The keys are in my coat pocket. The land mine is seven meters ahead of you and one meter to your right.

All location claims are relational - describing relationships between one thing and another.

Furthermore, when you are asked where something is, there is no law of nature that dictates what you choose to use as a reference point. Generally, when it comes to choosing a reference point (e.g., Colorado, the sun, my coat pocket, you), we choose what is conventional and convenient under the circumstances and given the context. An object that is seven meters away from you can be 400 kilometers away from me. Both statements are true, and neither is truer than the other. If I give its relationship to you rather than to me, I do so merely because expressing one relationship serves my interests more than the other.

There is an infinite number of ways in which I can describe the location of something. For example, we can say of the city of Denver that it is in Colorado, in the United States, on planet earth, about 30 miles southeast of Boulder, and where Brian lives for some specific individual named Brian. All of these are legitimate location claims. None of them are more legitimate than any other.

We do have standards - such as latitude and longitude - that we use in making location claims. We choose to use the equator as one axis, and a line going through the north and south poles and Greenwich, England as the other axis. Why Greenwich, England? There's a historical reason for it - but nothing that justifies the claim that it is the one true and correct base reference point for longitude. We just . . . decided.

If philosophers were to debate the one right and true reference point for all location claims, they would be wasting a lot of time.

However, in spite of these facts, nobody has any trouble including location claims in scientific publications or in investigating location claims scientifically. Location claims are taken to be as objectively true (or false) as any claim made in science.

In particular, the fact that a reference point is selected by custom and convenience is never used to question whether the relationship claim is true or false - or to challenge the claim that, if it is true, then it is true as a matter of fact and not as a matter of opinion.

All of these are true of relationships between states of affairs and desires as well.

We cannot tell the value of a state of affairs without describing its relationship to one or more desires. When we describe such a relationship, we describe what a specific person or set of people has reason to realize or prevent. It is only in virtue of having the requisite desires that one has the reasons to realize or prevent such a state of affairs. However, they have those reasons as a matter of fact. A person who ignores these facts cannot understand or explain what happens in the real world. There should be no problem investigating these facts scientifically.

There is no law of nature that dictates which desires are relevant in making such a claim. We can describe how a state of affairs stands in relationship to the desires of Uncle Joe, to the people of Atlanta, my pet rabbit Fluffy, or to those living near the coast. All of these relationships exist. Statements describing these relationships are all equally true. The truths that we talk about are the truths that we discover that we have reason to talk about.

When we choose a reference point, we choose it as a matter of convenience and convention - not because some law of nature dictates that particular use. It is because, "This is the relationship that we have decided to talk about. Other relationships exist, but we are talking about this relationship - in the same way that other rabbits exist, but we are talking about Fluffy."

There are some relationships that people have many and strong reasons to make a part of our conversation.

Take 'health' for example. This is a term used to describe changes in the functioning of the body or mind relative to the desires of the person whose body or mind is being evaluated. People, having reasons to avoid illness and injury, have many and strong reasons to make these relationships a topic of conversation. For convenience sake, we use the term 'health', 'illness', and 'injury' to describe these relationships.

Where these relationships exist, they exist as a matter of fact.

I will be arguing that the best use of moral terms uses them to describe relationships between malleable desires (or what some people call "intrinsic desires") and the other desires that the desires being evaluated can fulfill or thwart. Morality, understood this way, will require more than one person with malleable desires and the ability to mold those desires using tools such as reward, praise, condemnation, and punishment.

We are not there yet, so morality, in our simple universe, does not yet exist.

Our universe so far is one in which Alph, as a matter of fact, has a motivating reason to realize a state of affairs where "Alph is gathering stones" is true. It is a world where Alph sometimes has a reason to scatter stones (though not an intrinsic reason), just as he has a reason to know where the stones are, to avoid falling off cliffs, and to eat or drink that which allows him to realize a state in which he is gathering stones.

Gathering stones has value for Alph.

And that - within this hypothetical universe - is an objective fact.

XVI.   Realism about Values and "Literally True"

When a person claims to be a realist about value, they are often taken to mean that they take value properties to exist in the world independent of all mental states – that they are somehow a part of the very nature of that which is called "good" or "bad".

In this sense, the theory presented here is not a "realist" theory about value. Many would call it an "anti-realist" theory.

Yet, claiming that one is an anti-realist about value is typically understood as claiming that value, in some way, depends on the sentiments of the person making the judgment. Either the value-claim is made directly from the speaker's sentiments, or it is an expression of the sentiment that one would expect to have under some ideal conditions of perfect knowledge and sound reason.

In this sense, the theory presented here is not an "anti-realist" theory either.

I assert that this theory is a realist theory with respect to value, but it is not "realist" in the sense described above. It is a realist theory because value is real – it is a part of the world in which we live. However, it is not independent of mental states such as desire or emotion. Instead, it exists as a real-world relationship between objects of evaluation and desires. Value does not exist without desire. However, since desires are real, values are real.

We have another way of understanding value realism available to us. On this account, to be a realist about value, one must believe that propositions attributing value can be literally true.

This invites us to ask the question, "What does it take for a proposition attributing value to be literally true?"

On the account presented here, to say that S is good is to say that S is such as to fulfill the desires in question. This means that for the desires that make up the desires in question, expressed as, "Agent desires that P", that P is realized in S or S contributes to realizing P. Of course, such statements can be literally true.

In our imaginary universe, Alph values gathering stones. There is a way of understanding the proposition "Alph gathering stones is good" such that what it means is that a state of affairs in which Alph is gathering stones is a state of affairs that serves Alph's desire that he be gathering stones. Alph's desire to gather stones is the desire in question. When it has this meaning, the statement can be literally true. It is literally true that Alph has a desire such that the state of affairs in which he is gathering stones is a state of affairs that serves that desire.

When Alph has run out of stones to gather, scattering stone will serve his desire to be gathering stones. This accounts for its instrumental value. Knowledge of where to find stones, what to eat, what to drink, and how to avoid illness and injury serve his desire to gather stones as well.

Statements such as "Alph gathering stones is good" is an ambiguous statement. It need not relate the act of gathering stones to Alph's desire to gather stones. It may relate the act to some other desire, and it may or may not be such as to serve that desire.

Consider, by analogy, the statement, "The keys are on the table." Is this statement literally true? That depends. Which keys? Which table? We can know the answer to this question only from the context within which the statement "The keys are on the table" is uttered. However, once we know this context and can identify the keys and the table, we can determine if the statement is literally true.

The same applies to a statement like, "Alph gathering stones is good". A value claim can be true if it relates an object of evaluation to a set of desires. In this case, we know the state of affairs – it is the state of affairs in which Alph is gathering stones. However, we must refer to the context in which the statement is made to determine the relevant desires. When we know the relevant desires, we will know whether the state of affairs in which Alph is gathering stones is one that tends to serve those desires, either directly or indirectly. Knowing this allows us to know if the statement is literally true.

Everything in this example is real. The state of affairs is real. The desires are real. The relationship that exists between them is real. Therefore, the value is real.

Somebody may accuse me here of engaging in a bit of sophistry.

If I define a "leg" as an extremity on a dog not attached to the head, then the proposition, "A dog has five legs" becomes literally true. However, in common languages we would say that a dog has four legs and a tail. There is some point where the use of a term is so different from its common understanding that to claim that the theory accounts for the common understanding of the term is false.

Even though this objection has merit, the situation is more complex than even this allows. We can sometimes propose a new meaning of a term that corrects a previous error without being accused of anti-realism.

Consider, for example, that the term 'atom' once meant 'without parts.' The literal meaning of the term 'atom' was that it could not be divided into smaller elements. After scientists spent some time writing and talking about atoms of oxygen, gold, carbon, and the like, they came to the idea that what they were calling atoms had parts. They were made up of protons and (with rare exceptions) neutrons and elections. When scientists began to suggest that what they had been calling atoms had parts they were saying that original claims about atoms were not literally true. However, they were not suddenly being branded anti-realists about atoms.

The same can be said of the term "sunrise". One did not become an anti-realist about "sunrise" when one adopted the view that the earth spun on its access and "sunrise" was not, in fact, a case of the sun rising but the earth spinning to bring the sun into view for a given observer. One could continue to be a realist about sunrise, even though the term changed its meaning.

This leaves us with two ways in which to defend this theory as a form of value realism.

One way is to say that value terms include as a part of their regular meaning the claim that the objects of evaluation are such as to fulfill the desires in question. The other way is to say that, even though this is not a part of the current meaning, we can transition to this meaning as easily as chemists transitioned to a new definition of "atom" or people generally transitioned to a new definition of "sunrise". In the

same way that chemists continued to be realists about atoms, and people generally continued to be realists about sunrise, we can adopt this new definition of value and continue to be realists about value.

I wish to propose first that the "literal meaning" of value-laden terms does not, in fact, refer to mind-independent intrinsic values. It refers to reasons for action. Those reasons may spring from mind-independent intrinsic values (if they exist). However, they also spring from desires in that desires provide reasons for action. Consequently, relating an object of evaluation to one or more desires is a perfectly legitimate way of claiming that the agent has a reason to realize a particular state of affairs, and is a legitimate common use of value-laden terms.

As it turns out, desires are the only things that provide reasons for action. Consequently, all true value claims relate objects of evaluation to desires, while all value claims that refer to other types of reasons (e.g., intrinsic presciptivity or divine command) are false. However, all we need for value realism is a set of propositions using value terms that are true, and we have that much.

Yet, even if one wants to argue that this is a deviation from our standard usage, I would argue that "is such as to fulfill the desires in question" is so close to being a legitimate use of value terms that we can shift to it without losing realism in the same way chemists shifted to a new definition of 'atom' and people generally shifted to a new definition of 'sunrise'.

A third option is to say that values are not real. However, relationships between objects of evaluation and desires are real, and desires provide agents with motivating reasons to realize states of affairs in which the propositions that are the objects of their desires are true. We can make objectively true claims about what people have reasons to do or to avoid, and the question of whether we use value terms or some other terms will not alter these truths in any way.

Regardless of the stand one wants to take on "value realism", the entities under discussion in this paper are real.

XVII.   Summary of One Person with One Desire

I am about to add a second person to my imaginary community. Before I do, I would like to review the key points with respect to having one person (Alph) with one desire (to gather stones).

Desires are propositional attitudes. That is to say, they take as their object a proposition - a sentence capable of being true or false and adopt the attitude that the proposition is to be made or kept true.

A desire has motivational force directing the agent to make or keep true the proposition that is the object of the desire. So, for example, Alph's desire to gather stones motivates him to act to make or keep true the proposition, "I am gathering stones."

A desire "that I am gathering stones" is different from a desire that the stones be gathered. With respect to Alph's actual desire, having a large pile of stones is merely an unintended consequence of the desire to gather stones. It bears the same relationship to the desire to gather stones as pregnancy often has to the desire to have sex. Such a state of affairs might be brought about by acting so as to fulfill a desire, but it is not always wanted itself.

I am using Bernard Williams' account of what it means to have a reason.

A has reason to φ if and only if A has some motive which will be served or furthered by his φ

On this account, Alph has a reason to gather stones. He has a desire to gather stones that would be served by gathering stones.

If Alph runs out of stones to gather, then he has a reason to scatter stones, since that is the only way he can once again make the proposition, "I am gathering stones," true.

Alph also has a reason to avoid a crippling injury or death where these would prevent him from making or keeping true the proposition, "I am gathering stones." He has a reason to eat. That reason is not hunger or thirst, but the need to do so in order to continue to gather

These other goods - scattering stones if the supply of stones to gather runs out and avoiding crippling injury or death - are instrumental goods. They provide or preserve the means necessary to make or keep true the proposition, "I am gathering stones." They are not valued for their own sake, but for their usefulness in realize a state within which Alph is gathering stones, which Alph values for its own sake.

There is no intrinsic value. A state in which, "I am gathering stones" has value for Alph in virtue of his desire to gather stones. However, nothing in this state generates a reason for anybody else (if they should exist) to realize or preserve such a state. A reason to realize a state in which a desire is being fulfilled requires a desire that desires be fulfilled, or some other desire made true in a state where a desire is being fulfilled.

Not even Alph is motivated by desire fulfillment. Alph is motivated by a desire to gather stones, which motivates him to realize a state in which, "I am gathering stones" is true. It is, "I am gathering stones" that has value for Alph, not "My desire that I am gathering stones is fulfilled."

Alph, in this case, has no reason to enter a Nozickian experience machine that will stimulate his brain and feed him the illusion of gathering stones. Such a machine cannot make or keep true the proposition, "I am gathering stones." Similarly, parents who cares that their children are safe and happy cannot settle for an experience machine feeding them the illusion that their children are safe and happy. Their desire motivates him to seek the actual safety and happiness of their children.

If the reader thinks that Alph is wasting his life pursuing a meaningless end, then this would be because the reader is appealing to her own desires. The reader is saying, "I would not want to live that life." This may be true, but it is a separate issue. The fact that Alph is content gathering stones does not imply that the reader - with the reader's own set of desires - would be or even should be content with a life of gathering stones.

Nor does the reader's discontentment imply that Alph has a reason to shun a life of gathering stones. Alph has no such reason unless Alph has a desire that would be served by doing something other than gathering stones.

Value claims made in this world, such as Alph saying, "Gathering stones is good," can be literally and objectively true. Value, understood as a relationship between a state of affairs in which Alph is gathering stones and Alph's desire to be gathering stones, exists.

In this world consisting only of Alph with his desire to gather stones, there is no morality. Morality requires more than one agent.

So, now, I am going to drop a second agent into the world and describe what is true of this new state of affairs.

XVIII. Introducing Bett

So far, we have had a universe with one person (Alph) and one desire (to gather stones).

Recall that his desire to gather stones is a motivating reason for him to make or keep true the proposition "I am gathering stones."

Let us assume that there is a limited number of stones such that Alph must spend half of his time scattering stones so that he can gather stones again.

Now, I am going to plop a second person down on this world - Bett.

Bett has no desires. Bett is simply laying there on a slab of rock and, because Bett has no desires, he has no reason to do anything.

However, we are going to give Alph two syringes; one with a red serum and the other with a blue serum. We will inform Alph that the red serum will cause Bett to have a desire to gather stones – making Bett like Alph. The blue serum will cause Bett to have a desire to scatter stones.

I would like to repeat that Alph has only one desire – a desire to gather stones. Furthermore, we are using Williams' concept of reasons:

> There is a reason for A to φ iff there is some desire the satisfaction of which will be served by his φ-ing.

This means that Alph has a reason to give Bett the red serum if and only if this will serve his desire to gather stones. The same is true of injecting Bett with the blue serum.

Alph's desire to gather stones provides him with no reason to inject Alph with the red serum. That is to say, Alph has no reason to bring into the world another person with a desire to gather stones.

On the other hand, Alph's desire to gather stones does provide him with a reason to give Bett the blue serum, thus giving Bett a desire to scatter stones. By giving Bett a desire to scatter stones, Bett will be motivated to scatter stones – and will scatter stones so long as he can do so. Alph will no longer have to spend time scattering stones, just so that he can gather stones again.

I want to point out that Bett will not be scattering stones for the sake of giving Alph the opportunity to gather stones. Nor is it the case that Bett will be scattering stones because of empathy. Bett has no interest in Alph's welfare just as Alph, in giving Bett the blue serum, did so with no interest in Bett's welfare or the quality of Bett's life. Neither agent has an interest in cooperation for its own sake, no love of fairness, and no sense of duty or obligation. Alph's interest in Bett is limited to Bett's usefulness – Bett's ability to scatter stones that Alph can gather.

Nor is it the case Alph is giving Bett a moral sense or a belief that scattering stones has an intrinsic 'ought to be doneness' – that it is a duty. Alph only gives Bett a desire to scatter stones – and that is all that is needed.

Alph's desire to gather stones alone is enough to give him a reason to establish a cooperative system with Bett whereby Alph gathers the stones, and Bett scatters them.

Empathy, sympathy, communal interests, altruism, moral sentiments, or a moral sense are all entirely unnecessary when it comes to establishing a cooperative system. If they existed, they may be put to use, but they are not necessary. Morality can exist even where they are absent.

Those who equate the discovery of altruism or empathy or a sense of fairness or justice in humans with the discovery of the foundation of human morality are mistaken.

In other words, our community at this point contains

  (1) Two or more independent agents
  (2) Who engage in intentional actions (that is, act so as to fulfill their desires given their beliefs)
  (3) Whose desires can be molded by the actions of others – in this case, by Alph injecting Bett with a blue serum.

But we still have not created a moral system.

XIX.   Molding Desires

The one difference we can still see between the situation that I gave described above and that which was necessary to create a moral system is found in the method used to give Bett a desire to scatter stones. In the story so far, Alph injected Bett with a serum. In the story that follows, Alph will use rewards and punishments – with praise acting as a type of reward, and condemnation acting as a kind of punishment.

For the purposes of this discussion, we are simply going to say that reward and punishment works the same way as the injections did in the previous system. We have reason to take a great deal of interest in this system as it exists in the real world. However, those details are not important here. This will require a couple of additional changes.

The primary cause of difficulty at this point is that reward and punishment require that the agent have a pre-existing desire. To reward an agent is to provide him with something that fulfills a desire, whereas punishment thwarts a desire. Praise counts as a reward because of a desire for praise, and condemnation served as punishment because of an aversion to condemnation. Consequently, to introduce rewards and punishments we must give Bett some other desires. For example, we may give Bett a desire for pleasure and an aversion to pain, and a desire for praise and an aversion to condemnation.

For the sake of the current discussion, these facts merely add confusion. The only fact that I want to use at this point is that Alph can use rewards (such as praise) and punishments (such as condemnation) to cause Bett to value scattering stones for its own sake. Bett comes to value scattering stones in the same way that Alph values gathering stones – as an end or goal, and not merely as a means.

On this description, Alph has just as much reason to cause Bett to desire to scatter stones as he had when his method involved the use of an injected serum. From Alph's point of view, the only thing we have changed is the instrument to use in bringing about a state in which Bett desires to scatter stones.

We are also going to make it the case that Alph is like Bett in this respect. Bett can mold Alph's desires using the same tools of reward (praise) and punishment (condemnation).

We now have what we need to create a moral system.

XX. Universal Aversions

For ease of explanation, I am going to add one more complexity. I am going to give both Alph and Bett an aversion to pain – that is, a desire "that I not be in pain."

In the same way that Alph's desire to gather stoned gives him a reason to gather stones, and Bett's desire to scatter stones gives him a reason to scatter stones, they each have a reason to prevent the realization of states of affairs where the proposition, "I am in pain," is made or kept true.

This means that Alph has a reason to cause Bett to have aversions to things that would tend to result in Alph being in pain. Similarly, Bett has reason to cause Alph to have aversions to things that will tend to cause Bett to be in a state of pain. Chief among these would be an aversion to causing pain to others.

By our original assumptions, they would accomplish this by praising those go about the business of gathering or scattering stones in ways that did not risk causing pain to others, and condemning and punishing those who conducted their activities in ways that put others at risk.

This aversion to doing that which causes pain to others is an aversion that the people in this community (Alph and Bett) generally have reasons to promote universally – in all other members of the community – through the practice of praising those who avoid options that put others at risk of pain, and condemning and punishing those who act in ways that put others at risk of pain.

Once the people in this population acquire this aversion to causing pain, they will tend to prefer actions that (1) fulfill their individual desire to either gather or scatter stones, (2) fulfill their personal desire to avoid being in pain, and (3) fulfill their learned desire that they not inflict pain on others.

This, I would argue, is something that is so much like a moral system that, if we were to stumble on a jungle community engaged in exactly these practices, we would not hesitate to say that the beings in thus community view acts that tend to put others at risk of pain as immoral.

We would see a community in which people react to acts that tend to cause pain to others with condemnation and, perhaps, even punishment. They declare that what they condemn a given individual for doing is that which they would condemn any individual in that given situation for doing. They would treat the claim that the action deserves condemnation as an objectively true statement and, in defense of their claim, they would muster as their reasons for condemning and punishing the action, the pain caused to others. That pain is what gives others a reason to condemn or to punish the action. The prescription would be independent of what the agent they condemn wants to do. people ought not all individuals for doing.

Furthermore, they will assert that refraining from these types of actions that tend to cause pain to others is something that people ought to value for its own sake. Agents should not only be motivated by

a desire to obtain a reward such as praise, or to avoid punishment such as condemnation. Even if circumstances should put them in a situation in which they have a chance to cause pain to others without getting caught – without facing condemnation and punishment – they should be reluctant to do so. An aversion to types of actions that tend to cause pain to others would make them reluctant to do so, even in situations where they would not get caught.

Even if we were to expand our community to include 100 Alphs and 100 Betts, it would still be true that the people generally, in this population, have many and strong reasons to promote a universal aversion to types of actions that tend to cause pain to others. Their separate, individual aversions to their own pain provides them with these reasons.

In the sections that follow, I will show that we will find even more in this community that we judge in our community to be a part of morality. We will find them teaching moral lessons by holding the virtuous up as examples and making examples of villains. We will also catch them using stories and parables to teach moral lessons. We will see them building an emotive component into their moral claims. We would find them distinguishing between obligation, prohibition, and non-obligatory permission, as well as having the concept of a hero whose actions go above and beyond the call of duty. They will have the concepts of "excuse" and "*mens rea*"

Once again, I want to draw attention to the fact that this community has no empathy or sympathy, no dispositions towards cooperation, no evolved moral sense, and no moral intuitions of any kind. We could introduce these things, and as we do their moral institutions may change to take advantage of these traits, but they are not necessary.

Furthermore, what these people praise and condemn are actions (or types of actions) – they praise actions that avoid causing pain to others, and condemn actions that cause pain to others. However, the purpose of the praise and condemnation is to alter the desires and aversions of agents. The ultimate aim of their moral institutions is not the modification of actions, but the modification of desires and aversions.

XXI.   Learning Desires and Aversions

If we were limited in teaching desires and aversions by rewarding or punishing the specific individual who performed the action, our capacity to set up these desires and aversions would be severe limited. In effect, in order to teach a person to have an aversion to committing murder, that person would have to commit several murders and be punished for them. The same would be necessary for teaching him not to lie, steal, rape, and break promises.

Fortunately, we are not built that way. We have the capacity to learn from the experiences of others. When we publicly condemn one liar, others who discover the condemnation acquire something of the aversion to lying, and when we praise one charitable act, others who witness the praise also acquire something of the desire to be charitable. This capacity to make an example of people, or to hold them up as an example, significantly improves our ability to promote these universal desires and aversions.

Even when it comes to experiencing the rewards and punishments of others, we do not need to experience it directly. When Uncle Joe comes over for dinner and condemns the neighbor who woke him up early on Sunday cutting his lawn, this condemnation will tend to cause the people around the dinner table to acquire something of an aversion to disturbing the peace of others at such an

unreasonable hour. When the employee learns of the co-worker who was fired for lying about being sick when he went fishing, this helps to promote an aversion to lying.

In fact, the act of reward or punishment does not even have to be real. Uncle Joe's condemnation might simply be a story – or a lie he told because of his hatred for his neighbor (but which nobody knows to be a lie). It would still have the effect of promoting an aversion to disturbing others.

Along these same lines, one can create a parable or story about somebody who helped a stranger – one that contains an attitude of praise for those who would help others – to generate the same effect. One can create a story about a school bully that conveys an attitude of condemnation towards that person to help create a general aversion to bullying.

Because we are built to learn desires and aversions in these different ways, agents can learn the moral rules without actually committing and being punished for each specific moral wrong. By the time a child becomes an adult, she should have been well exposed to a lot of indirect praise and condemnation. Much of it comes from listening to parents and other adults praise and condemn each other. Some comes from stories with a moral lesson – from books, television, movies, and certain songs. Some comes from preachers, teachers, and group leaders. They will also experience some direct praise and condemnation – hopefully, in the latter case, limited to minor wrongs, both from authority figures and from peers. The net effect should be to adopt something near the standards of the culture.

However, I hasten to add that the desires and aversions taught within a culture can be significantly different from the desires and aversions the people within that culture have reason to teach others. What tends to thwart other desires may be different from what actually causes harm to others. What makes a society stronger and contributes to the well-being of its people may not be what the population believes makes it stronger and contributes to their well-being.

XXII.   Emotivism

There is a view of morality that says that utterances of right and wrong are not truth-bearing claims at all, that they are mere expressions of emotion. To say that theft is wrong is to say, in effect, "Theft, Boo!". To say that charity is good is to say, "Charity, Horrah!" Expressions of this type are neither true nor false.

I see no reason to claim that emotive claims are not truth-bearing. One can laugh at a joke as a way of expressing an attitude towards it. Yet, it is also true – as a matter of fact – that "that makes me laugh." One can turn up one's nose and make a look of disgust at the very thought of eating sushi, or one can say – honestly and truthfully – "I really hate the taste of sushi." One can be disgusted by a smell and say honestly, "I find that smell disgusting."

One really cannot get away from the idea that moral statements are truth-bearing propositions simply by associating them with an emotive response.

Another relevant fact about emotive utterances is that they can be combined with a truth-bearing statement. The emotivist invites us to see emotive utterances and truth-bearing propositions as mutually exclusive – something can be one or the other, but not both. Yet, they can be both.

If somebody shouts, "You stepped on my foot!" in a tone of anger and indignation, the statement itself is a truth-bearing statement. The proposition, "You stepped on my foot," is capable of being true or

false. At the same time, the tone in which it is spoken – in a tone of anger and condemnation – expresses the speaker's attitude towards this fact – or, more precisely, towards the person who stepped on her foot.

In the same way, a statement of the form, "People ought to keep their promises" can be both, at the same time, a truth-bearing proposition and an emotive expression. At the same time, it can report the fact, "people generally have many and strong reasons to promote an aversion to breaking promises by condemning those who do so," and express the condemnation that the truth-bearing proposition reports people have a reason to give. That condemnation also does the work of helping to promote the universal aversion to breaking promises. This type of statement, then, kills two birds with one stone.

In other words, the emotivists are correct in noting that there is an emotive component built within the use of most moral statements – either of praise or of condemnation (or neither in the case of non-obligatory permissions). Furthermore, this element exists on purpose – it serves a function. This is the tool that helps to create the universal desire or aversion that the truth-bearing component says that people generally have many and strong reasons to create.

In short, the explanation as to how the practices of this imaginary tribe mimics our morality includes the fact that their moral claims are both, at the same time, truth-bearing claims about what people generally gave many and strong reasons to promote using reward and punishment and, at the same time, embedding the emotive content – the praise of condemnation they claim people have reasons to deliver.

XXIII.   Right and Wrong Action

People generally divide actions into three moral categories; obligation, prohibition, and non-obligatory permission. We would find our people in our imaginary community using this same distinctions.

Specifically, I want to suggest that a right action is an action that a person with good desires (and lacking bad desires) would perform, a wrong action is an action that a person with good desires (and lacking bad desires) would not perform, and a non-obligatory permission is an action that a person with good desires (and lacking bad desires) might or might not perform depending on other interests.

The interests of people living in this community include the fact that people generally have many and strong reasons to promote an aversion to causing pain to others. It should not be surprising to also see such things as aversions to lying, breaking promises, destroying the property of others or taking their property without consent, assault, rape, and murder among those that people generally have many and strong reasons to promote – and to use condemnation and punishment to promote them.

The population has reasons to prefer to be surrounded by people who hate to do all these things – and to hate them for their own sake and not just because they do not find these things useful. This means we have reason to condemn those who do these types of things and, in some cases, to condemn and punish them quite harshly.

These people also have reason to be surrounded by people who like to help those who are in need, who will give a hand to a neighbor. This means that they generally have reasons to praise those who are kind and helpful and to condemn those who are cruel and selfish. They have reason to be surrounded by

people who like to make sure that they repay their debts – meaning that they have reason to praise those who repay debts and to condemn those who do not.

To say that people are obligated to repay their debts and keep their promises is simply to say that people generally have many and strong reasons to use praise and condemnation – rewards and punishment – to promote a universal desire to repay debts and keep promises.

To say that vandalism, theft, assault, rape, and murder are morally prohibited is to say that people generally have many and strong reasons to promote a universal aversion to vandalism, theft, assault, rape, and murder by praising those who refrain from these types of acts and condemning – sometimes by condemning very harshly by inflicting punishment – on those who perform these types of actions.

Note that X is wrong is not the relativist claim, "I don't like it when you do X," or the cultural relativist claim, "We don't like it when you do X!" or even the utilitarian claim, "Your doing X has bad consequences." The moral prohibition takes the form of the objectively true or false proposition, "People generally (not just me) have many and strong reasons to promote a universal aversion to doing X."

In addition to moral prohibitions and obligations, we have a large set of non-obligatory permissions where people have a liberty to choose among a set of options. These arise because there is a set of cases where people generally have no reason to demand that everybody have the same desire, and even some where people have reasons to prefer diversity. One example of this is evident in the example we have been using – a diversity of interests with respect to career choices.

In the same way that Alph, with a desire to gather stones, had a reason to prefer that Bett had a desires scatter stones, we have reasons to prefer that some people like teaching, others like engineering, still others like medicine, and still others enjoy construction. Because of this, the choice of a career falls into the category of "non-obligatory permission". People have a moral permission to become teachers or engineers, but no obligation.

The choice of what food to eat, what to wear, what to watch on television, and who to date – with some constraints – all fall into the realm of non-obligatory permission. People are allowed, even encouraged, to adopt a diverse set of likes and dislikes in these areas.

XXIV. Heroes and Supererogatory Actions

In addition to obligatory, prohibited, and non-obligatory permissible actions, we see another category of action – heroic or supererogatory actions. These are actions that go "above and beyond the call of duty".

When a person performs an ordinary duty, we tend to offer some measure of praise. When a person repays a debt, there is reason to give a word of thanks as a token of gratitude to help to encourage people to repay debts. We have reasons to praise the honest person for their honesty, and a kind person for ordinary acts of kindness.

However, in these normal every-day cases, we also have reason to condemn the failure to commit such an act. We have reason to condemn the person who does not pay a debt, or does not tell the truth, or who refuses to perform a simple kindness.

However, if the action is particularly kind – if the kindness goes far beyond what we think we have reason to expect from most people – we tend to give these people a great deal of praise, while refusing to condemn those who fail to perform such a heroic duty.

This applies, for example, to a person who gives a kidney to a stranger, or who jumps onto the tracks to lift a person out of the way of an oncoming train. It applies to the soldier or a first responder who performs a heroic duty beyond the heroic duties that such people perform every day.

We commonly use praise to promote a concern for the interest of others. Even the casual "thank you" for a kind deed is a way of promoting an interest in performing kind deeds within a community.

These actions cannot be expected to have the same effect on everybody. On some, they have a minimal effect and the target remains selfish and inconsiderate. Others, on the other hand, acquire a level of concern that is far greater than we can expect from most people.

We still have reason to promote this level of concern. We have reason to promote it in everybody – we just do not expect to e able to do so. When we do find it, we have reason to praise those who commit such acts. In fact, we have reason to give them a great deal of praise since we have reason to push all of society more in that direction.

We give such a person a medal – and we do so publicly at a ceremony and broadcast it as far as possible. We feature them on news casts in the hope of news of their heroic acts will inspire others to perform similar acts.

However, at the same time, we have little or no reason to condemn its absence. To condemn its absence is to condemn the bulk of the population – perhaps including ourselves. There is little to be accomplished with this level of condemnation.

Ultimately, supererogatory actions follow the same set of rules as obligatory, prohibited, and permissible non-obligatory actions. It is another category of actions defined by the degree to which it makes sense for people to use the tools of praise and condemnation to mold the sentiments of others.

XXV.   Mens Rea

The institutions of morality and law include a concepts of *mens rea* or "guilty mind". In order to demonstrate that somebody is guilty of a crime – or a moral wrongdoing – one needs to include proof that the accused had a particular mind-set.

For example, let us assume that Agent1 grabs a suitcase that is not his own off of the luggage carrousel at a busy airport and starts to walk away. Is he guilty of theft?

In order to answer this question, we need to know what was going on in the mind of the agent. Perhaps Agent1 was trying to take somebody else's luggage. But, on the other hand, maybe he simply thought the suitcase was his. Perhaps his suitcase looked exactly the same and he was simply careless about making sure that this one belong to him. Though carelessness is still culpable (for reasons I will get to later), it is not theft.

The difference between theft and negligence is all in the mind. We cannot tell by the physical act alone that the person picking up and walking off with the suitcase is a thief. To accuse him of theft, we need to also prove that he knew that he was not the owner of the suitcase and intended to claim the contents of

the suitcase (at least those that he found valuable) as his own. These are mental events – mental events that we refer to when we talk about *mens rea*

After he finds out these things, if he does not make an attempt to return the suitcase, we may accuse him of theft, but not before, even though, at that time, his physical actions could not be distinguished from those of a thief.

Another example involves the accusation of fraud. Is a person trying to sell a painting – arguing that it is a genuine Picasso painting when, in fact, it is not – guilty of fraud? This depends on whether he knew this was a fake (or claimed to a certainty he could not justify). Let us assume that he had just inherited the painting from his very wealthy grandmother, had an art expert appraise the painting, and had been told that it was a genuine Picasso painting. As he tried to sell it, a perspective customer had her own appraiser look at it and judged it to be a fake – and could prove it. Our agent in this case is not guilty of fraud. He does not have the *mens rea* necessary for guilt or moral condemnation.

Mens rea is also relevant on matters of moral praise and social rewards. Let us say that an agent rams a car that was about to run over a group of school children. Is he a hero? Let us assume that he had stayed up the night before working on a project and had fallen asleep at the wheel, and just happened to drive into the car that was about to ram the school children. He is not a hero. To be a hero, he would have had to have performed the act deliberately, for the sake of the children, while fully aware of the risks to himself. These mental attitudes are among the *mens rea* requirements for a heroic action.

If we hold that morality is substantially concerned with using rewards and punishment (including praise and condemnation) to mold desires and aversions – to promote a common attitude towards lying, fraud, theft, vandalism, assault, rape, murder, and the like – we can account for *mens rea* as an essential component of guilt.

In the case of the person taking the suitcase, *mens rea* tells us if the agent had the aversion to taking the property of others that people generally have reasons to promote using condemnation and punishment.

In the case of taking the suitcase, knowing that the suitcase belonged to somebody else shows that the agent lacks the level of aversion to taking the property of others without their consent. It shows an attitude that people generally have many and strong reasons to change using their tools of condemnation and punishment.

In the case of fraud, knowing that the painting was not genuine shows a lack of interest in dealing with others honestly. People generally have many and strong reasons to culturally promote an aversion to dishonest dealings with condemnation and punishment.

People generally have reasons to promote an interest in protecting children from harm, and to praise those who do so. However, our agent who fell asleep and accidentally protected some children from being run over does not displays an interest in protecting children from harm. Consequently, he does not show himself to have that interest that people generally have reason to promote with rewards and praise.

We can account for the mental elements of morality by accounting for the presence or absence of mental states that people generally have reasons to promote using rewards such as praise, or aversions

(e.g., to taking property without consent, to deception) people generally have reasons to promote using punishments such as condemnation.

XXVI. Negligence

A theory that argues that the focus of moral judgment is on mental states – in this case, on using reward and punishment to mold desires – needs to deal with the issue of negligence.

Negligence does not seem to require a bad mental state.

Let us assume that a farmer needs to haul hay from his hayfields to his cows. This involves loading up a truck with hay and driving it down a public road to a distant pasture where the cows are. The stack of hay bales he puts on the truck is a bit high, and he makes no effort to secure the load. As a result, while he is driving down the highway, a few of the bails fall off, directly in front of an oncoming car that hits the bales and crashes.

The farmer, in this case, is guilty of negligence and generally condemned – even punished – for failure to secure his load and for creating a hazard for others.

However, the agent did not intend anything bad.

This is different from the case of the agent who knowingly lies, or who refuses to pay a debt, or who throws a rock through somebody else's window. In those cases, we can identify an intent to deceive, to refuse to honor an obligation, or to cause harm. In the case of negligence, the farmer did not intend to cause the wreck. It was, instead, an unintended side effect of something the agent did not do – he did not secure the load on his pickup.

The moral crime here is the moral crime of negligence. The problem to address is now to account for negligence in a system that says that moral evaluation focuses on mental states.

We see the answer to this question in the simple description of negligence. Negligence shows a lack of concern (a "wonton disregard") for the well-being of others. What makes negligence morally wrong and worthy of condemnation is that the agent's aversion to being the cause of harm to others was not strong enough to motivate him to take precautions against a foreseeable risk that his actions created for others. A properly motivated agent, concerned to make sure that others do not suffer for his actions, and recognizing the easy-to-recognize risk of having hay bales fall off his pickup – would have taken two smaller loads or better secured the one larger load.

The purpose of condemnation in the case of negligence is to promote this aversion, not only in the person being condemned, but in society as a whole. It aims to make people generally take greater care when they act so that their acts do not produce undue risk for others.

Of course, people generally have many and strong reasons to promote this type of aversion, since people generally are the "others" who would otherwise be harmed by negligent actions. In our imaginary universe, Alph and Bett both gave reasons to promote a universal aversion to causing pain to others strong enough to cause them to pay attention to signs that a course of action creates such a risk, and to take action to avoid creating such a state.

XXVII. Excuses

A driver drives a car through a crowd of pedestrians – wounding or killing several of them. Our first instinct is to blame the person who performed the act. This is clearly a state of affairs that a person with good desires would have sought to avoid. Furthermore, it is a state that people generally can avoid. In light of the evidence, we have reason to suspect that the agent lacked aversions that people generally gave reasons to create using their tools of condemnation and punishment.

However, the agent brings to light the fact that the brakes and steering failed on the car. He had lost control. Even a person with the best of desires could not have avoided running over the pedestrians. Considering this evidence, we lose the justification for believing that the agent is somebody who we have reason to condemn. The agent has given us an excuse – the excuse of accident. In this case, the excuse works.

A man's daughter has promised to meet him for lunch. As time goes by, it becomes clear that she has not kept her promise. Generally, when a person does not keep a promise, people generally have reasons to condemn her. We have reason to prefer to be surrounded by people who keep their promises and that gives us reason to use our tools of praise and condemnation to create such a community. It gives us reason to condemn those who break promises.

However, we discover that the daughter, in this case, was driving through a tunnel when there was a serious accident. The daughter is a physician, and several people were seriously injured in the accident. She stops the car and renders what aid she can to those who were injured. Certainly, we have reason to want people to keep their promises, but not to the extent that they would ignore an accident such as this when they could help people. We have reason to prefer that people have an even stronger desire to help those they encounter in need of help than to keep a promise for a lunch engagement that can easily be rescheduled. Instead of condemning the daughter, we praise her, to promote a proper balance between the desire to keep promises and the desire to help strangers in desperate need.

Here, I want to note that we still expect the daughter to feel bad about missing the engagement and to offer her father an apology. The desire to keep promises does not vanish simply because it is outweighed. It is still there and, when it is frustrated, it produces regret. The offer of an apology acknowledges the legitimacy of the aversion to breaking a promise.

These two examples illustrate what an "excuse" is. It is a statement that breaks the implication from what appears at first glance to be something that a person with good desires would not have done (e.g., drive through a crowd, break a promise to meet somebody for lunch) and the conclusion that there is something in the agent's motives that is a proper object of condemnation. It may look that way at first, but the excuse tells us that this is not, in fact, the case.

XXVIII. Conclusion

In this paper, I attempted to show how it is possible to create a moral system with just a few common elements. All one needs is a community of individuals who are intentional agents – who seek to realize the most and strongest of their desires given their beliefs, and whose desires can be molded using the tools of reward and punishment, where praise and condemnation serve as a type of reward and punishment.

With this set of ingredients, one can create a moral system.

It is a system where objectively true moral statements refer to things in the real world. In the case of moral value, they refer to relationships between malleable desires and other desires. It recognizes that, where desires provide agents with the only reasons for intentional action, they provide agents with reason to promote some desires universally and inhibit others. This includes desires to help others in need and to repay debts, and aversions to lying, breaking promises, theft, vandalism, assault, rape, and murder.

One can create a moral system in which reward and punishment, as well as praise and condemnation, play a significant roll.

One can create a moral system that recognizes four categories of conduct; morally prohibited, non-obligatory permission, morally obligatory, and heroic or supererogatory action.

One can create a moral system in which excuses, mens rea, the moral crime of negligence, and the role that emotions play in moral utterances all make sense.

Because of its focus on the use of reward and punishment, including praise and condemnation, to mold desires – to promote desires that tend to fulfill other desires and aversions that tend to prevent the thwarting of other desires – I name this account "desirism". It is the view that desires are the focus of morality, and actions have value only insofar as they are signs of a creature with good desires.

It is a system that does not require empathy, sympathy, any type of evolved moral sense, moral intuitions, intrinsic moral values, categorical imperatives, or divine commands. Such things may contribute to a moral system, but they are not necessary. In terms of what this theory requires of the world, it requires little, and none of what it requires is considered even remotely exotic. This gives it a significant advantage over competing systems.