
MORAL RELATIVITY

ANTONY VAN DER MUDE

Copyright © 2016 Antony Van der Mude

VANDERMUDE.COM

All rights reserved. This book or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher except for the use of brief quotations in a book review.

First printing, September 2016

Contents

| | | |
|-----------|---|------------|
| 1 | Well-Being | 7 |
| 2 | Moral Relativity | 19 |
| 3 | The Moral Absolute | 39 |
| 4 | The Golden Rule | 51 |
| 5 | Moderation | 63 |
| 6 | Justice | 79 |
| 7 | Free Will | 93 |
| 8 | Logic | 107 |
| 9 | Reasoning | 133 |
| 10 | Deciding: Population and Probability | 153 |

Introduction

This is a book of metaphors.

The development of knowledge can be considered to have four stages. First, the use of metaphor and analogy, the dream castles in the air; formal theory, the image of the dream on paper; verification, putting the idea to the test; and finally implementation, making the dream come alive. This discussion takes us from metaphor into theory, with some minimal verification.

Metaphors are a basic way that the mind understands and applies knowledge. Where analysis breaks a situation down into parts and studies the interrelationships between these parts, metaphor preserves the overall structure of the situation, comparing it to the essential characteristics of other situations.

This book discusses theory, not experimental results. The goal here is to develop a set of principles that are central to most conceptions of what moral reasoning entails.

In this book a moral entity is not necessarily human. Objects such as governments, societies cultures dogs and cats ecologies Gaia and AI's can be moral. Even a paramecium can have a morality, but on a scale of 1 to 100, if a human is 100 a paramecium society has a complexity of 0.001 Many of the concepts involve seem to require a consciousness on a level that only the human individual is capable of. For example, the concept of empathy requires both a sense of self and also a sense of other as like them. This level of consciousness is not found in a newborn - it develops later in childhood. But dogs and other social animals do certainly act in an empathetic manner. Other living things or organizations may have correlates that function as a consciousness does where it is required for making moral decisions. A corporation may exhibit a behavior of reciprocity that takes the place of empathy, to show how corporations have moral behavior outside of the actions of the individuals that make it up.

Few references, ideas have come from society at large and not the com-

munity of moral philosophers. The goal is to redefine moral relativity, since the twentieth century concept of moral relativism by concluding that all moralities have equal weight leads to commonsense absurdities. Nazism and the Communist Kmer Rouge do not have a claim to moral correctness, in this redefinition of moral relativity.

Finally, a word about language. This book makes no argument or point that applies uniquely to one sex and not the other. . Therefore all language in this book is gender neutral, despite the limitations of the English language. This does not present a problem when referring to man or mankind as a whole. Unfortunately, individual pronouns such as 'he', 'she' 'him' and 'her', have no gender-neutral correlates, except for the dehumanizing 'it' and 'its'. Therefore I am stuck with either saying something that can be misconstrued as applying uniquely to one sex, applying only to nonhumans and objects, or making a grammatical error. I take the third choice as the least objectionable error, and use 'them' and 'their' where no plural is implied. I may actually use the plural where only a singular case can only be understood, a clearly grammatical error. I consider this the least objectionable alternative

Chapter 1

Well-Being

To live ethically, make moral choices based on your overall well-being; when another person is involved, temper that with empathy.

This is the basis of many ethical systems. The differences between these systems arise from how people define well-being and how they apply empathy. Most people begin with the ethical system of their parents - usually a system that is prevalent in the society at large. As time goes on, the system is modified by the person by the selection of variations in the original system or by the inclusion of ideas from other systems. Some people reject all given systems outright and try to start afresh. But we all give meaning to our lives in a way that makes our well-being as unique as our own identity. Likewise, each person is different in how much capacity they have for empathy and how they express it. Therefore, the meaning of these terms can be so different that the question arises whether there is an absolute sense of moral action or not.

Before we go further, I have to define my terms.

Morality is the study of appropriate behavior this is usually phrased in terms of what actions are good or bad.

Ethics are the practical aspects of morality. So to live ethically means that one is living by a moral code.

A choice is an action that can be good or bad, depending on the well-being of those affected by that choice.

A moral choice is a volitional decision whose primary aspect under determination is the goodness or badness of that choice.

Volition is the act of willing, choosing or deciding in a conscious manner. Volition requires judgment.

A moral code is a set of rules that help a person to make moral choices.

These definitions are open-ended: they are not given here as precise terms with an absolute meaning. The words "morality" and "ethics" are defined in terms of other words such as good and bad, that themselves need defining. The whole matter of definitions needs to be discussed in detail, but this will be done in a later chapter. What I shall do until then is to use the terms and definitions with sufficient rigor to make the analysis of morality possible, without bogging down in what are essentially technical details. But I shall appeal to the reader's intuition as the arbiter of ambiguity and source of detail.

Although actions are sometimes termed right or wrong in terms of their appropriateness, I will tend to use the terms good and bad. Note that there can be many good choices that cannot be considered moral choices by my definitions. Moral choices require a consciousness.

Also, if the characteristic of efficacy of action is being considered, good and bad can be considered to be equivalent to being effective or ineffective. Of course the final outcome of an action is only one aspect, although probably important one. Therefore good and effectiveness are not equivalent concepts.

So I began by saying that morality is based on well-being. Other moral systems are based on a principle of acting on the right motivations. Usually these motivations are the attempt to realize through one's actions some set of ideal principles. But even with this as a starting point for acting, the effect of these actions, if done in a moral manner, is to increase the well-being of individuals, societies and other entities. Therefore, although these actions can be specified in terms of the motivating principles, they can be analyzed and judged in terms of their well-being.

There have been other standards of morality that seemingly form a different basis. For example, morality has been defined as being imposed from above, such as from a god who oversees the universe. The moral code in this case is usually revealed to those expected to follow it through some supernatural event that is recorded for posterity. This type of morality may contain rules that are to be followed regardless of the well-being of the individual. It is usually implied, though, that the rules are for either the well-being of that god, or the well-being of the universe - or both. It is also implied that these rules are for the well-being of the individual that follows them, even sometimes when this is not apparent.

Morality has also been defined by what is good for society as a whole - to consider the most good for the most people. A moral code based on society's needs is usually not exclusive - that is, it is not expected that the good of society is the only determinant of the good. Instead, this is usually

a part of a moral code that also contains a set of rules that determine what is good for the individual alone.

I claim that any of these cases is also a form of well-being. The difference is that the context in which well-being is defined is larger. Both of these cases consider the well-being of the individual, but define some entity that is something larger than the individual and shift the focus of the determination of well-being outward to that larger entity. A metaphor for this shift can be seen when we consider that the well-being of the individual is greater than the well-being of the cells that make up the individual's body. An individual bacterium has only its own well-being at stake. But an individual cell gives up having only its own well-being as the basis of the goodness or badness of an action when it ceases being a cell all on its own and throws in its lot with cells that form a larger organism. Well-being is then defined over all of the cells in the body. Therefore, it is good to amputate a leg that is caught in a car wreck if that saves the body, even though the cells of the leg were lost. Typically, what is good for the body is good for the cell. But the human body lives each day with cells sacrificing themselves for the greater good.

But even though there may be a difference in the definition of well-being between the parts and the whole, a morality based on well-being must be a morality that is relative to the person, either as an individual entity or as the parts that make up the whole. So, to make a good choice, something must positively affect the well-being of a person. The question is at what level of organization the well-being applies to. There is certainly the well-being of the individual cells in the body. Most well-being preserves the basic life-force that is within each of our cells. But there are cases, such as the example of the loss of a leg, when the well-being of certain cells must be sacrificed for the well-being of the person. It can sometimes happen the other way - that the well-being of the organism as a whole must be sacrificed for the well-being of some or all of the parts. If the body is sick and needs to recover, this can mean a forced retirement until the healing process is completed, an action that may not be in the interests of the person as a person, but necessary for the person as a body.

Similarly, it happens that the well-being of the individual is set aside to maintain the well-being of the community. Sometimes it happens that the well-being of a whole community has been ignored in the interests of the state, which is an even larger community. It can even be argued that in some ways the well-being of all humanity can take second place to the well-being of Gaia, the ecosystem we are a part of. There is a whole hierarchy of entities, of parts and wholes, where the concept of person is centered at a particular place up the hierarchy. But there is a similar concept of personhood, a

wholeness that consists of a complex of parts at that particular level.

This leads to a second set of definitions

An entity is something that exists as a particular and discrete unit, having the property of well-being.

A person or an individual is an entity capable of taking volitional action in pursuit of its own well-being.

A community or complex is a collection of entities or persons.

Well-being is quantized by saying an entity is faring well or doing poorly based on good and bad actions, regardless of where they were moral actions or not. The terms virtue, vice or evil shall be reserved for morality; that is, I shall not assume an entity is evil just because it makes bad decisions.

If an individual acts in a manner to make the well-being of the entities affected by the action better, then the actor is said to be doing good. If the overall well-being is reduced, the actor is doing bad. Actually, it is most often the case that an action is considered good only if the well-being of all of the entities is improved. It can happen that a bad act is one in which many entities improve their well-being at the expense of one.

These definitions make a technical distinction that there are objects that are constituted in such a way that they have the property of well-being inherent in them, but they themselves are incapable of taking actions as an individual in support of this well-being. An example of this is a desert ecosystem where the conservation of water is necessary to the maintenance of the ecosystem, but the desert itself cannot be considered to volitionally act to preserve itself. This is an entity that is not considered a person.

A person is defined in a way analogous to the legal sense of a person. A legal person can be a corporation that is considered to act as a corporation in a way that establishes and preserves legal rights and duties that go beyond the legal rights and duties of the people who are part of the corporation. For the moral use of the term, a country can be considered a person also.

Is it possible for a machine to be considered to make a moral choice, such as an engine that turns itself off to avoid overheating? This may be a moral entity in the sense that there is a definition of well-being for the engine that includes not running while being overheated, but the problem is whether the machine is really acting volitionally or not. A moral action is considered to be volitional, in a way that leads one to describe the person as having free will. An electrical generator with a temperature shut off cannot be considered to have free will in any reasonable sense of the term, so consequently it is an entity, not a person.

I shall consider it perfectly reasonable for animals to make moral choices based on standards of well-being defined in terms of those animals and their

own sense or capacity for empathy, even those may be different or more limited than what a human is capable of. Although the ability to make a willful choice is not at the same level as what an average human is capable of, there is a sense of volition that can be applied to animals when discussing morality from a more inclusive perspective than is common. That is, there is no hard and fast dividing line that serves to separate the ability to make willful choices of the more advanced animals from that of the more limited humans. Morality must be discussed in terms of the normative behaviors of the collection of individuals under discussion. Although it is absurd to accuse a lion who kills a man of murder in a court of law, it could be meaningful to discuss morality of a lion who kills another lion in terms of the effect this has on the well-being of the pride of lions this individual is a member of, if it can be ascertained that the individual willfully made the choice to act.

To attempt to prove that well-being is the measure of morality, I would need to show that any other standard is just a special case of this generalized utilitarianism. In essence, this would mean that phrasing the definition of ethics in terms of well-being is universal, if the definition of well-being is broad enough. This seems to be a reasonable proposition. In effect, it is claiming that the notion of well-being has the essential property of being the result of good actions and that all good actions must bring about the well-being to at least one entity. If, starting from our intuitive notion of well-being, we show its identity as what is good to be tautological then we can consider well-being to be an axiom of morality.

This identity does not mean that what is good cannot have other universal properties besides the establishment of well-being. These other properties may be helpful for a richer understanding of morality. But well-being may be both necessary and sufficient for an essential understanding. If the equivalence of the two concepts is tautological, the establishment of a sense of well-being may be sufficient to establish the existence of what is good. Other properties may not be necessary.

I claim that it is not possible to find a case where some action is said to be good but there is no entity whose well-being is not positively changed. For example, the concept of delayed gratification is expressed in terms of the well-being it affords the individual in the long run, even though this is not obvious in the near term. Charity towards a stranger includes the good feelings that arise from the action for the giver besides the increased well-being for the recipient. Even if it is admitted that no increase in well-being will accrue to the giver directly, the increased well-being of society improves the well-being of the giver who is part of that society. If this is not enough,

the well-being of being part of God's plan can also be invoked. For that matter, anything that is good in God's eyes improves the well-being of God or the universe - there is no case where some good act diminishes God or the universe.

So, even though I could go on citing examples of how the concepts of goodness and well-being relate to one another, unfortunately, this relation cannot be proven. Even if every action that we conceive of results in the improved well-being of some entity, this does not mean that something that we currently cannot conceive of may not come along that would be considered good. We will have to take as a given that the improvement of well-being is equivalent to good. This principle is a thesis - a proposition to be accepted without proof.

Given that morality is defined in terms of well-being, there must be at least one entity whose well-being is in question. That is to say, some thing or action is never inherently good or bad.

The concept of well-being is not inherent in an object the way a physical property such as mass is, or in an action the way duration is. The mass of an object is the same for that object regardless of the environment the object is placed in and the mass of the various objects around it. But the morality of a slave owner in Biblical times, when slavery was a part of even the most advanced societies, is different from the morality of a slave-owner in the mid-1800's when most of the advanced societies had abolished the practice. It is instead a relationship between objects, or a consequence of actions of objects.

A physical property that is analogous to morality is friction. Friction comes from the interaction between two distinct objects, and the effect varies with the interactions of the different characteristics of the objects at the interface and the forces acting between them. There is also a difference in the friction between the objects if the two objects are just starting to move or have been sliding past each other for a while.

Just as friction is a function of the objects in contact with a given surface, morality is a function of the way objects are related to the entity whose well-being is under discussion. As an example, feeding a child is good, but feeding a child who is not hungry is not. This can actually be detrimental. A child who is not hungry but is force-fed can easily end up overweight, but also can suffer mentally from being forced to do something that is normally good in a bad way.

An all-encompassing entity from which to define goodness and badness is usually referred to as God. God may be considered in one of two ways, depending on whether God is considered to be all-encompassing or as a

separate entity from God's creation. That is, goodness, as viewed by God has the properties of being omniscient and all-encompassing. The second view is that God is an entity whose interpretation of goodness is only one of many such definitions, and the other views of goodness compete for attention with this interpretation.

Assuming God is conscious, or is an entity composed of volitional persons, the first view means that goodness, as experienced by any entity is just one aspect of the universal goodness of God. This means that the definition of the goodness of God simultaneously takes into account the well-being of all entities that exist in the universe. The second view was probably in the mind of Jonah when he tried to run away from the duty that God had called him to.

Let us illustrate the difference using an example. Assume that a person whose kidneys have failed would have their life saved if another person donates a kidney. For almost all people who hold to the first view, it would be believed that God's universal goodness would be in accord with donating the organ, even though this act would put the donor at some risk, both during the operation and if the single kidney is lost. But the nature of universal goodness means that this higher goodness takes all lesser contexts into account, so the act of donating is truly good for the donor too, despite the risk.

The second view of God seems to limit the deity - to make God into some sort of minor player in the universe. But this is not necessarily the case. In the second view, it is not to the good of the donor to donate the kidney, if there is nothing to be gained for that person to offset the loss. But still, the donation would be good as seen from the viewpoint of the Universe, because the loss to one person is more than offset by the gain in well-being to the other. This second case views the good of the individual and God's goodness as separate but unequal valuations, just as my household budget is separate from my government's budget but certainly unequal. Even so, the two views of the good interact and affect each other, even as these two budgets are affected.

If God does not exist, or is believed to be indifferent to good or bad, then the ground of action devolves on some other sort of individual. But it is still true that well-being is the goal, even with an indifferent universe. The universal requirement of well-being means that a good action to be taken by the individual entity, whatever that individual is, promotes the good of that individual.

Presuming that there are some unifying principles in the universe, it still may be possible that there is some unifying principle for well-being

that governs all entities. This principle can be postulated without having to consider the existence of a universal entity such as God as a prerequisite for its existence. This principle can be assumed to exist, for example, in the same way that gravity exists in the universe; as a force that affects the behavior and interactions of different entities in the universe which can be observed and measured locally for the most part without having to analyze or take into consideration the existence of a universal entity that gives rise to this universal force.

Instead of a universal power though, morality can also be viewed as a universal phenomenon or epiphenomenon of action the same way that consciousness is considered in relation to cognition.

In either case, the concept of well-being would have universal properties that enable everyone to measure the effect of an action, the same way friction has universal properties that make it a predictable and measurable property of physical action. But even though there are presumed to be unifying principles that govern the universe such as the law of gravity, this does not of necessity mean that well-being is as universal. Morality can still be defined without assuming a unifying principle for well-being. This limited viewpoint would consider that any definition of well-being would apply only to actions of entities only on this earth, say, or as applicable of to that part of the universe that human beings can significantly affect through their actions. That means that with only a limited viewpoint, morality would have to be totally redefined for an alien race, for example.

But morality must have some universe of discourse where general statements can be made if any general statements can be made at all. This universe may be limited in extent, but within that extent, if two situations are established and it can be shown that the difference between the situations is negligible in the way it affects the well-being of the entities involved, then the same rules or morality would lead to the same conclusions.

No definitive statement will be made here on the existence of God one way or the other, nor on the characteristics of God if such an entity exists. The theories about morality and the practice of ethics laid forth here are applicable regardless of the particular viewpoint of the question of God that each individual holds. Note that this does not in any way mean that the morality and ethics of a Christian, a pantheist, an atheist or a Buddhist are in any way the same. The question of God being a question of universals gives a profoundly different quality to the aims, judgments and practices of those who differ in their viewpoint of God. But there are still important characteristics of morality and ethics that transcend the differences that people have in their vision of God.

As we have mentioned, besides the individual actor, society can also be an actor. Just as the human body is made up of cells, society is made up of people. The cells in the body each have their own individual activities and behavior, but the behavior of the body as a whole cannot be predicted from studying the actions of the individual cells.

This is an explicit denial of extreme reductionism. Reductionism is the attempt to explain the whole as the sum of its parts. An example is to explain the biology of living things in terms of their chemical processes, where chemistry is ultimately reduced to physics and physical processes are reduced to the interactions of subatomic particles and physical laws. This is an ontological reduction, where ontology deals with existence or being. A related form of reductionism is epistemological reductionism, where the explanation of the activity at one level is best expressed by reference to the parts further down in the ontological hierarchy. Epistemology refers to the nature and origins of knowledge.

Western culture is pervaded with the idea of reductionism. Everything from mathematics, which is considered to be reducible to logic and set theory, to cooking, where the special qualities of a dish are often described in terms of the spices. We don't consider very satisfying an explanation of how things come to be unless we are told how the parts interacted to make it so. Extreme reductionism holds that this is the one sure way to knowledge.

Whether or not extreme reductionism is a valid way of looking at the world will be discussed in a later chapter, but until then, I shall not consider reductionism as a necessary prerequisite to understanding. That is, the actions of society and its well-being will be described in some cases as characteristics of society as a whole apart from the actions and well-beings of the individuals in society.

One problem that extreme reductionism has is choosing a particular level at which discourse takes place. After all, a human being is made up of cells. Therefore, the moral judgments of a person's actions could be considered to be the aggregate of the moral judgment made by each neuron and cell involved in the action. The well-being of the individual cells, summed together, will then explain the actions of the person. But at the next level down, the morality of the cell is based on the morality of the individual proteins that form the basis of cellular activity. We can also go down to the molecular level and the atomic level. But this is an absurd situation.

This is as much a question of perspective as it is considering how to analyze the action of the whole in terms of the parts. After all, if one presumes materialism, then there is nothing else making up the body other

than the individual cells and so on down to the physical particles. Extreme reductionism would say that there is nothing to be gained by talking about the actions of the aggregate when the actions are originated by the parts.

It is reasonable to say that the actions of society are made up of the individual actions of the people who make up the society. But it is both meaningful and helpful to consider the actions of the aggregate as an emergent behavior that does not occur at the individual level. Even if extreme reductionism were actually the case, the complex calculations involved to derive the societal properties from the individual actions of its citizens would make the whole process of discussing morality too complicated for any but the most trivial questions. Regardless of the truth of reductionism, we can arrive at a practical and more useful analysis by taking the individual, at whatever level, as the unit of discussion and considering the properties of society as emerging out of the nature of society or societies in a way that is above and beyond the properties of the individuals.

Moral action is an emergent behavior of the aggregate of parts that make up the individual. Emergent behavior means that a set of actions manifest themselves in an entity defined at a certain level of complexity, but will not be explained by reducing the entity into its constituent parts. It is not that we cannot give this explanation; we will not make this explanation if we are reasonably sure that this effort will not yield any deeper level of understanding. It is sometimes possible to judge the actions of individual cells as good or bad. For example, cell death is good, although not for the well-being of the individual cell. For another example, cancer is bad thing for a cell to do in terms of the person, although it is in the well-being of the individual cancer cell. But how meaningful this judgment is can actually be considered to be moot because free will is not involved.

At the other extreme there are certain good or bad actions that apply to whole ecosystems. One might ask if a choice results in the destruction of the whole habitat of an organism, has a moral choice been made? Since the well-being of the whole species is affected, certainly, even though this may not affect any person's well-being. For example, it is possible to destroy every last smallpox virus on earth. People currently know enough about genetics and probably enough about biology to reconstruct an existing smallpox virus. Due to the effect that smallpox has on people the destruction of the virus would definitely increase the well-being of humans. But in terms of the ecosystem, it could be argued that this destruction reduces the biological variability of the ecosystem as a whole, and therefore its well-being is negatively affected.

Therefore, the discussion of morality and ethics given here is framed in

terms of well-being as the single unifying principle that serves to unify the discussion of morality. This is itself a reductionist statement. Any entity that can consciously act on behalf of its own well-being or the well-being of others is a moral agent, regardless of whether the entity is human, animal, society or some universal principle such as God. Any entity whose well-being is at stake and can be affected by the actions of itself and others can be the recipient of moral actions. The universal ground of morality is free to be believed or interpreted in any of the different ways that humans have been known to believe, as long as the recognition is given that these views are individual views of the nature of the well-being of God or whatever universal property that substitutes for God. Finally, although the reduction of morality to well-being is a reductionist statement, it is not extreme in its reductionism, because the well-being of each individual entity emerges out of the nature of the entity as a whole and not the well-being of its constituent parts.

Chapter 2

Moral Relativity

People sometimes use the phrase "moral compass". The metaphor is quite apt. Ethical choices oft-times have to be made in light of the individual's unique goals, looking ahead to the ultimate end and not being deterred by what is around the immediate area.

The metaphor can be extended to imagine leading a moral life as sailing a ship. Living well is like setting a course and sticking to it, using that compass heading to prevent being blown off course by local wind or adrift with the current. If fate should cause the ship to go off course, reference to the charts and to the compass allows the ship to be corrected and the course reset. It happens to some of us that the goals change, or the realization occurs that it is not possible or very hard to reach the goal that was originally selected. Then the course changes, the compass reset and the individual heads off in another direction.

A ship in deep water is affected by wind and current, but other dangers arise close to shore. In that case the ship must be guided by a careful observation of what is immediately under the keel, with reference to the charts supplied by those who have been through there before. Each ship is at a different place, so the shoals are different for each. But the principle of sailing in shallow water is the same - make sure the water is deep enough for the ship to go through.

The metaphor illustrates how living ethically has often thought of, but like any metaphor it has its limits. One problem with the sailing metaphor is that it is built on a belief in moral absolutism. The compass is aligned to an absolute magnetic pole that guides each sailor no matter what the course of their individual life is. If a person gets into trouble the way a ship runs against the shoals, then the person has gone in the direction of immorality.

The shoals are fixed and absolute in some continuum of human misbehavior and an immoral action that brings that person to that place is directed away from the moral center. This would also equate the nautical chart to the sacred writings of a religion. The chart lays out the areas of danger to the person and warns against those actions that lead to those areas.

Footnote: In a 1994 address to the Acton Institute for the study of Religion and Liberty entitled 'The Necessity of Moral Absolutes in a Free Society', Justice Clarence Thomas says "Freedom did not mean that one could do exactly what one wanted. There had to be an understanding of right and wrong; of good and bad; of obligations; of responsibilities. These, among others were to provide the inner compass to navigate the vast oceans of a free society."

Religious teachings in the Judeo-Christian tradition tend to be like this. Activities such as lying, thievery or murder are expressly forbidden to all. Religious teachers set these laws out as absolute pronouncements from an all-powerful God that are meant to be followed by all and ignored at their peril. If the consequences of this wrong-doing does not come back to haunt the transgressor later, then certainly the penalty will be paid in the afterlife.

Moral absolutism is not limited to religions with a sacred book promulgated by a deity. Ayn Rand and her Objectivism defined a moral code that it is claimed can be defined axiomatically and can be applied through a rigorous application of logic. In her ethics, all actions that are life-sustaining are good and those that are not are evil. She particularly singles out altruism with its concept of sacrifice as denying life and thus being contrary to the positive virtue of rational selfishness. There is a clear demarcation between good and bad based on the recognition of objective human needs such that anyone who rationally analyzes the situation will come to the same conclusion. Her moral imperative is to "judge and be judged" - not to stand by and let actions that sacrifice life to go unanswered.

Other atheists and even some secular humanists consider moral absolutes to be a self-evident fact of existence. Subjectivism and Cultural Relativism do not lead to workable moralities, they claim, because there is no way to judge an action immoral if the person or culture believes it to be so. Although it is admitted that different cultures differ on the morality of certain actions, this does not mean that they differ on morality per se. They may differ instead on the facts that lead to making a judgment. For example, different groups of people differ on the morality of abortion. The difference between the groups is not necessarily exemplifying a difference in the underlying moral standard that questions if murder is wrong. Instead, what they differ on is the factual question of whether abortion is murder or

not. The standards themselves may be "self-evidently" true. This means that a moral standard such as "unnecessary suffering is wrong" should be self-evident. If someone believes that it is wrong, the burden of proof is on the disbeliever to find a counter-example.

An influential description of absolute morality came from C. S. Lewis. In "Mere Christianity", Lewis begins by pointing out that with most quarrels about right and wrong both sides accept the same standards of behavior. This standard is not rejected by either side, implying that it is universal. Those who deny the existence of an absolute moral law eventually get caught in an inconsistency.

Lewis suggests that one objection to the concept of an absolute moral law is that it is just an instinct, a type of herd instinct. But if two instincts are in conflict, the thing that judges between the two cannot be just another instinct itself. This higher law shows itself most strongly when it is called upon to suppress the stronger of two instincts and encourage the weaker. Also, if morality were an instinct, there should be some one impulse that we could call good. But every instinct is correct to follow in some cases and wrong to follow in others. Every instinct will eventually lead someone astray if followed at all costs.

A second objection is that morals are arbitrary social conventions. Although some social conventions are arbitrary, other social conventions can be absolute truths, the same way that the laws of arithmetic are absolute truths. Moral laws are in this second category. Since we can see changes in morality in time, changes for the better, we must be judging different moralities against an absolute standard that is greater than the moralities being compared. The fact that different societies differ in their moral standards demonstrates the presence of an absolute moral standard. The presence of a real New York allows two people to judge which of their impressions of New York to be truer. He argued that Christianity is judged to be preferable to the morality of Nazism by almost anyone because there is an ideal system that dictates those things that ought to take place.

Also, Lewis goes on, although there are moral differences between societies, the differences are minor. Although witches were put to death three centuries ago and are not now, both societies would agree that people who surrender to the forces of evil and do harm to others in bad ways should be morally condemned, even executed. The difference between then and now is a matter of fact, not of moral principle.

This argument for universality is mentioned by other thinkers. For instance, not only do we believe it morally wrong for the Nazis to perform medical experiments on concentration camp prisoners, we think everyone

should think it was wrong, including the Nazis themselves. This implies that certain actions are objectively good and others are objectively bad. If we were truly relativists then we would be willing to say that certain actions are acceptable under certain situations.

Before going further, I would like to make a distinction between two different forms of the absolute that arise out of the notion of absolute morality. The first type of absolute is that of an absolute origin for morality. This is an ideal position, from which all other standards are an imperfect deviation,

The second absolute is an absolute measure or standard. This absolute allows us to take two situations and compare them. An absolute standard is the measure by which we say that one situation is better than another or is the ethically the same. The standard is absolute in that our measure does not vary if in the two situations we change those features that are irrelevant to well-being. That is, abstracting the relevant issues in each situation, the situations will remain in relatively the same relationship if they are changed only by altering morally irrelevant conditions. For example, if it is agreed that buying a coat from a person with dollars is better than stealing the coat and paying nothing, this relationship still holds if it were a pair of shoes being paid for in rubles. This standard allows an observer to measure the degree of good or bad in a person or action by measuring the change in well-being brought about by that action.

Most everyone agrees that this second absolute exists; otherwise there would be no common basis for discussing morality of any kind. This does not mean that we do not measure the differences each in our own way. It is possible for each of us to judge the degree of goodness or badness differently. What is absolute is the agreement that if one situation is better than another, this relative ordering will not change with irrelevant changes.

Although the proponents of absolute morality can find problems with moral relativism, this does not mean that moral absolutism is free from problems itself. The first and most basic is the inability to pick an absolute center with unanimity. Beginning with the recognition of the great amount that is universal in the common experiences among all humans, it is possible to limit the bounds in which the absolute origin can be found. That is, the existence of common human instincts, while imperfect, does point out the general location of such a moral origin, if it exists.

But we cannot claim to fix this absolute center with absolute certainty without denying the fact of human fallibility. Although there could be an absolute moral center such as a moral code promulgated by a personal God, we cannot speak with certainty from that viewpoint without arrogance, usurping God. And there are many such viewpoints in the world that claim to be

the moral origin. Christianity is certainly preferable to Nazism. But which is preferable - Christianity or Islam?

Actually, it is possible to leave open the question of whether or not an absolute exists. Even if there is no absolute center, the fact that people share a common biological basis means that they may act as if the illusion of absolute coordinates were a fact. These absolute coordinates can be thought of as a sort of averaging over everybody's essential well-being.

But given this approximate center for human behavior, what is moral for human society may not be moral for intelligent bears or an intelligent anthill. If it can be shown that other biological systems have the freedom to act in terms of their welfare, the moral laws for these systems could be radically different from ours.

A second problem for moral absolutism arises when moral differences are dismissed as differences in fact instead of in moral principle. To say that the two sides on the abortion issue agree on the moral principles but differ on the fact of when life begins raises self-evident moral truths to become abstract quasi-logical axioms, that need facts for their instantiation. But is the start of life a fact? I would claim that both sides of the issue actually agree on the basic biological facts of fetal development; what they disagree on is a definition. And it could be argued that the definitions of certain situations form a basic part of one's moral code, different from the basic principles or morality. That is, both sides on the abortion argument agree on the same set of objective facts of biology, but their moral codes define that beginning of life differently.

Even granted the possibility of an absolute moral center, I am going to make the case for moral relativity as the simplest explanation of how moral judgments are defined and applied. To illustrate this, I am going to use the concept of relativity from physics as a metaphor.

There are three important types of relativity in physics: Classical relativity, Special relativity and General relativity. I am going to discuss Classical Relativity here, in particular the implications of the Copernican Revolution.

This story is one of the most well-known in physics. The way it is usually told, it begins with the observation back in antiquity that there are a number of heavenly bodies and planets that wander through the sky against the background of the fixed stars. They are the Sun, the Moon, and Mercury, Venus, Mars Jupiter and Saturn. They are known to remain in a bounded region of the sky centered on the ecliptic plane. It was believed that the planets were attached each to a sphere centered on the Earth. These embedded concentric spheres imparted motion to each other, making the planets move in a circular motion around the earth. The stars and the

Milky Way were fixed to an outmost sphere that was the furthest distance away.

The problem with this model is that the planets also exhibited a change in brightness at certain places in their orbits. Also, to make things worse, they were sometimes observed to move backwards in their orbits. This retrograde motion could not be explained with this simple model. Note that the fixed stars did not exhibit either of these behaviors. Except for the unusual events such as a nova, they remained the same, locked in a simple orbit that changed hardly at all year after year.

Around 150 BC, Ptolemy developed a more complicated system that would account for these discrepancies. The Ptolemaic system of describing the motions of the planets attached each planet to a circle called an epicycle, which was embedded, in the corresponding sphere, the deferents. Now, at the same time the deferents uniformly rotated around the Earth, the epicycle rotated around its center at a different rate of rotation. At different times during the rotation of the epicycle the planet would appear a little brighter or dimmer. Also, during the return motion of the epicycle, the motion would be greater than the motion of the deferents and the planet would appear from the Earth to move backwards.

This was not enough to completely explain the observed motion, though. A closer approximation could be reached by one of two ways: embedding other epicycles in the epicycles, or by offsetting the epicycle in the deferents, so the center of the epicycle was now not the center around which the epicycle rotated.

This geocentric notion of the universe was replaced by the heliocentric universe pioneered by Copernicus in 1543. In this model the planets rotate in circles around the Sun. The Earth itself becomes one of these planets, with the Moon rotating around the Earth.

It is important to note that the Copernican system is better than the Ptolemaic system not because the Copernican system is right and the Ptolemaic system wrong, but because the Copernican system is simpler than the Ptolemaic system. Actually, the Copernican system as originally conceived - as circular orbits- still requires a system of epicycles to be consistent with observation. Both systems, given an infinite number of epicycles on epicycles, each one much smaller than the last, in the limit reduces to the elliptical orbits that the planets actually move in. But the Copernican system, especially with the introduction of elliptical orbits has the simplest and more elegant explanation.

Another important concept is that this new way of viewing planetary motion gave rise to classical relativity. The Earth-centric system as conceived

by the ancients postulated that the heavens existed in an ideal condition away from the imperfection of Earth. Newton's laws of motion established a calculus of motion applicable both to the heavens and the Earth, such that calculations of relative motion were shown to be the same regardless of the choice of origin.

In classical relativity, the calculations of position and motion of different objects can be computed relative to any arbitrary origin. Given three objects A, B and C, the location of C as seen from B is the same as the position of the C as seen from A minus the position of B as seen from A. Now if all three of these objects are in motion relative to each other, the calculation of the motion of C relative to B is simply the difference of the motion of C relative to A minus the motion of B relative to A. A necessary precondition for relativistic calculations to be valid is that the measures of distance and time duration are absolute regardless of which origin is chosen to make the calculations.

This calculation of relative motion can be extended to calculations about the dynamics and forces of objects with mass, as long as the origin and the method of calculating distance and time defines an inertial frame of reference. An inertial frame of reference is one in which Newton's First Law of Motion holds. That is, given an object that is moving (or not) in an inertial frame of reference, if nothing comes along to interfere with that motion, the object will maintain that motion for all time.

It is possible to define a non-inertial frame of reference. For example, we could define our frame of reference to be centered on a rocket ship that is moving away from the earth at a constant acceleration. How this happens is not important - it could be scooping up the fuel to do this from interstellar space. By using this rocket as a frame of reference means that any object moving with a constant velocity appears to be accelerating away from the rocket ship. Using this frame of reference, basic techniques of measuring relative motion are rendered invalid. Attempting to make calculations about the motion of other bodies as if it were an inertial frame of reference would lead to inconsistencies.

The embrace of a relativistic framework does not in itself deny the existence of an absolute origin to the universe. For all we know, the universe could truly have a center based, for example, on the midpoint of the Royal Observatory in Greenwich England, and all motion of physical bodies is around that point in some complicated dance. For a more reasonable example, if all of the galaxies in the universe had a spherical distribution, where the bulk of the mass fell within a certain radius, and all of these galaxies were rotating around the center of that spherical distribution, then there

would be an absolute center for the universe.

Footnote: in General Relativity, this is Godel's solution of the Einstein field equations 1949.

I will show that moral relativity works because it is a cleaner explanation of moral action. That is to say, using an inappropriate absolute frame of reference leads to the explanatory equivalent of epicycles to explain the same phenomenon that can be described in a simpler manner by a more appropriate choice of a moral center. Even if there is an absolute moral framework, if this moral center is sufficiently different from the context of the observers, participants and society in which the particular situation occurs, then it is better to abandon that moral absolute in this case and use the relative moral framework appropriate for that situation.

What evidence is there for preferring moral relativity? The first is lack of a definitive majority of people agreeing on a single moral standard. Although there are many things in common with most moral systems, they differ significantly in the details, such as moral codes delimiting appropriate family structures, or the importance and degree of truthfulness required to be considered honest. Although from the similarities between human moral systems we could infer a moral absolute, no individual or group can claim to know this center to any degree of certainty without being challenged by the majority of people.

The argument for the existence of an absolute moral center is challenged by the existence of a significant number of major moral codes in the world represented by the major religions, and by the fact that these religions invariably fracture into different sects each of which begin to move their moral codes away from the center of that religion. The attempt to claim a moral absolute in spite of this fracturing ends up having problems of a nature similar to the epicyclical hypothesis. I will show this with two of these attempts to claim moral absoluteness that I discussed earlier.

The first attempt to show that despite the existence of sects there still is a moral center was based on a claim that these groups differ not in their moral principles but in differences of fact. If that were the major cause of the differences, then sects would arise due to these differences of fact, holding their principles the same. Under this interpretation, for example, the split between Sunni and Shiite Moslems is simply one of a difference between them on question of the existence of the hidden imam. This can be stated as if it is a difference in fact, but this difference also accompanies a number of differences in belief that magnifies the differences in practice between the two groups above and beyond the fact of the existence of such a person.

There have been a number of Christian groups that have formed on the

basis of a difference in fact, such as the historical claims of the Mormons or the claim of the Christian Scientist that health problems are matters of the spirit. But each of these splits is also accompanied by changes in moral principles and practice not explained only as differences in fact.

But granted that some of these differences arise from factual disagreements, there are just as many splits caused by a reinterpretation of basic moral sources, such as those Christians who consider the Bible to be a spiritual truth but not necessarily a factual record, or those who "speak where the Bible speaks and are silent where the Bible is silent." Thus ascribing a factual difference to moral differences is insufficient to account for sectarian differences. A relative interpretation that postulates group fragmentation as a difference in moral judgments relative to a precursor group explains things better.

The second attempt to show that there is a moral center was C. S. Lewis's claim that the existence of different moral codes shows that there is a moral absolute which is revealed through the process of comparing them: "If your moral ideas can be truer, and those of the Nazis less true, there must be something - some Real Morality - for them to be true about." This argument conflates the two concepts of a moral absolute - that of a moral center and that of a moral standard.

Most people consider their own moral codes to be either the best or close to the best. That is, they consider themselves to be able to identify the moral center as being close to their beliefs, even though they admit to not being there themselves. Therefore the comparison of Christianity to be better than Nazism will hold for a Christian just because of this relativization of the individual moral viewpoint to the moral center of the prevailing moral code. If most everybody shares what is basically a similar moral code and shares a moral standard, then for most people the Christian morality will look better than the Nazis - except to a Nazi. By the time a Nazi decides that the Nazi moral code is less true than the Christian code, they have already shifted their moral center away from Nazism.

It is certainly possible that the relative coherence of moral codes for different religions does seem to imply the possibility of a moral center, but this is not absolutely proved. One problem is that the moral codes are lumped together into major religions and philosophies with a few outliers, instead of being a uniformly varying whole. It is as if instead of us all living on this Earth with an easily determinable center, we were all living on an agglomerated collection of separate rocks kind of floating around together. When each person who believes in a moral absolute is asked to pick the center, they pick the center of the particular chunk they are a part of: the

Christian defines the center as Christianity, the Buddhist as Buddhism and so on. But it could just as easily be that if there were a true moral center, it would be some sort of average over all moral traditions - a center that lies outside of any one religion as currently practiced. In any case, the attempt to compute such a moral center would be difficult, if not impossible, and certainly contentious. It would be better not to even bother, since a moral relativism fits the facts better and is more practical.

This choice of the moral center as the religion of the person making the choice is the standard simplifying assumption of relativistic reasoning - it is easier to make sense of and make decisions about a situation if we choose ourselves as the center. This is termed the reference point of the observer. Choosing oneself as the center is more of a simplifying assumption rather than a claim that the observer is at the absolute center. For example, a Baptist could function very well morally, thinking in terms of Baptist concepts of right and wrong even though God were a Methodist. In fact, since people are fallible, the claim that a person's moral code is centered on the moral absolute is guaranteed to be wrong in some respect.

Another sign that morality is relative is that most past attempts to derive an absolute morality has only come up with a morality that is applicable only to humans. That is to say, attempts to identify the moral center have failed because these moral centers are not universal for all entities - they have no way of handling non human morality. Human morality cannot be uniformly applied to other animals even if limited to the more intelligent mammals who bear some behavioral relationship to us.

A comparison of other mammals shows differences in sexual mores, wars with other groups, ownership of property, territorial disputes, infanticide, population and resource pressures and other things. If they had a moral code, the code would be different according to the behaviors. In many of these aspects, higher mammals show limited free will, but free will nonetheless. They are capable of choosing an appropriate number of mates, whether to fight or take flight, whether to claim territory or surrender it, to take care of or kill infants and react cooperatively or in competition. They are driven mostly by instinct but the details are subject to individual and collective will. These moral codes cannot be integrated into a generally accepted core of human moral beliefs and have much of substance left.

The remainder of this chapter will discuss the implications of moral relativism. Besides dispelling some misconceptions, we will show how relativism allows us to see morality as it works in the real world, instead of having the inconvenient details buried by the attempt to fit actual morality to some perfect moral absolute.

I have been discussing morality as being measured from a center, either absolute or relative. This origin measures some space that represents entities and their actions. What does it look like? Of the two metaphors used in this chapter, the image of a sailing vessel and that of the motions of the planets, the first metaphor is probably closer to the image of a moral space. Considering all estimations of the well-being of an entity as a single value of being well-off or poor and considering time as the second dimension, the moral space becomes a two dimensional chart divided into two regions; the well half of the chart and the poor half. There is a boundary line between them that some moralities consider clean-cut. In practice this boundary line is a gray area - a stripe of varying thickness running between the two regions. Although the boundary is usually thought of as drawn between good and bad, that determination is a secondary consequence of the well-being of the entity. To decide if an action is good or bad, we must first determine whether as a consequence of this action the entities involved are better or worse off than before.

Entities are represented as a line going up the chart in time. Some timelines veer from the well-off to the poor-off region and back, but some timelines can be seen to be mostly in one or the other region. Moral actions are events that move these entities in one direction or another. They represent forces that interact with the entities. Given some moral act, if in general the entities affected by the action are taking directions that are toward the poor side, the action is bad; otherwise it is good.

In this, the simplest way of imagining the moral space, the boundary separating well from ill goes straight up the center of the chart, with the origin, whether an absolute center true for every valid chart, or a center that is relative for different ways of calculating well-being, would be located at some defined time but centered in the middle of the gray area. The region of well-being then would have positive values relative to the center, the poor region would be negative. God, being thought of as an infinite good or an infinite state of well-being, would be the infinity off one edge of the chart.

The difference between an absolute morality and a relative one can be seen in the changes of the entity lines. In an absolute morality, every chart would have the same trajectories for the entities, the only possible differences perhaps being when time is presumed to start and how to assign a number to well or poor. So different charts would start at different times and appear thicker or thinner, but that's all. The charts can be shifted or expanded to make them fit.

On the other hand, a relative chart would have the same regions for well or poor, and the same straight line going up the center, but every chart

would show different time lines for the entities depending on how well or ill is measured for each action at that time for that entity. The remeasuring of the well-being of an entity does not change the boundary between good and bad; it moves the position of the entity.

Ayn Rand's Objectivism, since it equates the good with being enhancing to life tends to look like the two dimensional chart. But there are many dimensions to well-being, some of which cannot have a number assigned to them. For example, a person's financial worth can be considered one numeric measure of that person's well-being, but their medical health cannot be reduced to a number, or even a vector of numbers. To have a more realistic portrayal of the moral space, we need to go to multiple dimensions. In this case, the regions on the chart become some sort of volume of space, and the line separating good from bad becomes a complex surface. The only requirement is that given any two points, one in the region of well-being and the other in the poor region, the difference to go from ill to well must be positive in at least one dimension.

In a multi-dimensional space, some of the dimensions are qualitative, not quantitative. This means that it is difficult to get a measure on a single entity for a single event in time. Even granting the existence of an absolute standard of measurement (as opposed to an absolute origin) it is hard to pinpoint exactly where something is in terms of its well-being. That means that even for an absolute moral standard of measurement, a measure cannot be reliably made on a single event alone. Except for the most egregious outliers, any single event is independent of morality. A moral standard can best compare between two events, especially those who differ in some qualitative feature. It is difficult to discuss morality of a single situation, but much easier to discuss differences. An example is to point out that for both the pro-choice and the anti-abortion viewpoints, a live birth is preferable to an abortion. The two sides may differ on the moral outcome of these two actions, but they will both agree on the relative difference between them.

So we have a moral space that is comprised of a number of different dimensions by which we measure well-being. Actions we term good or bad are seen to be forces that move the position of the entities around in the space. Unlike physical forces, good or bad actions do not usually act to gradually change the well-being of the entities involved. An action typically takes an entity from one point of well-being to another point located at the same time, but some distance away depending on the loss or gain of well-being and which dimensions changed. Also, unlike physical trajectories of objects like billiard balls, the entities do not have to be touching to interact. So a single moral act can make a number of entities dispersed through the

moral space act in quite different ways.

The moral judgment of an act as good or bad is then a two-stage process. First, the entities affected by the action must be identified and a determination must be made as to how each entity's well-being was changed. Second, these individual measurements must be combined to make an overall determination. Just as a multi-dimensional determination of well-being must have at least one dimension positive to say that the entity is better off, there must be at least one entity whose welfare is improved to judge an action as good.

If an action could be considered to act on a single symbolic entity that represents the cumulative well-being of all the entities involved in a moral act, then good and bad can be drawn and measured in the moral space as well or ill for that representative entity. Given this composite position both before and as a consequence of the action, well or ill is the line drawn between the two positions, with a good act being one with a positive direction from beginning state to consequence and bad as negative. The line does not even have to be in the gray area to determine this; goodness makes the world relatively better, and badness makes it relatively worse.

Most moral codes do not handle all of the complexities that real life is able to manifest itself. Instead, we usually consider only changes in the one dimension of our measure of well-being that has had the greatest effect from an action. This single measure is used to determine whose lot has improved or gotten worse. We also tend to focus on the entity that had the greatest change and base our measure of good or bad on whether the person most affected did better or worse. In this judgment, we are usually quicker to focus on loss than gain, probably because our normal expectation is that in most human interactions, both parties come out the better. Therefore, any loss in well-being is the exception and is indicative of an act that is not good.

Thus, if there are two people, one who buys a loaf of bread from the other, we tend to view that transaction as good because they have both received in value something that to them was better than what they gave up, so they both shifted slightly to the better side. On the other hand, if the loaf was stolen, we judge the action as bad because one person lost, even though the other gained. Even if the person stole from someone who has an abundance of bread and used that stolen loaf to feed a starving child (which means that the decrease in well-being for the one with abundance was more than offset by the increase in good) this action is at best condoned, but never thought of as virtuous.

Footnote: unless the thief is Robin Hood and the victim is the Sheriff

of Nottingham.

One advantage of looking at morality in this way is to show how the moral viewpoints of different moral origins are affected by perspective. Like the perspective of real objects in space, two moralities that are close together have the most pronounced disputes with actions whose effects are in the same range as the differences between the moralities. Where one Protestant denomination would disagree strongly with another about whether salvation occurs at the moment of confession of sins or at the moment of baptism, both would view with equal disdain the emphasis of the Buddhist on the elimination of suffering without regard to salvation at all. At the time of the Protestant reformation, it was accepted that to burn an apostate at the stake was the lesser evil, compared to the alternative of letting this heresy continue to be promulgated, thereby putting countless souls at stake. To a Buddhist, practicing the ways of non-violence even to the point of vegetarianism, this would be seen as just another killing. But all would agree that the actions of Genghis Khan and the Mongols were beyond the pale.

Instead of having greater differences between moralities at greater degrees of good and bad, then, we tend to have the differences more pronounced at the range of difference similar to the differences in origin. These differences can even appear to change an action from good to bad in the same way the planets move in retrograde motion, but only for actions whose moral codes measure these actions near to the gray area between right and wrong. When actions are far away from the center, then like the fixed stars, the motion is viewed as the same for all. Different Protestant faiths may differ on the moment of salvation, but once this is passed, the length of the lifespan of the saved is not an issue. Therefore both would consider a cure for cancer as a good because the resulting change in lifespan only affects when, not if the person will be saved. But for a morality that views overpopulation as the fundamental problem and a Malthusian collapse as the ultimate ill, such a cure may not be considered an unalloyed good, but just another pressure placed on an overburdened Earth.

It has been claimed that, with a relative morality, one can admit the possibility that for any action no matter how reprehensible, a moral code can be constructed that will admit it as a virtuous act. This is looking at each event in isolation, not relative to the goodness or badness of other events. Of course, if a relative morality is unbounded in its choice of origin, this is possible. The fact that a shift of perspective is possible is given as an argument that a relative morality leads to inconsistencies because of course no true morality would admit to a whole range of reprehensible actions that

could be allowable in some moral relativity.

This is no more a refutation of moral relativity than the equivalent astronomical observation that given any three points, it is possible for a mass to pass through these three points in an orbit around the sun. The fact that there is a whole region of points far from the ecliptic plane where no planet or asteroid passes does not make classical relativity absurd.

The problem goes away because the selection of the relative origin only serves to give an arbitrary name to a particular event. But this does not change the relationship between this event and all the other events in the moral space. In comparing two events, we can still judge one event to be more or less bad than another, even if we have arbitrarily defined genocide as a virtue. This affects the utility of our definition of virtue, not the way in which we run our lives. By comparison to the other possible choices we can make we can still steer ourselves to be more and more good, regardless of what value is assigned.

As for the argument that this means that anything is acceptable under a morally relative framework, I will point out that this is in effect claiming that for any two given situations, there is at least one choice of a relative origin where the relative location of the two situations is reversed. This is as false as claiming that under a relative interpretation of celestial motion, there is a viewpoint from which the Moon can be determined to be over the North Pole, or the distance of Mars from the Sun is less than the distance from the Sun to the Earth.

The choice of an appropriate origin in a relative morality makes a big difference in how easy it is to make moral judgments. The choice of origin can be a choice of focus - it can be individual, group, society, culture, mankind or Gaia. Once this gross level is decided on, a finer distinction can be made. For instance, different Protestant denominations could value differently a particular exchange between two people. This is a matter of context. The first choice is an 'order of magnitude' choice. Choosing the origin in this case scales the measurement of the moral standard so that the important features of the actions of interest become visible.

The choice of origin should be reasonably close to the types of entities involved. It would be hard to resolve a moral problem between two members of two similar Baptist groups, if taken from the perspective of the moral obligations of countries. Similarly, it makes no sense to judge the moral actions of the United States with the same criterion you would use to judge your neighbor next door, unless you could find a meaningful translation of the actions of the body politic into the actions of the person. Things such as killing or stealing do translate, but a quality that applies only to groups

such as political cohesiveness would not make much sense to consider in terms of a single individual.

The choice of origin as a reference point can be relative to one of the entities involved but this does not mean that any one of the entities has an equal right to be considered or that each entity's origin is valid. To go back to the metaphor of classical relativity, choosing some moral codes are like choosing a non-inertial frame of reference. Using this frame of reference leads to difficulty in analyzing the well-being of others who do not share the same frame of reference, and may even lead to inconsistencies that show that the choice of this origin to be invalid. A classic case of this is the choice of a moral frame of reference to be a person who considers themselves to be privileged in ways that no other person is. That would lead them to consider any action good which increases their well-being regardless of the effect on others. This person may certainly be making moral determinations, and in their perspective the determinations could be internally consistent, but their conclusions would not translate into any other frame of reference without inconsistencies. Such a person will have to do without the freely given cooperation of other people; no moral consensus can be reached that includes that person's limited viewpoint.

One of the questions to ask is whether the choice of origin allows us to make valid and consistent determinations of the well-being of all of the entities involved in the situation. As a further example, consider making moral judgments from the standpoint of the Nazi government. If there were nobody but Nazis in the world it would be impossible to declare this frame of reference invalid. But there are more than Nazis in the world. Using this frame of reference to measure the well-being of other people and other governments leads to inaccurate or inconsistent results according to the absolute measure of well-being. This is not just a matter of taste. For some group to hold a Nazi viewpoint of morality requires holding logically inconsistent propositions that lead others to conclude the morality to be flawed, such as the believed superiority of a given race despite the inability to objectively prove such superiority or to even be able to define who is a member of that race or not.

The use of a morally relative framework is required when moral practice is not limited to be within a particular society with a generally accepted moral tradition, such as when two or more disparate traditions interact. When two different religious traditions have a dialog, they start from a basis of ecumenism that takes the beliefs in common between them and ignores the differences. The unilateral imposition of some aspect of a moral standard is not greeted with acceptance by the society being imposed on;

this is even the case for a weaker society that has to accept the imposition. These types of dialogs typically result in an agreement of a limited nature, covering only the practices that are in common between the two cultures. If there was a single absolute moral standard it could be argued that each of the cultures, both being correct but only seeing part of the truth, would be combined by adding together the practices that each alone has instead of subtracting off the differences.

One of the biggest problems with trying to use an absolute frame of reference for morality is the problem of time. An absolute origin as it is usually conceived, implies that the moral code is unchanging. But in actuality, as time goes by, moral codes change and they often change in similar ways. This shows that moral codes stay comparatively the same relative to a moving frame of reference.

Let's give a pair of examples comparing Christianity, Judaism and Islam. In the case of slavery, this was an accepted part of society in all three religions. In Christianity, the religion even urged the slaves to obey their masters. In Islam, masters were admonished to treat their slaves fairly. By the Industrial Revolution slavery began to be abolished entirely and currently slavery is found in only a handful of the most backward societies. Despite the permissiveness of the religions, as time went on the societies imposed more and more prohibitions on the practice. It may be that the belief in an absolute moral code made it difficult for groups such as the Southern Baptist Convention to give up the idea that slavery was acceptable.

In the case of adultery, in the Jewish tradition, it is punished by stoning. Christianity, coming later than Judaism, eliminated the practice, since the Gospels indicated that adultery was grounds for divorce. Although Islam reinstated the practice, in the more advanced Islamic countries, society eventually eliminated it.

In both of these examples, claims that the moral code is absolutely fixed in time are contradicted by the actual practice of the adherents of the religion. Later religions supplanted the older teachings with more relevant practices, and they themselves changed in step with the advancement in the societies that came in time. Even supposedly fixed and unmoving moralities are seen to shift in time, as if they have an inherent moral force pulling them towards a greater possible well-being.

Some of the changes in ethics have led to moral precepts covering situations that were not even foreseen by the original religions because they involve technologies that weren't even invented at the time. An example of this change through novelty is the condemnation that parents would receive if their children were not vaccinated for infectious diseases if the vaccine

were available and an epidemic were raging. Although there are moral admonitions for parents to take care of their children, the definition of care is extended as technology opens other possibilities.

Stealing property is a crime prohibited in practically all moral codes. This is another case where the moral code changes as the definition of property expands. In the modern world, such thefts as those of air rights and stolen data have no precedent in the original moral codes and must be specified by analogy to the older moral code, abstracting principles that seem to be generally applicable out of more concrete injunctions in the original laws.

More recently, the development of birth control has created profound changes in the morality of sex. These changes are taking place in similar ways in the ethical practices of different cultures as the changes in technology are introduced.

Although Judaism, Christianity and Islam usually make a claim that their moralities are absolute in the sense that their moral codes are fixed at the time of their revelations, there are some moralities that do not have this limitation. One of the most important in the Judeo-Christian-Islamic tradition is that of the Baha'i faith. They believe in a concept of successive revelation where the moral code is fixed for an age by one of a successive number of prophets. Each of these revelations extends and perfects the moral understanding of the previous age in step with the increased understanding of the people and their capability to receive these revelations. This admits that each moral code is an approximation to an ideal, in effect pulling a current relative moral center closer to the absolute moral center that only God can comprehend.

So, despite the claims that only an absolute morality is workable, we can see that morality in actual fact is relative. We have an illusion of an absolute only when moral choices are limited in time to the present and to a particular homogeneous group. When trying to make moral decisions where nontrivial differences are involved, we need to choose a relative moral framework in which the well-being of the entities involved are easily represented. This does not mean that we average the frame of references together to come up with a composite, something not even possible if one is to judge the effect of a given set of actions on the well-being of an individual, group or country. The test of having an adequate choice of a moral center is three-fold. First, it must be valid in its conclusions; it cannot allow for inconsistencies that lead to contradictory evaluations. Second, the choice of a framework must allow for the distinctions that are drawn to be visible. This requires a choice of a framework close enough to the entities involved so that their differences

are visible but not exaggerated out of proportion. Finally, the choice should be one that simplifies the judgments and simplifies the comparisons of relative well-being. If the conceptual equivalent of epicycles is required to be introduced to make the theoretical framework match reality to the accuracy required to make a decision, then the choice of framework is impractical.

We used the metaphor of classical relativity to illustrate how morality is actually judged in a relative manner. The metaphor worked even though real space is different from a space of moral actions, because some of the principles governing the two are the same. That is, in both spaces we are making an assumption that there is some universal principle at work that is underlying our attempts to measure events in these spaces. The concepts of universality and measurement are general principles, not just in physics. These ideas, when first recognized as basic conceptual principles, led to the spread of the scientific method to the study of many fields.

Of course, every metaphor has its limits. Universality means that two events are relatively the same if the changes between the two events are not relevant to measuring the effect. That is, the example of buying versus stealing is invariant if the relative distinctions are preserved between buying, stealing and the moral codes under whose judgments the comparisons are made. But in classical relativity, events are invariant under both translational and rotational transformations. That is, if the events differ by moving from one point to another, then the events do not vary in their effect, and if the direction of action is rotated, this also means that they do not vary. Although it is possible to translate a moral action from one region to another and still have its moral judgment remain the same in the same moral framework, there is no analogue to rotation. Rotation works in real space because there is no special distinction between the three dimensions of space. But it is not possible to equate moral dimensions, such as those of health, freedom, happiness and knowledge, for example.

To repeat, the necessity of making judgments in a relative moral framework does not deny the existence of an absolute moral center. But regardless of relative or absolute origin, there must exist an absolute moral standard that allows for comparisons to be drawn between the different reference frames, or there can be no way to have a discourse. The characteristics of the moral standard form the topic of the next chapter.

Chapter 3

The Moral Absolute

The absolute standard of morality is the capacity for empathy.

Well-being, since it is a multidimensional measure, is a difficult thing to quantify. With some features being quantitative and others qualitative, it cannot be simply reduced to a single measure. Therefore, to judge an action as good or bad is usually a relative statement, unless there is an unequivocal measure that overrides all of the other considerations. There is also the problem of our own perspective. For example, for most of us, we consider ourselves as good as most other people, regardless of whether we consider mankind to be mostly good or mostly bad. This makes it difficult to stand as a distant observer apart from an action and to judge with any reliability. We cannot impartially measure, the way we would measure the physical properties of an object in space. We need empathy.

We have shown that all actions are relative to a moral reference frame. This reference frame is the local context of the entities involved in the action. But laws are written from an absolute perspective, sometimes by invoking God. The question is, are there any real absolutes? It could be argued that life or death is an absolute. The result of walking off a high cliff - that should be an absolute. But it is not to the person whose life is at stake - that person is dead. To the grieving next of kin, the effect may depend, for example, on whether the body was recovered or not, and there the question is more relative. Therefore even the condemnation of someone a murderer is not a simple absolute. The judgment must take into account the values of the people who remain and the society they are a part of.

So how can I claim that there is a standard that is morally absolute?

The moral absolute is a measure of well-being and thus good and bad. To be an absolute standard, it must be, by and large, the same value for

everybody. Due to the nature of well-being, it cannot be quantitative, since there are too many features of well-being are qualities that cannot be turned into values. Any attempt to pick a selection of values of the moral dimensions and attempt to claim this as the standard yardstick is bound to lead to failure - no matter how well chosen, a significant minority will disagree with any such choice. Thus the moral standard is a limit value; a measure of maximal well-being that everyone can agree on.

Consider any moral action that involves more than one entity. There must then be at least two measures of well-being involved. Then how can one get any number of observers to agree on how to measure the consequences of the action? Well, any measurement done using the origin chosen by any one of the observers is liable to be challenged by any of the other observers. So their moral framework cannot be used. Also, if we tried to use the moral frame of reference of the first entity as the origin to measure the well-being of the second entity, we would again have a problem because every observer would have a different perspective on the relative differences between the two entities. The only way to achieve a uniform measure is to use the frame of reference of the entity involved in an action for that entity's well-being and no other.

The absolute moral standard then must be the maximum well-being possible for an entity using that entity's moral reference frame. This measure is what the entity uses to determine what is best for it, if that entity is capable of some degree of introspection. Even if the particular entity has no sense of what its ideal well-being is, thinking beings such as humans can make reference to the ideal in their deliberations. So even if the entity itself has an unconscious striving for its well-being, there can be an external entity that can make moral judgments using the moral standard.

Well-being as an absolute measure arises from the ability of an entity to act to enhance its own well-being, though. - it is not simply a consequence of life. If the Gaia hypothesis is true, the ecosystem acts in its own well-being without being considered a living organism. It may not be able to make reference to any conception of its ideal well-being but it can still take action. It may be possible to equate the two and say moral action is the definition of life. This is not what is usually thought of as living, but it might lead to a different and fruitful viewpoint of what life is.

This absolute standard is a usually unobtainable measure. That is, even the entity itself usually cannot measure it completely. Most entities, conscious or not, have in them an imperfect sense of this ideal yardstick as a reference point with which to measure their actual well-being. This yardstick may be limited to only a few dimensions of well-being, if the entity has

a limited sense of what is good for it. Even in these dimensions, the entity may not be aware of all that it is capable of. Also, the entity's understanding may be inconsistent or faulty instead of simply incomplete. This means that it is possible for some observers to have a better conception of the well-being of some person than that person does. That would mean that the empathy of an observer could be a better guide than the entity's self-knowledge. But for most entities, in reality both self-knowledge and empathy are less than perfect.

To take an example of the moral standard, consider some moral action that involves a Christian and an atheist. One is committed to the belief that people possess an immortal soul, and their measure of well-being must contain that as an aspect. The atheist does not hold that view. There cannot be any consistent moral judgments made by the two parties if the atheist considers the welfare of the Christian without reference to the other's soul. Likewise, the Christian will not be able to have a correct view of the situation if the question of the atheist's soul was a factor in the matter. To achieve a true empathy, each party must assume the other's viewpoint, as well as have an understanding of the other's situation. The atheist can understand and take into account the Christian's need for salvation even to the point of knowing that the Christian believes that the atheist's soul is at risk. But this concern is part of the Christian's sense of well-being - it has no bearing on the well-being of a committed atheist.

This standard does not mean that empathy has to be limited to an erroneous understanding, though. With empathy the atheist can even have a better sense than the Christian of the Christian's well-being, for example with a better knowledge of human psychology or even a deeper knowledge of the relevant theology than the Christian possesses.

This moral standard is not a fixed quantity. As time goes on, the maximum well-being can change. There are two types of change that affect any person. The first is how they act to improve or worsen the capacity for well-being of themselves and others. This is not an improvement in well-being per se. Instead, this change comes about through an increase in self-knowledge. The second change is external. These changes can also cause a shift in the moral framework. As the capability to achieve well-being changes, the way well-being is determined changes to account for this increased capability.

These changes can be illustrated through the capacity for happiness. Although a dimension in which overall well-being is measured, happiness itself can have a number of qualities. An internal change in the capacity for happiness, then, would be the case of a person who never left the place they were born until later in life. Assume that the first time they left home, they

find that they enjoy traveling. This newfound enjoyment has increased that person's capacity for happiness.

A case of an external change has come recently through the creation of the Internet. The happiness (or sadness) that comes through interpersonal communications has been qualitatively expanded through the use of this medium.

Note that both of these cases broaden the degree of happiness possible. They do not add a new dimension to well-being that did not exist before. This is extremely rare - it may not have happened that a totally new dimension has been created since the use of money led to wealth as a measure of well-being.

I would like to say a little about justice in the context of moral standards. Justice, of course, is the attempt to return equity to moral acts. That is, if some act has resulted in the gain of one at the cost of the loss of another, justice attempts to right this imbalance. Ideally, one would want to act to cancel out the effect of an bad act, but it is quite often impossible to undo the effect of a loss of well-being, restoring a person to the point before that loss occurred. Instead, many acts of justice try at least to eliminate any gain made by the bad act. This is easier to accomplish than to create a gain to offset the loss - it is a fact of life that it is easier to destroy than create. But somehow it is rarely considered that undoing the loss is equivalent to undoing the gain. That is, it may be justice if an embezzler is fined by the court though the victim gets nothing, but it would not be as satisfactory if the court lets the criminal go unpunished but have the victim be given restitution, even though the two are the same in terms of righting the balance.

There is also the concept of social justice. How is this different from justice as is usually understood; is there a need for the concept at all?

Social justice is defined by the World Bank as "equality of opportunities for well-being, both in and among generations of people." They consider social justice to be an attempt to create sustainable development in at least three areas: economic, social and environmental. Development is sustainable if it "meets the needs of the present without compromising the ability of future generations to meet their own needs."

The difference between simple justice and social justice is that there does not necessarily have to be a wrongdoer for social justice to apply. In the case of simple justice, there is a determination that a loss of well-being occurred as a result of an act usually accompanied with the gain of well-being of the perpetrator. In the case of social justice, the imbalance of well-being between different groups does not result from the consequences of a moral

act, because no single such act can be found. Instead the determination is made by looking at the difference between the well-being of the parties being compared at a given time, regardless of the actions involved.

This comparison is usually thought of as a simple comparison of well-being between or among the parties involved. This works if the parties use the same frame of reference for determining well-being. But it leads to inapplicable results if the frame of reference is different. For example, a comparison of the well-being of a person in a first-world country compared to that of a third-world country would show the need for justice for the third-world people until the day that third-world country becomes a first-world country. The person in the first-world country will cry foul about this, especially if the claim can be made that the effort of that person contributed to that country becoming and remaining in the first world. But the third world member can also claim the same degree of effort into bettering themselves, except that the effort does not show equivalent results.

The way out of this impasse is to base social justice not on a comparison of well-being relative to the action in which the parties were affected, but by comparing the well-being of the parties in terms of the ideal well-being they are capable of at that time. This analysis provides a comparison of differentials between where the party is to where they can be; a difference value that I have pointed out is easier to analyze and more reliable than an absolute value.

A contentious issue involving justice is the question of redress for past wrongs. Although it would seem that justice is the same now matter how far back the transgression occurred, the effect of the transgression is diluted by time. This is due to the fact that the ill-gotten gains of a bad act and the resultant loss of the injured party are the prime components of change in well-being in the here and now, but if it occurred in the past, subsequent events will have a larger effect on where the person is today. There comes a time that restitution - beyond a simple acknowledgement that something bad had been done - is itself unjust, for it stakes a claim to benefits gained by other good efforts of the transgressor after the transgression occurred. The gains realized by the transgression may have magnified the benefits of any efforts after the transgression, but that does not given the injured party claim to all future benefit.

This dilution of effect seldom plays out in a single person's lifetime, but it is important to factor in between the generations. No one would seriously claim restitution, for instance, from the current citizens of Mongolia for the actions of Genghis Khan. Even a generation can make a difference. It is very difficult to justly address the claims of descendents whose parents had their

property confiscated against the rights of the descendants who inherited that property. Both sides have a legitimate right to the property and it would be impossible to come up with a consistent policy on problems of this sort that would handle every case.

Although the well-being of an individual is an absolute value, determined from that person's own sense of what well-being is, determining the good is done by comparison to the frame of reference in which morality is being defined. This frame of reference is determined in large part to the degree of applicability the moral code extends to; that is, the field or region of applicability. Specific things such as alcohol consumption for example are local to a country or group. As a frame of reference changes to include more people or other entities, the larger the field becomes and the more moral judgments are limited to more general issues. So it is possible to make judgments for or against polygamy in certain groups, but not for all of humanity, yet it is possible to judge the more general characteristics of family life, such as a condemnation of the abandonment of young children.

So, we have found that the absolute in morality is in the capacity for empathy: the ability to consider the well-being of others in terms of their definition of well-being.

This determination changes the consideration of a person's acts to be relatively good or bad as the field of applicability changes. This change, besides having the effect of washing out differences in the interests of generality, is also related to the fact that good and bad are in part determined by comparison to a reference value drawn by summing over the entities in the field.

Let us use the metaphor of a forest. A forest is made up of trees. Every tree is striving for sunlight as people strive for good. A tree does not grow forever, trying to reach an unobtainable perfection. Instead the tree grows to its appropriate height as we strive for an appropriate level of good. We are all defeated if we try to eliminate all our faults; the best we can accomplish is to be good enough.

But any given forest has different conditions that allow for a certain level of development of the trees in the forest. This means that to label a tree as a small or large tree, a comparison is made with reference to the trees that are its neighbors. A tree of a given height may be great or small in one forest, but fit right in the middle in another. The height relative to sea level is not important often, unless one is trying to consider characteristics, such as air pressure, that are measured from that general referent. Thus, the closer the neighborhood, the finer tuned the measure.

Thus, a moral judgment is not relative to the individual - it is a measure

whose upper/lower bounds tighten as the neighborhood contracts. It is almost meaningless to an individual alone, and can only be given in terms of a society. A moral judgment would make sense in an individual context only to a loner, or for acts that do not involve another person.

Problems arise when making moral judgments in the presence of change. This can often lead to paradoxical situations. They often occur because different people undergoing change at different speeds will have different notions of parity. That is, the differences that arise when considering different characteristics of well-being for comparison become more complicated when the reference value for that characteristic changes in time enough to be noticeable.

Consider for example, two people who come from the same village, but one leaves for an advanced culture and the other stays behind. Regardless of who rejoins who's society, the one who returns will seem less moral. This is caused by the fact that it rarely happens that one society is "better" than other, despite the fact that one may have a greater level of well-being. In going from one society to another, there is a change from one norm to another. Acting in the same manner after leaving one culture and going to the other would lead to a different judgment against those actions because the norm is different. For example, if a person is coming from a place with a higher standard of living and attempts to live at that same standard, it would make that person appear selfish to the poorer neighbors. On the other hand, if a person were coming from a lower standard of living, their lower giving to charity would appear to be selfish to the more affluent neighbors. These changes are ameliorated by increasing the frame of reference to include both societies, and are resolved by the newcomer changing behavior to fit the prevailing society.

An even more curious paradox comes from the fact that the more societies adhere to their moral code, the more likely the society is to function efficiently and thus inspire changes. This is due to a faster increase in well-being leads to a faster increase in human potential. A society that does not have a functioning morality is more likely to stay mired at a certain level of development or at best proceed slowly. A society that is more moral, due to bringing out the best in people, will create out new potentials and even new problems, which have to be addressed by changes in the moral code. With the novel ways that well-being can be increased, there are also new opportunities for action, both good and bad. These must be identified and dealt with.

This new identification is not just limited only to new moral situations, though. Like any field of knowledge, the more that is known about the field,

the more that people will become aware of and address the ramifications and previously unrecognized subtleties. This applies also to moral acts. This is more likely to happen in a society where moral values are better understood and followed.

The final moral paradox has to do with a frame of reference that moves with time, remaining forever in the here and now. This seems intuitively to be a sensible thing to have, since it acknowledges that moral actions are best judged in the present, and a moral frame of reference that stays in the present will be more likely to provide judgments unqualified by the need to include conditions that are made necessary by referring back to a moral code centered on one era in history.

This choice of a moving frame of reference involves both the need to keep one's moral code up to date and to keep a constant perspective. For example, the admonition that society should "plan for seven generations" is such a moral imperative. It provides a metric upon which the effects of moral actions are judged - provide for enough time for the repercussions of the act to dampen out.

The distinction between a fixed and moving frame of reference can best be appreciated with unrepeatable actions. Common cases of this involve the exploitation of unrenovable resources, or the creation or destruction of something that results in permanent changes. For example, the production of oil results in a permanent loss of this resource for future generations. If it is believed that the amount of oil is available for many dozens of generations but the moral center is fixed on the point at which oil was first utilized, then a case could be made for the unrestricted use of this resource in the context of three or four generations.

But recognizing a moral center that is always in the present does two things. First it forces the judgment of the wisdom of exploitation to take into account the fact that every successive generation will be faced with increasing scarcity. Second, being centered in the ever-changing present also forces the recognition that the very concept of exploitation is not a fixed property. As the generations change, newer uses come up for that resource that are totally unknown. The moral ramifications of how to handle nuclear waste are a case in point. The effort is now currently underway to make inaccessible to future generations what we now consider a dangerous form of garbage, but could in the future become more prized than jewels.

But this simple consideration of a moving frame of reference does not lead to any paradoxes. They arise from having a coordinate system that moves in time along with the generations in such a way that the generations collapse into one.

Consider an example of a forester. The forester has children, and the single source of income for this family is the growth and harvesting of trees in a forest. Assume there are two adults and four children and that any trees cut down in this generation will not be replaced by harvestable trees for another 20 years - the time it takes for a generation to pass.

How many trees to cut down in the next 20 years? The first choice is to cut down half of them. This will support the family until the children are grown up. The children then will have the same number of trees available to them, and the size of the forest is stable over time. But this does not take into account the fact that there are twice as many children as adults. It is reasonable to expect that all four children will marry and start their own families, but only half of them will stay to inherit the forest. The other half will leave to partake in the family business of the families they married into. That half will marry and stay could seem to indicate that the forester should only cut down a third of the trees, so that there are two thirds of the trees available for the next generation, which is twice as large.

But what about the next generation? If there are 6 times as many trees as are necessary to sustain a family at a subsistence level, then the grandchildren, presuming that the size of the families remains the same, will not have enough to live on. This situation may be held back for a generation or two depending on how large the forest is, but there is a limit that will eventually be reached regardless. This is the kind of argument that led Malthus to claim that eventually there comes a time when overpopulation causes a collapse in population due to a scarcity of resources.

A way out of this impasse is to consider every generation to be collapsed into a single archetypal family. Considered this way, it is possible to reason in the following manner. It is possible to subdivide the forest only a finite number of times so, except for only the initial generations that are populating a new territory, this is not an option as a way to provide for a an archetypal family. That means, as far as the family is concerned, only one family can take over in the next generation from this one. Due to the consideration that a family is composed of two adult parents, only two of the children can be provided for as eventual owners. The others must be provided for in another way.

The only conclusion to draw in this case is the realization that if there is only a fixed amount of resources, effort must be made in each generation to increase the resources. This can be done in a variety of ways: First, the parent can buy more land. This is also equivalent to earning enough in this generation so that the children could make that purchase. A second alternative is to create new resources. There are two ways this can happen.

One is to discover new uses that the forest can be put to that would increase its yield, such as growing mushrooms, or hunting game. The other way is to transform it, such as turning it into farmland. These pair of second alternatives requires creativity, because finding new uses for a resource or transforming the character of the resource both require thinking in new ways. This requires education. Again, it does not matter who is educated, the adult or the child. When viewed through the concept of a single archetypal family, the effect is the same.

The paradox lies in this collapse of many generations into a single one. For example, since there is a certain satisfaction in being able to give to one's children, there is usually an emphasis placed on sacrificing for this generation so that the next generation can be better off. This is due to the fact that the increase in well-being of the child is worth more than the increase in well-being of the adult because both generations get something out of the children being provided for. But the calculus of delayed gratification comes out differently when collapsed to a steady state. In that case, the benefits given to the children are equivalent to the benefits given to the parent, since those benefits were the ones that the grandparents had work for. The way to determine what to save for your children and spend for yourself is to find a balance. Obviously, both generations' needs must be met, but after this is done, where is the balance point? It may not be possible to rationally choose what to save or spend if the two balance out over all choices in the acceptable range. If 60% of the parent's income meets the needs of the family, then taking all of the rest and leaving the children none is the same as leaving all the rest and taking none.

The way out of this impasse is to sometimes consider the generations as separate, even though, in the previous example of the forester, this collapse helped to reason through to the need to educate both adults and children as being essential to sustainability. In this second example, we could split the excess between the generations and give an equal share to each in the family, but this balance point is essentially arbitrary. If we go back to considering the transaction as going from one generation to another, the balance point that took into account the extra pleasure in caring for one's children is the more reasonable one.

So we have determined that the absolute in morality is the capacity for empathy. This capacity allows each of us to see and measure the amount of well-being around us in the same way. Justice is the attempt to right the imbalances of well-being that can be seen around us. This can be a relative measure based on the context of the injustice, but it can also make reference to the maximum attainable well-being, as is done with social justice. The

redress of injustice can be tempered by time, though. We have also found that the measurement of relative well-being depends on the context in which the measurement is taken.

We have also discussed some paradoxical situations associated with the attempt to measure well-being. Since we determine what is good or bad in terms of our local context, the stranger is usually found to be wanting. Another paradox comes from the changes that come in a moral society that make the morality change faster to keep up with the changes. Finally, we have a paradox in how to measure the change between generations. Do we look at the generations as a sequence, or do we fold that change into an eternal now?

Chapter 4

The Golden Rule

The golden rule is a direct consequence of empathy being the moral absolute.

We have established as a basic thesis that for a moral act to be good, that act must work to promote well-being. These acts must be more than just random; morality has a volitional component to it that presupposes free will. Without this presupposition morality is nothing more than a descriptive body of knowledge. That is, if all actions are randomly or deterministically generated, falling into either the good or bad category based on a measurement of their outcome, then there is nothing more that can be done except to catalogue and describe the different outcomes as a result of these unwilled processes. But we usually demand more out of morality than that. It must then be translated into ethics, the prescriptive body of knowledge that tells the individual what to do.

An analogy can be drawn between basic science such as physics or chemistry and engineering. In the basic sciences there is no consideration of volition - the planets move around the sun and hydrogen and oxygen combine to form water by their nature. There is no need to discuss how they came to be that way - if these fields were expressed as logical syllogisms, the preconditions would be given as arbitrary postulates.

But engineering implies volition. For a satellite to trace an orbit around the earth, someone must have designed the satellite, filled its tanks with liquid oxygen and ignited it at a given time and place. The results of these volitional actions are the forces generated by the chemical reaction leading to the orbit the satellite assumes. The engineer has the free will to design that satellite according to engineering principles or not, but the success or failure of the mission depends on the successful application of this free will.

The question of free will in engineering is not as vexing due to the objec-

tivity with which the results can be measured. Making the wrong calculations will lead to the satellite ending up in the wrong place - an objectively verifiable event. In morality, we instead find that the outcome of volition is determined by the measurement of the well-being of the entities affected by the action. To have any hope at all of having a morality that is universal in its understanding and application, it must be necessary that each entity's well-being is defined in terms of itself. This results in a measurement that is subjective no matter how hard we try, since it is ultimately measured in terms of the ultimate well-being of the entity - due to the fact that each of us is unique - so no all-inclusive standard can be applied.

It could be argued that each person might be characterized as a complex, or vector, of quantities and qualities that express the person's well-being, and that an objective standard could be derived by a complete description of the dimensions that make up well-being. But this would only work for entities that are essentially nonsentient. The inclusion of sentience makes the characterization subjective, almost by definition. It is more than a simple definition, though. We shall see in a later chapter that the implications of modern logic mean that any objective measure is bound to fail, due to the will of the thinking person.

So given that ethics consists of the prescriptive injunctions that point the way to moral action, the question is what basic principle can be given as the overarching framework for moral action. The answer is clear. To act in a way that is good, strive for the greatest increase in well-being. The greatest increase in well-being is defined relative to the ultimate well-being the entity is capable of. This maximum is measured in reference to that entity's definition of well-being. Therefore to be able to volitionally act in a way that maximizes the good, act with the other entity's definition of their well-being in mind. This is the Golden Rule. Usually expressed as "Do unto others as they you would have them do unto you", this formulation results in the precept "Take each action with thought of the best interests of each entity as that individual sees it."

This is why, in most major religions, some expression of the Golden Rule is a primary virtue, sharing this primacy with that religion's conception of the ultimate. This is one of the two commandments of Christianity, but it is also a basic tenet of Islam, where this basic force is expressed as the Principle of Equality. In Buddhism, where the existence of suffering forms the ground of being, with enlightenment we can achieve empathy through the four boundless states of loving kindness, compassion, sympathetic joy and equanimity. In Taoism, empathy is also a primary virtue, expressed as compassion. The three treasures for the Taoist are compassion, balance and

humility.

Even without volition, descriptive morality shows that the most successful society is one that practices the Golden Rule. This can be expressed as a simple series of logical statements. The first proposition to show is that, given any two societies, one of which is more moral than the other, the more moral one works as if each moral act of each individual serves to minimize the overall difference between the aggregate well-being of all of those entities affected by an action and the well-being they are capable of. The second proposition observes that, regardless of how the aggregate well-being is calculated, the individual components of that aggregate are determined in reference to the maximal capacity for well-being of each entity, and that each of these calculations are done independently of the other, since the individual definitions of well-being do not overlap. This is true not because each entity functions independently of the others (they don't) but even allowing for the interactions, the definition of well-being is centered on that entity alone, incorporating those interactions as part of its unique definition. Finally, we can show the main result that, given any two societies, one of which is more moral than the other, the more moral works as if each moral act of each individual is operating to minimize the difference between the entity's definition of maximal well-being and that entity's current state of well-being. Thus the more moral society functions in accordance with the perceived well-being of each individual, whether the individuals are perceived to be volitional or not.

This descriptive sense of morality can be used as the viewpoint by which to create and describe a morality, but the lack of an assumption of free will leads to a somewhat contorted expression of the results. Instead of describing the functioning of each individual as a behavior that manifests certain moral properties, it is simpler and more direct to assume free will exists and talk about the functioning of the individual as volitional in nature. Although the question of the nature of free will shall be addressed in a later chapter, for this reason we shall now assume that free will exists.

The determination of the overall goodness or badness of an act depends on how the judgment combines the changes in well-being of the participants into an overall aggregate or whole. This determination is part of the relative moral framework in which the judgment is made. It hardly ever suffices that the combination is a simple combination where the well-being of each entity involved is equally considered with the rest. This is especially false for those moralities based on the concept of a God whose wishes and desires are sometimes in opposition to the wishes and desires of lesser beings. In that case, the difference could assume an almost infinite order of magnitude.

The only way that a pure equality would work is to reduce morality to a set of operations between separate entities, denying of the possibility of any aggregate well-being, and where the potential well-being of each entity is adjusted so that all appear equal. This approach is too simplistic to capture all of the complexities or moral action between people, groups, nations, animals, plants and ecosystems.

Another extreme case of an aggregate function is the equally simple morality of an absolute selfishness. That is, regardless of the well-being that other entities possess or are capable of, these values have no consideration in the determining of what is moral. Only the well-being of the actor and that actor's potential count for anything - the rest is zero.

The first objection to this type of a moral aggregation is that it does not take into account the fact that each individual operates as part of a whole world and that this aggregate is the same whether the whole world exists or not. To make an absolute selfishness work, then, is to assume a type of solipsism. Another way to accomplish the same thing is to admit that the outside world does exist, but the well-being of the other entities becomes components in the measurement of the well-being of the individual's selfish needs, so that, having been taken into account at the level of well-being, need not be given a nonzero estimation for the purposes of aggregation. This in effect means that the outside world only matters in terms of how it helps or hinders one's own well-being.

This type of self-centeredness is, when considered from a purely analytic standpoint, certainly a potentially functional way of operating. But to be truly functional it results in an analysis whose conclusions end up no different than a more conventional morality. The only real difference between the two is merely a matter of definition - what is the individual? The solipsistic definition is really doing nothing more than equating the individual with the universe. This usually creates a need to find some sort of description or means of thinking about what the individual formerly was - that person whose well-being now encompasses everything. This often translates to a separation of the individual and the outside world into parts of this greater individual, where the "will" or some "controller" determines the nature of reality and the rest of the world is just perception or illusion under its control.

The self-centered viewpoint that admits the existence of the world but ignores everything not related to the selfish needs of the individual is little different. Instead of expanding the definition of the individual to be that of all reality, it actually truncates reality to fit the individual. That is, all aspects of all entities in the world that do not impinge on the well-being

of that individual simply do not exist anymore, because they do not play a part in the evaluation of the morality of an action.

Because the differences between a pure selfishness and a more conventional morality are mostly definitional in nature, there is abstractly no difference in the functioning of an effective morality using these viewpoints or a more conventional aggregate. In practical terms, though, the difference may be large. These differences hark back to the discussion of why morality actually functions from a relative standpoint. The inappropriate choice of a moral framework makes it very difficult to effectively compute the well-being of the entities involved in the action.

In an absolute selfishness, these entities do not go away: their moral coordinates are transformed so that they are viewed strictly from the viewpoint of the individual, meaning that any determination of the individual's well-being requires a non-trivial coordinate transformation to be successfully incorporated into this new, much more complicated measurement.

The attempt to completely perform this transformation ends in failure, usually in one of two ways. The first and most obvious is that the person who functions from a purely selfish viewpoint makes bad decisions which eventually catch up with them. This happens when the person just does not take into account the needs of other people into the moral equation. The consequences of this omission leads to the other people acting in a way that negatively impacts that person's well-being in a direction that was unforeseen or ignored. The second mode of failure is that the individual chooses to limit the interactions with the rest of the world to what is manageable. This is usually done by focusing only on considerations of what is in it for them. In reality, it could potentially be true that the whole of what is good for the other entity can be translated into the viewpoint of what is good for the individual, but this transformation to the selfish viewpoint is too complex to carry out in a limited time. In effect, pure selfishness just puts back the epicycles of the absolute frame of reference, because it is an absolute frame of reference centered on the individual.

In opposition to pure selfishness is a pure altruism. Altruism in its purest sense - taking action with no thought of gain for the self - is simply the aggregation of the well-being of all parties involved, where the well-being of the individual is scored as zero. A pure altruism could lead to an overall maximization, if it could be proved that any thought of self would lead to a gain far outweighed by the loss of the aggregate well-being of a number of other entities effected by the action, even if their individual losses are small. But this attitude, even though popular in many periods of civilization and among many groups is ultimately suspect.

First, it ignores the very real need for well-being that every entity possesses. If the well-being of the actor is not taken into account by the actor itself, the responsibility for that individual's well-being is shifted to other individuals, if it is not to be simply left to chance. But this is not efficient: typically we know ourselves best, so selfishness is more certain and also more efficacious.

It may be that a person is considered altruistic for the simple reason that their needs are few. It can be argued, though, that this person is practicing the virtue of simplicity and is just as selfish and non-altruistic as everyone else. If the lack of desire for needs were actually altruistic, then the normal person would be more altruistic than someone whose health or handicaps place extra burdens on others.

Second, although implied balance can be true - that the thoughts of oneself take away from the concern of others - this can be extended to those other entities as well. There are potentially a myriad of other entities whose well-being might come into play in almost any circumstance. That is, if everything is interconnected, any action in reality affects every entity in the universe. The reduction of the scope of an action to a select number of individuals is the recognition that although there can be an effect on everything for every action, the practical effect is nil for almost every entity. But if there is an infinitesimal effect upon every entity and all of those entities are considered equally, it is possible to imagine that the overall effect of an action is very large but the inclusion of any entity is vanishingly small. To make it possible for any action to be evaluated, it is usually taken as a prerequisite that there is a method of aggregation where the effect of an action on an entity drops off in inverse proportion as the size of the context of the action increases to include these entities, eventually going to zero at some point. This is done, if for no better reason than to make judgment practical. This is in contrast to the aggregation function of the pure altruist, though which starts with its origin at zero.

Instead of a pure altruism, the more practical approach to morality is a rational self-interest. That is, each entity should value themselves as slightly higher than the other entities in their neighborhood, if simply for the fact that they are more capable of discerning what is in their best interest. If everyone were operating in their best interest using their own self-knowledge, the overall efficiency of society is bound to increase since everyone is functioning as an expert in what they know best. This also extends to their nearest neighbors. Having knowledge of those closest to oneself, simply due to the fact of familiarity, means that one can take their well-being into consideration at a higher level of understanding than that of a stranger. This

aggregation naturally decreases as the context expands.

Since the most basic precept of moral action is the Golden Rule, which urges the actor to take into account the viewpoint of well-being of the other entities, consequently the most basic bad action involves dehumanization. Dehumanization is simply defining the well-being of a human as something other than is best for that person. The most naked and obvious type of dehumanization is to consider other human beings as lesser creatures. This is an overall downgrading of that person's worth that is unwarranted between all people of equal potentiality.

Racism is usually justified in terms of potentiality. It is usually justified by the racist with a rationalization for considering that member of another group as having fewer capabilities. But this type of distinction has been consistently shown to be wrong. Due to the flexible nature of mankind, any claim to a limitation on one group or another will be shown to be false when enough counter-examples are found.

The question arises whether it is ever possible to value one person or group over another. For example, consider a moral act involving a saint or a person with exceptional talents who might bring great benefits to mankind. This act may be judged with extra emphasis because it involves that person's well-being. This person is valued over others, since that brings about a greater well-being for everyone else affected by that action. This type of reasoning can be extended to groups that may have special talents or are organized to handle situations where their abilities are of exceptional value.

It is certainly reasonable that we can value a saint higher, but this decision must be done on a case by case basis. Dehumanization comes from attempting to make general prescriptions that do not take into account the particular qualities of each case. Sometimes we grant too much consideration to an exceptional person; more than they are actually worth. Even for a saint that moral consideration fails to bring the expected results more often than we'd like. We also can hold special group in too high an esteem and the populace at large suffers as the result. Again, this is due to having the gain in our special consideration outweighed by the decrease in well-being of the ordinary people.

Besides devaluing people in an overall estimation, dehumanizing can occur if one does not value another entity in terms of all the dimensions that entity considers important to its well-being. For example, to act to benefit a person may not be a help if the action, for example, just benefits that person's happiness, when in this case the person considers their financial state as equally important. But complications arise from the fact that even though the well-being of an entity is defined in terms of that entity, they

themselves may not know what is best for them. Therefore, a determination of what is best for a person must be made with a careful distinction between what the actor feels is in the best interests of the recipient - in terms of what is important to the actor - versus what the actor is able to determine is a dimension that is truly important to the recipient's well-being.

This set of distinctions applies to justice as well. Justice is too often applied from the point of view of the well-being of society. Since justice is the redress of a loss of well-being of an immoral act, justice must be applied by taking into account the definition of justice that the victim has. If a penalty is to be applied to the transgressor, the penalty that society demands may be different from that which the victim believes is just; both viewpoints need to be taken into account. Also, any redress must take into account the well-being of the perpetrator. Justice has often been seen as dehumanizing because of this belief that it must be blind. Justice must be blind in the sense that everyone must be considered equal in terms of the worth that is attached to each person. But a justice that does not take into account the particular situation of each person, victim or perpetrator, is stunted in its ability to respond to the needs that it is required to meet.

Social justice is especially difficult in this regard. Simple justice is usually in response to an action where the effect on well-being is clearly a consequence of this action. Still, a humane response must take into account the actual and potential well-beings of the parties involved, a group whose identity is often obvious. Social justice is a comparative statement where no single time or set of actors can be immediately determined, but still an injustice has happened. Because social justice is in response to situational problems where there may be no transgressor, the act of dehumanization can sometimes come more readily. One of the most common cases is that, in the absence of a real transgressor, one is invented. It happens quite often that in a case of social justice, the group who is better off is blamed for the poverty of the other group, even if they did not act with malice toward the other group.

Dehumanization reveals itself sometimes in the way that the well-being of friends and strangers are differently valued. This can be considered a question of aggregation the way selfishness and altruism is, but with a function that takes different people into account at a different rate, depending on how close they are to the individual. Obviously, each entity's well-being can not be considered equally when there are a large number involved. It would seem reasonable that the consideration given to each entity drops off in proportion to the amount of effect that action has on the entity. But it is common that the effect on an individual is discounted much quicker the

more they are considered a stranger.

An example of this is the injunction to "Buy American" - that is, consider the worth of a job of a fellow countryman more important than the job of a person in a foreign country. Although it is admirable to take care of your own, this does not mean that one should consider more valuable the well-being of one's family, group, city, state or country above that of another. It is good to consider that the effect of an action has a stronger effect on those close to you, but it is not good to claim that these people are more valuable. An unqualified injunction to buy products made nationally is nothing more than an economic nepotism. It is important to recognize that the foreign worker has just as much right to a job as the local worker.

These two considerations are often confused. It is right and natural to consider most important the well-being of one's family and friends. There is where one's primary responsibilities lie and one has the greatest effect. Because of this, it is right to value the benefit of one's actions on one's family over the benefit given to a stranger. But this increased valuation does not mean that these people are inherently more valuable, even though they are more valuable to you. Everyone should be considered as equal in their inherent worth. That is why it is important to take into consideration the country one is a part of but to remember it is part of the larger world. It is important to strike a balance between regionalism and globalism.

The heart of the Golden Rule is to treat each person with empathy. This is not an easy thing to accomplish. It often happens that one honestly believes that one is being empathetic to the other party, but what is actually happening is that one mentally makes the other into a person who is actually a mental copy of oneself. This is another form of dehumanization because it does not recognize the values and beliefs of the person - it just takes the outer appearance of the other and ignores what that person really is. For example, for a Christian to consider an atheist as just another soul to be saved is dehumanizing, because it ignores a basic and very fundamental part of that atheist's belief system. Similarly the demand of the humanist that the theist acts in the social sphere without involving their spiritual beliefs about the relationship between God and Man negates an equally important part of the theist's nature. There must be parity here between the two groups.

Sometimes this type of dehumanization is rationalized by claiming that the value system of the other person is flawed and that one's own value system is "right" in some deeper way. Too often this is not real concern - it is arrogance. The imposition of part of the observer's belief system in the determination of what is in the other's best interest must always be done

with a great deal of humility. It may be true that the other's person's self-knowledge or sense of values is flawed, but this in no way means that their values are wrong. True empathy comes when one is able to take the other person as they are, using your personal knowledge to augment and complete their view of themselves, not to replace that view.

One place where this comes up currently is in the treatment of patients at the end of life. Too often palliative care orders are ignored in the attempt to sustain a life that the person does not want sustained. This sometimes is due to the care giver considering their own value system as more important than the patient, although sometimes it comes out of a sincere belief that this is the best for the patient. But it is ultimately dehumanizing in that it does not truly value the patient and the way that patient defines what is important to them.

Another form of dehumanization is the attempt to sanitize one's viewpoint of the other - that is, to see the other person without seeing their flaws. Of course, the opposite is equally true. It is dehumanizing to turn a person into a saint or a devil when everybody's nature contains both imperfections and good qualities that keep them from being an absolute symbol. Besides flaws of character, there are also flaws of understanding. To consider another person as if they are aware of all of the facts or know the full implications of their beliefs does not admit the full humanity of that person. Each of us constantly make errors of science or religion - not allowing them the possibility of error and making allowance for that does not show empathy to the fallible nature of humanity.

Like the balance between globalism and regionalism, a truly empathetic concern for others is a balance between the recognition of how each person is essentially the same and yet how we are all different. There is no simple prescription here. A true empathy must balance both similarity and difference in determining the well-being of others and deciding on what course of action to take. With the idea in mind that the best action is the one that maximizes well-being, the balance actually comes out to be a balance between knowledge and ignorance.

If one has perfect knowledge about each entity involved in an action, then the well-being of the entities can be easily maximized by taking their individual characteristics into account. But we never work from perfect knowledge. We are functioning in a subjective state ourselves, trying to determine the subjective state of the other individuals who we are interacting with. In absence of knowledge, then, we make reference to our own experience in similar situations and with our knowledge of people who have similar characteristics. This ability to generalize makes up for our ignorance.

Once we have learned the particulars, we can act from knowledge. But our knowledge can never be perfect.

Although an empathetic understanding of other humans is hard, harder still is the attempt to have empathy for that which is not human at all. Although the difficulty of being empathetic in our treatment of animals is the obvious example of this, we must also be empathetic when we consider the well-being of society and groups. All of these have needs that are totally different from that of a human being. An animal's well-being might be tightly bound to responding to the instincts that drive it, for example. Societies have needs such as peace, health, and happiness that may be similar to these characteristics in the individuals that make up the society, but what these characteristics means to a society will have completely different parameters than in a human. The whole in this case is not the sum of its parts. Acting in ways that are in the best interest of a forest or river is also an act of empathy, even though it is not possible for a person to "put oneself in their place" when it comes to an ecosystem. The attempt to do this though, has led to some of the more transcendent and spiritual feelings that mankind is capable of, affirming a oneness with all of life and the universe. It is necessary though, to translate these experiences into action, so that the things we do to the world around us take into account the needs of our world in and of itself.

I will end with an illustration of the different ways in which empathy can manifest itself by considering vegetarianism. Although there are many vegetarians who are so for purely dietary reasons, I am considering those who are vegetarian out of a moral belief. It is quite often the case that vegetarians consider each living creature to have the same right to life and well-being as a human, in effect establishing an equivalence that transcends the bounds that reserve special concern to just humans. In aggregating the well-being of different entities, then, an equal weight is given to living things regardless of the difference in mental and physical capacities that they are capable of.

Quite often the difference between the vegetarian and the meat-eater is expressed as a difference in this relative valuation. The meat-eater is often thought of as giving a consideration to each living thing that distinguishes between the levels of awareness that they are capable of, valuing a human over an animal. Some vegetarians may consider both human and animal as equal, but there are also vegetarians who recognize that a human with its greater capacity may be more highly valued, but not highly enough to subject these creatures to suffering to meet a human dietary need. In either case though, the vegetarian is thought of as one who values the life of an

animal higher than a meat eater would.

There are those who would consider meat-eating as acceptable though, even though they may equate the value of an animal to that of a human. The difference is that the meat eater considers that each human or animal is acting in accordance with their own nature. That is, the right for a human to eat meat is granted in the same way that the wolf has a right to eat meat. This type of equivalence recognizes that it is no virtue to abstain from eating animals such as rabbits and mice whose reproductive capacities are geared to an equilibrium where their members are eaten before they reach reproductive maturity. This is a recognition of the basic balances of nature. In the case of domesticated animals such as chicken and pigs, these animals gain from the caretaking of the humans who eat them. This bargain can be struck with a respect for life that is just as profound as that of the vegetarian.

This chapter has derived the Golden Rule from the absolute moral standard of empathy and discussed the implications in applying the rule. We have discussed both the need to consider the well-being of an entity in terms of the definition of well-being that the entity uses, and we have also compared the different ways in which the well-being of the entities involved in an action aggregate to determine the morality of an action. Besides selfishness and altruism, there are ramifications in how to value the welfare of others such as nationalism. In determining the well-being of an entity the ways in which dehumanization occurs are explored.

Dehumanization is the most basic kind of moral error, but not the only one. The next chapter shall consider that each action has to strive a balance between cost and benefit. Many moral errors come from imbalances of this kind.

Chapter 5

Moderation

In the application of moral principles, moderation is the key.

There are many people who do not drink alcohol. Abstinence will certainly ensure that one does not overdo, not having to face the effects in the morning after a couple of bottles of wine the night before. But wine is not all bad - medical research has shown that a glass of wine a day is good. The inability to sustain moderation is the problem.

Moderation also comes from the toleration of fallibility: the recognition that the real world does not allow the existence of perfect states or qualities. This is especially true for well-being. Since there are a myriad of qualities that go into a full determination of well-being and they are often interrelated, there is no way of achieving the highest possible condition for even one of these qualities.

As an example of this, consider the tradeoff between freedom and justice. Absolute justice can be considered to have been accomplished when every action that has an imbalance between the parties involved in that action is rectified in order to restore that balance. Absolute freedom, on the other hand is the ability to take action unfettered by any restraints placed upon those actions by other persons. To be able to achieve an absolute justice in an imperfect world, each action would have to be answered with some sort of rectification, since no act can be perfectly balanced. But this means that every act would be constrained and there would be no freedom at all. On the other hand, an absolute freedom would mean that no act, no matter how unjust, would be constrained. The best place to be is where the amount of freedom and the amount of justice is maximized. This balance point is not necessarily at the 50-50 level. Given some measure of how much freedom and justice we can give up (or attain) it may be possible to have 95% of the

possible justice by giving up 10% of one's freedom, or 10% of justice gives one 92% of one's freedom. In that case, the best choice may be to settle with 95% justice and 90% freedom.

Morality has similar properties, although there is no equivalent to Planck's constant to measure the effect. Given any two dimensions of well-being, the attempt to maximize one of these will, once we pass a certain point, ultimately lead to a decrease in the other. Once we pass the balance point for these two dimensions, the combined measure of well-being will decrease, eventually dropping off towards zero.

This trade-off occurs both for the person and for society. The attempt to maximize the health of a person will, if taken too far, lead to a decrease in that person's happiness, sense of security and wealth. A society that is secure against enemies, internal and external, gets out of balance in terms of its honesty, happiness and efficiency. Even the effort of striking a balance itself, being a separate (though hardly independent) dimension of well-being, cannot be maximized. That is, the effort involved in "hill-climbing" serves to lower the hill. It is better to expend a certain amount of energy to "get close". Although a higher hill may be out there, it won't be as high if you strive to reach it.

One case where this tradeoff is seen in its simplest and most abstract form is the two forms of the Golden Rule. "Do unto others as you would have them do unto you" is the positive injunction. Its negative is "first, do no harm". The trade-off between the two is usually a probabilistic one. That is, if one does not allow for a reasonable risk of something harmful happening in order to accomplish a good, the available choices of action are circumscribed to the most limited positive gains. Neither the ultra-conservative path nor actions that recklessly endanger the other person are called for. To function in a way that maximizes the chance of the best well-being, a moderate risk should be accepted, or even encouraged.

There is a balance between good and harm in most situations, which is not just the result of risk. Although the injunction "first, do no harm" is often applied to the field of medicine, both medications and surgery incorporate both health and harm in their effects. The cut of the surgeon's knife destroys health flesh before the illness is treated. The benefits of the medication come with the side-effects that must be accepted. Although the benefit and loss have some randomness, this is secondary. Primary is the fact that almost all medical treatments have both benefit and harm. Moderation demands that both be taken into account.

Because it is usually recognized that there can be no gain without some loss, the two different injunctions are actually implicit ways of evaluating

the negative effects higher than the positive or vice versa. "Do unto others" works as a positive precept because, in the daily affairs of mankind, doing a good thing for someone else has little risk of a negative outcome. Stepping aside for someone who is obviously in a hurry does not make one that much later, but can certainly help that person who is late for an appointment. On the other hand, in the medical sphere where harm can mean a loss of function or even the loss of life, "First, do no harm" explicitly mentions the benefits of good outcomes, which sometimes can be, at best, limited, versus the costs of bad outcomes, which can sometimes be disastrous.

In general, every moral action is a maximization problem of some kind. People are urged to live moderately. This advice applies to the single person alone, but it applies just as well to any action, no matter how many and what type of entities are involved. But just as often people are urged to search the highest quality. Yes, but at what cost? The advice to strive for the best will work only in cases where the unspoken simplifying assumption is close to true: when the costs of each alternative are the same.

This simplification is found also in epistemology - the philosophy of knowledge. Epistemology is a necessary condition for the types of moral action being discussed here, since one cannot act to maximize well-being unless, one knows what well-being is and how to improve on it. This requires knowledge and the ability to learn.

The simplifying assumption that the costs of each choice are similar underlies one of the basic concepts of epistemology - Ockham's razor. The razor as it was originally expressed was that "entities are not to be multiplied beyond necessity". This means that one should not make more assumptions than the minimum needed. This principle of parsimony urges us, when trying to model some process, to choose the simplest model from a set of equivalent models. This principle was the reason why the Copernican was superior to the Ptolemaic planetary system and to a third competing system by Tycho Brahe. This third system postulated that the planets circled the sun, but the sun circled a fixed earth. Even though the three systems are computationally equivalent, the Copernican system had fewer entities in the form of epicycles.

The problem with Ockham's Razor is that in real life the models are hardly ever equivalent in explanatory power. Given a list of observations to explain, a simple listing of all of the observations is the most accurate representation of what has happened. Unfortunately it has no predictive power, since it claims that there are no other cases than the ones we have seen. This prediction is probably wrong. At the other extreme - the model that says that all observations are possible - is easily the simplest model. It also

has no predictive power, since it claims that any other possible observation can happen.

Somewhere between the most detailed and the simplest model can be found a wide range of models. Some of these models make simplifying assumptions that enable certain details in the observations to be ignored, thus producing a simpler explanation than the complete listing. If the details that were ignored are extraneous, then the accuracy is the same as the complete listing and Ockham's Razor applies - the model where the extraneous details are left out is preferable. But it is more likely that the model, due to its simplifications, gives up some accuracy, so that the cases it claims is in the set it is trying to model is not quite the set of observations, but a list of items that resemble the observations with minor, perhaps random differences. Other simplifications can also be done on this model, resulting in a series of simpler models of decreasing accuracy.

This process can go in the other direction. Starting with the model that claims anything is acceptable, restrictions can be placed that removes whole classes of data from consideration as part of the set to be modeled. This process can even eliminate all of the observed data, replacing it with a set of archetypal values that are close. These predictive values may each be slightly different from the real observations, but the claimed values may have statistical properties that match the statistical properties of the observed data, something that may be more important than the ability to model each data value exactly.

This tradeoff can be purely considered in terms of the descriptive power of the models - the ability to summarize the data that has already been seen. But it is usually considered even more important that the model has predictive power. The two abilities are usually related, if the ability to descriptively summarize the observed data is due to the model's being able to capture some underlying principle that governs how the observations came to be in the past which works just as well in the future.

An example of this progression of models is the series of astronomical theories created by Kepler, Newton and Einstein. Kepler's model postulated the motion of the planets in elliptical orbits. Newton postulated a general force of gravity that predicted the motion of the planets just as well, but was able to predict the motions of comets, which were not explained by Kepler. Newtonian physics was also applicable to objects outside of the solar system, such as other galaxies. Einstein's General Theory of Relativity was able to explain the precession of the orbit of Mercury, which was a deviation from observation that Newton's theory could not explain. Einstein's model was also able to predict the existence of gravitational lensing, a phenomenon

that was only observed recently.

In the tradeoff between simplicity and accuracy, sometimes it is not necessary to consider the most accurate model. In calculating the orbit of man-made satellites, it is sufficient to perform the calculations using Newtonian gravity, but it is not necessary to take into account general relativistic effects.

The tradeoff between simplicity and accuracy implies a means to calculate the tradeoff. What is an acceptable accuracy? What is the cost to make a more precise calculation?

If the accuracy and costs can be translated into probabilistic terms, then techniques such as Bayesian models of learning can be used. Bayes theorem says that the probability of one event times the conditional probability of a second event (given that the first event occurred) is the same value as the probability of the second event times the conditional probability of the first event assuming that the second event has occurred: $P(A)*P(B|A)=P(B)*P(A|B)$.

For a made-up example, assume that a person has a 25% chance of having lung cancer if that person smokes, and the percentage of people who smoke is 15%. Assume that if the person does not smoke, then the chance of getting lung cancer is 2%. Then out of 2000 people, 34 of the 1700 nonsmokers and 75 of the 300 smokers will get lung cancer. This gives a probability of getting lung cancer at 109/2000. Given the 109 cases of lung cancer, the chances that one of these people is a smoker are 75/109. The probability of getting lung cancer times the probability of being a smoker if you have lung cancer is $P(LC)*P(S|LC)=(109/2000)*(75/109)=75/2000$. The probability of being a smoker times the probability of having lung cancer if you smoke is $P(S)*P(LC|S)=(300/2000)*(75/300)=75/2000$.

Bayesian theory is applied to learning theory by computing the chance of generating a given model out of the universe of possible models by a random hypothesis generator. This is the first event. The second event is the data itself: how likely is the data set, assuming the model is the right one? A measure of how good a model is at explaining the data is the probability of the model given the data: $P(M|D)$. Since we have the equation $P(D)*P(M|D)=P(M)*P(D|M)$, the value $P(M|D)$ equals $P(M)*P(D|M)/P(D)$. That is the probability of the model times the probability of the data given the model divided by the probability of generating the data set out of the universe of possible data set by a random data set generator. Given two models M and N, we compare the values $P(M)*P(D|M)/P(D)$ against $P(N)*P(D|N)/P(D)$. Since the denominator $P(D)$ is the same in both cases, it can be ignored. So model M is better

than model N if the value $P(M)*P(D—M)$ is more than $P(N)*P(D—N)$.

Jerome Feldman later generalized this technique by replacing Bayesian statistics by a calculation that involves two functions, the cost of the complexity of an hypothesis $C(M)$ and the error rate of the hypothesis $E(M—D)$. As the complexity of the models increases, $C(M)$ increases. Similarly, the error rate increases as the models get worse in their predictive ability. These two values are combined, but not necessarily by multiplication. Whatever combination function X is used, it must be increasing in each argument. That is, if A is greater than B , then for any C , AXC is greater than BXC and if D is greater than E then for any F , FXD is greater than FXE . Therefore, given a pair of models M and N , M is the best model if the value $C(M) X E(M—D)$ is less than $C(N) X E(N—D)$. This means that the best hypothesis to express a set of data is the one which has the smallest combination of the complexity and error rate.

This way of coming up with the best model for a set of data was also developed independently by Ayn Rand as the Epistemological Razor: "Concepts are not to be multiplied beyond necessity - the corollary of which is: nor are they to be integrated in disregard of necessity". In this case, the more the concepts are multiplied, the more complicated the model. Integration simplifies the model. Rand recognizes that Ockham's Razor is one-sided. If the model is kept simple, it is usually at the cost of accepting incompleteness in the explanation that lead to erroneous results. The correct balance is maintained as a tradeoff between the integration of the concepts and the necessity to model the data accurately.

These concepts have their application in morality in cost-benefit analysis. Every action that can be taken has a certain cost to the actor that reduces some part of their well-being. This cost must be taken into account when considering the benefit of the proposed action. Sometimes the most advantageous alternative comes at too high a price.

The benefits and benefits of any moral act can be more complicated to determine as the number of dimensions of well-being are increased. The study of different alternatives to poverty can consider actions that benefit the poor's financial state, but can also affect their educational level, their happiness and self-esteem, their health, the risks of failure and even how well they fit into society. But the costs of the alternatives may be more than the money that would be spent to address the problem. The complexity of the solution can delay the time it takes to implement it, the amount of expertise required for the implementers, the number of people involved and the changes that it imposes on society as a whole. Some techniques may not even be directly comparable - their comparison might only be made by

reference to a third approach, which could be a combination of the two. Even if the combined approach fares worse than the two pure approaches, it might reveal which method is comparatively better. So, in terms of the Epistemological Razor, considerations should not be ignored in disregard to necessity.

But when we make choices that are not the best possible in a particular case, there are usually one of two reasons for this: the first is not having good evaluation functions to guide one's choice, and the second is putting an inappropriate effort into finding and implementing the best choice. Both of these types of failures can either be due to a well-meaning but unfortunate selection of methodology to make the choice, or could be due to a premeditated choice to select a method that stints in its effort to choose the best result. Whether or not the action is morally questionable depends on the degree of volition involved and the effort made to arrive at a methodology that works and to apply that methodology.

Consider a person who decides to rob a bank. The simple thing to say is that they are making a bad moral decision, but it is more interesting to ask what they were thinking of. Typically, their thought process involves only a few of the entities involved in their decision - themselves, who win, and the bank, that loses. Their evaluation does not adequately judge the effort that society will make to try to catch them. Such a person usually also does not usually judge the chances of success very well, believing that the odds of getting caught are low if they succeed. This makes them more likely to try again: the typical criminal does not stop at one bank robbery. Also, they do not usually consider the consequences of becoming an outlaw. Even if not caught, there is a chance that they could be identified, making it hard for them to fit into society. They also have the problem of accounting for their new found wealth. A more sober assessment of the costs and benefits of their actions alone should be enough to show that crime does not pay. Leaving aside any desire to lead a righteous, upstanding life a rational self-interest would argue against bank robbery as a profession.

Now consider another person who is not considered generous. A lack of charity is not a crime - certain people are just not giving by nature. It is not possible to make a blanket claim that everyone gains through giving to others. Some people, such as those who grew up with never enough would feel a keen sense of loss if they were required to give of what was theirs to someone less fortunate. But with most people, the gains that result from a reasonable amount of charity are worth more to them than that amount of the money spent on themselves. It is important, though, to be aware of one's own sense of values. If it turns out that being viewed as a miser

were actually part of a person's sense of well-being, then a well-functioning ability to make moral decisions should be able to identify this as the choices are being made, not just as a later consequence as that person finds out how they are viewed by the community.

Although I am arguing that a moral injunction to be charitable is not a universal law that applies equally to all persons, it is important to note that, since humans are social creatures, the welfare of those around us is important to our finding ourselves in a society that we find comfortable. Therefore, although not reaching unanimity, the need to be generous is almost universal. This is another case of having an appropriate evaluation function. Although moral injunctions such as the need to be charitable could be stated with conditions such as "don't bother if it hurts you to do so and you get nothing out of it", this encompasses such a small number of exceptions to the general rule that the Epistemological Razor defaults to Ockham's original rule and these other qualifications are not added because they are not necessary - they add a complication that to most people is not helpful.

This analysis applies to societies also. During the Great Depression, the New Deal was instituted because society decided that the gain in overall welfare of the populace was worth the cost in extra taxes even though that might result in a slower recovery. Societies make these types of tradeoffs all the time. The decision to go to war is usually an analysis that includes other alternatives. The society believes that these alternatives just do not yield a benefit, or result in a net loss, whereas a war, it is believed, can be won at a reasonable cost. These debates are usually framed in very emotional terms, with different sides arguing for different conclusions. In these cases, the factions do not agree on a method of evaluating the costs and benefits. To be able to reach any understanding, it is important to lay bare these differences in values.

One of society's ongoing debates concerns the different approaches to helping the needy. Some people look at the children in need and urge that every family be given the resources it requires to raise each child in at least a minimum acceptable manner. Other groups look at this welfare and warn that it leads to a loss of independence and initiative in the families who receive these gifts. Both sides have a tendency to overgeneralize, making claims that seem to imply that the needy are all poor helpless victims, or they are all people who have lost the will go earn a living for themselves. The Epistemological Razor demands that to reach a meeting of mind on this issue, both sides must expand the number of identifiable groups to that which is necessary to account for the different cases that make up the

population of the poor and needy. There are certainly children in need whose parents, despite their best efforts, are not able to care for them, as well as those whose parents are not fulfilling their obligations. It is not helpful to try to argue the existence of either group - they both exist. Also, those who are mentally ill or incapacitated, and those who have fallen upon hard times. The successful evaluation of this problem must identify the different classes up to the level of necessity, in that each class that requires a different method to address their situation has to be determined before the solution to their problem is implemented. This means that no group or viewpoint is correct by itself. They are right for those they have correctly identified but wrong for all those who have not been.

This construction of appropriate ways of evaluating moral actions is the primary task of moral living. How many dimensions of well-being to consider? How many entities to consider - who must be included and who can be left out? What degree of aggregation?

The second kind of errors that we make are caused by our failure to put appropriate effort into finding the best choice, followed by inappropriate effort towards implementing our choice. Even if we had the best possible evaluation function to allow us to make discriminating decisions that accurately measure costs and benefits, we still make poor decisions. These are operational errors.

Some operational errors are due to the fact that even making the decision involves costs, in terms of the time and effort involved in the decision making. Sometimes this tradeoff is simply due to impatience, or a lack of ability on our part that makes decision-making hard. Sometimes it is willful decision to cut corners with the expectation that this will not adversely affect us. Sometimes it is simply that we cannot spare the resources of time and effort to apply to decision making, having to use our efforts in implementing a solution even though we know that this choice is far from the best.

Sometimes this lack of effort in looking for a good decision comes from temperament. Some people are just the type who are comfortable with what has worked for them before. This is not associated with a particular viewpoint. Both the stodgy conservative and the knee-jerk liberal are people who rely on the same approach that has worked for them in the past.

Examples of this type of behavior can be easily seen in the decisions that people make about their own welfare. These are not usually considered moral choices because they usually affect only the person who takes the action, but in the sense that moral actions affect well-being, and the well-being of that single person is at stake, these are moral choices also.

Examples abound. One is the tendency to quickly turn to vitamins

instead of more nutritious food. This is an error of seeking out the simplest answer to the question of eating right instead of putting effort into choosing a diet that is healthy. Another common situation is where individuals choose a field of study or even a career based on an incomplete or superficial analysis of what makes them happy or what they are good at. Sometimes a lifetime commitment is made on the basis of a subjective observation that they are good at certain subjects at school and good money can be made in a job using those skills.

The most common way in which people take too little time and effort in making good decisions is the tendency to follow the crowd with fads. This is the situation where each individual considers that the number of people making the same decision means that this decision is right for them without stopping to analyze deeply enough their own unique situation. To be sure, the decision to follow a current fad may be simply because this action is the right thing to do for that person, but until it became a fad the person was not aware of it. Sometimes following a fad may be appropriate, simply since doing so the individual will fit in, accommodating to society in a way that is important for that person to fit in, or conversely, it is just not worth the effort to be different. But quite often it is a lack of imagination, suppressing the person's individuality and choosing the common course. This trades the simple answer for what could be a more satisfactory answer.

The opposite of not making enough effort to reach a decision is making too much effort. Sometimes this is due to a psychological dysfunction in the individual: they tend towards anxiety, or are obsessive by nature. This causes them to delay and agonize over decisions that would come easily to a normal person. It can, though, come about not through an inner turmoil but instead a mistaken impression that if one puts out an extra effort, this will result in a choice that has more quality. Up to a point this is true. But there is a point of diminishing returns. The amount of effort should be comparable to the differential gain that the effort yields.

The problem of putting in the right amount of effort is that of moderation. To determine the appropriate amount, the first approximation is always to apply oneself to the same extent as others are doing. After all, people in similar circumstances should expend the same amount of time and effort making up their minds, all things being equal. After this starting point, the effort can be adjusted according to individual results. If a person is not making good choices or expending too much effort making them compared to the value of the outcome, then that person should try to adjust the amount of effort that they put into decision making. It is a curious paradox that even this second-guessing can be done with too little or too much effort

and requires a balance.

The other type of operational error is a failure to put enough effort into implementing a choice once it has been made. Even the best decision suffers from not putting enough into it. Obviously, trying halfheartedly will result in an outcome that is not as effective as the analysis led us to expect. But also many moral choices are being simultaneously made by many members of a community. Evil often comes not from making decisions that we know to be wrong, but from a more banal reason: we put less effort into our choices expecting the effort of others to make up the slack.

This second type of moral failing can be seen in many ways. For example, it is common to see dishes of extra pennies beside cash registers that are available for people to use for change for purchases that are a penny or two above a round number. These dishes usually work because the value of the pennies and the nuisance of having them are not great enough to result in people taking out more, on average, than they put in. On the other hand, unless the community involved is especially unselfish, a similar dish of quarters in a laundromat to be used for those caught short when putting in that final load to dry, will eventually be left empty. Although people consider themselves well-meaning, the cost of replenishing the dish when the person has a few quarters to spare is too high to offset the gain of taking one to finish the load.

This lack of effort can be dehumanizing in its effects. Although there was great evil in the decisions of the Nazi government in setting up and implementing the extermination of European Jewry, the lack of effort of each individual German in opposing these actions leveraged the effort of the evil few.

Sometimes the result can be more mundane. The rate of overweight and obesity in this country is growing. Losing weight is a combination of eating right and exercising more. It requires effort both to maintain vigilance in making the right eating decisions and to put a vigorous enough effort into a workout to produce results.

A social problem where insufficient effort is put into implementation of what is known to be adequate solutions is the care and treatment of the mentally ill. In the last few generations, the discovery of effective medication has emptied the mental hospitals of all but the most difficult or violent cases. But too often these people have been turned out into the streets to fend for themselves. The medications they have been given go only part way, but society has lacked the will to provide the resources we require to address the problem adequately.

Is it possible to have a case where too much effort has been put into

a solution? Typically this happens when a poor decision has been made but due to pride or willfulness the decision is pressed forward well beyond the time when it would be more appropriate to step back and reassess. It has been experimentally determined in psychology that this ability to appropriately give up on a failing decision is a trait in human beings. We tend to value the effort that we have put into a poor choice in and of itself. This makes it harder to admit defeat and go back and think about a better way of tackling the problem.

One pervasive failing that is both a combination of insufficient contemplation of the better solution and an insufficient effort to implement that solution is in poor parenting. A common problem is the case where the parent wants a certain type of behavior changed - for a child to do something they ought to or to stop doing something they should not be doing. Instead of finding an effective way to make the change happen, the parent simply orders the child to change. This is accompanied with threats of punishment if the command is not obeyed. There are very few cases where giving a command under threat of punishment is the optimum way to effect change. It is certainly the easiest on the parent - it requires no thought to come up with. It is much harder to think about why the behavior is happening in the first place and to decide how to correct the situation in a more effective way. Compared to an analytical approach to parenting, giving commands the easiest action to take.

The analysis of the situation sometimes shows that the problem comes directly from the parent's behavior in the first place. Children are influenced by the behavior of the people around them. Ordering a child not to eat unhealthfully is ineffective if the adults do not do this themselves. There is a special paradox in ordering a child not to use force or threats against other children. The very act of commanding this behavior belies the command.

The ability to arrive at a good approach to parenting requires more observational skills than many parents have or take the time to use. A child cannot be told to eat everything on their plate if the amount they are given is more than is appropriate for them. Sometimes standards of behavior are established that are beyond the ability of the child to meet. It is essential to know what a child is capable of before making expectations.

The moral failings that have been discussed here can be well-meaning but ultimately incorrect choices or they can be result of deliberate evil. What is the difference? Evil is usually considered to have a volitional component - a deliberate choice to do what one knows is wrong. Much of the evil in the world comes from acts of dehumanization as we covered in the previous chapter - the deliberate turning away from the Golden Rule. What

deliberate acts of immoderation can be considered as acts of evil?

One of the most prominent ways that evil comes into being is from acts of destruction instead of construction. To tear down is always easier than building up. Both actions are great equalizers. If one person is better off than another, this may have come about by luck, a better starting point, or just a greater capacity or ability. Given a better start in life or the benefits of good luck later on, the advantage may be established and remain throughout the different people's lifetime. Given a difference in abilities or looks, the difference may even widen through no fault of the person left behind. Since the chance of reaching equality is unlikely, the evil person turns to destruction to force those ahead back down.

Sometimes this inequality does not result in acts of destruction, but in giving up. Defeatism is not bad, but it is the absence of good. Sometimes just giving up is the right choice to make in certain cases, when one has done all one could. But this must be selective - it cannot apply to all of life, if there is enough energy for more. But every life comes to an end, and that time, it is appropriate to let go. But for someone in the prime of life actively embracing surrender cannot be considered virtuous.

The other way that evil comes is when either the benefits are unequally distributed or the costs. When the benefits are more than a person's fair share, such as in acts of thievery, swindling or embezzlement, the nature of the evil is usually obvious. It is often less obvious when the act of evil is an unequal distribution of the costs. One can sometimes tell whether a person did not pull their load due to laziness or incapacity, at least if the effort is made in the open with that person's previous capability to compare to. This is a lack of effort in implementation. What is especially hard to determine is when there has been a lack of effort in deciding on an appropriate course of action. This can be due to a deliberate mental laziness or an unwillingness to do the work to collect the facts to make a good decision. But since this mostly occurs within a person's head, there is little evidence to base a judgment on.

The difference between good and bad is between symbiosis and parasitism. In many cases an act of evil results both in unequal benefits and unequal effort. Since the basic equation of moderation is the tradeoff between costs and benefits, these types of inequalities are in effect the same because the difference is the same.

Most religious traditions classify the sins and failings that human nature is capable of. I will use the Buddhist three poisons to illustrate how these sins are either a consequence of dehumanization or a lack of moderation. This shows that although the classification here and the Buddhist classification

are different, they seem to cover the same failings

The first poison is termed aversion. It is basically dehumanization in its causes. Aversion leads to failings such as hatred, fear, hostility, envy and jealousy. It is caused by people's failure to follow the golden rule and to treat others the same way that they expect to be treated.

The second poison is that of ignorance or delusion. This leads to failings such as sloth, foolishness, doubt, confusion and boredom. This can be easily seen to be caused a lack of effort into making and carrying decisions out. The effort to dispel ignorance is amply rewarded in increased effectiveness and improved results.

The last poison is desire and ill will. The resultant evils are lust, anger, greed, pride and arrogance. Initially caused by dehumanization, it is the easy choice of destruction that leads to these failings. Sometimes instead of destruction, it comes from the unwillingness to put in the effort required to get the desired result.

Despite our best efforts, problems constantly arise in making moral choices. The question is whether there is some way to improve our ability to make moral choices. A more detailed analysis of this question will have to wait for a later chapter, when it will be discussed as part of reasoning. But we can make some preliminary remarks.

First there is the need to avoid complacency. Although we try hard to make good decisions, the fact that we ourselves and the world around us are constantly changing means that there never will be a fixed, unchanging moral code, or even a set of methods to come up with the best moral code for a particular time.

Second, moral decisions are made through a combination of careful observation, logical reasoning and gut feeling. Similarly, to improve the decisions one makes, we need reliable data about our past efforts and the willingness to use both reasoning and intuition to make things better.

Third, we learn from our mistakes. The fact that we are correct time and again strengthens our confidence in our methods, but only our mistakes give us the knowledge we need to fix the imperfections we inevitably have. Sometimes too much success makes us overconfident, and that makes it harder to change when push comes to shove.

Finally, perfection is impossible. Strive as we might to reach a golden mean, doing too much or doing too little is part of the imperfectness of life. We must also recognize that we have to err on both sides to be in balance. A person who chronically errs on one side or the other is guaranteed to be doing worse than the one who balances their mistakes. In learning to do better, we will have this problem also. It is unavoidable that in correcting

our mistakes we will sometimes incorporate into our new methods some special circumstances from our past that are not applicable anymore, or conversely, we will leave some detail out that would make our decisions even better. It is best that we don't make too big a fuss about it and recognize that the ideal is a figment of our imagination.

This chapter has been about the virtue of moderation. It is a basic fact of life that all things suffer from too much or too little. This is even true when trying to do good. Attempting to achieve a state of perfection in any one aspect of our lives means that other aspects will suffer. In our moral choices, we must strive to follow the Golden Rule but at the same time avoiding doing harm to others through our efforts.

In moral decision making, our analysis should not be directed towards simplicity in our choices, but instead by the use of the Epistemological Razor - a balance between simplicity and necessity. This can be done by looking at the cost our actions will demand versus the benefits that accrue.

In making these types of tradeoffs we will make mistakes both in making the decisions and implementing them. These mistakes can either be innocent errors or they can be the result of deliberate destructiveness or pushing the costs onto others and taking more than our share. Most sins and failings come down to these kinds of errors, if they are not acts of dehumanization.

Although it is possible to learn to make better decisions, learning is as prone to error as the decisions themselves, since moderation is a balancing act, not a state of grace.

Chapter 6

Justice

Justice arises naturally out of the human condition, but it is a very weak form of justice.

In a number of the world's great religions, one finds the concept of heaven and hell. The basic idea is that there is a place that you go after you die where your good deeds are rewarded if you are good and your bad deeds are punished if you are bad. Of course there are differences between faiths. Certain branches of Christianity believe that going to heaven depends solely on the acknowledgement of Christ as Savior - this is the only determinant of whether a person will escape punishment in Hell or not. The Restorationist branch of Universalist Christianity believes that those who are sent to Hell for punishment are eventually restored to Heaven when the punishment is over. Some religions, such as the Baha'i, believe that Heaven and Hell are not separate places, but some sort of continuum that delineates the distance of the soul from God. Even the concept of reincarnation is sometimes expressed with a component of retribution for sins and reward for good: the form of the next incarnation is dependent upon one's behavior in the past.

All of these speculations meet a basic human need that there is justice in the world - if not in this life, then in the next. Recent experiments have shown that people are willing to give up something if it helps ensure that transgressors are punished. But too often it seems that bad people can get away with murder, figuratively and even literally, while no good deed goes unpunished. The question arises if there really is justice in the world at all. It is possible to show that, depending on the nature of how people respond to the actions of others, justice arises naturally for any social creatures, but it is very imperfect - probabilistic in character and very attenuated, so that not only do the sometimes innocent suffer for the sins of the wicked, but

even their victims suffer.

Because the natural forces of justice are so imperfect, humans impose their own justice upon transgressors. This has met with different levels of success. For example, sometimes a focus on justice in individual cases can lead to an unjust society.

Justice was described in Chapter 3 as a force that returns equity to moral acts. Justice is the attempt to restore a balance in well-being that had been disturbed. In most cases of justice, this imbalance can be seen to be simply due to the fact that one entity's well-being is improved at the expense of another. In cases of social justice though, the imbalance is due to different groups progressing at different rates.

The metaphors that help to describe how justice comes about naturally are the concepts of conservation and feedback. Conservation of some property is observed when that property remains the same in some system even though there have been other changes. In physics there are a number of conservation laws for different properties such as momentum, charge, spin and so on.

Conservation laws are defined and confirmed by observation and experiment. Some hypothesized laws that seemed to make sense are shown not to hold, such as the conservation of parity in nuclear physics which was shown to be false when nuclear reactions were recorded where the number of cases of one parity did not match that of the other.

In the physics of motion, inertia is the result of conservation of momentum. Inertia is the tendency of a body in motion to remain in motion unless acted on by an outside force. This is the first of Newton's laws of motion. This conservation law was not known to the early Greek philosophers such as Aristotle, who believed instead that objects composed of certain elements would tend to seek their level. The level of each of the elements was the following: earth was the lowest, followed by water. Above that was air and the highest was fire. This concept meant that motion of an object that was taken it out of its element would naturally change and eventually reverse itself until the object would get back to the level which is appropriate for it. For example, a rock, thrown up in the air would return back to the level of other rocks, composed of earth. A bag full of air, opened below the surface of the water, would bubble up until it reached the level of the atmosphere. This erroneous belief is in itself a conservation law, where the elemental level of an object is conserved.

Aristotle's law of motion tried to equate two different properties: motion and density. The four elements that the Greeks recognized were actually four states of matter. Earth is the solid phase of matter, water is the liquid phase,

air is the gaseous phase, and fire, having enough energy to remove some of the electrons from the electron shells, is a plasma phase. Matter at normal temperatures tends to have different densities depending on the phase it is in. Solids such as iron or quartz tend to be denser than liquids such as water, which has more density than the nitrogen in the air. Fire, being at a higher temperature, is even less dense than air. There are exceptions, though. It is well known that the solid form of water is less dense than the liquid, so ice cubes float.

What this view of physics missed was that the density of matter is changeable - it is a property that is not conserved. Instead it was believed that matter was composed of these four elements; as if the densities were mixed together in certain proportion but they did not change.

A conservation law that addressed the composition of matter that finally supplanted the Greek view was encapsulated in the periodic table - a theory of different elements, each element being an atom with a different number of protons that give the elements their mass, and whose electrons are organized into electron shells which give them their chemical properties. This conservation law applies to chemical reactions: that is, situations where the nucleus of each atom is unaffected. The conservation law states that if you count the number of atoms of each element, where an element is an atom with a certain number of protons in its nucleus, the count of the number of atoms of each element are the same. Besides this basic conservation law there is a second conservation law concerning chemical reactions. This law states that the overall count of the number of electrons in all of the atoms, before and after is the same. These laws are the law of conservation of electron charge.

A basic part of physics is to define the properties of matter and energy. We can see in this example that physics has sometimes begun with a concept of some property that is an inherent, fixed, and immutable property of matter, where different quantities of this property lead to the differences we observe around us. A deeper understanding, though, leads us to postulate some other underlying property that can be distributed through a system, but is the same after a redistribution.

Besides conservation we also need to consider feedback. Feedback is when the output of a process turns around and influences the inputs of the process. When a positive change in output is fed back to become an even more positive change then there is a situation of positive feedback. We also have positive feedback when a negative change becomes a cumulatively more negative change, because the negative quantity has become greater. If, on the other hand, a negative change in the outputs is fed back to become a

positive change and a positive change becomes a negative change, then we have negative feedback.

Although in human interaction, the act of positive feedback is considered good, this it cannot be maintained for long outside of restricted circumstances. Positive feedback is usually thought of as an input to a person praising what they have done, with the goal of having them do even better. This cannot be taken too far. If the person, in their attempt to do even better can get to a point where they work too hard to attain the next level, this is ultimately detrimental. So positive feedback is effective only within certain limits.

Two situations in which positive feedback result in uncontrolled behavior are population explosions and monetary inflation. The negative effects of positive feedback can be seen in the genesis of wars. Sometime a negative action is responded to with an even more negative act. This leads to more and more negative acts which finally explode in war. At its most benign, positive feedback of this type will not increase the degree of the response, but instead result in a reduction in interaction between the parties, which will feed back to cause a resultant further reduction in interaction, which, instead of leading to an explosion, will eventually result in no interaction at all.

Negative feedback is preferable if a stable situation is desired. This type of feedback is built into many automatic systems, both mechanical and electronic. It can be found in cruise controls and autopilots, bathroom toilets, ovens and air conditioners. Negative feedback can be applied to make a system stay at a preset limit, the way an air conditioner is set by setting a thermometer.

There can be systems with negative feedback that have no absolute fixed setpoint. These types of situations are found in nature between predators and prey. An increase in the number of deer, for instance, can lead to an increase in the number of wolves who eat them, which can lead to a decrease in the deer population if there are more deer eaten by wolves than are being born. This will lead to a reduction in the deer population which leads to some of the wolves starving. Although it would seem that there can be a certain fixed level that would apply to a particular pair of predators and prey, the negative feedback can be such that there can be a range of possible stable situations, not just a single level.

Conservation laws and feedback both function on the degree of well-being in the world. The conservation laws in human behavior are not as absolute as they are in physics, though - conservation means the properties are approximately the same before and after a moral act occurs. But, aver-

aging over many such acts, given a certain level of well-being, the resultant level of well-being will be similar to what it was before. This level cannot be specified exactly because humans are more complicated than physical objects and are therefore affected by changes in the environment that would cause changes in well-being even if no moral act takes place. As to moral acts, there may be general principles that are in common between similar acts, but when it comes to the details each act is as unique as the players involved.

This variability makes it hard to determine if there is a conservation law at all. Human nature being what it is, we might postulate certain innate properties of humanity, which, although they are manifested slightly differently in different people, are essentially the same within certain bounds. On the other hand, what might appear to be an innate property could be mutable, but bound by a conservation law that make predictions about the limits of change. Of course, the third possibility is that there is no such constancy at all in human nature - we have been observing people in only a limited number of circumstances due to the limits of this earth and our human history, and that this supposed innate property could manifest itself in people in different places in almost infinite ways.

To address the question of a conservation law we first note that historically, human nature has been variously viewed as essentially virtuous or evil depending on the philosophical tradition of the viewer. This is a claim that humans naturally function so that they maintain a certain level of well-being in society. Since humans tend to act consciously when they consider actions that are good and bad, a virtuous person is one who tends to make good decisions. An evil person is judged to be that way because their decisions tend to be bad - certainly for others and sometimes for themselves. This does equate good and bad with virtue and evil as if there is some sort of perfect volition at work. Whether or not this is actually the case will not be discussed here. Instead we will frame things in terms of good and bad, leaving out whether the person made moral or immoral decisions.

One of these viewpoints is the Christian concept of original sin. This is the belief that humans are essentially bad in comparison to what is attainable from the viewpoint of the Almighty. This is usually described in terms of the human capacity to do things that result in a lowering of well-being of others. This leads to a condition where society functions in a state of malaise.

A different viewpoint is that which sees the world as a good place. That is, despite the bad actions of people at various times, the world itself is essentially a place of well-being where although the capacity of pleasure and

happiness can fail to be realized in certain times and places, the world will in time right itself.

Both of these viewpoints look at the world as having these elemental properties of goodness and badness. That is, goodness and badness are innate properties with certain fixed quantities the same way we see the world as made up of atoms whose elemental properties are fixed. Besides the obvious difference that one viewpoint sees mankind as bad and the other sees the world as good, there is an interesting difference in how these properties are stated. Original sin is given as a capacity of people to function in a certain manner - a property of action. The other viewpoint is given as a capacity of people to exist in a certain state or condition of well-being.

Original sin states that people have the capacity to do bad, but this statement is made in terms of people's falling short of the glory of God. That is to say, people are certainly capable of doing good at certain times, and in fact do. The problem is that people do bad things, and in the Christian viewpoint, people tend to do more bad things than they do good. But the question is whether the two balance out or not. Even if less good acts are done and more bad acts take place, each good act may have a greater positive effect than a number of bad actions that each has a smaller effect. In that case, even though there are more bad actions, the cumulative effect is that the world is a better and better place as time goes on and both viewpoints are simultaneously correct.

Those Christian who take the darkest view of original sin would disagree with this estimation. They tend to believe that the number and degree of the bad actions tend to outweigh the good. This leads them to consider the world as having fallen away from its initial state of goodness. At worst, the world can even be seen to be worse and worse the more we go along.

This belief in an innate capacity of human action has both a nature and a nurture component. The nature component is genetically determined - it is the degree of goodness a person is naturally capable of. The nurture component is the capacity for doing good or bad that society imposes upon the person. This may be different for different societies.

The same question that applies to individuals - whether they are innately bad or good - can also be applied to human societies. The innate capacity of the society is institutionally determined, creating a certain climate of morality that is handed down through the generations, preserving a certain basic capacity to societies even as they change. There is a capacity for change, just as individuals can be better or worse than the innate capacity for virtue or vice that humanity has as a basic characteristic. But a concept of an innate capacity for morality in a society would lead one to claim that

society can only be perfected so far, or that society can degenerate only to a certain level below which it would cease to function as a society.

If the capacity for good or bad in the world we live in is considered as an innate property of our environment, it does not mean that good things are constantly coming our way to such an extent that they outweighs the bad. If that were so, our lives would eventually be raised to such a height of ecstasy that we could not contain our happiness. Instead we have a question of balance. What might come to us might bring us sorrow or pain, but for most of us, most of the time, what comes to us sustains us sufficiently that life goes on. Because life goes on most of the time and that requires a preponderance of good, it is reasonable to conclude that the world is basically good.

These two views of the innate qualities of mankind and the world may be true simultaneously. If either viewpoint is wrong, it may be wrong in either or both of two ways. The first way is that the claim of that particular innate quantity is wrong - no such quantity can be found that is consistent over all of human nature. If it is not the case that people are intrinsically bad in their actions, it may appear that this is the case due to the world being an intrinsically bad place. This would mean one innate property is true and the other false. Or they could both be wrong and that no property this type holds at all - humanity is neither good nor bad, and neither is the world.

The second type of error is that there is an innate quantity, but that the quantity being claimed is in error. That is mankind could be essentially good or the world is essentially bad. Whether this is an error or not is really a question of perspective. The Christian viewpoint of mankind's essentially bad is obviously true if the standard of comparison is a god capable of being perfectly good. But this viewpoint does not actually say very much about human nature. It may make a statement about humankind's eventual salvation, but it does not really describe how well people function in the world. Yes, we fall short of a standard of perfection, but are we good enough? Again, this is a matter of perspective. We can be considered to be good enough because humanity has thrived enough to populate the world. Or one can establish an attainable ideal of human functioning that is represented by exemplars of human behavior whose actions lead to the best level of human well-being that is realistically attainable and ask how close humanity comes to this level. This may lead to a conclusion that says humanity is falling short of this ideal, but is not inherently bad - just not as good as we could be. This perspective has practical results, because a measurement of this type can point the way to improving the human

condition.

If no such property can be said to hold, then humanity is neither innately good or innately bad and so for the world. This would mean that the degree of good or bad in our actions has an arbitrary setpoint that depends on the situation that we find ourselves in. If this were so, then the typical human, if born and raised in a utopia, would function all their lives as a paragon of virtue, and if they were unfortunate enough to live in some dysfunctional society, they would rely on their baser nature to get by. If this is the case instead of having some innate level of functioning, then there are some observational predictions that can be made that would be able to distinguish between the two cases. If there were an innate level of good or bad in humanity, then a person, placed into a utopia would, on the average revert back to a normal state of imperfection, and if placed in some negative situation, would tend to rise above that state, leading to an eventual correction. These arguments apply equally to the existence of an innate property in the world, where eventually every Garden of Eden reverts back to a more normal state, and some inhospitable desert eventually gets transformed into a place where life can be sustained.

Although no definitive answer is given here as to the innate quantities of goodness and badness, we can still show that there is a conservation law at work in human behavior. The law simply states that humans have a tendency to respond to good or bad actions with a similar level of good or bad actions themselves. This conservation law has been amply demonstrated by experimentation: for example, people who were abused as children tend to grow up to be child abusers themselves. Even without experimental verification though, this law can be derived from first principles.

If there is an innate capacity for good or bad, then people would tend to act in a way consistent with that innate capacity. Different people would tend to have different innate capacities, though. If there were no conservation law, then there would be no difference in a person's response between being treated well or ill. Within the variability of human action, a person would tend to respond at about the same level, consistent with their capacity.

But this implies a mechanistic functioning of human beings that is inconsistent with human nature. This type of functioning would be seen in organisms that are driven by their instincts. This extreme of unresponsiveness does not even occur with other mammals. Once there is a capacity for learning, there is a response to actions that incorporates the levels of goodness and badness of the actions that the organism has been subjected to along with other details of the action.

Even if an innate tendency towards a certain level of good or bad were functioning, this would not totally override, but only influence the degree of goodness or badness in the response. This is a consequence of the conditions we discussed in the chapter on the moral absolute. That is, there is an absolute standard of well-being, but it is determined by finding the maximum possible well-being attainable. This must be so, because as time goes on, human ability has worked to make it possible to attain well-being in more and more possible ways. Therefore, if the innate capacity of humanity limited the response of humans to respond to novel ways of moral functioning, we would be limited in the capacity of well-being we would be able to learn.

For example, issues of well-being around money and its exchange are not innate in humanity. Wealth as a measure of well-being came along later in human civilization, and has been constantly changing as we have gone along. If we are limited in our response to novel situations, we would not be able to develop concepts such as charity around money. Instead we have developed such concepts once money entered into human society. Actions such as giving alms or embezzling money entered the human repertoire and these behaviors were passed from person to person as they were experienced. This does not mean that a victim of a holdup turns into a thief themselves. But it does mean that if the degree of thievery increases in the world, there is a basic conservation of action that left to itself, preserves this activity at a certain level.

The existence of this approximate conservation law is further shown to exist due to the fact that there is no absolute cutoff between a sufficient and an insufficient level of well-being. That is, as human history has gone on, the level of what is acceptable has changed. In the distant past, a subsistence level of food, a certain level of health, the comfort of ones housing and clothes may have defined a good life, but this level of comfort would not be acceptable in advanced societies of today. This lack of an absolute cutoff between well and ill implies that the level of acceptable functioning we have in modern society is different from that of the past and also different from less developed societies, and therefore we respond to the actions of those around us with a level of good or bad that maintains the level of well-being in the society we find ourselves.

This conservation of action leads to a feedback situation. Our outputs are influenced by our inputs, and this in turn influences the outputs of the people around us. There does not appear to be an innate character of positive or negative feedback in human nature, though. Different types of situations lead to different forms of feedback. Some feedback conditions lead to stable situations, others lead to cases where things get out of hand, until

all society breaks down. But although this breakdown can happen, it is not due to random action - it can usually be shown to be a continuous change in the situation caused by the feedback in the system.

This feedback system naturally leads to a certain level of justice in society. If someone takes an action that is worse than the typical actions of other people in the local environment, the conservation law predicts a tendency for the recipients of that action to respond in kind. This will affect others in the immediate environment, and eventually lead to a change in level of well-being in the society at large. Eventually this will come back to the initiator of the action, unless that person has left the environment entirely.

This level of justice is certainly imperfect. This means that the victims of an injustice suffer most, and that the perpetrator only indirectly. To achieve a higher level of justice requires human intervention.

If it is the case that humanity has a certain innate level of good or bad then this conservation law functions within the bounds set by this innate functioning. That is, if a person is treated well, they will respond in kind up to the maximum level of well-being that they are capable of. Similarly, ill treatment may be returned with ill-treatment, but not below a level that defines a basic humanity. Inside these bounds, there may be at different times a different setpoint established in accordance with the conservation of good or bad, but outside that range, any reaction will act only to preserve functioning within the operational range. What that means is that in a good society, the person who is even better will not be rewarded for this behavior because the other people are not capable of it. In a bad society, a person who is truly debased will go unanswered.

One important consequence of this conservation law is that one a society changes for the better or worse in terms of its level of functioning, the tendency is for that society to naturally stay at that level instead of reverting back to the previous level. A similar behavior is seen in probability theory in the arctangent law.

Consider a situation where someone is flipping a coin. There is a 50-50 probability that the coin will come up heads as often as it comes up tails, so that if you count the heads versus tails, they will come out about equal. But what happens if they don't? For example, what is the expectation if after 100 flips, there are 55 heads and 45 tails? Although intuition might say that after 100 more flips, there is just as much chance of coming out with more tails than heads than the other way around, probability theory shows that once there are more heads, the tendency will be that it will remain true that there will be more heads than tails. The arctangent law says "quit while

you're behind": playing longer will not erase the setbacks of the past.

This type of behavior comes about when a something is learned without reference to behavior in the distant past. This can happen in periods where there is a rebellion or revolution from the past, or when the social situation appears novel enough that the old ways do not apply. In that case a different way of doing things will be taken up and copied by people who have come in contact with this new behavior or have heard about it.

When a new behavior is introduced where this novelty is unwelcome, people will have a tendency to revert back to the past. This is a negative feedback situation, where some change will be met with a reaction that has less of that change and some of the previous behavior. The effect of this is to reduce the level of change until society gets back to functioning the way it did before.

It often happens that a change has both good effects and bad effects. An example of this is the tradeoff between individuality and conformity in society. Individuality has the positive effect of personal self-expression, but has the negative effect that others have to make more effort to adapt to all these differences. Conformity allows us to know what to expect of ourselves and others, but at the cost of having to suppress parts of ourselves. When there is a situation in which changes can lead to two different states of well-being, approximately equal in their degree of personal satisfaction but different in their composition, society tends to veer between the two states in a kind of pendulum effect.

This crude form of justice described here plays itself out in ways that can be surprisingly long-lasting. The biblical phrase "the sins of the fathers shall be visited upon the sons, even to the third and fourth generation" is an example of this effect at work. Some of the social strife that occurred in England during the last century is this principle in action. The confiscatory levels of income tax on upper income levels was in part due to children of the labor class acting against the children of the ruling class in a way that embodied the resentment left behind from generations of injustice by the upper class in years past.

This effect of spreading the consequences of injustice to others in the local environment can affect the victims unduly, especially when the perpetrators are not to be found. It has been said that some of the problems in Russia in the last few generations is due to the Russian character never having fully recovered from the negative effects of the Mongol invasion of the 1200's.

This effect has an important consequence in our actions. If it happens that we are the recipients of an action that was bad for us, one of the first questions we must ask ourselves is what we have done that might have

sparked such an action. If we ourselves were not the initiator of an action that may have resulted in this type of response, then we need to ask who in our immediate environment might have caused it. Our search needs to go back in time as far as required to match the duration of the behavior in the others. This needs to be looked at very carefully and ruled out before we can claim to be a victim of injustice.

The justice being described here is a form of retribution. Instead of trying to right wrongs, it just returns to an evildoer the same treatment that had been afforded to others. This form of "just desserts" also returns to those who do good a measure of the benefits they gave to others. Retribution alone is not considered sufficient for an adequate sense of justice. Justice requires a redress of the imbalance: a determination of inequities in human behavior and a restoration of equity when they happen. Justice is not served when the victims suffer more than the perpetrators. What is necessary is that justice restore to the victims what has been lost and ensure that the perpetrator not profit from an unjust action.

To attain a just society, it is best to function in accordance with the natural feedback of good and bad in human society. A perfect justice is neither attainable nor desirable, since feedback systems with an immediate feedback lead to instabilities. What is preferable is that there is a certain amount of lag in the response. This is known as hysteresis, and leads to a functioning that is overall much smoother and efficient.

Besides a time lag, there should also be a lessening of the response in relation to the degree of the injustice. Due to the complexity of human nature, absolute equity is impossible. It is impossible to restore things to the level that existed before. Because of this, so it is better to round down than up. That is, when taking an action to redress an imbalance, do not try to make things whole again, but err on the side of doing a little less. Otherwise, even though some individual cases may even out exactly, there is a possibility that if the differences are too often rounded up, the impression will be created that a further injustice was created in this overcompensation. This would lead to a situation of positive feedback, which often leads to disaster.

The advisability of a lag in the application of justice implies that it is possible for justice to be too swift. Every action has its consequences, including actions of justice. The consequences affect the future behavior of the wrongdoer who has been subjected to this redress and the victim. But the action also has consequences to the society around them, in a degree relative to the closeness of those to the parties in the action. Without having time to consider these effects, the beneficial effects of justice may be

worse than expected.

In the spirit of moderation, it is possible to take too long, though. The point was brought up in Chapter 3 that as time goes on, the incorporation of past injustice into the prevailing functioning of society dilutes the effect of the injustice with subsequent events and makes it impossible to restore things to that previous level. The balance must be struck between justice too swift and justice too late.

Applying justice with dampening effects leads to mercy. Mercy is sometimes thought of as valuable because showing good to a wrongdoer leads to that wrongdoer doing good in the future. This is using the conservation effect that is the basis of this chapter to bring about a permanent change. A second reason for mercy is that mercy has a dampening effect on feedback, leading to an effective and maintainable system.

Justice in human society is often associated with the application of laws that carry punishments such as fines or jail. The penalties are imposed to change the balance between costs and benefits that exist for a particular person or group of people to better reflect the costs and benefits of this action as seen from the perspective of society as a whole. The intention in many cases is to prevent the individual from taking without giving back. This is the same thing as adding a governor or a thermostat as a feedback mechanism in a physical system.

Feedback enters into the type of punishment handed out to criminals. The tradition of sentencing has been driven by two goals: to punish an offender in proportion to the degree of the crime and also to provide a correction. Correction is the attempt to make it possible for the criminals to improve their lives so they will not be led back into crime. There is also a third factor that comes into play - a prisoner in jail simply has no chance to commit a crime against a member of the general public. These different effects come out of the application of justice at two different levels - justice for the individuals involved and justice for society as a whole.

The goal of the individuals is to redress the imbalance that the act creates. The goal of society is to ensure that the society runs smoothly. A first-time offender can have a certain tendency towards wrong-doing which can be made better or worse by the experience of jail. If a correction system works effectively, the tendency is lessened. It has historically been true that in a system of punishment, the tendency toward recidivism is increased. Historically, a correctional system has shorter sentences. The two effects tend to function in different directions. A system of corrections puts ex-convicts out on the street sooner, but with less tendency to commit crimes. A system of punishment creates people with higher recidivism rates, but

having not had the opportunity to commit more crimes. As the criminal is older, both of these factors are attenuated. The older criminal is less like to go back to crime, and the sentence of an older criminal is a larger fraction of that person's life span.

This chapter applied the concepts of conservation and feedback to morality to understand the mechanism of justice. We determined that there is a natural conservation of good and bad in the way people interact with each other, even though there may be an innate level of functioning or an innate level in the world at large. This leads to a feedback system that establishes a crude type of retribution in the world. For humanity to improve on this, it is advisable to establish a system of justice whose feedback mechanisms are tempered so that the whole system remains stable.

Chapter 7

Free Will

Moral choices require volition which is an emergent property of thought, but the determination of whether an entity has free will is a relative characterization, not an absolute one.

Every action has a measure of goodness to it if the welfare of some entity can be affected by it. But morality is not meaningful to a rock falling downhill. If falling downhill causes the rock to shatter into small pieces, then the action of rolling downhill is not in the well-being of the rock qua rock, since it now has ceased to exist. In terms of the constituent parts of the rock, the action is neutral.

Morality comes into being when there is conscious choice involved. A rock falling downhill is not making a moral choice; it is acting in its own nature.

Living creatures - plants and animals - can make choices. That is, there is some internal representation of the entity's well-being that is referred to in the determination of the action to take in response to the current environment.

A choice was initially defined an action that can be good or bad, depending on the well-being of any entity affected by that choice. This definition needs to be refined, since the action of rolling downhill in a way that causes a rock to shatter is a choice by this definition. To call this a choice by this rock stretches the definition of choice beyond what is reasonable.

The definition will be extended by requiring that the entity performing the action must have access to some representation of the well-being of some entity (not necessarily itself) affected by the action. This representation is used in decisions that serve as a precondition for the action. This means that there is some measurable condition that derives to some characterization of

well-being for an entity, such that a change in this condition would result in a different action. The entity is said to react to this condition. In a simple creature such as a sunflower, the action of turning towards the sun is a result of the determination of where the best sunlight is, a condition that is a measure of the plant's well-being. In that sense, the sunflower chooses to turn towards the sun, or it reacts to the sun moving across the sky.

Note that the word choice is quite often used to mean a decision made by conscious deliberation. This is not the sense that is being used here. If a conscious choice is being considered, terms such as judgment or volition will be used to describe it. Choice by itself refers to actions that are simpler or more basic - all they have to do is affect well-being with a determination of that well-being part of the preconditions.

For simple plants and animals, these conditions exist as measurable quantities and configurations of the animal, but the representation of well-being and the resultant choices that are made are qualitatively different from the decisions of more complex animals. Many animals have a nervous system, part of which is reserved for the representation of the well-being of that organism. Pleasure and pain are terms used to identify some of these neural representations.

Animals choose to react to pleasure and pain, but only the most advanced creatures use judgment. Judgment requires that the organism not only identify the pleasure or pain that the organism is experiencing, it also requires the organism identify to some extent the alternative actions that are available to it. The sunflower does not use any judgment to choose to turn towards the sun - the action is automatic. Similarly it is unlikely that a fish uses judgment regarding whether to school or not.

When we get to the mammals, the situation is not so clear-cut. It is obvious that almost all normal humans use judgment. But the question of whether a dog makes a judgment when choosing the time to have its dinner is a question that would take a psychologist to answer. A conditioned response such as salivation takes no judgment on the part of the animal. But there are other choices that an animal can make that could arguably be said to include judgment - that is the animal's decision involves not only a determination of well-being, but also a determination of the expected change in well-being an action would result in.

An example of this is learned helplessness. A dog placed in a cage with a barrier in the middle where either of the two floors can be electrified, will soon learn to jump to the other side to avoid a shock. By itself, this could be considered to be simply a conditioned response that requires no judgment. But if both sides are electrified, then this action becomes useless

and the animal stops jumping. This results in the animal's exhibiting signs of distress. Again, if the animal just stopped jumping, this could be argued that the animal is capable of choice without judgment - it chooses to stop jumping because this wastes energy without helping. But the resultant distress the animal experiences makes a compelling case that the animal has some notion of its own condition, a necessary component of the process by which it forms judgments.

This is the basis of intensional action. Intensionality requires both reflection and adjustment. The adjustment for the dog occurred when the situation changed and the animal found that the avoidance was not possible anymore. Intensionality requires a certain amount of self-consciousness - the ability to reflect on the state that oneself is in and to make these adjustments based on that reflection.

As far as we can tell, though, humans have a qualitatively greater level of consciousness than other creatures, and therefore a greater degree of volition. Where does this come from? Is there some special organ, or part of the brain that gives rise to this heightened consciousness or is it some emergent phenomena of thought?

Although there are most likely significant differences in human brain structure leading to the uniquely human level of consciousness, and thus a physiological explanation of consciousness may be achievable, I would argue that the experience of consciousness that most humans share is an emergent property of the mammalian brain.

First, I would make the case that consciousness comes from the ability to learn. The proof of this is not difficult. It is only necessary to point out that all people do not start out being born conscious. We tend to have conscious self-awareness by some time around age three or so. At this point there either must have developed some complex of neurons that caused self-awareness to arise, leading to a qualitatively different mode of functioning in the brain, or that we have accumulated a certain body of thoughts by learning them. Even if there is such a neural complex, it would appear at the same time that language acquisition, abstract thoughts and ideas and other traits arise in the child. These behaviors and ideas are learned. Actually, self-awareness is useless if one can't learn, because there would be no change in behavior. Therefore learning is necessary for the development of consciousness.

Without speculating about physical changes, the way consciousness arises can be thought of by using the metaphor of a nuclear reactor. A pile of uranium contains atoms of fissionable material. These atoms, after having absorbed a suitably condition neutron go into an unstable state and split

in two, generating more neutrons. There is a qualitatively different state between criticality and subcriticality. The subcritical pile is inert, benign. A critical pile is hotter, more radioactive, and contains larger amounts of non-uranium atoms. It has characteristics that were too diffuse to be a part of the subcritical mass.

This is probably what consciousness is all about. Any collection of neurons capable of registering pleasure and pain has a subcritical consciousness, in a way that inert materials are not capable of. But the mammalian brain has a sufficiently high degree of thought processes that give rise to thoughts, which in turn learn other thoughts, just as one fission gives rise to two or more others in close enough proximity to suitably conditioned neutrons.

Animals have a wide range of consciousness, but probably even the most advanced mammals outside of the primates are subcritical in their thought processes. Therefore their level of consciousness is qualitatively different from people.

This way of looking at consciousness may help explain how consciousness fades away when going to sleep and disappears in deep sleep or moments when one is zoned out. People who are sleepwalking can commit good or bad actions but are not morally responsible because they had no powers of volition. Even if their unconscious states come about because certain parts of their brain are shut down, this may mean that they are only capable of a series of thoughts that remain subcritical using this nuclear reaction metaphor.

Although an interesting metaphor, defining the point at where consciousness begins is difficult. The attempt to pin down this point may not be useful, though. Most creatures are well above or well below the line. Even a sleeping person, who goes from a waking state into a deep sleep, then into a dreaming state, then back into a deep sleep spends most of their time either conscious or unconscious and very little time in the transition between the two. A two to three year old human is typically the only creature who lives at the line for long. It is possible that a chimp can be at this level, also. It has also been speculated that dolphins can achieve this also, but it would be hard to give a definitive answer.

With these views of choice and consciousness, we can now look at free will. Free will has traditionally been considered to be an all-or-nothing proposition: either an entity (usually human) has free will or it does not. There are a number of types of free will that have been considered, typically characterized by whether the choices of the entity in question are nondeterministic or not. For instance, some consider humans to have free will because the choices they make have an essentially nondeterministic,

unpredictable component. There is an essentially random character to the choices. Others consider an idea of free will that is compatible with an essentially deterministic view of the choices that entities make. What both views have in common is that the entity does not employ a simple mechanistic thought process. Even for the believer in compatibilism - a free will based on a deterministic thought process - that thought process is usually non-trivial, involving considerations of the immediate environment, internal mental states and future expectations.

I shall instead argue that free will is a relative judgment call. As such it depends on both the viewer and the entity being viewed. It is not an absolute determination. The only reason free will has appeared to be absolute was because the only entities who have considered the question of free will - that we are aware of - are our fellow humans. This makes it appear absolute, because there is no relative frame of reference except the variation among people.

Although the field of Artificial Intelligence has repeatedly met with failure, there is, to our knowledge, no theoretical reason that explains why there can never be an electromechanical computer that could attain or surpass human intelligence. Some progress has been made, but so far there has not been a theoretical design that has led to a computer capable of thinking on the same level as a human over any but a very narrow range of abilities. Despite this failure, the rest of this discussion will, for the sake of argument and some examples, presume that such a machine is possible.

I first point out that, with this relative viewpoint of free will, there is no free will if there is no observer, which implies a certain level of consciousness. The consciousness is required to provide a reference point by which to judge the range of choices made by the entity under observation to the range of choices made by the observer. Therefore, the question of free will does not come up for plants and simply animals. They are incapable of evaluating free will - it does not enter their mind, such as it is.

We humans have free will relative to the viewpoint of other humans. Expressed in the alternative manner, every organism has free will as far as it can determine, if it is capable of making that determination. Given that you yourself have free will, you look at most other humans and determine that they have free will also. There are exceptional cases, but they are determined by acts of introspection and empathy. That is, a person in a vegetative state may be considered to have no free will because observation of that person compared to introspection of oneself leads one to conclude that the ability to make significant choices is profoundly impaired.

If an entity is completely deterministic to an observer, then that entity

has no free will. This is a basic precondition the absence of which leads to a judgment that denies free will. It is necessary to show that the most humans cannot have their basic actions be reliably predicted by a human observer, even if their actions are completely deterministic, and thus to other humans, the typical human has free will. The argument is based on a computational complexity argument.

When a person views an object, either living or a machine, this observer abstracts a certain level of behavior from the object's actions. That is, if a human watches a dog play a game of catch, the person does not describe the actions of the dog in all its minutia such as the precise leg movements the dog makes. Instead, the observer abstracts the condition that the dog is chasing the ball.

These abstractions characterize the object as a member of a set of similar objects. I can say that a human has certain characteristics that make a human different from other animals. But we claim that humans have free will because there is no set of abstractions that serve to characterize the ability of humans to make choices except in the most superficial of ways.

This failure to abstract the essential qualities of choice is not necessarily due to a stochastic component of the choices that people make. It can come entirely out of deterministic processes, although they are incredibly complex. Recent work in chaos theory shows that many non-linear systems exhibit chaotic behavior very dependent on initial conditions. The gross qualities of this behavior can sometimes be characterized in general terms, but the detailed moment-by-moment behavior may appear unpredictable. And in some chaotic regions even the gross properties are beyond analysis. Some of these systems are termed "trap-door" functions: they are very easy to compute in one direction, but incredibly difficult to compute their inverse.

Human choices are type of deterministic behavior. The choices that a typical person may make can show an overall regularity from day to day, although that can disappear when that person goes through life-changing events. The choice function can be considered to be a trap-door function. The individual choices may be the result of an easily-explained rationale, but the attempt to arrive at that rationale from the simple listing of the choices made can be impossible to compute efficiently.

Any attempt to make such an analysis is bound to fail due to the capability of humans to anticipate the reactions of other humans to their activities and adjust their actions in reaction to this anticipation, if a conflict in goals or desired outcome exists between them. As we will discuss in a later chapter, this ability of humans leads to an essential inability to decide what a person will do in a certain circumstance.

To abstract the qualities required to make predictions of what a person will do is often beyond human analysis. It may be reasonably predicted that this person will have a meal sometime in the next day or so, but unless the person's diet is forced to be limited, the particular menu that day may be off the mark. The appearance of free will arises when the most efficient way to make such a prediction is through the use of simulation - that is, create a copy of the person and sit down and watch what the copy will do. Unfortunately, it is not currently possible to create such a detailed, usable copy of a person with our current state of knowledge. So at present, we can only observe the original and see what that person will do. If the most efficient way to predict an entity's behavior and the choices they make, then that entity has free will.

This view of free will implicitly requires an evaluation that is relative to the entity making the determination. For a human or a computer as powerful as a human, a lizard may have no free will, but another human does. Whether or not a dog has free will is debatable. A dog is able, to a limited extent, anticipate what a human may do in reaction to what the dog does, but this level of anticipation is limited both in time and degree of sophistication, limits that humans do not possess. For a theoretical computer more powerful than the mental processes of most mammals, a dog may have no free will but a human has. For a computer with godlike powers, even a human has no free will.

This relative view of free will clarifies the relationship between God and man in terms of free will. To a human, another human has free will, but to an omniscient god a human does not. This is true even if humanity's actions are not willed by an all-powerful deity. Since free will is a matter of evaluation, omniscience alone is enough to lead to a lack of free will in God's eyes.

The typically absolute interpretation of free will naturally leads to a misapplication of the consequences of divine predestination found in some Protestant denominations. The presumption that an omniscient god can perfectly predict the actions of any given human says nothing about the way in which a human should function. This is due to the fact that predestination and a lack of free will is obvious for an omniscient god, but the reduced capacity of human reasoning leads to a presumption of free will and the necessity of making appropriate moral choices.

This relative interpretation of the concept of free will is compatibilist - that is, even if there is a recognition that the world is essentially deterministic, free will is compatible with this determinism. Even though there may be a non-deterministic component to the universe, this nondeterminism does

not play a significant role in the typical choices that humans make.

The opposite of compatibilism can be addressed with this relative model. The incompatible "requires that there is a plurality of futures open to me consistent with the past (and laws of nature) being just as they were." [Stanford encyclopedia of philosophy] The incompatibilist viewpoint requires two assumptions: first that there exists a basic indeterminacy in reality, and second, that this indeterminacy plays a nontrivial part in the decision making process of entities such as humans. Otherwise, incompatibilism is an illusion brought about by the impossibility of identifying all of the deterministic processes. The relative view of free will implies an identification of an entity viewing another as if that viewed entity has a plurality of futures even though no such thing actually exists. This plurality is not inherent in the nature of reality; instead it comes from the inability of the observer to reliably rule out significant alternative outcomes. As the depth of analysis of the entity's future possibilities increases, an essential determinism leads to the eventual elimination of alternate futures as possible until just one is left. That is why the determination of free will disappears as the complexity of analysis increases.

This complexity also increases within humans. As we learn more and more about psychology and sociology, we identify less and less situations in which an individual is practicing free will. Although the potential exists that a state can be reached in which no human could be judged to have free will, in practice this is likely to be unobtainable.

One argument against a purely deterministic view of nature is from the discoveries in the twentieth century that the laws of physics in the smallest scale exhibits and essential indeterminacy that can only be expressed using probability functions. The best known expression of this is the Heisenberg Uncertainty principle of Quantum mechanics. Quantum mechanics could supply an indeterminacy that is potentially an important influence on human decision making.

But it does not appear that way when looked at carefully. Quantum mechanical processes, even in a world controlled by chaos theory - which is profoundly affected by minuscule changes in initial conditions - affects the range of situations available for a human to make moral choices, but does not appear to be an important factor in the way a human makes those choices. It appears that all of the probabilistic quantum fluctuations may operate on such a low level that they do not enter into the day-to-day decision making processes of the typical human. Rational mental processes seem to follow deterministic logical rules, that are either internally consistent, or if they are not have a rationalization that can be described by a determinable

sylogism. Even emotional mental processes have a distinct physiological process that can be described in terms of neurons firing in certain regions of the brain, where the quantum fluctuations more often than not cancel out. Therefore human decision making tends to be composed of reasons and desires that have an almost discrete objective existence that is as little conditioned on probabilistic events as the chance at a word on the printed page could, through quantum mechanical processes alone, rearrange themselves into a piece of gibberish. For example, the operation of today's computers is potentially affected by quantum mechanical events, but the computers are designed so that this is minimized. It also seems likely that these effects do not play a part in human reasoning. Therefore, although quantum mechanical processes can make minute changes to the possible choices that a human can make, in general, it is unlikely that these processes significantly affect outcomes.

Besides noting that there is a significant difference in the judgment of the amount of free will an entity has based on whether the observer and the observed are qualitatively different entities, such as a lizard, a dog, a human or God, it is always possible to draw finer distinctions even when the comparison is between entities of the same type. These finer distinction can come from the relative intellectual capacity of the observed entity versus the observer. An extreme case is between that of an adult and a child, but a distinction can also be drawn between the capacity for free will between an educated person and an uneducated one, or expressed in terms of a person's life experience instead of their years of schooling, between a wise person and an unwise one.

The other difference in judgment is of course in the determination of the situation that the observed entity is found relative to the situation of the observer. It is obvious that the determination of the amount of free will an entity possesses is different if the situation that the entity is in is understood with a deeper level of knowledge of its implications by the observer than the observed. This is always a factor in culpability. This determination of free will can change depending on the situation of the observer or the situation of the observed. There is an essential symmetry changing the determination of free will such that a change in the observed has a corresponding possible shift in the situation of the observer.

The more educated or wiser, a person becomes, the finer the distinctions can be made. A deeper appreciation of the complexities of the situation leads to a change in the determination of whether the actor acted with free will or not and is therefore culpable. Most often this increased knowledge of the situation leads to a determination that the actor had less free will than

a more superficial analysis shows, but this is not always the case. It can also happen that a deeper analysis can show that the entity actually did have alternatives available that the entity was not aware of. This means that the observer would declare that the observed acted with free will even though the observed would claim that the action was the only alternative open at the time.

The determination of an entity's free will is not just an abstract thing. The notion of free will is important because it affects the way that individuals react to the actions that an entity takes. The reactions that other people have to an entity's actions sometimes take the form of rewards and punishments, but there are other reactions that involve the notion of free will besides this. For example, the expectation that a person can learn from their actions is dependent on the degree of free will the entity is thought to possess. Therefore, the amount of teaching given to this entity and the depth of this teaching, such as expressing this knowledge in terms of motivations depends on the amount of free will the entity is judged to have³, especially since free will is partly determined by the entity's capacity to reason with self-awareness. This is true both for humans and for teachable animals, such as most mammals.

But the question of rewards and punishments is usually the starting point for determining how to react to another's actions in light of their capacity for free will. An unfortunate outcome can be excused or even met with sympathy if the action is deemed to be beyond that person's capacity to alter. On the other hand, the determination that a person acted with willful disregard to what is right will not only cause a harsher punishment to be meted out, but the overall judgment of that person's character will suffer as a result.

The determination of free will is considered in light of the number of options available to a person and the ability of that person to choose amongst those options. With an absolute notion of free will, this determination is complicated enough. With a relative notion of free will, not only does the mental capacity of the actor come into consideration, but also the mental capacity of the observer making the judgment. This two body problem leads to a complicated dance of reasoning that can lead to different results, especially as the mental capacity of most living creatures is constantly in a state of flux.

But because free will is relative, systems of rewards and punishments are relative also. Punishment has two main purposes - deterrence and retribution both of these are dependent on a person's degree of complicity and this depends on their volition. Even if complicity can be objectively determined,

volition never is.

One way to reduce the complexity of making these judgments is by reference to a standard. Standards still exist even when the determinations are relative. It was already pointed out that in the determination of what is good, there is an absolute standard even though what is good is relative. This absolute standard is expressed in the Golden Rule and is derived from an application of empathy. The same process applies to judgment. The reaction to a person's actions, the reward and punishment, requires empathy on the part of the person doing the judging to determine what capacity for free will the actor had in that case. But this is not entirely enough because every person's capacity for free will is dependent upon their situation. So the capacity for reasoning empathetically must take into consideration the typical situation the actor is in, not the situation of the person doing the judging. One way to objectify this is to postulate some hypothetically typical person, and compare the actions of that actor to this hypothetical.

The starting point of this evaluation is for the person making the judgment to consider themselves a typical person, and then make the judgment based on a sense of empathy. This requires that the observer take into account the difference in the situation that the observer and the observer find themselves, especially since such judgments are usually made based on events that have happened in the past. This means that the situations differ greatly, since the observer knows the actual outcome when the actor probably was not able to predict this with any reliability. It takes a fair degree of empathy for the observer to ignore these differences and to put themselves in the same place as the actor with the same lack of knowledge.

This comparison to the typical person forms the bedrock of society's laws. But comparison to the individual is the basis of most people's judgment of others. This is not a good basis of defining and applying the law, though.. This is because our comparisons are usually distorted by the high self-image that the typical person holds of themselves. This means that empathy must not be given in comparison to the observer of a moral act, but made on the basis of contemplating a third, theoretical construct - the image of the typical person.

When, in the determination of the responsibility that a person has for their actions - whether that person be punished for an evil act or rewarded for being virtuous, the point must come when free will is presumed, if the reward or punishment is a response sufficient in itself - that is, the response is made for no other purpose than to reward or punish. This determination can be made by comparing this person to the typical person. If a typical person would have acted more benignly in the same situation, then the person being

judged would be considered to have done something wrong. If the situation was such that the typical person would have done no better, regardless of the bad consequences of the act, then the person is judged not to be responsible. We also declare someone not responsible if the outcome was predetermined - this happens for people such as children or with impaired mental capacity.

If the application of reward and punishment is done for the purpose of correcting future behavior, then it appears that free will is not as important a consideration in whether or not to mete out that response. In this case, the determination of free will is not as important as the determination of whether the future behavior of the actor will be changed in a way that is desired by the observer due to the reward or punishment. If the recipient of the reward or punishment is deemed to possess free will, this means that the observer is less likely to successfully predict that the application of this corrective action will lead to the desired result.

One unfortunate effect of this indeterminacy in the face of free will is that the person who metes out the judgment may tend to over-emphasize the degree of the reward or, especially the punishment in the light of free will, because the result is less predictable. That is, knowing that the actor had limited free will and knowing the effect on the future behavior of the punishment, the punishment will be more finely calibrated due to the degree of knowledge. In ignorance, an attempt will be made to further guarantee success on correcting the behavior of the actor by increasing the degree of punishment.

That question arises then, to what degree does one punish a transgressor who grew up in a deprived household? There is often a tendency to over-estimate the degree of culpability of such a person, or an overestimation of the effect of punishment or reward on future behavior. This is due to the tendency to neglect the difference in situation that the observer came from relative to that of the transgressor. To make an appropriate response, this decision is made in terms of the average person - whether they would grow up in the same situation would make the same transgressions also. This is not to excuse the bad behavior of someone who grew up in a bad situation from the consequences of their actions when they ought to know better. But it does argue that a person of privilege, who is given the opportunity to truly know the difference between right and wrong, should by rights be judged more harshly than the person who does not have the opportunity to see those differences.

One other consideration that must be made is the recognition that there are different levels of entities, and that each of these levels can be judged on their own merits. That is, the situation that an individual finds themselves

is no more than the characteristics of that part of the society in which they were brought up. It is reasonable to grant a certain amount of free will to a society as a whole, since that society manifests a certain level of volition that results in an emergent behavior that is different than the individual behaviors of the people that make up the society. That is, the volitional actions of the individuals can lead to a qualitatively different volition for the society that aggregates these actions. This means that it is possible to judge and reward or punish a society independent of the judgment of the individual in the society.

This observation leads to the conclusion that praise or blame for an individual's actions can be apportioned between the individual and the society of which they are a part. An example of this is the matter of racial prejudice. If a society as a whole is prejudiced against a particular group of people, it is reasonable to condemn the society more than the individual, although both must share the blame to some extent.

It is interesting to note that if there is a god that is all loving and omniscient, and the rewards and punishments are done for no other purpose than reward or punishment itself, then that god can certainly create a heaven or hell to separate the sheep from the goats. On the other hand, if the purpose of reward and punishment is to correct, then the all-loving nature of that god implies that god must be a universalist.

We have predicated the ability to make moral choices on the existence of consciousness. It is this consciousness that makes the moral choice volitional. The question arises though, whether these choices are predetermined or due to free will. We have seen that free will is not an absolute measure - it is instead a relative judgment based on the different mental capacities of the observer making the determination of whether free will exists or not relative to the mental capacities of the observed entity whose will is being judged. This determination must also take into account the situation in which the entity finds themselves. The degree of free will determines the culpability of the entity making a moral choice - an entity that can be judged not to be able to make a free choice cannot be held to blame by the entity making the judgment. Finally, we note that the blame can be shared between individual entities and the larger entity that defines the situation the individual finds themselves.

Chapter 8

Logic

Logic is the primary way that humans reason, but it has its limits.

Up to now, we have talked about some of the basic principles of morality, Morality, is the knowledge of good and bad. The remainder of this book will discuss ethics - the methods by which moral decisions are made.

A recurrent debate in the field of ethics has been to what degree, if any, moral decisions are subject to logical reason. Reason is the way that people make decisions. The formalization of these methods of reasoning is the study of logic. In contrast to the application of logical reasoning to matters of good and evil, there is the emotional response to our moral choices. The emotions appear not to follow the laws of logic, creating a dynamic of their own, a different method of reasoning.

This chapter will go into some detail about logic and formal methods of computation. Some of the more advanced results in logic, discoveries of the last century, have an important bearing on the power and limits of logic.

What is logic? Like the lever and inclined plane to physics, it is a simple and fundamental piece of machinery for the intellect. Most intellectual processes are reducible to logic just as many dissimilar machines are reducible to the inclined plane. But just as machines can be especially tailored to optimally perform certain tasks, there are different logics and formal calculi, each fulfilling a different purpose. And just as there are other machines besides an inclined plane, logic is only one of the tools that the intellect possesses.

Logic essentially is a calculus of discourse - it is a way of talking about things. Logic is defined by its rules of syntax and rules of inference. Syntax rules define the structure of the sentences. They provide the method of determining if a statement is well-formed; whether the sentence structure

has no errors. The syntax can also determine the structure of the overall discourse and it defines a space of sentences.

Logic as we know it was first systematized by Aristotle. He proved a detailed analysis of logical syllogisms. To give a hoary example found in standard logic textbooks such as Copi:

All humans are mortal Socrates is human Therefore Socrates is mortal

Aristotle showed that when the concepts in the sentences are replaced by abstract symbols along with their logical connectives, we can analyze logical arguments in and of themselves without recourse to the specific facts at hand. This work was later extended in the nineteenth and twentieth centuries by people such as Boole, Frege and Russell and Whitehead.

Logic requires a formal language that maps statements to some notion of truth and falsity. Besides the simple notion of logical connectives such as and

and , or, ... implication

implies negation and equality =, there is also a need for a definition of logical predicates - properties of individual objects and their interrelationships that are true for some selection of individual objects and some properties but false for others. Such formal languages can be represented as a Herbrand universe, where various interpretations of the objects and properties form different models of the formal system.

The first step in logical reasoning is to formalize the situation at hand in terms of the applicable objects and the predicates that describe their interrelationships. This can be a source of disagreement. Some aspects of definitions and their use will be discussed in the next chapter. But the formalization of a statement of fact into a logical notation requires a clear and consistent use of terminology and a clear understanding of the assumptions that are made when people convert real world facts into logical notation.

After the statements are defined, logic then provides rules of inference. The inferential rules show how derive conclusions from axioms and premises. For example, if A is true and B is false, then A OR B (A \vee B) can be inferred, but A AND B (A

and B) cannot. If B must be true whenever A is true, but B can be either true or false if A is not true, then A implies B: A

implies B. If both A

implies B and B

implies A, then A if and only if B, given as A IFF B (A

implies B)

There are two different forms of logic, the propositional logic and the predicate logic. Propositional logic is the rules and methods of reasoning

for basic propositions, made up of simple symbols A, B, C, \dots connected by AND, OR, IMPLIES, IFF and NOT. The rules of propositional logic are relatively simple. The situation becomes more complex with in the predicate calculus. There, the goal is to reason over universes of objects, finding out what is true for all individuals x , or if there exists at least one y . This leads to the addition of predicates such as Ax and Byz , for example where A may or may not be true for all x and only certain pairs of values assigned to variables y and z can have the relationship B . Then we need to consider universal and existential quantification: $\hat{A}x Hx$ implies $\hat{E}y Hy$ and x implies y and Myx , which under some Herbrand universe be interpreted to mean: For all objects x , if x is a human (Hx) then there exists a y , where y is also human (Hy), the humans x and y are not identical, and y is the mother of x (Mxy).

The truth or falsity of logical statements is usually determined in light of certain premises: Given that A implies B and B implies C , it must be true that A implies C . A

implies B B
implies C — A
implies C

This form of presenting a logical argument can be packaged together by making a conjunction of the premises and implying the conclusion: the assumption that both A implies B and that B implies C are both true further implies that A implies C is true. $((A$ implies $B)$ and $(B$ implies $C))$ implies $(A$ implies $C)$.

For a logical argument to be true, it must be true under every model that serves as an interpretation of the objects and predicates. There are a number of sets of methodical rules that can be applied to determine if a logical argument is valid. One of the simplest is to package the logical argument as a single proposition. If the logical argument is true, then this statement is a tautology. Therefore, there can be no assignment of truth or falsity to the propositions A, B and C where the statement is true.

It is actually simpler to negate the statement and show that every assignment of truth or falsity to the negated statement leads to a contradiction.

We can represent all these assignments through a device known as a 'truth tree' where statements connected by an AND are replaced by listing the statements one below the other, and statements connected by and OR are fanned out side by side, forming a tree. See Richard C. Jeffrey 'Formal Logic: Its Scope and Limits' for a detailed description.

For example, assume the negation: that ((A implies B) and (B implies C)) implies (A implies C) is false. Now if X implies Y is false, it must be true that X is true and Y is false. Conversely, if X implies Y is true then either X is false and Y is true.

So, assuming ((A implies B) and (B implies C)) implies (A implies C) is false, it must be true that A implies B is true, and B implies C is true, and A implies C is false - that is (A implies C)

We write these cases this way: A implies B B implies C (A implies C)

If A implies B is true, then either A is false or B is true. We replace A implies B by two paths, one containing A and the other containing B: B implies C (A implies C) A B

Similarly, we replace B implies C: (A implies C) B C A B A B X

The 'X' marks a branch of the tree that contains the two contradictory propositions B and B.

Finally, for A implies C to be false, then both A must be true and C false. Replace (A

implies C) by these two statements:

A C B C A B A B X X X X

Now all four branches have a contradiction: in the first branch, A and A, in the second B and B, in the third A and A again, and in the fourth C and C.

All propositional syllogisms can be proven true or false this way: if you list the premises and the negation of the conclusion, then apply these rules to expand the argument into a tree, the syllogism is true if and only if every single branch has a contradiction.

The use of truth tables is extended to the predicate calculus with universal and existential quantification. If a universally quantified statement is false, then its existentially quantified negation is true: $\neg \hat{x} (...Px...) \text{ iff } \hat{E}x (...Px...)$. A similar rule is true for existential quantification.

The goal for using the truth tree method is to replace each universally and existentially quantified statement by a statement where the quantified variables x,y,z are replaced by constants a,b,c. The process is defined as follows:

If an existentially quantified statement $\hat{E}x (...Px...)$ appears in the truth tree, then for each branch of the tree, replace it by $(...Pa...)$, where a is a constant that has not yet appeared in that branch of the tree.

If a universally quantified statement $\hat{A}x (...Px...)$ appears in the truth tree, then for any constant a,b,c, etc that appears in the path of the tree, the statements $(...Pa...)$, $(...Pb...)$, $(...Pc...)$ can be added to the tree. If there are no constants in that path already, choose a constant a. Do not remove the universally quantified statement, though, as it can be used again as new constants are added.

There are also a couple of rules for handling equality ($x=y$). The first is that any branch that contains the statement $(a=a)$ for some constant a contains a contradiction. If a positive statement $a=b$ appears, then it is possible to add statements in which a ...b... appears with the statement with the constant a substituted: ...a...

The Godel Completeness Theorem shows that with a systematic application of these rules, any true theorem in the Predicate Calculus can be proven to be true. That is why it is referred to as the predicate calculus - calculus means a process of reasoning by a formal symbolic computation. A major aspect of logic is that the procedure is a way of describing recipes, that when followed as given will yield absolutely the same results. This shows that formal logic is a model of effective computation - it is a mechanical, cookbook method to find every true statement. An effective computational procedure is fully determined - there is no appeal to magic or a supernatural

force.

Two other methods of reasoning are Turing Machines and the General Recursive Functions. A Turing machine consists of a head positioned over a potentially infinite tape of symbols. The head can be in any of a finite number of states. Each position on the tape can hold one of a finite number of symbols, or blank, and the head can read and write those symbols. The program for the Turing Machine is a state table, where the rows of the table are labeled with the indexes of the states that the Turing Machine can be in and the columns of the table are headed by the symbols that can be written on the tape, including the blank. If the Turing machine is in state i and reads symbol a , then a lookup in the table shows the next action the Turing Machine can take: it can overwrite the symbol a with a new symbol b , it can move Left or Right, and it can go into a new state j : table entry $\langle i, a, j \rangle$ is the triple $\langle b, M, j \rangle$, where M is either L or R. It is acceptable for a and b to be the same, or i and j to be the same, and either a or b or both can be a blank. The tape is presumed to begin with only a finite number of non-blank symbols written on it. There is a unique initial state for a given Turing Machine, state 0. If the Turing Machine goes back to the start state at any point of its computation, it is said to have halted. If it never reenters the start state then it never halts. If the Turing Machine halts, then it accepts the input sequence. To give an example, if the number 19 is given as a binary sequence 10011 and Turing Machine $TM(x)$ halts if only if x is a prime number, then $TM(10011)$ halts. Turing machines can also be considered as functions. When the Turing Machine halts on input x , inspect the symbols left on the tape. This can be considered as the output y , thus $TM(x)=y$.

It is possible to consider a nondeterministic Turing Machine that can have more than one alternative move from a given pair of state and symbol. In that case there is an exponentially growing set of sequences of states and tape symbols that can arise from a single input. In this case, we express these sequences as a branching tree of alternatives. If one branch halts, the Turing Machine accepts the input. If the Turing Machine is considered to be a function, choose the shortest sequence and look at what is left on the tape at the end of the sequence for the output.

Due to the simplicity of the Turing Machine, the rules for computing the truth of a logical statement can be rather involved, since it takes a lot of these simple instructions to do anything of significance. Even simple arithmetical operations such as adding one to a binary number can take a handful of states to express. To take a simple example, a nondeterministic Turing Machine can determine if there is an assignment of truth or falsity to

a set of Boolean statements in conjunctive normal form (equations of OR's linked by AND's). A representative set of equations is: (A ... B) and (A ... C ... D) and (B ... C) and (C ... D) This is satisfied if the truth assignments are A=T, B=F, C=F and D=T.

To make the construction of the Turing Machine simple, we shall assume that the input has a very simple syntax. If there are n statements and m variables in total, then the input is n statements each m characters in length, with each statement separated by a '#' from the previous one. Character i of statement j is a 'X' if variable i does not appear in statement j, a 'P' if variable i appears but is not negated, and 'N' if it is negated. So (A ... B) and (A ... C ... D) and (B ... C) and (C ... D) is given on the input tape to the Turing machine as ...#PPXX#NXNN#XPNX#XXPP...

Where '.' represents a spot on the tape that is blank.

The Nondeterministic Turing Machine has 10 states, numbered 0 (the start state) through 9 (a reject state). The tape begins with the symbols #,X,P,N on the tape and the Turing Machine can also read and write the symbols a,b,c,d,e,f,T,S. To save space in the following table, if a set of productions start in the same state n and go into the same state m and move in the same direction M, for a set of symbols x,y,z, where the machine replaces x by r, y by s and z by t, then I will write the three productions n,x implies m,r,M n,y implies m,s,M n,z implies m,t,M as the contraction: n,x,y,z implies m,r,s,tM The only nondeterministic state is state 1. The state 1 productions will be written out without a contraction.

Here is the full Turing Machine 0,# implies 1,#R (i.e. the Machine in the start state 0 reads the leftmost # and goes into state 1 and moves right, leaving # unchanged) 1,X implies 1,a,R 1,X implies 1,b,R 1,P implies 1,c,R 1,P implies 1,d,R 1,N implies 1,e,R 1,N implies 1,f,R 1# implies 2,#,L 2,a,b,c,d,e,f implies 2,a,b,c,d,e,f,L 2,#,T,S implies #,T,S3,R 3,a,b,c,d,e,f

implies 4,S,S,T,S,S,T,R 4,a,b,c,d,e,f,#
 implies 4,a,b,c,d,e,f,#,R 5,a,b,c,d,e,f,#
 implies 5,a,b,c,d,e,f,#,R 4,X,P,N
 implies 2,a,c,e,L 5,X,P,N
 implies 2,b,d,f,L 4,.
 implies 6,..L 5,.
 implies 6,..L 6,a,b,c,d,e,f
 implies 6,S,S,T,S,S,T,L 6,#,S,T
 implies 6,#,S,T,R 6,.
 implies 7,..L 7,S
 implies 7,S,L 7,T
 implies 8,S,L 7,#
 implies 9,#,L 8,S,T
 implies 8S,T,L 8,#
 implies 7,#,L 9,#,S,T,.
 implies 9,#,S,T,..L 7,.
 implies 0,..L (halt state)

Another useful formal system is what came to be known as the primitive recursive functions. (The following description of recursive functions is quoted from Rogers Theory of Recursive Functions and Effective Computation):

The primitive recursive functions are a set of functions satisfying the following rules: (1) all constant functions $f(x_1, \dots, x_n) = c$ (2) the successor function $f(x) = x + 1$ (3) all identity functions $f(x_1, \dots, x_n) = x_i$ (4) if f is a primitive recursive function of m variables and g_1, \dots, g_m are primitive recursive functions of n variables, then the composition of these functions is a primitive recursive function: $h(x_1, \dots, x_n) = f(g_1(x_1, \dots, x_n), \dots, g_m(x_1, \dots, x_n))$ (5) if f is a primitive recursive function of $n+1$ variables and g is a primitive recursive function $n-1$ variables, then h is a primitive recursive function, where $h(0, x_2, \dots, x_n) = g(x_2, \dots, x_n)$ and $h(y+1, x_2, \dots, x_n) = f(y, h(y, x_2, \dots, x_n), x_2, \dots, x_n)$

Although the primitive recursive functions seem to be general, it was found that they were not general enough to express Ackerman's function: $A(0, 0, y) = y$ $A(0, x+1, y) = A(0, x, y) + 1$ $A(1, 0, y) = 0$ $A(z+2, 0, y) = 1$ $A(z+1, x+1, y) = A(z, A(z+1, x, y), y)$

This function is a generalized exponential, a function where the first argument represents an index of a class of functions of the second two arguments: $A(0, x, y) = y + x$ $A(1, x, y) = y * x$... $A(z+1, x, y) =$ the result of applying y to itself $x - 1$ times under the z operation

The fact that clearly defined functions existed that could not be defined as primitive recursive functions led to their being classified as primitive and

a more general class of functions defined instead. The General Recursive functions are defined by Kleene as a computation of a finite sequence of equations. The definition of a function P is a set of recursive equations. A computation is a finite sequence of equations, beginning with P , where each equation after P is obtained from the preceding equations either by the substitution of a numerical expression for a variable symbol throughout an equation or by the use of one equation to substitute "equals for equals" at any occurrence in a second equation or by the evaluation of an instance of the successor function $\lambda x[x+1]$. In P , we allow auxiliary function symbols in addition to the main function symbol in whose evaluation we are interested. Thus the set of equations $f(0) = 0$ $g(x) = f(x)+1$ $f(x+1) = g(x)+1$ with f as the main function and g as an auxiliary function symbol, that computes the function $\lambda x[2x]$. Although this definition allows for ambiguity in choosing the computation sequence, the procedure can be made effective and unambiguous by defining a way of enumerating all acceptable sequences and choosing the first complete sequence in the enumeration that comes along.

It is also possible to express the General Recursive Function as a simple extension of the primitive recursive functions by the mu theorem. If $f(x, z_1, \dots, z_n)$ is a partial function with $n+1$ arguments, then the mu function implies μx is the partial function with arguments (z_1, \dots, z_n) that returns as its answer the smallest value x such that $f(x, z_1, \dots, z_n) = 0$. If no such x exists, then $\mu x f$ is not defined for the case z_1, \dots, z_n .

The mu theorem is expressed as follows. Given any general recursive function g of n variables that can be computed by Turing Machine TM_i , it can be defined as the mu function applied to a primitive recursive function based on f : $g(x_1, \dots, x_n) = \mu z [z = \langle y, t \rangle \text{ and } v = \langle x_1, \dots, x_n \rangle \text{ and } f(i, y, t, v) = 0]$, that is, $g(x_1, \dots, x_n)$ equals z , where the value of z is the smallest value such that $f(i, y, t, v) = 0$.

There have been a number of models of effective computation defined besides formal logic: General Recursive functions Church's lambda calculus $\lambda x [\dots x \dots]$ Post's correspondence problem Turing machines

Each one of these has been shown to be equivalent to the others. That is to say, the elementary computational steps taken in each of these methods of computation can be simulated step by step in any of the others.

For example, the predicate calculus of formal logic can be converted into a Turing Machine. Any logical statement can be built up out of the connectives

and \vee ,

implies \rightarrow , $=$ and parenthesis $()$, where any number of variables and predicates can be represented using numbers: x_{111} is the 3rd variable symbol and

P111111 is the 6th Predicate symbol. A single space can be used between statements. This is the form that the logical statements can be written on the tape. The Turing Machine then produces the proof the using the well-known rules that a person uses.

It is also possible to express any General Recursive function as a Turing Machine. It can also be proved that a single step of a Turing Machine can be expressed as a primitive recursive function. From this we can define a primitive recursive function $f(i,y,t,x)$ that returns 0 if and only if $TMi(x)=y$ in t steps.

Because Turing machines are so simple, they often form the basis of the canonical (or standard) definition of this essential set of effective procedures. Each Turing Machine, for example can be assigned a number, known as its Godel Number. The method of doing that is simple: since a Turing Machine is defined by its state table and each table entry $\langle i,a \rangle$ is the triple $\langle b,M,j \rangle$, each entry in the Turing Machine's state table is a five tuple $\langle i,a,b,M,j \rangle$. Now each individual state i and each unique tape symbol a can be assigned a number. The two movements L and R can be assigned the numbers 1 and 2. Then the five tuple is five numbers $\langle i,a,b,M,j \rangle$. Each one of these tuples can be converted into a sequence of 0's and 1's of the following form: 'i' number 1's followed by a 0 followed by 'a' number 1's followed by a 0 followed by 'b' number 1's followed by a 0 followed by either by 10 for $M=L$ or 110 for $M=R$ followed by a 0 followed by 'j' number 1's followed by a 0. Stringing together all of these five-tuples gives a very large unique binary number for each Turing Machine.

Now, given any Turing Machine converted into a binary number, assume that symbol number 1 is a blank and symbol number 2 is the letter a. Then let us consider all of the inputs to the Turing machine that consist only of strings of a's - the rest of the tape is blank. If the Turing Machine is started in state number 1, assume that the Turing Machine can run until it reaches state number 1 again, at which point it halts and stops running. Otherwise, it goes into an infinite loop, or it writes an infinitely long string of symbols onto the tape and never halts.

We can now define sets of numbers the following way: given Turing Machine i , the Recursively Enumerable set W_i is the set of all numbers x such that, given an input tape with x letter a's written on it, Turing Machine starts in state 0 and halts. If there are only blanks on the tape, then the number x is zero, so zero can be included in the set W_i also.

Therefore we have the following equivalences First order predicate calculus is equivalent to Turing Machines Turing Machines are equivalent to the General Recursive functions Turing Machines are equivalent to the Re-

cursively Enumerable sets. Thus First order predicate calculus is equivalent to the Recursively Enumerable Sets.

[Effective procedures]

The notion of an effective procedure can be represented in any of the equivalent models. Alternatively, classes of effective computations have always been shown to be either some transformation of the class of Recursively Enumerable sets, or some proper subset of this class.

These proofs that different types of models for effective procedures are equivalent have led to the claim known as Church's Thesis: there is an essential set of effective procedures that can be represented in any number of equivalent ways. Any class of functions that is greater than this essential set must have some non-effective, non-cookbook step in their computation that cannot be reduced to a step-by-step deterministic method.

So logic is just one of the methods of performing effective reasoning. It is, though, the most common and useful of these methods. Originally formalized by Aristotle and further refined by Boole and Frege, logic is still today the most important tools for arriving at truth.

All computer programming languages are effective procedures Some employ some of these formalisms as an idiom. One language in particular, Prolog, is actually composed of logical statements, where a computation follows the process of logical deduction. Another language is Lisp, which styles itself more like the partial recursive functions.

[Logic has limits]

Logic is a very important tool for reasoning. But it has its limits. As a practical matter, there is nothing special in the operations of logic themselves. These operations form a concise cookbook of instructions. What makes reason special comes from the objects of contemplation - the concepts in the real world that get cast into the logical syllogism as axioms, and the meanings attached to the predicates and terms.

This leads to the first limit. Using only the real world as the source of the concepts we subject to logical manipulation limits the range of our logical conclusions.

To quote from Alfred North Whitehead (From Religion in the Making):
"Any proof which commences with the consideration of the character of the actual world cannot rise above the actuality of this world. ... By considering the world we can find all the factors required by the total metaphysical situation; but we cannot discover anything not included in this totality of actual fact, and yet explanatory of it."

Therefore, since there are no observable infinities in the real world, whole branches of mathematics that rely on transfinite concepts cannot be derived

from logical reasoning applied to real world objects alone. This may seem not to be very limiting, except that differential and integral calculus are in this category, and a large part of modern technology is built on a basis of reasoning based on calculus. So a large number of very practical applications are based on the postulates that certain ideas that are not observable are presumed to have some ideal existence.

A second shortcoming is that the relation to logic as a tool and the real world as described by logic is murky. One problem is that the representations of the real world that logic leads us to can distort our understanding of the world. For example, in recent generations, it has been fashionable to consider the real world as quantized - made up of discrete individual units - and therefore amenable to direct representation as logical constants. This seems to make this discreteness to be a real characteristic of nature, but it could just be an artifact of the way we represent reality.

A representation is partly a matter of convenience. The original use of logic was to express geometrical arguments, where, although the objects being manipulated - lines, points and circles - were discrete unities, they could not be decomposed in any fundamental level to quanta. But even here, these representations led to trouble - they gave rise to the paradoxes of Parmenides and Xeno, such as the Achilles and the hare. It took the creation of the calculus to resolve these paradoxes adequately.

A third shortcoming is that the map is not the territory. To be useable, logic can only abstract a limited set of aspects of reality at any given time. This means that there are an infinite number of aspects that must be left out. Giving the driving directions to the Empire State Building and back is not equivalent from actually going there. It only abstracts one dimension of the experience - that of the way of making the trip there. It does not discuss the view, for example. You could add that dimension, and others, but then the dimensionality becomes infinite and you cannot even begin to reason about it.

[Logic is hard to apply to a deep level]

Another limit to logic is that it is hard to rigorously apply. According to C. Hartley Rogers, humans seem to be limited to three alterations of quantifiers - a syllogism in which a proposition of the form $\hat{A}x\hat{E}y\hat{A}z\hat{E}v Pxyzv$ - seems to be beyond the capabilities of most people.

This incapacity to analyze logical syllogisms with too much complexity is another strong reason why humans deal with the emergent properties of things instead of breaking them down to their basic components. As an example, take a living organism. Presuming no supernatural aspects, an organism can be completely defined, for all intents and purposes, by the

location and arrangement of its chemical components in space and time. This is what logical analysis does to any item of discourse - it breaks it down into its parts.

It is mostly practical for simple things. Logical theorems work on things like proteins, but not human beings, represented as assemblages of chemical molecules. With our ability to experience these emergent properties, what we perceive with our senses is more real than what we perceive with reason.

The final limit to logic is that logic is just a means of discourse. Logic does not impart meaning to what one writes, any more than a paintbrush imparts meaning to a canvas. The meaning is imparted in the identification and interpretation of the properties used in the logical syllogism. Logic itself does works for any interpretation given in some Herbrand Universe. The wonder is that such a simple tool works at all, and can be applied to so much. But it is just one tool, which is why we sometimes forget that there are others.

[Logic is not the only tool]

What alternatives are there to logic? The most obvious is action: pure action, without analysis. Although certain aspects of action at certain times can be expressed in logic, the totality of action cannot be expressed in logic. A second alternative is the recognition of gestalts - experiencing things as an unanalyzed totality. This is the passive side of action.

An further alternative to logic is storytelling and other art forms. The arts can be reduced to logic, but this would be an error. In logic a story this is anecdotal evidence - usually considered a problem. But in the artistic sense the anecdote reveals more than a logical analysis does - casting it into logic loses the essence.

Related to both action and storytelling is learning by example. This is a combination of both. The Taoists such as Chuaung Tsu knew that the essence of doing is not reducible to analysis. This fact is still understood in education. A chemistry textbook can never take the place of a chemistry lab experiment. Even a computer simulation runs the risk, through its limited representation of the events it is simulating, is missing important aspects that are not built into the simulation.

Despite its limitations, though, logic works wonders. For example, at the level of DNA and proteins, the discrete logical formalism works admirably for the understanding the basic building blocks of life. This bottom up approach can be used for many aspects of understanding, as long as we recognize its limits.

[RE sets and the Incompleteness Theorem]

To understand modern logic, it is necessary to understand the notion

of the incompleteness of logic. We shall demonstrate this by discussing the logical paradoxes.

The most ancient of the logical paradoxes is the Liar Paradox. This is attributed to the philosopher Epimenides: 1. Epimenides is a Cretan 2. Epimenides says 'All Cretans are liars'

If we assume that a liar always lies, then if all Cretans are liars and Epimenides is a Cretan, we can conclude that Epimenides is a liar, thus this statement is a lie. Actually, then, the opposite is a true statement: 'some Cretans sometimes tell the truth.'

As originally presented, the pair of statements does not lead to a paradox. But we can sharpen the statement so that the paradox is obvious. Eubulides Paradox is the statement 'This statement is false.' Assume this statement true. Then it follows that the statement is false - a contradiction. On the other hand, assume the statement false. Then we can conclude that the statement being false is false - i.e., the statement is true. Also a contradiction.

There are a number of self-referential paradoxes that can be easily expressed in mathematics. A classic is the Russell Paradox. Consider sets of elements. Sets of sets are known as classes. Consider the class of all sets that do not contain themselves as a member. Call this class S. If S does not contain S as a member, then S is a member of S by definition. This is a contradiction. So S cannot be in the class S, which also leads to a contradiction.

[Axiom sets for religion and morality]

In the greatest proof in formal logic in the twentieth century, Kurt Godel was able to formalize Eubulides Paradox from the axioms of arithmetic and the formal notion of general recursive functions. As mentioned before, the equivalence of the general recursive functions and formal logic means that if you are able to assign a number x to every possible logical statement, there is a function $fn(x)$ that returns 1 if and only if the logical statement x is true. This function is given the number n (called a Godel number) in some enumeration of all of the general recursive functions. This function fn must be sound - that is, if $fn(x)=1$, then logical statement x must be true. Godel first proved the Completeness theorem - the function fn exists and can be given a particular formal description (that is, the Godel number can be explicitly stated) such that if for every logical statement x that is true, $fn(x)=1$. It is complete for all true statements.

The function is partial recursive, though. That means, that there are some cases where $fn(y)$ does not compute a value at all. One way of thinking about it is to consider the function fn represented as a Turing machine,

TMm, where TMm, given input x on its tape will, if x is true, go back into state q_0 with its tape only having a single symbol 1 on the tape and the rest blank. Therefore $fn(x)=1$ iff $TMm(x)=1$. But there will be some cases of logical statements y where TMm just goes on and on without arriving at back at state q_0 ever - it does not halt. Of course, this does not happen all the time. There are many cases y where it halts with a tape containing only blanks, or containing some tape symbols other than a single 1. That is $fn(y)\neq 1$ iff $TMm(y)\neq 1$. In those cases y is provably false.

The main result was Godel's Incompleteness Theorem. The Incompleteness theorem constructed a logical statement number a , which states the following: if $fn(a)=v$ then $v\neq 1$. That is to say, if function $f(a)$ returns a value then that value is 1. Assume then, that logical statement a is true. By the Completeness theorem $fn(a)=1$. But that means a is false, because it states the opposite. Assume that statement a is false, then. We thus have two alternatives: either $fn(a)$ does not return a value or if $fn(a)=v$ it must be true that $v\neq 1$.

Assume that if $f(a)$ converges to the value v , and $v\neq 1$. Then again by the Completeness theorem, a must be false, or else $fn(a)$ would have returned 1. But this means that statement a is false, contradicting the assumption, which was just the statement of a . Therefore, the statement must be false, but $fn(a)$ does not return an answer. Therefore, the proof method expressed by function fn is incomplete.

Actually every such proof method must be incomplete, since there was no special characteristic of the procedure fn that made it fail in the case of statement a . Every formal method fy that is sound and complete must have a logical statement z such that $fy(z)$ does not return a value. Turing later showed a related result: There is no such Turing Machine TMx such that for every value y , If Turing Machine y is presented with its own number y on the tape, then if TM_y halts on y and returns a number - any number at all - $TMx(y)$ halts and returns 1, otherwise $TMx(y)$ halts and returns a value $v\neq 1$.

The proof just defines an Turing Machine TM_y that calls TMx as a subroutine. Given some input z , TM_y simulates TMx with a simulated input of z . Then, if the simulation of TMx halts with a 1 on the simulated tape then TM_y goes into an infinite loop of some kind - it doesn't really matter as long as it doesn't ever reach the halt state. If the simulation of TMx halts with any other value than 1 on the simulated tape, TM_y cleans up the tape, overwriting it with blanks and halts.

Now consider what happens if TM_y is given the value y as an input. It passes the value y to the simulated TMx and watches what happens. If

TMx halts with a 1 on the tape, then by the definition of TMx, this means that TMy with input y halts in some finite time. But by the definition of TMy, since TMx halted with a 1, TMy then goes into an infinite loop. So TMx should never have halted with a one in the first place. On the other hand, if the simulation of TMx halts on input y with something other than a one, then that means that TMy on input y goes into an infinite loop. But actually, since TMy observed the simulation of TMx halt with a value other than 1, it cleaned up the tape and halted. So TMx could not have returned a value other than 1 and be correct.

Therefore, in this case, TMx must not halt. This is certainly permissible. Functions such as TMx can be constructed such that for all z where TMz(z) halts in some finite time the function TMx will halt with a 1 on its tape and for many cases where TMz(z) never halts, then TMx can halt with a value not equal to 1. But there must be an infinite number of cases where TMz(z) never halts and TMx(z) never halts either. At the very least, they are the infinite number of TMy functions that have one of the infinitely possible ways of simulating TMx and do the opposite.

A further result related to the Incompleteness Theorem is the Recursion theorem. The Recursion Theorem in its most basic form is a technical fixed point theorem that says that for any way of defining the family of effective procedures such as logic, general recursive functions, Turing Machines or any other method, there exists some function f_x in that formalism that, if presented with its own representation in that formalism - that is, its own number x - it will recognize itself and return that value x . That is, it is a fixed point, where $f_x(x)=x$. This is a formal notion of what it means to be conscious of oneself on a very rudimentary level. Actually, some of the most successful computer viruses spread this way, by getting access to their own computer program and copying it on other computers.

A related result is that any property of finite objects is trivial by Rice's theorem. Let any property be given. If this property is true of all numbers or false for all numbers, then it is trivial. But assume instead that the property is it is black and white, black for some objects and white for others. Then there are two cases, one true, one false. Assume that in every case it is decidable that the property can be determined for that object. Now pick two objects, one white and the other black. Redefine the two cases so that the halting problem is solved if the property is decidable. This decidable property will then allow us to solve the halting problem.

The Incompleteness Theorem and its related results have profound implications for morality. They show that you cannot logically prove all true moral statements, especially moral problems that involve consciousness as

an important feature or contain self referential moral statements related to the generational tradeoff that was presented in an earlier chapter.

For example, let us assume that someone there is a group that claims that using their holy book you can rationally decide what is the maximum state of happiness attainable for each person. Well, perhaps in an ideal world, such a methodology could be applied to everyone, and everyone would achieve their optimal level of happiness. The attainment of these levels may require an allocation of scarce resources that may result in some tradeoffs. For example, it is probably unlikely that everyone can live in Hawaii who wants to, or if they do, it would lose the character that would make it so desirable.

But this perfect philosophy is incomplete in the face of evil. Assume that there exists a person x whose goal is to prevent the maximum happiness of some other person y . If the existence of this perfect philosophy is generally known, then person x can avail themselves of that philosophy, determine what would make some person y happiest, and if it involves some scarce resource, move to acquire that resource before person y , thereby thwarting them. This would then make that ultimate state unattainable. Now person y , anticipating such a setback, could also use the same philosophy to determine the nature of their ultimate theoretical happiness and what would endanger its realization. They then have the possibility of readjusting their goals to something more attainable. But person x could also anticipate this readjustment and account for it also. This kind of infinite regress is what makes the Halting Problem unsolvable.

Note that all that is required to end up in this untenable state is a philosophical system that claims to be effectively deterministic and universal, but that there are finite resources involved that put limits on the process, and that there is at least one actor who is trying to accomplish a negative. Since the second and third conditions are quite common in real life, they preclude many cases where people make claims to universal certainty in their philosophical systems. A consequence of the Incompleteness theorem and its corollaries is that all sacred books cannot be used as axiom systems and be considered divinely perfect at the same time. This means that in general one cannot work from a holy book as if it were an axiom set. There will always be cases where the knowledge in this book is incomplete or incorrect (that is, unsound).

This fallibility means that all sacred books are written by fallible and incomplete men. They can point to or indicate a perfect or divine nature, but they can never embody it.

There is no perfect and objective measure of well-being, since that would

solve the halting problem. It also means that we take our axioms on faith - they can not be completely proved. This objection can be removed by claiming a metalogic. This metalogic can be a methodology from which we build our own system. This system must be open to error and interpretation and in a continual state of development. It is also true that there is no foundation we all agree on - we all start from an arbitrary place. We hypothesize an axiom set and then learn from our mistakes and refine our axioms. But there is no single absolute morality - each moral code is relative to our own situation.

[Learning and immune sets]

This means that a moral code is incomplete error-prone and subject to revision. This naturally leads to taking a look at the formal properties of learning. I will discuss these properties through the use of Kolmogorov complexity.

Learning can be formalism by considering the problem of learning to be one of being presented with a set of data and trying to find a theory that fits that data. Although most interesting problems are open ended - the number of potential cases to learn is unlimited - the information presented at any given time is always finite. That is, it can be written down on a sheet of paper. Let us put this data into a computer readable form. Then the expression of the data set is reduced to bits and bytes. It now becomes a finite binary sequence - a list of zeros and ones - of some finite length.

It is of most interest that the theory learned to explain the data be effective. It should be capable of being expressed as a logical expression, or better yet, as a Turing Machine that computes the data set. Since we are not concerning ourselves with function pairs, but instead of a finite binary sequence of digits that is the computer representation of the data set, we can consider that an acceptable theory to explain the data is a Turing machine that, beginning with an empty tape, runs until it halts and leaves the sequence representing the data set on the tape when it halts. We will also reduce this Turing Machine to a state table represented in a computer memory as another sequence of zeros and ones.

One of the most basic rules of learning is Ockham's razor - when deciding between two theories that each present the data, choose the simplest. Leaving alone the question of the computing time for two competing Turing Machines, we will simply choose the Turing Machine with the shortest state table as preferred over any other Turing Machine whose state table is longer than this one.

Therefore, in general, learning problems reduce to the search for the smallest Turing Machine that computes the data set that you wish to learn.

When systematically searching through a space of models, such as Turing Machines, it is obvious that if a smaller model fits, it is easier to find. There is not necessarily a correlation between the size of the data set and the model that explains it. This means that certain data represents phenomena that is easy to learn - that is, has a small Turing Machine to compute the data set - due to inherent regularities in the data set. For example, a sequence of a million ones is very easy to learn. All it requires is a Turing Machine that, starting with a blank tape, writes the number one million as a counter on the tape, and then performs the following loop: go to the end of the tape and write a one, then come back and subtract one from the counter until you reach zero. At that point, erase the zero and halt, leaving a million ones on the tape. Similarly, the first million digits of pi can be computed by a Turing Machine that is not much more complicated. Instead of writing a one at each step, what it writes is the next digit of pi.

Per Martin-Lof showed that if something is hard to learn, it is a data set whose digit sequence passes all generally known tests of randomness. In effect, it is a random number. A random number is defined to be a data sequence where the shortest representation of a Turing Machine that computes the number is as long as itself, up to a constant. Note that every data sequence has a corresponding Turing Machine that computes it. This Turing machine just has the digit sequence written in its state table. All the Turing machine does is, for an arbitrary digit sequence of a hundred zero and one digits, has one hundred states. At state 50, the Turing Machine writes the fiftieth digit of the sequence, moves one to the right and goes into state fifty one. At state one hundred the Turing Machine simply halts. So every data set has a theory to explain it, although many data sets are random. That is, the shortest expression of the data is simply the listing of the data in another form: as a state table for a Turing Machine.

It is easy to see that virtually all data sets are random sequences. For every data sequence length n , there are 2^n data sequences of that length. For example, for length 3, there are 2^3 or 8 data sequences of length three. There are about as many Turing Machines of this length or shorter. Actually, the sum is $2^n - 1$. But of these shorter Turing Machines, most of them are copies of other Turing Machines of the same length or shorter. These redundant Turing Machines just renumber the states of the copies or contain dummy states that don't do anything. So most data sequences are random.

Learning at its simplest is a brute force generate and test method that tries to find the simplest model to fit the data. That is how evolution 'learned' how to live in more and more inhospitable environments, such as

on land, in the air or in the arctic. Although there sometimes are techniques, such as drawing analogies to other phenomena that can be used to shorten learning times, in those cases where we are learning something truly novel, we have to fall back on this king on exhaustive search to discover something truly new.

Kolmogorov complexity is the study of the complexity of data sets in terms of the shortest algorithm to represent them. It seems somewhat artificial to equate learning with finding the shortest Turing Machine to compute a data set, but it actually is a reasonable paradigm for learning. Church's Thesis, the belief that all effective methods have essentially the same computational power means that we can use Turing Machines as a method of representing any other kind of effective procedure. That is, we are using Turing Machines as a way of representing the concepts used some other arbitrary formalism. Of course, it is also possible to use a formalism that simply names the data set but does not try to compute it, but this type of formalism is of limited usefulness. Learning by and large provides a method of effectively working with the information at hand and this usually implies some sort of procedure. So the representation of learning as finding a Turing Machine to compute the data set is an adequate, if abstract way of thinking about learning.

One surprising result that falls out of Kolmogorov Complexity is that for any given set of axioms there are only a finite number of cases for which it is possible to prove that we have the shortest explanation for the data. The proof goes like this: assume that there are an infinite number of proofs of the form 'TM_x is the shortest representation for data set y'. Since formal logic has a set of simple rules, it is always possible for any axiom set to create a machine that prints out all of the true theorems for those axioms. Create a machine from this one that only prints out the true theorems in that form. From this machine create a machine TMA(z) that, given a number z, enumerates the z'th case of TM_x is the shortest representation for data set y' and then runs a simulation of TM_x, producing value y: TMA(z)=y. Since not all x's and y's are included in this listing, the values of x and y increase faster than their enumeration index z. So eventually, every value x is smaller than the fixed value a for TMA plus the index z, and so TMA(z) is a shorter representation for data set y, than TM_x itself. If we have the shortest explanation for the data, then this explanation is irreducibly complex. Thus irreducible complexity is provable only for a finite number of concepts.

Because of this argument, the set of pairs {x,y} where TM_x is the shortest representation for data set y' cannot be computed by any effective procedure. This is known as an immune set: an infinite set such that every attempted

means to calculate the set effectively will be in error in an infinite number of places. In a profound sense, this set is unknowable.

Although this set is immune, and thus incomputable, it is easily learnable in the limit. It is done in a very simple way that mirrors the process of scientific discovery by an ever expanding group of scientists. Given some data set that we wish to be explained, we simply farm out more and more possible theories to explain the data to an ever larger group of people, each of which tries the theory assigned to them until they get an answer. Some in this group are assigned a theory, represented as a Turing Machine that does not halt. They will wait futilely for an answer. Others will arrive at the wrong output for their theory and must try another. But for those that match, we simply use Ockham's razor and choose the simplest theory. We just don't know when we'll have found the best if it ever happens that one of our candidate models, shorter than the current best, never halts.

These theoretical results discuss how to find a theory that expresses a finite data set. But learning is mostly prized for its ability to generalize. When the finite data set is only the list of phenomena we have seen so far out of an unlimited and unbounded number of cases, the best theory is the one that successfully predicts the specifics of any future phenomena based on what we have seen already. It is not so much the accuracy of the model in summarizing the past as how it prepares us going forward.

It is generally known that we learn from our mistakes. Too often, our success in applying what we learn does not prepare us better, only give us a false complacency that dissipates when we come across something completely original. Seen in this way, learning is like managing to walk down a dark hallway. We bounce off the walls when we learn from our mistakes, and this teaches us to aim for the center as our mistakes become rarer and rarer. In a universe whose ultimate reality is immune to any final theory of everything, we may find that it is not possible to come up with a completely satisfactory theory to predict a particular set of phenomena. The best that we may do is to come up with some procedure to identify and avoid our mistakes. In that case the best that we can do, to learn proper behavior is to delimit the bounds of acceptable, by setting the boundaries that mark off the space of the unacceptable. This gives a deeper meaning to Justice Potter Stewart's oft quoted dictum that he may not be able to define pornography, but 'I know it when I see it.'

[FOOTNOTE: Movie Day at the Supreme Court "I Know It When I See It" by Judith Silver, Esq. In 1964, Justice Potter Stewart tried to explain "hard-core" pornography, (legally synonymous with obscenity), "It is possible to read the Court's opinion in *Roth v. United States* and *Alberts*

v. California, 354 U.S: 476, in a variety of ways. In saying this, I imply no criticism of the Court, which in those cases was faced with the task of trying to define what may be indefinable. I have reached the conclusion, which I think is confirmed at least by negative implications in the Court's decisions since Roth and Alberts, that under the First and Fourteenth Amendments criminal laws in this area are constitutionally limited to hard-core pornography. I shall not today attempt further to define the kinds of material I understand to be embraced within that shorthand description, and perhaps I could never succeed in intelligibly doing so. But I know it when I see it, and the motion picture involved in this case is not that." *Jacobellis v. Ohio*, 378 U.S. 184, 197 (1964) END FOOTNOTE]

This may be the essence of learning proper behavior in society. Because the limits of human behavior are constantly expanding as the possibilities available to us are expanding, the best we can expect is to be able to quantify what is unacceptable. That is, to know what is good and beautiful may not be possible - the best we can hope for is to have seen enough pornography to know it when we see it again. As for the rest, it is acceptable if it does not cross our boundaries.

Besides learning how to predict behavior, we are also learning in the presence of noise. It is seldom that the things we are called upon to learn are expressed as a perfect set of data tables, otherwise probability and statistics would not have become such important tools. Even the most fundamental phenomena of physics, chemistry and astronomy have to deal with noise in the data.

The existence of noise and randomness presents us with the question: can reality be reasonably expressed as a set of integers? After all, we have reduced formal logic and computation down to numbers, which, although unlimited in number, are each finitely bounded and perfect. There does not seem to be any justification for this. Yet it does seem to work. We have learned how to function effectively in the world with the help of computers that convert the world down to bits and bytes, expressing reality in practical ways that allow us to some sense of mastery.

Turing's original paper on the Halting Problem addresses this problem. As he says 'the justification lies in the fact that the human memory is necessarily limited. That is to say, the functioning of the human mind requires the representation and manipulation of information in terms of a finite number of elements, each of which can be in one of a number of finite states: 'We may compare a man in the process of computing a real number to a machine which is only capable of a finite number of conditions q1, q2, ..., qR.' This means that even if the quantity is in some sense unlimited in

the degree of its specification, for the purposes of human computation this expression must be truncated at some finite place and the before human cognition can begin to work on it.

[Overfitting is inevitable]

Although we can use numbers, with their deterministic character to represent a world built out of random processes, randomness does incorporate itself in our attempts to learn about the world. This is the phenomenon of overfitting. If data is presented to a learning system and that data contains a certain amount of randomness, then there is a tendency for the learning system to see some of that randomness as a pattern and incorporate that pattern into the model. This happens for any kind of learning system, not just humans. Computers will overfit noise into their models, and so will animals, also. It has been observed for laboratory mice that have been put in Skinner Boxes where, they are given a reward for pushing a lever. If there is a random reinforcement, that is, the reward does not come every time they push the lever but randomly for every second, third or fourth time or so, they will incorporate random behaviors, such as turning round to the left before pushing the lever or other such motions. This is like the stories of baseball players wearing their 'lucky socks', or the mistaken belief of basketball players that they sometimes develop 'hot hands'. Some methods of prediction are nothing but overfitting, with no meaningful patterns, such as the field of astrology, which claims to have empirical models of future behavior built assuming the influence of distant planets on the person.

Although instances of overfitting can be refuted, they are impossible to avoid. One problem is that humans always look for explanations and they will be satisfied with bogus ones, if there is nothing better. A case in point was the attribution of stomach ulcers to diet and emotional tension. This was not refuted on its lack of merit alone. It took the substitution of an alternative model of infection by *h. pylori* before the previous model was finally discredited. We seem to prefer an inadequate explanation to none at all.

[The world is immune]

But this is what may be the ultimate nature of reality. The universe itself is just as likely to be expressible as an immune set as it is to be expressed as an effective procedure. It does appear that nature is amenable to our explanations, though. Newton's Laws of Motion and the laws of chemistry, for example, have given us great control over the world around us. But the immune nature of the universe may present itself as an onion with an infinite number of layers. Once we peel back one layer and express it in the closed form of a natural law, this merely reveals a new layer, a fine structure

that was not fully explained by this model. In that way, Newton's laws were supplanted by Einstein's, and in the future, a new model will correct it.

We are probably not even starting from the outer layer to begin with. Outside of all the layers of explanation, there is a new level of generality that extends our already known theories to larger domains. But this may lead to an even larger generality. Since the world is immune, there is likely no Theory of Everything that will last more than a generation. The fact that every effective model is infinitely wrong gives us eventually a new body of exceptions that need explaining that provide a new layer of generality, or a new level of subtlety to our calculations.

Some evidence for the indeterminacy of the universe comes from the nondeterministic processes of quantum theory. Although this is usually thought of as a completely random process, it may be the case that the universe is deterministic but immune. This would make these random events random in the sense of Kolmogorov Complexity, in that their expression cannot be summarized in a neat theory. They are not so much random events, but instead events incapable of being summarized.

In an immune universe, moral decision making is also immune. This happens in one of two ways. First, as we learn more and more about the world, our theories get more and more complicated and so do our capabilities. This means that we shall never come to an ultimate answer in determining what is the good. What looks to one generation as an obvious determination that choice A is preferred over Choice B, to an later generation who can see the world with a deeper level of sophistication, this may be true only for the most obvious cases. This is not to say that future generations will overturn what we know now to be right and wrong any more than they will say that every action has an equal and opposite reaction is false. But they will be able to point out that in certain cases, there is a subtlety to the situation that means that what looked like an obvious choice of one thing over another does not really apply, because in this case a previously undiscovered law applies.

Besides the immunity of the outside world to explanation as a single closed set of laws, we ourselves are probably immune also. This means that there is an underlying indeterminacy to our actions, This indeterminacy can arise even if we were completely deterministic machines, expressible as recursively enumerable functions, as long as we are able to tap into the indeterminacy of an immune universe. This indeterminacy is sometimes mistakenly considered to be the basis of free will, but is nothing more than an underlying unpredictability. In point of fact, we tend as a general rule to strive to eliminate these random processes from our decision making. They

tend to be a hindrance more than they are an essential part of how we think and act.

[Mysticism]

This indeterminacy can lead to a mystical view of the universe, even in the face of the efficacy of rational processes. Mysticism and logicism are not incompatible - it is perfectly reasonable to be a rational mystic.

The notion of mysticism comes from a word meaning hidden. There are two ways mysticism enters. The first way is via consciousness - the self-awareness expressed in concepts such as Godel Incompleteness Theorem and the Recursion Theorem. The second way is via the inherent complexity of an immune universe. The essential capability of the universe to achieve self-awareness allows for the ability of creatures to arise that will learn and grow in understanding, which outstrips the ability to understand, since this very understanding can result in future learning. The second level of mysticism leads to an infinite depth of levels of complexity that require an infinite level of understanding that reveals more and more of the complexity of the universe, but always leaves an infinite depth of still hidden complexity waiting to be revealed.

Chapter 9

Reasoning

A workable methodology for arriving at moral precepts requires more than logic.

The goal of this chapter is to formalize and define how to reason about what is the good. Much of moral reasoning is involved with the search for the rules to determine what is good, then deriving consequences of these rules for certain situations. Essentially, this is the process of deriving conclusions from a set of axioms. But we have seen that the limits of logic are found in its undecidability. This comes out of the ability of logic to be self-referential, leading to infinite loops or self-contradictions. This is a problem especially when the entities being reasoned about are themselves capable of complex reasoning, especially the ability to learn from the actions of others and to respond to those actions in ways that aim for different goals.

It is worthwhile to retain, as much as possible, an objective, scientific approach to morality, because it has worked so well in other fields. In fields of knowledge such as physics and chemistry, and even in areas of biology that look at the workings of life not involving human psychology, formal logic is adequate for reaching conclusions about almost all problems of interest. This is true for the hard sciences, but also for applied sciences, too. Besides reasoning about what is observable in physics and chemistry, logic helps to determine what is healthful, in medicine, and what works and is practical in engineering.

But once the ability to consciously reason is put into the mix, paradoxes can arise that make it hard to reach workable conclusions on the basis of logic alone. These self-referential paradoxes can have moral implications. I am indebted to the book *Paradoxes from A to Z* by Michael Clark for some of the examples in this chapter and the next.

To give a simple example of a logical paradox, consider the Barber Paradox, first expressed by Bertrand Russell. A village passes a law that every resident male who does not shave himself can only be shaved by the village barber, who must also be male and also a resident. The barber also cannot shave anyone who shaves himself. If the barber does not shave himself, then no other male in the village can by the first part of the law. But the barber cannot shave himself, because that leads to breaking the second part of the law. Therefore the barber must either break the law, not shave, or be shaved by someone who is not a male resident of the village. It is, of course possible to add amendments to the law that forbid other groups to shave the barber or not allow him to go unshaved, eventually reaching an absurdity. Eventually, the restrictions lead to a situation where the law refers to no admissible situation. Although there are arcane ways to get around this problem, in effect, making it a pseudoparadox - the situation such as this can have moral implications.

Can these types of paradoxes arise in morality? It is certainly possible. A practical example of a paradox is the way a placebo works. It is known that placebos, pills containing no medication, can actually cure people. People with psychological disorders, for example, respond positively to placebos more than one quarter of the time. The only problem is, they work only if you believe in them. Should I believe that it cures me or not?

A related moral problem would arise for a person who believes in voodoo. If a person claims that they have put a curse on someone, is it morally wrong if the person who claims to have put the curse doesn't really believe the claim? After all, the belief of the person who is the source of the curse is not the source of the injury. As in the placebo example, it is the belief of the recipient of the curse that determines whether the curse is effective or not, and that is not under the control of the person making the claim. Therefore they could be held responsible for an action that is not under their control, since that belief is the product of the mind of the person who believes they are being cursed. There is a case to be made that there is no moral wrong if the statement that there is a curse was made in jest, but taken seriously by the person who was the object of the jest.

This is not a strictly logical paradox, though. It may possible to create a true logical paradox by creating a situation where, for example, there are two stock-picking computer programs that go head to head for the same stock, where each program is able to get full access to the programs and data available to the other machine. Given that a buy order for the stock will result in an increase in the price of the stock, it is to the advantage of the two machines to go second. In reality neither would go at all if both

machines were to simulate the computations of the other as a means to make the decision. So a true logical paradox is hard to achieve in reality.

But still, the rise of pseudo-paradoxes is troublesome enough in themselves. For example, consider the case of a legal guardian taking power over a person. This guardian is often required to protect that person in general, but one of the things being protected is the ability of the person to do as they please. The actions of the guardian in themselves violate this protection, especially if the actions of the guardian are to optimize the person's freedom of choice.

Another example comes from those religions where humility is a virtue. Someone who claims to know the will of god is a sinner because they are guilty of hubris. But someone who does not know the will of god is a sinner because they do not know how to act with righteousness. Either way, it is impossible to completely avoid sin.

In general, the determination of what is moral or not is based on the determination of what is best for the well-being of the entity involved. But due to the fact that moral entities such as humans are conscious, evaluating well-being must take that into account. But we have shown that these moral arguments are subject to paradoxes due to this self-awareness of the entities means that there is no objective measure of well-being. In essence a pure, objective logical approach to well-being must address an unlimited variety of recursive function to compute well-being. To claim a complete measure of well-being that would, in effect, solve the halting problem, or give the appearance of that violation.

Because these logical paradoxes can arise it means that a universal pronouncement, such as 'Thou shalt not kill' can immediately be contradicted by a counter example, no matter what it means to kill. A typical response to this conundrum is to postulate that the absolute truth exists at a higher level, a "meta-level" that contains the non-contradicted essence of the principle. This metalogic could be a methodology from which we build our own moral system relative to our own well-being. This system must be open to error and interpretation and in a continual state of development. If the universe were immune, this would allow us to build a moral system in a context of a learning system that transcends the limitations of a purely recursive-theoretical process.

But this metalogic would lead to a morality that is relative to each entity, once the individual's well-being is superimposed upon it. This implies that there would be no practical foundation we all agree on, only an abstract higher level. In contrast to the absolute morality of Western philosophy and religion, the ensuing metalogic is more like Taoism - we all start from

an arbitrary place, trying to view the Tao in different ways, as best we can. In this methodology, we hypothesize an axiom set to reason our way to decisions, and then learn from our mistakes and refine our axioms. But there is no single absolute morality - each moral code is relative to our own situation.

The viewpoint of a limited consistency then brings us to the view that this attempt has not brought us higher; instead it has pushed us to a different place in the conceptual web, where a new set of inconsistencies immediately arise. Therefore, there is no ultimate level - there is just an infinite series of successive reformulations of principles that incrementally differ from the previous one, *ad infinitum*. Eventually, you just have to stay with the principle 'Thou shalt not kill' as is, contradictions and all.

Because there is no absolute truth to the practice of morality then it is sometimes unavoidable to speak in contradictions. This is the essence of a Taoist way - the recognition that any duality is inaccurate. Thus a universal pronouncement, such as 'Thou shalt not kill' can immediately be contradicted by a counter example. This means that universality is not absolute, at least in the way it is observed, analyzed and applied, and that all concepts are to some extent, oxymorons. This leads to the intellectual equivalent of the golden mean. No philosophy is absolutely correct, but since people have many characteristics in common, it is possible to define an evaluation of well-being and an ethics that uses it that is by and large correct. No philosophy can be completely correct, but there are many that are close.

Along with the observation that the ability to learn avoids the problems of a fixed absolute system, the presumption of an adaptable relative morality makes the process of developing a moral viewpoint adaptable itself. This leads to a flexibility that is actually a source of strength. A philosophical system that claims absolute consistency is fragile. An example is the moral philosophy of the more extreme followers of Ayn Rand - accept it all or reject it. This leads to a brittle system. The same problem can arise with Roman Catholicism and the insistence of Papal infallibility when making a pronouncement *ex cathedra*. An absolute moral statement of this kind runs the risk of demanding that a believer must be excommunicated even though they agree with 99% of the dogma. As a metaphor of the differing way the Taoist views truth, the example of bamboo is presented. Compared to a tall, strong tree, bamboo bends with the storm, but a tree eventually falls.

Another strength of departing from a strict requirement of consistency is that it allows for the compartmentalization of reasoning domains. This allows for reasoning in a context. The conventional opinion is that this is

undesirable, but consistency limited to a context can actually be a virtue. Typically, emotions and subjectivity provide the context to make relative choices. This changes the focus to what is relevant or most important in a particular context. As an example when the lack of compartmentalization can lead to trouble, consider the problem of dogmatic reasoning. When someone has dogma, they will force all new facts into the dogma. This happens because the person has extended a context with limited consistency into an inappropriate universal outlook. In effect, the relativistic fallacy is addressed by deciding that living with fallacies, if handled in a common sense manner, is not so bad, after all. Indeed, in real life, that is the norm.

[Law of the excluded middle is impractical]

This brings us to the Law of the Excluded Middle. The law states that in logical reasoning, there are no gray areas in deciding. Something is either true or false. The identity of an object is fixed - A is A, not B. I would claim that the Aristotelian law of the excluded middle, although a good first approximation, is not an absolute basis for understanding reality.

To continue the use of Taoism as an example of an alternative method of looking at reasoning, the Taoist knows that it is impossible to see good without evil - humans are always seeing in opposites. But this belies the essential nature of reality. Good and evil is not a duality inherent in the universe itself. Instead, they are an artifact of the way we reason. Such dualities come about because our neurons work this way. They either fire or they don't.

It is not hard to imagine why the Law of the Excluded Middle is so pervasive: logic is a tool of the mind to make sense of the world that has been remarkably successful. Therefore the mechanism by which logic is applied is superimposed upon reality. Since all reasoning is mediated through this mechanism, it appears to be an essential aspect of reality itself.

This does not deny that there do exist absolute laws or rules. For example, there are absolute laws like "1+1=2". They are absolute because they exist in an absolute universe built around the laws of logic and mathematics. But that does not mean that this world actually exists. This absolute universe is a mental construct. It is part of the universe of logical statements that gives rise to paradoxes. We can admit the existence of absolutes and therefore come up against paradoxes. Even the postulation of an absolute rule denying all absolutes has an exception in the absolute world of absolutes. Since this absolute rule exists in this world of absolutes, it becomes an example of Russell's paradox. In reality 1+1=2 is true up to Planck's constant due to the uncertainty principle. Every instance in reality of an instantiation of the notion of a quantity is open to the natural uncertainty

of the universe.

If the universe is made of integers, then the law of the excluded middle follows. If it is an integer, n-tuple or whatever then it has a boundary. This boundary forces everything into a state of being that is either that integer or not. It even tries to force the real numbers into some infinite representation of a finite sequence of digits, allowing for a countable number of real numbers, such as pi, to be represented by a finite algorithm with an infinite computation. Turing's representation of functions works because you name some function up to a finite number of states. But each naming has to stop at a level - the process could go on forever, but unlike the expansion of pi, each new rule makes the function more unpredictable.

An attempt to address this incompatibility leads to the development of the dialectic. Dialectic is a natural psychological process. It starts with the false assumption that the Excluded Middle actually exists. The synthesis of thesis and antithesis happens when people finally realize that neither theory is completely correct and truth is somewhere in between.

In a certain sense, quantum mechanics seems to imply the existence of integers. In this case, the world has an indeterminate existence until it is measured, then it collapses into one of a finite number of states. But even this attempt is transient. Further observations and interactions with the world bring the uncertainties back. Randomness seems to imply more than quantumness. It makes the uncertainty the base state of reality, and a measurement into some transient attempt to pin this randomness down, an attempt that lasts until the next interaction. We do have the ability to set up situations such as the example of Schrodinger's cat, that is alive or dead depending on a single quantum event, but that is an outlier - a fixed even in a world of more or less probabilistic events - the exception, not the rule.

Instead of demanding that reality conform to the laws of thought, we can admit that logic may not necessarily be universally applicable. An example of where this type of reasoning has entered Western thought, consider the Complementarity Principle of the Copenhagen School of Physics. Complementarity shows that some phenomena, such as the electron, manifest themselves in ways that are mutually exclusive forms of existence such as a wave and a particle. Instead of forcing the physical phenomena to be absolutely one manifestation or the other, possibly relegating the alternative manifestation to an inferior role, complementarity maintains that the two are equally valid and one should not be preferred over the other in any universal sense; the preference applies to the context of the phenomenon.

Complementarity can be also true of the identification of moral characteristics. Any action can be any combination of good and evil, right and

wrong, helpful and destructive, and in the real, nonideal, world, every action is all of the above if you look long enough. The Taoist would argue that it is better to think of the person as part of a river of experience. The flow is sometimes good, and sometimes evil, but any direction is dependent on the flow of river and the goal the person is trying to achieve.

The ultimate lack of universality of the Law of the Excluded Middle may be especially relevant if in fact the world is immune as was discussed in the previous chapter. The truth may be that no matter how many laws we have and how well they work, reality is ultimately unknowable.

Related to the law of the excluded middle is the law of identity. Having some property as either true or false makes us identify an object with these properties and exclude other objects where the association of one or more of these properties is false. Thus the law of Identity gives rise to definitions. Definitions of words are in turn defined by other words. This gives rise to semantic nets. A semantic net is a collection of facts and concepts that are related to other concepts. For example, the fact 'a whale is a marine mammal' is a fact in which the object whale is related to the concept mammal by an ISA link, with a link to the concept marine that indicates that it is a qualifier of this ISA link. The concepts mammal and marine each have links to other concepts, such as water and animal.

Seen this way, definitions are operations in a network of words and concepts. Definitions are the local contexts in the web. As new definitions are being acquired it is impossible to have a perfectly consistent knowledge space a priori. They must be checked.

The attempt to preserve the Law of the Excluded Middle sometimes leads to taking excessive pains to preserve an absolute consistency. This has two avoidable costs. The first is that the attempt to preserve consistency forces an inflexibility upon the conceptual system that makes it difficult, if not impossible to incorporate new or changed knowledge that contradicts some aspect of the cognitive structure that has already been established. The second cost is the simple expenditure of time and effort that it takes to preserve this consistency. As the number of concepts grows linearly, the number of interconnections that must be checked grows exponentially. Eventually, there comes a point where this effort must be curtailed or the system cannot learn in a reasonable time.

In computer science the breakdown of this effort arises in the creation and maintenance of semantic nets. The addition of each new fact requires comparing it to each other fact in the net. This is simple for the second concept, but the one millionth concept requires a cost a million times greater. For example, if the fact that mammals have four legs has already been en-

tered into the semantic net, consistency would have to be re-established by making an exception for marine mammals, or by indicating that whales only have vestigial legs. This effort can be limited by only requiring consistency to the objects in the semantic net that are within a fixed number of linkages between one fact and the next. The consistency between facts about whales, mammals and animals may be required to preserve consistency, but consistency with facts about plants and bacteria may be waived, even though certain properties between bacteria and whales may present consistency problems. Both bacteria and whales can have individual facts attached to them relating to how they react in the presence of oxygen and each could be a statement inconsistent with the other, but a general statement about how animals react to oxygen might be consistent with both. If the number of connections between one object and its neighbors is limited, then the requirements on consistency can be bounded, even though the number of concepts grows without bound.

This points out the fact that although formal definitions are a useful tool, they cannot be absolutely applicable. Formal definitions cannot take the place of common usage, because common usage enshrines implicitly the recognition of what consistencies are important. If a consistency is not established explicitly, it does not necessarily mean it is not important, but precedence must be given to the explicit consistencies over the implicit. This means that ultimately we have to use a utilitarian definition for our concepts over one that is of formal rigor.

This is an important distinction to make. The creation and use of knowledge is a practical art, subject to usage and not required to preserve a mathematical rigor. The rigor of abstract mathematical concepts are a luxury that can be allowed in the creation of a formal representation of knowledge about the real world, just as a map is a formal representation of a territory. But the map can never be the territory - it can only be an abstraction of it. Inconsistencies can arise from the selection of what details are abstracted away in the creation of the formal map.

Letting go of the claim to universality of formal logical validity opens us up to belief systems that contain contradictions - a situation in logic that leads to the complete breakdown of the system. How can one prove the correctness of a claim that violates the excluded middle, thus leading to these kinds of contradictions? The answer is, you can't - not logically. One has to use a non-logical method, or to extend or alter logic in some way.

A simple fix is to preserve logic for the most part, but limiting the context of logical conclusions. In that case, for every use of logic we must define a local context and require soundness in that context. This means that any

conclusion that is drawn is not absolute but relative to a particular universe of discourse. This can reduce the problem of paradoxes, especially by limiting the universe of discourse in time as well as space. In that case, reference to a single entity can be qualified by a conditional such as 'knowing then what I know now' where what you know can be determined well enough to draw the necessary conclusion. Most paradoxes implicitly assume some sort of infinite regression, a situation that has not been shown to exist in reality. Providing some sort of limit that removes infinities from consideration, if done carefully, eliminates most paradoxes.

Most other paradoxes come about by presuming a limit, although an undefined one. Also, most countable infinities, that is, infinities involving every number that can be potentially written down can be bounded by existential quantification - that is, instead of saying 'for all t ' you can often reformulate this as 'there exists an s such that for all t less than s '. Let's consider how to convert an infinite regression into this other form.

Consider Zeno's paradox of Achilles and the tortoise. The tortoise gets a head start of 8 feet. In one second, Achilles covers 50 feet and the tortoise has barely moved. A half second later, Achilles has covered 2 feet. Dividing the interval in half, the infinite regression goes forever, without the sum ever amounting to the starting distance from Achilles to the tortoise. But given any finite interval, there is a specified time that Achilles crosses that interval, and for any such interval, there is an approximate time at which Achilles passes the tortoise, which serves as an upper bound to this time. Now the question is to determine the smallest such interval. Of course, if there is no limit on the smallest time, this infinite regress has been replaced by a fruitless search to find an infinitely small infinitesimal.

In most realistic cases, there is a recognizable limit. For this example, the length covered by a photon in Planck's constant is the shortest time and distance. For human decision making, the shortest reasonable time is often the time it takes for a neuron to fire. For the longest time the current age of the universe is acceptable, or twice the typical human lifespan when thinking of an individual. This makes most infinities of interest into finite universes.

But this bounding can lead to unacceptably large spaces, even though they are finite. For example, it is possible to calculate the number of ideas every possible person on earth could ever have, given the size and age of the earth, and different physical characteristics of every typical human. In these cases, we may have to resort to heuristics to find contradictions and resolve them.

It is not even necessary to rely strictly on logic. People use other methods

of decision making, such as a reliance on intuition or the emotions. Emotional reasoning does not require the excluded middle. Emotional reactions to people and objects overlap - they are not mutually exclusive. Emotions have their limits, though. They lead to simple, unverifiable results, which can be shown to have limited validity as later experience comes along. This is because the emotions are built on past personal experience that may or may not be generalizable to future experience. Emotional reasoning must be extended by experimentation and learning to moderate the limits of past experience. Evolution built emotions into us because of their survival value, but provided the rational abilities to ride above them, a later development that extends and increases the applicability of our first impressions.

An analysis of the limits of logic does not necessarily lead to modified formal methods of logic such as paraconsistency and dialetheism. There is a difference between the attempt to change the formal rules of logic and the recognition of the limits of logic. To use an analogy, the creation of paraconsistent logic is like modifying a simple tool to make it useful under more situations. It appears that the ultimate goal of paraconsistent logic is to modify the rules of logic in an attempt to better match the way the world works. In effect, it takes the map and tries to make it more like the territory. But this can only go so far, then the map becomes too unwieldy and loses its value as a map.

It is preferable instead to recognize that, since logic is a mental construct, there is no guarantee that it is universally applicable to the world around us, and that any attempt to make it so will lead to failure. Instead, it is better to qualify the use of logic by trying to identify its limits. In those situations where the Law of the Excluded Middle is too restrictive, logical reasoning can bring us part way to a solution. Sometimes it takes more than one tool to finish a job.

[Statistics reasoning as an alternative]

For an objective method of reaching conclusions that can be an alternative to a purely rational reasoning, without the problems of emotional reasoning, probably the best method known to date is the use of statistical analysis of populations of similar cases. There are a variety of reasons why using statistics to make ethical judgments is worthwhile. For example, even though individual cases are subject to incompleteness problems, this is often moderated in populations because individual actions tend to average out. Even when individuals are acting in a manner that makes the individual analysis difficult, the statistical analysis of populations can bring out trends. Other advantages of statistics come from the regularity of populations. The Central Limit Theorem allows us to use the Standard Normal distribution

in the limit, even though the analysis of smaller populations can be more complicated. Even cases where the statistics cannot be effectively analyzed in detail, the statistics can usually exhibit some noneffective but approximate order. This helps us in the formation and modification of hypotheses. As time goes on, the probability that the confidence of the hypothesis tends towards certainty will itself tend toward certainty.

Statistics are especially helpful where the Law of the Excluded Middle is not practicable, even though the analysis of probability distributions is almost always done with the tools of standard truth-valued logic. Fuzzy logic is one of the formalisms that tries to marry statistical probabilities to logic.

Because the world is so complicated, you can't be absolutely right or wrong unless you are omniscient. This means that a practical analysis cannot classify things absolutely. Thus, something like statistics are needed to add approximations to the analysis and reach conclusions that, although they are not absolute still have practical validity.

An example of this in physics is the three body problem. Even in well-known fields such as gravitation, approximate analysis is necessary. Although the two body problem has a simple elliptical solution, adding just a single third body makes the problem unsolvable in a closed form, except for the simplest of cases, such as the Lagrange points.

Three body problems arise in morality as often as they do in physics, but in an even more complicated multidimensional form. Since human beings have controlling forces and interests that do not solve simple equations like the inverse square law, it is impossible to determine most common day-to-day moral positions from first principles. It is necessary to either make simplifying assumptions, or observational studies or simulations. Thus the need for statistics.

But even this is not enough to be usefully right or wrong. A practical use of statistics requires that one balance the cost of a false positive condition versus a false negative. These costs in turn are offset by the benefits for the correct positive and negative choices. A false positive is an erroneous determination that an object has a certain property when it does not. A false negative is the opposite error. In this case it has been missed that the entity possesses that quality. To give an example of a medical diagnosis, a false positive is a test that shows a patient has signs cancer although there is no cancer present. A false negative is to miss a potentially dangerous cancer tumor.

An advantage of statistics is that it provides a measurable frame of reference to a moral relativity. Statistics triangulates a moral relativity.

Although taking many measurements does not make an absolute morality, it does instead determine the relative frame of reference of the moral act. It both quantizes the frequency of a moral action and the cost of it.

Another advantage of a statistical approach is that it gives a way of handling the inevitable gray areas that arise in practical ethics. Logical analysis does not have clean way of handling what is known as the 'heap paradox'. Given a heap of sand made up of 1,000,000 grains, if you remove one grain of sand, you still have a heap. But a single grain of sand is not a heap. Not even ten grains of sand are a heap. At what point does a heap of sand become a heap? This decision is made in the context of how a heap is defined in this case. One can then set two limits. A number of grains of sand above the upper limit will certainly be considered a heap of sand. A number of grains of sand below the lower limit will certainly not be considered a heap of sand. In between, we can assign a probability, or just be vague and say that this is 'possibly' a heap.

Statistical reasoning brings its own paradoxes also. One of the paradoxes is Simpson's paradox. Consider the case of two hospitals, where A and B have differing cure rates for diseases X and Y. Assume that 350 patients with disease X were admitted to hospital A and 7 of them are not cured. That is a 2% fatality rate. Out of 650 patients with disease Y, 15 of them die. Although there are twice as many fatalities, almost twice as many patients with disease Y die, so the fatality rate is 2.3%, slightly worse than that of disease X. On the other hand, of 650 cases of X admitted to hospital B, 40 die - a 6.5% fatality rate. Of 350 cases of Y, 28 die - a fatality rate of 8%, which is the worst of all. It appears that no matter which hospital a patient with disease Y is admitted to, they fare worse. But consider the fatality rates over both hospitals. In both cases there were 1000 cases of X and 1000 cases of Y admitted in total. But the total fatality rate for disease X is $7 + 40 = 47$, compared to the total fatalities disease Y of $15+28=43$. When both hospitals are compared together, patients with disease Y come out better. How can this be?

The answer is that the patients with disease X are being admitted to the hospital with the higher overall fatality rate. In hospital B, 1000 people are admitted, and 68 of them die. In the case of hospital A, only 22 fatalities are recorded. The lower overall fatality rate for disease Y reflects the fact that they are going to the hospital with the better outcome. This may be a reflection of the better quality of the care, or that the people being admitted to hospital B are sicker than the patients in hospital A. There just isn't enough information available to explain the discrepancy. But it brings up the point that statistics can be incomplete, and this can lead to

paradoxical situations.

Another type of statistical paradox comes from the problem of false positive errors versus false negatives in statistical analysis. Clark refers to this as the Xenophobic paradox. I will use his medical example of false diagnosis. Consider the case where a disease occurs in 10% of the population and can be identified with 80% accuracy using a particular diagnostic test. If a population is tested for the disease, and an individual is told that the test shows that they have the disease. It turns out that there is only about a 30% chance that this is in fact true. Out of 100 people being tested, 90 are disease free, but since the test has a 20% error rate, there are 18 false positives. Of the 10 people with the disease, there are 2 false negatives where the disease is not identified and 8 cases of the disease that the test. So there are 18+8 positive tests, but since the false positives outweigh the true positives, only 8/26 or about 30% need to be treated. This is not a problem with statistics per se, but with the incapacity for human intuition to correctly judge the chance of the false positive being confused for what are actually rare events.

The paradox of the ravens is another case of a statistical paradox involving confirmation. To claim that all ravens are black as a scientific hypothesis, we would look for confirmations in the population of ravens to support or disconfirm the hypothesis. The claim logically translates to 'Nothing that is not black is not a raven. Therefore the observation of a white swan would logically confirm the hypothesis that nothing that is not black is a raven, but it does not say anything about ravens. The problem seems to arise from the fact that there is a much larger space of things that are nonravens than ravens, so the observation of a nonblack nonraven says very little about ravens. It does give some confirmation, but the amount of confirmation is vanishingly small because the sample size of nonravens is so much larger than the set of ravens. In a sense, this is related to the heap paradox. It takes a heap of observations of nonravens to balance any direct observations of ravens, so that it does not appear that confirming the hypothesis by looking at nonravens is very productive. It is possible to say that one has looked at the universe of all things and found that every case of a nonblack object was also a nonraven is an inefficient way to test the hypothesis. It takes a lot of observations of nonravens to make a heap of nonravens that gives a similar confirmatory weigh to a set of observations of a bunch of ravens.

Some statistical paradoxes can be downright strange, even when they are explained. Consider a case where there are three adjacent regions that are being explored for oil. It is known that the oil is down there, but not well enough to determine if it is under region A, B or C. Assume you choose

to drill in region A, but before you start, a competitor drills in region B and comes up dry. Does it make sense then to switch drilling to region C?

Paradoxically, it does make sense to switch to C. Region A had only a $1/3$ chance of coming up with oil. But, knowing that region B came up empty, choosing to switch to C now gives you a 50-50 chance of getting the oil. Therefore, you have a higher probability of success.

The trick to understanding this paradox is in the nature of what it means to switch. Consider that, instead of just switching to C, you instead decide to switch based on the flip of a fair coin. Then you have a 50-50 chance of switching, not to C, but actually to your original choice A. The act of flipping the coin and choosing on the basis of the switch is the same as switching to C, but half the time your switch takes you back to A. But in this case, you have made a second choice of A, this time with a better likelihood of success. The paradox is resolved by identifying the psychological quirk that switching from A to A is not usually considered a switch, but in this case it actually is an act that changes your chances.

There are other statistical paradoxes, some of which are esoteric. Bertrand's chord paradox points out the problem of sampling. Even though you are randomly selecting samples from a population, there may be a bias in the way the samples are selected that can lead to vastly different results.

It seems counter productive to leave pure logic as the predominant reasoning tool because of the natural paradoxes that arise in its application to moral reasoning, but to advocate as an alternative the use of probability to reason with, when it has its own set of paradoxes, some of which are equally intractable. The answer is that we are not leaving logic behind so much as adding statistical reasoning to our toolchest. Sometimes logic is sufficient for arriving at an acceptable answer to a moral problem. But, due to the problems given here, we cannot rely solely on logic to arrive at an answer. On the other hand, the weaknesses of statistical reasoning are sometimes amenable to a logical analysis. Therefore, we are at our best advantage when we look at a given moral problem with an eye to which tool would give the best answer and pick accordingly. People have a tendency to create false universals. In this case, no one tool works for every case. Sometimes we choose one and sometimes the other.

[Definitions]

To return to the topic of definitions, probability theory gives us a different way of looking at how definitions are created and used. In a formal system of absolutes, definitions are the axiom set that forms the starting point of reasoning. In an absolute sense, the definition consists of the properties that the concept contains or does not contain. Instead, using a statistical

approach, definitions are given by consensus usage. Seen as a statistical process, the meanings of words are operational and based on common usage - they are not absolute. Meaning comes from observable phenomena. Observations are not black and white.

Because of this, definitions grow and change as knowledge grows and changes, but within limits. The definition attached to a word is dependent on the population that uses the word. In fact, the definition can change for different populations. Definitions can have no absolute definition outside of the population using the term. This definition has limits. It is common for people to have different meanings attached to the same words. Some of those differences can lead to violent disagreements. These disagreements must be solved democratically. An individual or group may shift a definition, but cannot go too far or others will not agree to follow that change

Keith Stanovich in his book 'How to Think Straight About Psychology', that inspired and informed some of the thinking in this chapter and the next, points out that the creation and use of definitions in scientific reasoning fits this statistical approach. He mentions that meaning changes as the operational definition changes, since this changes the relation from the definition to the observable phenomena. This means that definitions depend on consensus, because this is required for sharing observations.

A consensus approach means that the definition contains the properties that are deemed most often to occur. This set of properties form a boundary in which the objects fitting the definition tend to fit inside the boundary of properties. If a property does fit to some extent or is completely absent, the object may still fit the definition if the rest of the properties fit the consensus. A definition is therefore useful in the degree in which the boundary is such that the objects all fit well.

Since any axiom set is incomplete, definitions must grow and change. A definition is thus some sort of generalization. It starts from pattern matching. Logically this type of definition is self-referential. As you learn more the definition changes. This extends the set of items matching the pattern as the generalization changes. Some items can later be dropped.

Definitions are, and always have been a social construct. They have meaning relative to the society that created and uses the definition. As the society changes, the members of that society use the words in different ways, thereby changing the definitions by usage. This allows for the meaning of words to change in small increments as time goes on. But if a single person uses a word in a significantly different manner, that person will not be understood unless a coherent subgroup is built around that changed usage. Once that subgroup grows either in size or influence, the new definition will

become, sometimes suddenly, the current meaning of the word.

It has often been remarked that philosophical disputes quite often are disputes of definitions rather than something more substantive. The dispute may claim to be about different ways that a certain group should operate in a certain situation, but because the definitions are different, the applicable groups could be different. It might even be possible for the different parties in the dispute to agree on the rightness of a certain action once they have agreed on the same group. But since the definitions are different, the underlying agreement about actions might be obscured.

One way to resolve these differences is to distinguish between the qualities denoted by the definition from the qualities expressed in the word's connotation. The distinction between the qualities of an object that considered for it to fit a definition versus the other qualities that it connotes is well known. Connotations that have been associated with a definition can have both the problem of not being sufficiently universal or only applicable at the present time. They are subject to the problems of overfitting in learning. We must be aware of the emotional content of definitions. The emotions recognize and associate properties with the definition they denote, but the emotions act as the selector in a search space - a preliminary choice that the more rational parts of the brain must judge, upon reflection - to be worthy of keeping or discarded. The emotions also pull in a framework of related things that carry a similar emotional load.

But attempting to reduce a definition to a formal construct removes the definition of its power. The practical use of definitions cannot be reduced to a formal logical game. The reason we define things is in fact to provide that framework for us to react to the thing defined. Therefore, the emotional content cannot be dispensed with. It just must be used judiciously

[An answer to Hume]

A statistical approach to moral reasoning also has a bearing on Hume's observation that one cannot derive an 'ought' from an 'is'. Hume's argument is that morality has a deontological basis - morality is not supposed to be contingent upon the achieving of any type of goal. Therefore rationality does not rule the passions - you cannot derive an 'ought' from an 'is'.

Hume's approach to morality seems to lead to a moral code that is essentially arbitrary. He says that "Where a passion is neither founded on false suppositions, nor chuses means insufficient for the end, the understanding can neither justify nor condemn it. It is not contrary to reason to prefer the destruction of the whole world to the scratching of my finger. It is not contrary to reason for me to chuse my total ruin, to prevent the least uneasiness of an Indian or person wholly unknown to me. It is as little contrary

to reason to prefer even my own acknowledged lesser good to my greater, and have a more ardent affection for the former than the latter.”

Hume goes on to say ”Reason is, and ought only to be the slave of the passions, and can never pretend to any other office than to serve and obey them.” He is right - rationality does not rule the passions. In fact, it is dependent on them for the source of the objects of reason. But then we must ask the question where do the passions come from. Both humans and animals have passions. But passions are the driving force behind any entity with the power to make decisions. It could also be argued that an Artificial Intelligence could be imbued with passions, even a simple one. The source of the passions in a human or animal is in their genetic makeup. In a computer program, they are part of the goals that the system is programmed to accomplish. That is, the passions make up an entity’s nature. Without passions, an entity would have no character at all - the entity may process sensation the same way it processes food, but nothing would come of them.

But where does what we ought to do come from? Does what we ought come from the realm of what God intends - that, some external agent that defines the goal of the entity? It does not matter that we ought we do what is right because we want to. We also do what’s wrong because we want to. The oughts arise out of our evolutionary nature - out of our programming. The male peacock ought to strut. We ought to raise a finger if it improves our survival. Sociobiology shows ought can be social.

What this means is that what we ought to do actually does come from what is, but the ‘is’ that is important is the nature of the entity for which the question is asked. The oughts arise out of the nature of the entity and can be determined by an inspection of the makeup of the entity - its nature and the needs that it requires. These needs are somewhat different for each individual, but can be given generally by a statistical analysis of the population. They can be further refined by looking at populations with similar qualities to that of the entity in question. Therefore, what an entity ought to do comes from the nature of what it is. This determination is made not just by the analysis of the individual in isolation, but as part of a population of similar entities.

This also addresses the mind-body problem. There is no such thing as a pure mentality - there is no pure mind. Similarly, for any entity capable of moral action, there is no pure body. Mind is not an epiphenomenon, but emergent behavior. It is not a separate duality, but a process that arises out of and is embodied in a physical manifestation. It is not directly observable, in the sense that physical properties are observable. It is inferable by the use of empathy - the analogous reasoning that compares the nature

of the observer to the observed. Here again, a statistical approach can be useful. Generalization over a set of populations can arrive at the essential characteristics of the mind.

[What is the good?]

The matter of definitions brings us back to the question of what is the good. The good was considered as almost axiomatic in the first chapter. This is so, but the good is different in different times and different places. More important than being axiomatic it should have measurable dimensions. What is considered to be the best for the well-being of an individual must be cast in a form that a statistical analysis can be performed.

This does not entirely remove qualitative features, such as asking whether a person is happy or not. These qualities cannot be meaningfully be quantified, but they can be compared to previous states to give a simple better-than or less-than answer to the subjective judgment. These types of comparisons are easily capable of analysis either by logical means or by statistical analysis.

But it is not possible to give an absolute definition of well-being that applies to all time. Definitions are reached by consensus and change when the consensus changes. Five hundred years ago, the state of one's immortal soul was more important than the condition of the body. This has changed. With the changes in technology since the Industrial revolution, there are a variety of factors than go into an estimation of well-being that never existed before, and there will be more to come that we cannot imagine today. It is not possible to abstract a core set of qualities that are timeless, because we are not omniscient. As knowledge increases, the dimensions by which well-being is measured will increase also and any attempt to collapse this growth in dimensions into a number of basic qualities runs the risk of coming up with a set of qualities that are increasingly irrelevant. It is like the construction of language. The letters of the alphabet had an original meaning when they were first given, but as the language developed, the meanings became irrelevant. And there is no way to come up with a fixed number of qualities in the way that phonemes describe the words of a language, because human welfare is not predefined the same way as the vocal tract is.

Again, it is important to remember that this statement of well-being must be used in a relative sense. Measures of well-being that are relative to recent developments, such as access to electricity are changeable and may be irrelevant when electricity is superceded or universally available. Other measure of well-being, such as happiness or satisfaction are more timeless, but are measured differently at different times. Basic parameters such as life and death, or even expected lifespan have a more timeless component. But

even this is capable of change, for example, if a true immortality is achieved.

Since we discussed in the previous chapter the mathematical aspects of learning, we can apply the concept of learning to the problem of the relative applicability of definitions of the good. Is it possible to create a moral code that enumerates over all moral codes and looks for improvements or generalizations that increase the applicability of the moral code to more times and places? The answer is that this can increase the boundaries of applicability, but will never reach the level of absolute generalization. Assume the universe is immune. Then the procedure will not halt on an absolute rule since the actors themselves are learning also.

In summary, the measure of well-being used to define the good changes over time and distance relative to the local consensus. No dimension of this measurement is absolutely timeless, but some have a greater degree of universality than others. But it is not true that the more timeless the dimension is, the more important it is. The relative importance is determined by the consensus across the dimensions of space and time one is comparing well-being. And some of the qualities applicable only to this region of interest may be very important indeed.

Chapter 10

Deciding: Population and Probability

Even though morality is relative, it is objective, so actions can be empirically tested to determine what is moral or not.

The final chapter will discuss how to apply statistical reasoning to ethics and some discussion of some of the important moral issues of today. The material here is heavily influenced by the book "How to Think Straight About Psychology" by Keith Stanovich. It forms the basis of the first part of this chapter.

One movement in moral philosophy in the last century or so is the concept of moral non-naturalism - that morality cannot be reduced to natural properties. It appears that this type of morality entails a belief that there are qualities of morality that are not addressable by naturalistic means. They form part of the human psyche, since they act as the motivating forces for people to perform moral acts, or at least give a moral perspective to potential acts. But this implies that psychology must have a non-natural component; psychology cannot be explained naturally. Although this is perhaps possible, it begs the question. Perhaps there are qualities that a naturalistic explanation cannot explain, but the postulation of such characteristics leads to the creation of a set of new dimensions in which these qualities can be said to hold, which can be addressed through an extension of naturalism. We are then led back to a naturalistic morality.

If any one comes up with a theory of morality that makes claims about what is good or bad, it must be capable of experimental verification. Otherwise, there is no connection between the morality and what is good. Claims that cannot be verified are arbitrary, and arbitrary claims cannot be rea-

soned about.

One thing that we strive for when building a conceptual framework for a body of knowledge is the principle of connectivity - every new theory must account for all old data as well as the new. This makes the overall theory into a semantic network where the interrelationships can be traced from one area to another, or alternatively from the general to the specific and back. That does not mean, though, that the new theory need be a gradual development from previous theories or even that the theory contain a single monolithic set of concepts. Kuhn's identification of paradigm shifts in scientific development shows that the new theories can be radically different from the previous ones. A theory can also have different methods of reasoning at different levels of abstraction or in different specialties and still retain connectivity. The science of biology has connectivity even though the principles of biology cannot all be reduced to molecular biology. The methodologies in ecology of populations, in evolution or taxonomy all retain their own qualities. The addition of new knowledge in areas such as molecular biology augments these methodologies but do not supercede them.

It is important that these theories are more than a set of practical applications or observations. In building a moral framework, it is necessary to have a conceptual structure of definitions and theories that tie the framework together. Personal experience is not sufficient to lead to a fundamental understanding of the world. This is true in every field of thought. Before the laws of motion were discovered. Even common everyday events were misinterpreted. For example, in his *Physics*, Aristotle makes the claim that "a weight which is twice as great will fall from the same height in half the time."

Common preconceptions lead to erroneous conclusions about even basic physical laws. An example that Stanovich mentions is where people are asked to reason about the situation where a person is twirling a weight above their head and then let go. The intuition of the majority of people lead them to predict that the weight will move directly away from the person, instead of following the tangent at that part of the circle the weight was following. With Newton's *Principia Mathematica*, the laws of motion were abstracted to an extent that we could make accurate predictions, and were not led astray as often by mistakes in our intuition. The intuitions lead us even further astray with statistical reasoning, especially because the laws of probability are so new that they have not yet become part of the basic human culture the way the basic laws of science have. The laws of probability are recent - barely more than a century old in most cases even though games of chance are old. The intuitions we have built around these games will take a

long time to be superceded.

Any theory has two parts to it - a theory of knowing and a theory of doing. This is known as procedural versus declarative knowledge. The declarative knowledge tells us what we know about the field. The procedural knowledge tells us what to do in the field. Procedural knowledge cannot be reduced to declarative knowledge or vice versa. Even in something as abstract as mathematics, there is a vast body of knowledge of the definitions and theorems in each specialty of mathematics, but there is also the often unformalized, but still equally important, body of knowledge and techniques that tell us how to reason about these concepts. For example, formal logic has a number of results about reducibility relationship between problems because one of the most successful techniques to arrive at a new result is to reduce an unknown problem to a variant of an already known problem and to discuss the differences. But this can lead to the case of the drunk looking for a set of keys under the streetlight. The problems that are solved are the problems that the techniques allow us to be solved. This illustrates the importance of ensuring that procedural knowledge has a learning component. We must always be willing to try new things when we come up to the limits of our declarative knowledge. One of the hallmarks of an expert practitioner of a field is that they can maintain the balance between knowing theory from the books and knowing when it fails and it is time to try new things.

Almost all fields of knowledge except the purely theoretical (such as mathematics) have to be able to reason under uncertainty. This means that there must be a statistical component to the reasoning. Although this statistical reasoning is sometimes due to an inherent uncertainty underlying the forces involved, this is mostly not the case. When building a body of knowledge in a field it is important to recognize that the analysis of something due to chance does not mean that it is completely random. It may simply be that we choose to remove certain deterministic components from the analysis because they are irrelevant and treat them as random factors. Also, it is important to recognize that, quantum mechanics to the contrary, situations that are treated as random means that they are currently indeterminate not that they are forever indeterminable.

The basic claim in this book is that an absolute morality is insufficient to reason effectively about the world, but that a moral relativism that does not incorporate moral absolutes cannot achieve a coherent theory. In creating a methodology of ethics, it is necessary to come up with testable predictions to distinguish between moral relativity versus moral absolutism and pure relativism. When we get to specific moral situations later in the chapter, we

will provide specific cases to measurable predictions.

One family of testable predictions of ethical questions that distinguish between the correctness of moral relativism versus the other two theories is the following: the degree of relativism to a moral choice (for example, bigamy versus abortion versus murder) is correlated with the distance of the choice from the absolute best. This requires that we look at the underlying conditions that make societies different, to tease out a measurable metric to quantify the degree of relativism. In an absolute morality no such distinctions would apply. Similarly, a pure relativism would lead to the incapacity to distinguish between degrees of transgression at all. A moral relativism instead tries to capture differences between cultures and situations to find relative degrees of moral acceptability.

This is somewhat of an overstatement, though. Actually, a moral absolutism has a distinction that there is more leeway the less severe the transgression. But there still is a sense of rightness and wrongness to an action that tries to establish a bright line between right and wrong. Tolerance of deviation from the absolute moral standard then does not mean that there is a gray area as much as it means that the effort to correct the deviation can be proportional to the degree of injury the transgression involves.

Instead of a fine line between right and wrong of the moral absolutist, moral relativism recognizes that there is a gray area that is proportional to the distance from the moral absolute. This sets up a metric on the characteristics of the cultures being compared based on the relevance of their differences to moral choices, rather than, for example, differences based on climate, population density or access to resources.

How does one compare? By negotiating a standard. What standard? One that can be agreed on by all parties. What if you cannot agree on a standard? You cannot compare. For example, a religious fundamentalist requires absolute consistency with their sacred scripture, whereas a secular humanist requires an adherence to logic or rationality. There is no standard of comparison. The differences in the conceptual basis between the two parties means that if they both hold to their basic principles, even if one compromises on derivative characteristics - there is still no basis for comparison. Both parties must use their own standard of comparison. But does this mean that we cannot judge what is right or wrong? Yes, you can, by using your own standard. This is meaningless to the other parties, though.

To some degree, a comparison is possible, even without an agreement on a standard. Every religion believes that their moral code leads to better behavior in the people who practice it, and therefore a better world. These positive changes (or negative changes, in the case where their moral codes

are not followed) can be measured and compared analytically using their own standards. On the other hand, the correctness of sacred scripture can be judged by looking at how well the sacred scripture's picture of a just and good society compares to the actual society. Unfortunately, the degree of symmetry is violated in that there is less chance of compromise on the admissibility of religious faith in the comparison. If the recognition of this principle is demanded, it appears likely that strict rationalists will have doubts about the comparison, even though the belief that rationality holds is in itself a form of faith.

Before getting to specific situations, it is necessary to talk about the methodology of moral reasoning. The claim from the last chapter is that this reasoning heavily involves the use of statistics. One thing that a statistical analysis would de-emphasize is reasoning about 'lifeboat' cases. These are situations that are extreme or artificial that are presented to illustrate or reason about a particular point.

The lifeboat case also finds use in the physical sciences. There it is known as the gedanken experiment, or thought experiment. Its purpose is to reason about a particular situation in a way to arrive at a way of looking at the world - a methodology, or an example of the essential reasoning underlying a set of related phenomena. But it is used only as a conceptual framework. It plays no further part in the validation of the theory. The solution of the gedanken experiment does not validate the theory itself, even though it may give confidence to the theoretician in its applicability.

Similarly, in moral reasoning, the analysis and solution of a lifeboat problem has no real bearing on the general validity of a moral principle. To make a good moral judgment, you need to fit the case into a population of similar cases to determine what is best. For a given unique case, no matter how well reasoned, the choice is arbitrary due to the uniqueness of the case. The lifeboat example resolves itself through statistics. This demonstrates the degree of applicability of the example to the universe of human experience.

Instead of stopping at the archetypical example, it is necessary to find enough cases to make a statistical sample and base our verification on that. If enough cases cannot be found, then we have to relax our conditions to allow for a larger population and see if the more general situation can be validated instead. If there are too few cases where the determining factors are in effect you can make no moral judgment of any assurance.

For example, think of the case of the runaway train, sometimes used in moral analysis as an archetypical case. Take a hypothetical situation of a runaway train bearing down on 10 people. A moral agent, someone capa-

ble of making a decision affecting the outcome, is standing near a railroad switch. If this agent switches the train to another track it kills one person standing there and saves the ten. Studies have been made on how people would choose in this situation. These studies do give valuable insight into how human psychology operates in making decisions of this nature. But ultimately it may not give as much insight as one would like into the moral principles involved and their outcome.

It is more important instead to do studies of actual cases of this nature. Some of the questions we can answer are: did the actions that were taken save lives? Were the lives negatively affected? What happened to the person who pulled the switch or not? How did they feel about themselves? How did others treat them? These answers give insights into the real effects on well-being that are involved in cases of this kind. They also help to indicate the nature of the discrepancies between human reasoning and the theoretical calculus. People seem unable to switch the train even if it saves lives. This may be due to a built in part of human psychology that implicitly recognizes that this kind of sacrifice does not actually yield gains in the real world. Although the suspicion is that humans are not making an optimal choice and that real cases arise where the taking of few lives saves the lives of many, a study of a population of such cases is required to verify this.

But there are also problems that arise with statistical reasoning instead of reasoning from cases. People have a variety of misconceptions about statistics. These are not just statistical paradoxes, but cases in which human intuition has not been organized in a way to handle statistical reasoning very well. Coincidence is a basic part of life, but people do not handle it well. Humans feel the need to explain coincidence by some nonrandom mechanism than consider it due to chance, or due to a sampling from a population of different, related events.

This tendency for rationalization comes out of our basic need to be effective in dealing with the world. We depend on our emotions and perceptions to identify and react to what is going on. They make up the type of outlook that each of us has. Our outlook determines how we reason with the facts. It also means that we ascribe a deterministic functionality to our responses to these observations. This can happen if the facts at hand have simple causes that come from a single source. But facts are not simple monads. This leads to inconsistency due to the context of the event and the degree to which our interpretation addresses this context. We make decisions confident in a reliable outcome, but the underlying variability thwarts this confidence.

One aspect of human expectations where this deterministic viewpoint

goes awry is in the determination of how deserving a person is of the consequences of their actions or the justice of what befalls them. People have a tendency to believe in a just world hypothesis. If some random negative event occurs, people tend to ascribe to the victim some blame. All too often, the world is not just. Misfortune befalls people through no fault of their own. Someone may have taken reasonable precautions and still disaster struck. Someone may have gotten lucky to a degree greater than they deserve. Not every bad thing is a sign of failure. Not every success is a mark of virtue.

This phenomenon is expressed in the moral luck paradox. If our makeup is determined by chance combinations of our genes along with the random events of our lives, how can we be held morally responsible at all? Is it our fault if we end up drowning in a dangerous river? We were not guaranteed to have drowned by entering the river. The chance of each drowning is low, but still we can make the argument that it's our own fault because we ignored the risk. One famous example of this was Ayn Rand's conclusion that it was all right to smoke cigarettes because the consequences of getting lung cancer were at best statistically determined. She was forced to admit the error of her reasoning after she herself became diagnosed and had to admit that her actions had led to the outcome. A resolution of the paradox is not in the relative effects of luck versus choice but by looking over whole populations. Is it the fault of the inner city gang member that they became a criminal when that choice was determined by the bad luck of their situation? We can best give an answer by inspecting a population of average people in the same situation. By comparing a particular person's choice versus other people's choices can we arrive at an answer, and that answer is not necessarily one way or the other. If, given the fact that a significant proportion of average people placed in the same situation goes bad, it is unreasonable to ascribe blame to each individual. If instead, it is the exception rather than the rule can we reasonably blame them. This avoids the problem of deciding how much free will is at stake. It also points out the solution to the problem from society's point of view. If we are faced with a situation where it is unreasonable to ascribe blame to the individuals, then the solution is to fix the situation as a whole than to concentrate on changing the actions of the individuals. A moral problem that is seen over a population should not be transformed into a collection of individual cases in a false attempt to apply reducibility. Some problems are fixed by changing individuals, some by changing society as a whole.

Humans have a tendency to explain random events through overfitting. Some moral choices don't have not moral outcomes, they are random events. Randomness does not mean unmeasurable, it does though mean we have not

measured it. What we must guard against though, is basing our conclusions from too small a sampling of a population to warrant the degree of confidence we ascribe to a conclusion.

Although stories and fables are used to teach moral principles, we must not be taken in by anecdotal evidence. This is just another example of reasoning from cases. One or two archetypal cases prove nothing by themselves. One of the problems is the problem of self-selection. These archetypal cases are chosen because they are special in some way, but their specialness may be unrelated to the underlying principles involved in the moral choices. They can also give the illusion of causality where no such thing happened. To violate this principle by supplying an example, there is the example of the story of Richard Feynman (recounted in "Surely You're Joking, Mr. Feynman") where he had a premonition that his grandmother was dead. Right after that premonition, there was a telephone call to his frat house. He recalls clearly that the call was for another roommate, and had nothing to do with his grandmother, who, it turned out, wasn't dead at all.

The incident points out the fact that most events are not premonitions at all, and that it takes an unusual mind such as Feynman's to point this out. The average person only remembers the premonitions that come true and forgets the myriad of negative ones. So we end up believing in premonitions because we self-select the ones that come true, while the ones that are false or unverified tend to be forgotten in short order.

Statistical evidence tends to be overwhelmed by personal anecdote and vivid evidence. This varies from person to person. Some people have skeptical minds and demand a higher level of independent verification than others. Some people are very good at statistical reasoning, but that appears to be the exception. Instead most people seem to overfit their experience, and have the unfortunate tendency to retain these beliefs even when disconfirming evidence is presented. At the extreme are the authoritarian personalities, that is, those people who rely on authority figures to judge them. This is actually the personal anecdote magnified.

Once it is accepted that there is randomness in the world, you actually make better predictions if you accept error. One example given by Stanovich was a psychological study where the subjects were asked to predict which of two lights flashes on when the underlying situation is random and 70% of the time light 1 is lit. If you decide on some strategy that ends choosing light 1 70% of the time, in the same proportion as the light comes on, then you are correct $.7 \times .7 + .3 \times .3 = 58\%$ of the time. If you continue to try different strategies to predict an underlying cause for the randomness, the fact that there is none means that at best you end up trending to a random strategy.

Your attempts to model the behaviors of the lights will result in you choosing light 1 70% of the time just as if you had a random strategy, giving you the 58% success rate by the false assumption of an underlying causality. Instead, it is better to choose a second strategy: choose light 1 always. This gives up any pretense of finding a pattern to the variability and just goes with the most likely outcome. In this case you are then correct 70% of the time.

Stanovich points out that in situations where randomness is a significant factor, if you give up the hope of predicting then you gain better accuracy. This is termed actuarial prediction. It has been shown that actuarial prediction over populations is more accurate than case prediction, trying to analyze a case in detail -clinical prediction is inferior to actuarial prediction. Paul Meehl in his 1954 book "Clinical Versus Statistical Prediction", showed that actuarial models almost always are better than the judgments of human experts.

Because the just world hypothesis is demonstrably false, this applies to morality. Trout and Bishop, in an essay on the relationship of actuarial prediction to philosophy of science, point out a possible explanation for why actuarial prediction is better. It is known as the "broken leg problem". Given that an actuarial formula accurately predicts a person's weekly movie attendance, the fact that this person has just suffered a broken leg is enough to jettison the formula in the near term. That is, outside causes not part of the statistical model, such as the presumed working of a just world, will cause us to discard the model. Unfortunately, experts given this leeway find more exceptional cases (such as broken legs) than the statistics would indicate. In the final analysis, it is better not to overrule the actuarial analysis, because the attempt leads overall to worse predictions.

Besides the introduction of statistical analysis into the analysis of moral problems, we also have the problem that human well-being is a multidimensional feature space. This leads to two problems. Each action can have an effect in different dimensions of well being, and that there can be multiple events that lead to a particular situation. Humans just don't do well with this. As Stanovich points out, multiple causation is ignored in cases of pre-existing bias - when we believe that there is a causative agent, we tend to think it is the major causative agent and that there are no interactions.

A multidimensional feature space makes it difficult to tease out these multiple causations. This means that it is often necessary to perform testing of a theory with nonrandom data. If another nonrandom sample does not have the same response, the difference between the samples may tease out underlying causes. The failure to have the same response does not invalidate the first set, but only shows its limits. This appears to be a way of

reintroducing the lifeboat cases to moral reasoning in the guise of objectivity. It is not. The lifeboat example is the attempt to use anecdotal evidence in place of statistical reasoning. The use of nonrandom data is the attempt to control the effect of randomness in a study without claiming that randomness does not exist. If done correctly, it allows for the random factors to be easily incorporated into the final model. Done incorrectly, it leads to erroneous conclusions due to sampling error.

Finally, it is important to note that population and the effects of probability over large populations affect ethics. Take for example the problem of atmospheric pollution. Pollution for one person can be ignored - but for many, exactly the same behavior can be evil. This means that a morality based on some sort of pure libertarian principles will not work because they do not allow for a population effect. This points out that a social ethic - an ethics of society as a whole - cannot be reduced to an ethics of individuals. In some cases the moral judgment of an action can only be based on a determination over a population. In this case, the individual act is not and cannot be judged to be good or bad in itself. Probability determines good and bad for the typical case of an event of this type. In that case something may be bad sometimes but not always.

In conclusion: statistics is a good method of making morality relative. Logical reasoning can help us triangulate the relative situation, but probability helps us to fix the bounds of applicability. By relying on the study of local populations to determine what is good or bad, morality is defined relative to the society from which the statistics are drawn. Despite the differences, they also preserve the absolute metric. Using the standard of what is best for each individual, statistics may show different outcomes in different places.

For example, the restrictions placed on membership in many religious communities lead to conditions that the general society does not accept. But these conditions have their advantages. The article "Why Strict Churches are Stronger" by Laurence Iannaccone points out that some of these restrictions keep out people who are not as committed and would act as freeloaders on the society. This leads to the maintenance of a more cohesive and committed community, which gives more satisfaction to those members who choose to live with the restrictions. If an objective statistical survey of the personal satisfaction of the members of such a community reveals that they are happiest living under the restriction of their group, then that should be considered the moral choice for that time and place.

Next, let us look at some specific cases. We shall not attempt to reach definitive conclusions on ethical situations, though, since this is a book on

the principles of morality and ethics, a quantitative study of ethical behavior will of necessity involve a detailed analysis of the psychological and sociological literature to tease out the factors that apply in a given situation and their interrelationships, a time-consuming task. Our goal here is more limited. Instead, we will consider ways to test the correctness of the basic theme of this book - that morality applied in a relative manner rather than an absolute one leads to an overall increase in well-being.

Some characteristics that indicate an absolute morality is preferable is that the boundary between acceptable and unacceptable behavior is, except for obvious conditions, largely independent of individual circumstances and cultures. This tends to lead to a situation where the boundary between good and bad is well-defined and the similar for different societies.

A relative morality is harder to quantify. Here, the morality of an action is best measured by comparison to alternative actions. This comparison is relative to the degree of difference between these actions. Moral absolutes, such as the Golden Rule can be applied as a further comparison. In this case, the alternative actions can be compared by their relationship to the absolute. Therefore, a moral absolute can probably be the best model if the evidence shows a uniform applicability in all societies, whereas a more relative moral choice differs from society to society.

Instead of a single yes-or-no comparison of absolute versus relative morality, there needs to a relationship drawn between the universality of moral values over societal and individual characteristics and the effect on well-being for different choices. This means that the results of transgressing more absolute values such as murder have greater consequences than violating a more relative standard such as a prohibition against polygamy. Or for example, the seriousness of a lie determines the degree of transgression of the lie.

But absolute morality makes a distinction between degrees of transgression also. Stealing a loaf of bread is a minor crime compared to massive fraud or armed robbery. Shooting someone in the head is worse than punching someone in the nose. So the method of comparing the claims of absolute versus relative applications of morality cannot just look at the comparison of degree of transgression versus consequence. Both type of morality claim this relationship exists. The difference is that in an absolute morality, the relationship between transgression and consequence is the same over different societies and individuals, whereas for a relative morality, as the degree of transgression decreases, the variability in consequence increases, and this variability is due to differences between societies and people.

This means that the tests that determine whether an absolute or relative

morality is better are statistical tests over many moral problems, individual characteristics and cultures. So in the following analysis of various moral topics, we will attempt to define what an absolute and a relative moral outlook to a given situation would manifest themselves, and to provide some testable predictions or studies that would show which of the two moral outlooks is a better model for this problem.

[Corporal Punishment]

We begin with a moral injunction that comes out of the Bible: "spare the rod and spoil the child." James Dobson, founder of Focus on the Family is famous for his book "Dare to Discipline" where he advocates spanking as part of raising children.

Curiously enough, in the discussion of corporal punishment, the religious conservatives are on the side of moral relativism, while the religious liberals talk in terms of moral absolutes.

For example, Dobson is quoted as saying "Some people (particularly those who are opposed to spanking in the first place) believe that the use of a neutral object in discipline is tantamount to child abuse. I understand their concern, especially in cases where a parent believes 'might makes right' or loses her temper and harms the child. That is why adults must always maintain a balance between love and control, regardless of the method by which they administer disciplinary action." The advice is replete with warnings that each child is different. Dobson warns, "Be sure the child gets the message while being careful not to go too far." He warns that spanking may not work on some kids such as a child with attention deficit/hyperactivity disorder. On the other hand he qualifies the degree of spanking if a child is strong-willed. Although he warns against inconsistent discipline, he notes that the key to effective discipline is in knowing your child. Dr. Dobson says "the only way to raise children correctly is to understand each boy or girl as an individual and to design parenting techniques to fit the needs and characteristics of that particular child."

The liberals tend to be more absolute in their condemnation of corporal punishment. They point to studies that show significant correlations between spanking in childhood and psychiatric disorders and substance abuse problems in adulthood. They also point out that people who have been abused as a child tend to be abusers themselves, and that this result applies to corporal punishment also.

There are many factors to consider in determining the difference between corporal punishment and abuse, and the degree of corporal punishment. Some factors are the frequency, the anger displayed by the parent, the method of punishment, the psychological makeup of the child, and the

parenting style. Collapsing these dimensions into a single measurement of degree of punishment is not easy.

In summary, both sides agree that beyond a certain point, corporal punishment leads to abuse and abuse is detrimental. The difference lies in the range of mild spanking (however that is defined). The anti-spanking group would make the claim that there is a straight, constantly increasing correlation between the degree of punishment and the negative effects on the child. In contrast, the pro-spanking people make the claim that there is an inflection point, and that if there is no corporal punishment at all, the negative effects on the child again increase. This may be a case where it is not necessary to compare differences between culture and personalities to find the difference between a relative and an absolute measure, although this can lead to further insights. Simply looking for this inflection point is sufficient.

[Capital punishment]

Capital punishment is a response to moral situations that has a built-in absolute stance to it. Although there are some gray areas between death and life, such as the degrees of brain death, great care is taken in the application of capital punishment to ensure that these situations do not arise. If capital punishment is applied, it is applied in such a way that the person it is applied to is dead and stays dead.

People on both sides of this issue discuss it in absolute terms. Those against capital punishment argue that it is never acceptable. Those who are for capital punishment refer to it as the ultimate penalty, often citing the Biblical injunction of "an eye for an eye."

A morally relative viewpoint of capital punishment would look at the tradeoff between the lives saved by deterring crime versus the lives taken by execution, especially the execution of innocent people. One famous case of this deterrence was the execution of the Romanian dictator Nicolae Ceausescu. His execution caused the collapse of a resistance movement by his secret police, which was in the middle of a bloody rampage. Many lives were saved because the resistance collapsed in the aftermath of seeing him executed on national television.

The situation with capital punishment may never be resolved if it is just looked in terms of two competing standards of absolute morality. Both sides make rational arguments for their claims, based on a set of axiomatic principles. Both sides claim that with their given standards the world is a better place. The proponents argue that if capital punishment is applied, the deterrence (both for the criminal and others) results in a decreased murder rate. The opponents argue that no such deterrence effect is seen,

and besides, a certain number of people are executed in error, when in fact they are innocent of the crime. There is some evidence that this is alarmingly common, not a rare miscarriage of justice.

Although this debate is often expressed in absolute moral terms, capital punishment in most modern societies is applied using the principles of moral relativity. Instead of executing every criminal convicted of killing someone, there are various degrees of transgression, such as murder, manslaughter and criminally negligent homicide, with the imposition of capital punishment based on mitigating circumstances, such as age, mental competency, facts about the victim and the heinousness of the act.

Since the arguments for and against capital punishment are absolute in nature, the question arises whether it is possible to give a relative justification for capital punishment, more in line with the way it is applied. The relative argument would have a number of parameters: the application of the punishment on the irretrievably depraved, the possibility of being applied in error, and of course the question of the tradeoff between the positive and negative effects of capital punishment on the overall well-being. It may well be that a case can be made that reconciles the two sides of the issue, making capital punishment legal, but rare.

[Homosexuality and adultery]

Homosexuality has become accepted in modern times. This is an application of moral relativity where the changes that have developed in the human condition over time have led to corresponding changes in moral standards. Currently, the discussion of gay marriage is an example of the further social change in these definitions and social mores.

The sacred books of the Abrahamic traditions inveigh against a number of behaviors that had serious negative consequences. Homosexuality and adultery were probably banned because they led to the spread of sexually transmitted diseases. In the same manner, meat such as pork and shellfish very also banned because of the danger of disease. If this was the underlying reason why they were proscribed, it can actually be argued that there is nothing inherently wrong with these types of sexual activity at all. The actual cause of disease is promiscuity, no matter how it manifests itself.

It is fair to ask what the difference is in the well-being of the individual and the well-being of society since these changes in social mores have been in place. For the individual, the practice of adultery and homosexuality leads to positives, such as the increase in sexual pleasure for the acts alone, but there are also negatives, such as the disapproval of other people. The effects on society also have negatives and positives. Besides the negative in the spread of disease due to unsafe sex there is the positive effects of bringing

out into the open activities that have always been part of sexual activity but can now be dealt with in an open manner.

Whether a relative or absolute morality is most applicable here depends on how much or how little things have changed as society and its mores have changed. An absolute morality would say that regardless of the changes in society over time, these behaviors should still be proscribed. A relative morality would claim that allowing what was formerly forbidden results in a net positive effect on well-being, if the negative consequences are mitigated. The analysis is complicated by the fact that society has changed through time, so any analysis must take this into account. The relativist can make the claim that allowing the change in mores has not led to a society that is any worse off by any objective standard of well-being. The absolutist would have to show that society is worse off for allowing these changes, a claim that, in fact, the absolute moralists do make. In this case, though, the null hypothesis, that there is no measurable degradation due to the changes in morality favors the relativist argument.

[Zero tolerance]

The term 'zero tolerance' has been used in the last few years to describe the reaction of authorities to behaviors that are unacceptable. This has come into vogue with a resurgence of the attitude that absolute morality is correct and that moral relativism is somehow wrong. A moral relativist would argue that this attitude is immoderate in a negative way - that zero tolerance is intolerance and an inappropriate application of moral norms.

A study of the effects of applying zero tolerance as a general principle would make a good testing ground to compare the effectiveness of absolute morality versus relative morality. Absolute moralists would claim that the application of zero tolerance policies will, in most cases, lead to a better overall outcome for individuals and society than the alternative of providing leeway, such as allowing first time offenders the benefit of the doubt. Moral relativists would argue that once we move away from extreme transgressions such as murder, a reasonable allowance for mitigating circumstances and first time offenders does not lead to an overall negative outcome for society as a whole, and will even allow the individual transgressor an opportunity for self-correction that would benefit both themselves and society.

Related to this difference in attitude is the changed notion in society towards sin and repentance. In American society in the nineteenth century, the notion of rehabilitation came to be the governing attitude for law enforcement. Instead of jails, prisons were termed penitentiaries, and instead of punishment, the offender was supposed to have undergone a process of correction and rehabilitation. In recent generations, attitudes have moved

away from this. Prisons are now a place for punishment, where prisoners are no longer coddled and given luxuries. Even training courses are eliminated along with television or exercise equipment, and hard labor and chain gangs substituted in their place.

This response has extended the absolute attitude towards right and wrong to an equally absolute attitude towards a person's character. Less effort is made to reintegrate the person who has served their time back into society. Instead, the prison sentences are longer so that the transgressor is removed from society entirely. With this attitude, punishment is more important than redemption, and more often redemption is not even considered an option. Character is conceded fixed instead of leaving open the possibility of a change of character.

A moral relativist would tailor the incarceration of a prisoner based on the actions of the prisoner. Instead of fixed sentences, moral relativism is implicit in the establishment of the concept of parole. The elimination of parole and the return to fixed sentences marks another turn away from moral relativism to an absolute morality. Whether society is better off for it can be seen in studying the effects of these changes of attitude on society and on the criminals behind bars or the recidivism rate.

[Animal rights]

One of the moral movements of this generation is that of animal rights, led by notable moral philosophers like Peter Singer. The goal of the movement is to grant certain rights to animals that are similar to the rights of humans. Considering that animals are killed for food and used in scientific experimentation, along with the fact that many animals, especially mammals, feel pain, they should be granted rights that respect their need to avoid suffering.

The argument is based on a logical analysis that says that since rights are based on the needs and desires of humans and that animals have similar needs and desires, animals should be granted similar rights. The argument is based on the observation that the simple fact of being human is not a reasonable distinction for granting rights or denying them. Just as the denial of rights to slaves or women was eventually shown not to have logical validity, the same denial to animals is unjustifiable.

The argument for animal rights relies on absolute consistency and drawing clear demarcations. The granting of rights is dependent on the purpose and protections that those rights are for, not due to some artificial distinction. There are no statistical arguments necessary.

The problem with this argument is that it does not satisfy most people. Even given the logical argument, it does not get to heart of why we feel there

is a distinction between human and nonhuman. There is an emotional bias against animal equality that goes beyond the logical argument itself that makes it difficult to accept this equality. Of course, there was an equally emotional bias against equal rights for slaves and women that in retrospect is now considered rationalization. But the question still exists - perhaps there is some justification in this case for the bias against animal rights.

But it is undeniable that animals have rights, too. These rights are represented by humans, though, since animals cannot bring them to our attention. But the existence of groups such as the Society for Prevention of Cruelty to Animals exist because humans have for a long time recognized that animals suffer and have a right not to. These organizations arose out of a basic empathy that people felt towards domestic animals such as horses, dogs and cats. These rights, though, are attenuated (the way that well-being is considered aggregated in Chapter 4) by the fact that humans are recognized to have a higher level of awareness than animals do.

In the distinction between absolute and relative morality, the traditional approach to animal rights as practiced by groups such as the SPCA is more relative than the attempt to grant absolute rights. Which approach is better can be resolved by studying the effectiveness of the two approaches. If the current relative morality is the less successful approach, then the granting of animal rights will lead to a greater level of well-being for both humans and animals. But if the attempt to impose these rights leads to contentiousness that works against the ultimate goal of the reduction of suffering, then the more gradual relative approach to animal rights that has developed through the last century or so is preferable.

[Abortion]

In the abortion debate, both the anti-abortion and the pro-abortion advocates discuss the morality of abortion in absolute terms. But the majority of people are moral relativists.

For the anti-abortion movement the issue is at what point a fetus becomes a human - where it is immoral to take an innocent human life. For many, the start of life is defined to be at conception. Therefore, any abortion is murder. But this simple moral absolute leads to paradoxes in its application.

Many people on the anti-abortion side of the debate rely on Biblical references to bolster their case that life begins at conception. But there is no absolute biblical statement in the Christian or Jewish holy books that directly discusses abortion or the moment when life begins. They must base their argument on a logical inference drawn from other statements that do not address the matter of abortion directly.

It is argued by some that life begins when the egg is fertilized and by others when the egg is implanted. Although the point of fertilization is a clean distinction, common usage shows that this is not the point that is regarded as the start of life by most people. If it were so, a fertile woman who had sex during the month but had her period would be regarded as having had a life which died, but this is not how people see things, except for the extremely doctrinaire.

One question that naturally arises is the question of how temporal is conservatism. Many of the people who claim that life begins at conception tend to consider themselves conservative. If that is so, then how far back do accepted beliefs have to go before they are acceptable? For a Christian belief system, is a century adequate, or must it be two thousand years? The belief that life began at conception came only after the creation of the microscope during the Enlightenment. Before that, life began when the womb was quickened. Therefore, the claim that life begins at conception is a recent development.

Quickening is still the standard that the average person uses to determine whether abortion is acceptable or not. This distinction, though, creates a gray area. It is obvious that life has not begun in the first trimester, but it certainly has begun for a healthy fetus in the last trimester. The middle trimester is the question. Some people make an argument about the existence of a human life based on neurological development of the fetus, but this is also ahistorical. The stages of development were only known in the last century or two and do not carry much credence with the average person in making the decision on when to define a fetus as a human being.

Abortion is an example of the heap paradox. If 10,000 grains of rice all piled up form a heap but one grain of rice is not a heap, how to determine that intermediate stage where a few grains become a heap? Similarly, when does a fetus become a child? One way to phrase the question is that of probability. A zygote one day old has .00000 chance of surviving; a fetus eight months and twenty nine days old is .99999 chance of surviving. A fetus becomes a person when its chance of surviving becomes reasonably greater than zero.

A moral relativity is one that acknowledges that there can be competing rights that are in opposition to each other, and that neither right is absolutely attainable without the complete negation of the competing right. This situation holds in this case where it is necessary to balance mothers right to choice to carry the child versus the potential child's life. The definition of murder being the unlawful taking of a human life, the implication of a legal abortion is the recognition that this is not murder. A secondary

implication is that the fetus is not a human life until the third trimester, probably because it is incapable of life on its own. It is generally regarded that abortion in the last trimester is killing and should not be allowed. But for the majority of people, abortion in the first trimester, even if it is killing, is not killing a human being.

One of the best arguments for a relative morality is the fact that there are less abortions where it is legal. From the standpoint of moral relativity, the actions that may lead to the greatest well-being is to keep abortion legal, but to actively work to make it rare, by providing alternatives such as adoption. In this case, then determination of whether a relative or absolute morality applies would be to undertake a comparative sociological study of societies that differ in terms of whether abortion is legal or not. If the claim that legality actually leads to a reduction in abortions is consistently verified over a large number of societies, then a relative morality would bypass the whole debate of when life begins and establish a compromise where the ultimate goal of reducing the use of abortion is the goal. The goal can be shared in common regardless of the absolute definitions either side holds.

[Group responsibility]

If an organization commits a wrongful act, then how does the organization and its members bear the responsibility? The responsibility must be shared between the organization and its members. If you can't evade responsibility, there are no other alternatives.

This question explicitly recognizes the fact that social organizations can exhibit emergent behavior. Even a group of well-meaning individuals can act together as a group with an outcome that is morally unacceptable. One of the places that this can happen is in social services organizations, especially one that has been starved of the resources to adequately fulfill its mandate. Every member of the organization could be trying hard to do the best they can, but the organization could magnify any errors through a lack of oversight or self correction.

Ideally, if an organization does something wrong, justice demands that the people responsible for these misdeeds are punished. But this is not so clear-cut. Someone can take an action in good faith, but not have the information to make a moral judgment. Presumably the executive who gave the order does have this information. It is theoretically possible but difficult to come up with an actual case where no one has the knowledge to judge an action wrong but the organization as a whole has an organizational "awareness". Since an organization does not have a brain, this awareness would only be manifest through and informational flow or a set of procedures. If this could actually happen, it might even be hard for an outsider to recognize

that some wrong action was being committed.

But if a morally wrong action takes place and the people who took the action and those in authority take responsibility there still can be a residue of blame for the organization as a whole and its members. There have been cases in this century where governments and whole societies have had to deal with this. If no person or persons come forward to take the blame, the price paid by the organization can be severe and long-lasting.

The concept of having an organization and its members as a whole pay for the transgressions of some part of the group underlies the motives of terrorist actions. Terrorism, except for a small number of pathological individuals, is justified by the perpetrators as punishment against the group as a whole, including the innocent and even those members of the group who might have even objected to the particular action.

The terrorist response is an unfortunate and an extreme type of reaction in an absolute moralistic context. Since evil acts were judged to have taken place, all of the members of the group responsible are judged to have equal responsibility. This absolute moral response underlies the equal treatment in the law of accomplices to a crime such as murder, even for those who did not pull the trigger, such as the drivers of getaway cars.

A relative morality acknowledges differing responsibilities of the group members. Responsibility is apportioned to the individuals depending on their degree of involvement. In the case of an action where individual responsibility cannot be determined, it is still possible to apportion blame differently to different members of a group by considering each individual's overall contribution to the group in activities that this group has performed.

There is also a difference between absolute morality and relative morality in terms of how long a sanction of a group should be applied for transgressions in the past. An absolute morality that does not recognize degrees of blame has a difficult time of letting go of the past, since there is no mechanism for this to happen. Too often, blame continues up to the point where the original transgression is almost forgotten, at which point the group abruptly goes from being reproached to being held blameless.

A moral relativity must take into account whether any sanctions against the group or its members lead to an overall positive increase in well-being. This can be as a deterrent effect or by restitution of the injured parties. But the application of sanctions is taken relative to the individual's responsibility for the action and the amount of time that has passed. Also, the organization of the group must be taken into account. Although the citizens of a society with a corrupt government can be held responsible for the actions of that country, the responsibility must be tempered if the citizen was a not a direct

member of that government or its institutions. Also, blame is much less for a citizen of an authoritarian country than for a democratic country whose citizens voted for this wrong course of action.

[Aggression and deterrence]

In Michael Clark's "Paradoxes From A to Z", two paradoxes about aggression and deterrence are presented, that have differing responses for absolute and relative morality. The first paradox is a variant of Prisoner's Dilemma.

Two sides make a disarmament agreement, which is easy to evade by hiding the weapons. There are four outcomes, ranked in terms of what is best for us: 1. We are armed and they disarm: they are at our mercy, so we are safe. Call this state AD 2. We are both disarmed. Then we are at peace, unless we arm in the future. Call this DD 3. We are both armed: Then a war is possible, but neither of us has the advantage. Call this AA 4. We are disarmed and they are armed. Then we are at their mercy.

So, we have the following situation: if we break the agreement and arm ourselves then either case AD or AA applies. They are at our mercy, or we are at risk of war. If we stick to the agreement and disarm, then either DD or DA apply. Neither outcome is preferable to the two cases if we cheat and stay armed.

This is a basic part of game theory, where the situation is expressed as a game in which either player can win or lose or tie.

There are two ways around the problem. The first is to recognize that in real life, the situation does not lead to immediate disaster - in effect, there are multiple time steps. If the other side breaks the agreement, we have war. But we usually have time to rearm ourselves before war breaks out. The situation with nuclear weapons was a classic case of the prisoner's dilemma that could not be resolved by repeated trials. That is why the United States has refused to disarm, and why it has been untenable for the United States to demand that other countries do what we are unwilling to do ourselves.

Before going to the other solution, it is time to introduce the second paradox - the paradox of deterrence. If you deter an enemy by threatening retaliation, which you know you will not want to carry out, can you make the threat, knowing you don't really mean it? At its most benign, it is a lie. Again, this has come up in the nuclear conflict - the threat of mutually assured destruction. In real life, during the Cuban Missile crisis, the solution was found by the two powers taking their rational self-interest to heart and working out the problem instead of resorting to the threat. But the situation can certainly arise in real life where one or the other adversary is just not acting rationally. This led to the First World War with the result that the

countries of Europe ceased to be the Great Powers they were before the threat was carried out and the war started.

In both of these paradoxes, the way out of the problem is usually to expand the parameters of the situation to incorporate other conditions that change the nature of the game. This is most successfully done by transforming the situation into one where it is much less likely for there to be a cases where there is a winner and a loser or both lose. If, for example, in the case of the Prisoner's dilemma, where disarmament leads to a net positive for the power that disarms, this can translate into further turns of the game into situations where it is in the best interests of both parties to disarm. For example, the costs of maintaining a powerful military are high. The positive outcome of disarming is that the resources of the nation are turned to positive endeavors, such as trade. This leads to a situation for the other party where starting a war with the former adversary is not a win, but actually a loss due to the reduction in trade.

The relationship of absolute versus relative morality to these paradoxes turns on the flexibility of defining the rules of the game. An absolute morality has the danger of defining these problems in terms of absolute rules, making it hard to find a way out. A relative morality redefines the rules of the game. One way of doing this is to redefine the meaning of well-being. Since well-being is not a single one-dimensional function, shifting the different ways of computing well-being changes the nature of the game. Of course, as the consequences of the game become more absolute, a relative morality becomes more like an absolute morality. In the case of nuclear war, it is more difficult for a relative morality to redefine the consequences.

An historical analysis of actual cases of these nature can determine which is the more successful approach. If history shows that these situations have been most satisfactorily resolved by redefining the rules of the game, a relative morality is preferable. If it can be shown that usually one or the other side wins according to the defined rules of the game, then an absolute morality is more successful.

[Welfare and socialism]

The political definitions of conservative and liberal mostly consist of connotations. Essentially, a liberal is someone who is open to new ideas. A conservative is someone who prefers the tried and true. In the United States among the main connotations the words liberal and conservative have is that implies that these two attitudes to change result in a different point of view on social welfare.

Welfare comes down to how much selfishness is tolerable. You can completely focus of your own family and friends and let others fend for them-

selves. Or one can consider that, since we all partake in the benefits of society we also have obligations to contribute to the greater good.

Ayn Rand became famous for espousing the virtue of selfishness. Relative to her era, she was correct. But the differences between selfishness and altruism are not absolute. Like the debate between the effects on human development of nature versus nurture, the two sides are both wrong when they take out extreme positions. In the era in which she spoke, the world had an ascendancy of socialism. It was the predominant morality. People like her came along to challenge this position because it had gone too far. The pendulum has swung, perhaps too far now the other way. Even the Democratic Administration of Bill Clinton pronounced that they would end welfare as we know it, and social security is under attack. With the collapse of communism, the reliance on Objectivism to show the way to a better world has lost its relevance.

But in the current era the pendulum has swung further the other way. For example, Christianity is at its heart a communistic religion. The Acts of the Apostles talk about giving, each according to their ability and taking, each according to their need. Christ also talked about the necessity of charity. A major force for socialism in the Nineteenth century was the Christian Church. Some of the earliest communes were established by devout Christian sects. Currently, the most vocal parts of Christianity are espousing a social virtue that is incompatible with what is taught in their holy book.

But in this era we are witnessing the effects of too much selfishness. It is remarkable that in the last twenty years it has been Republican administrations that have run up the largest governmental debt. This is due to a selfishness that demands services of a government but is not willing to pay its fair share. The lack of socialized medicine in the United States leads to an overall more expensive and less effective medical outcome. For example, the rate of infant mortality in the United States is worse than many other developed countries to the extent that even Cuba fares better.

But the heart of socialism is the financial assistance given to the less fortunate by the government, and here the results are mixed. Although the United States gives much less in terms of its percentage of wealth in welfare compared to other countries, its individual citizens give much more than other citizens or other countries. But citizens are unwilling to give up social security and for good reason. When Social Security was first established, the majority of senior citizens lived in poverty. That is not the case now. The system works. It is just a question of the relative balance.

Given the question "how much welfare is enough" a relative morality can give no absolute answer. The appropriate level is dependent upon the

situation, and the situation always changes. Therefore, even presented with the same situation, such as a response to a natural disaster, or helping out a person down on their luck, the degree of altruism is different in different eras depending on changes such as the development of education and technology. In some ways technological changes enable people to help themselves. But in other ways, the complexity of technology requires more outside assistance to enable the unfortunate to rise to the level of current society.

An absolute morality would attempt to provide fixed rules, such as the concept of tithing. Of course, the rules in an absolute morality are not as simple as a single percentage - different situations demand different responses. But an absolute morality would claim that there are universal rules that hold over all times, places and situations.

Neither liberals nor conservatives can claim to be exclusively absolute or relative in their approach to welfare. Although one can draw a caricature of a conservative having an absolute approach to welfare, it is equally likely that the liberals frame their attitude in absolute terms also. Unlike the debate on abortion, the discussion of welfare in absolute terms is across the whole population, not just the strongest advocates either for or against the welfare state. But when it comes to individual action, it usually happens that the group is defined by absolute rules, but individual welfare is relative to the situation. So even though both sides talk about welfare and codify it in absolute terms, they become relativists when given the opportunity to deliver it as they see fit.

If it were possible to develop an absolute morality of welfare, it could become more quantitative an endeavor. We could measure effectiveness and apply optimal strategies. But if the moral basis of welfare is inherently relative, it must be invented anew as situations change. One way of determining which type of morality applies would be to look at the effects as the pendulum swings from too much or too little welfare. If there is an inherent stability and predictability to the effects of these swings, with similar initial conditions, then an absolute morality would be in force. If the changes are chaotic, a relative morality applies.

[The Golden Rule]

An absolute golden rule cannot be defined in a way that applies to every single person uniformly because that is subjective. But it can be defined statistically for a society.

There are many ways to formulate the golden rule. An analysis can determine both the quantitative aspects and the qualitative of a formulation. A major aspect of the golden rule is the principle of reciprocity involved. A statistical analysis can determine whether an absolute equality is best over

a range of exchanges, or whether something less or more is best in certain circumstances.

The qualitative characteristics of the golden rule can be studied also. Is it better to formulate the golden rule as a positive "do unto others" or as a negative "do no harm". This would be an observational study. Different communities phrase it in different ways. The statement of the rule in a given society can be matched to how the rule is applied.

Although individual applications of the Golden Rule are by necessity relative to the people involved, if there is some objective absolute standard for the Golden Rule, it would show itself in some sort of invariant characteristics that manifest themselves in a variety of different societies. It would be interesting to see if the Golden Rule leads to an absolute sets of common behaviors, or is relative to each society, or even each community.

[Morality and God]

The practice of ethics goes on whether or not a person believes in a God or not. Belief or disbelief changes morality but does not negate it. The atheist and the theist are both as likely to lead a moral life. Different viewpoints of the ultimate lead to different conclusions about good or bad actions.

But how does one compare the effects of different theologies on morality? Without a rigorous proof of the correctness of one or the other viewpoint of the universal, it is not possible to judge and accept or discard a particular theological viewpoint based on theoretical analysis alone. In any case, it is probably impossible to give a rigorous proof of a universal, except for those that contain obvious contradictions. The best one can do is to make an empirical evaluation of which viewpoint leads to the best overall outcome in terms of observable well-being.

The definition of good and bad must be quantifiable - if they involve salvation, they cannot be measured. They can, though, include things that might be believed to be associated with salvation. This is generally true of concepts - the measurable properties can change as they get refined.

Since we have no direct knowledge of the afterlife, the only measurable properties of salvation involve the individual in this life and the effect on the living after the individual is dead. For example, it is impossible to quantify the effects of following God's laws on the chance of eternal salvation, since this question cannot be answered either way. Pascal's Wager is a motive for believing, but it has no direct bearing on how to live a moral life. Pascal's Wager cannot even help us indicate which God to believe in and thus which morality to follow. But we can quantify the effect of belief in salvation on how a person lives their life and interacts with others. We can also quantify

the effect that this person's example affected the behavior of others.

Can morality be applied to prove the existence of God? It would appear that such a proof cannot be done, because, if God existed, morality would be defined in terms of God-like absolutes. This would make any argument circular. But this is no barrier to a proof, since the measurable characteristics of the morality that is associated with a particular definition of God may distinguish that belief structure as being superior in terms of morality. In effect, there could be an Intelligent Design argument for morality.

Whether a relative morality or an absolute one functions better does not necessarily mean that a God exists or not. Although many theists believe in an absolute morality, this is often inferred as part of the dogma of that religion, when in fact no such claim can be directly found. Instead, the proof that certain unique and specific practices are especially related to a good life may provide a compelling argument for belief in God, regardless of whether these specific practices are absolute or relative in their nature. Although existence of a God implies the existence of a set of absolute standards, the application of these standards do not have to be absolute also.

[The meaning of life]

Even if the basic nature of humanity could lead to a universal definition of the meaning of life, each of us is unique enough so that we have to define our own within those parameters. The meaning of life is expressed in different religious traditions in different ways. The Christian, the atheist, the nihilist, the Buddhist - each has a different description of the basic meaning of our existence.

The sources that these traditions use are varied. Meaning comes from our nature. Meaning comes from our experience. Meaning comes from the facts of life, and our existence on Earth. Some traditions use the nature of God to define meaning, even to the point of making that axiomatic and deriving the other sources as secondary. This makes free will complicated, in that free will is constrained by how the sense of what is meaningful gives a perspective to their available choices.

Just as the Golden Rule can be shown to be an absolute rule in a relative morality, it is possible to state an absolute, universally applicable meaning of life that can be applied relatively to each person capable of free will: the purpose in life is what you believe that helps you decide what to do. This purpose gives meaning. So, meaning is what makes it possible to make the decisions in your life.

Even for someone who believes that there is some external entity or force in the universe that provides it with a meaning, this meaning only gives purpose to a person's life through that person's decisions. An entity

that exists without deciding is probably non-existent. The feelings that a person has are in reaction to the environment and are not the cause of ultimate meaning because they are just there and cannot be avoided. Only in making decisions does the actor consider what is meaningful. Therefore, meaning drives purpose and this drives decisions.

What helps a person decide can come from many sources. If a person belongs to a community of faith that has an explicit statement of the meaning of life and the purpose that it gives to living, then this statement helps the believer decide what to do. If a person is a complete nihilist, then even the thought of suicide is driven by a purpose - the identity of the essential nihilism of life.

For the atheist or the pantheist, the meaning of life comes from within and is an active give and take between an identification of the facets of the essential nature of the individual. Meaning can come from the most trivial acts such as what to choose to eat for a meal, or the most universal desires such as how Gaia is evolving and to what state it should reach. Meaning and purpose develop a multifaceted richness to their definition. They create a conversation between action and its consequences, both good and bad. They juxtapose the essential nature of the individual against the nature of the world that individual is an actor in, no matter how big or small the degree of this person's actions. Meaning also has dimensions of long term and short term characteristics. Not all of what makes life meaningful need to focus on the big picture. The immediate experience, even a unique experience, never to be repeated again can be the result of meaning and purpose that exists only for the moment.

Meaning can also change as life goes on. To each season and stage of life, the parts of meaning can ebb and flow, shifting as different factors are less important in the decisions that affect the future. Meaning changes as the individual learns about themselves and learns about life. As this understanding changes, meaning does too. As responsibilities and demands change in life, the different aspects of meaning come and go, and as each chapter in a person's life is closed. Some parts of the meaning are left behind with it.

Of course, people can make decisions in their lives that can lead to a reduction in well-being of themselves and others as well as an increase. But this does not mean that meaning is only in the positive choices that the person makes. It is possible to conclude that bad decisions imply that this individual's meaning and purpose are maladjusted, and that redefining their meaning will lead them to a better state of existence.

Meaning and purpose, like the Golden Rule, may not be provable, but

given as axiomatic. But like the Golden Rule, it is possible to experimentally determine whether the meaning of life is absolute or relative. If, through an empirical study of what drives different entities to make choices in their lives leads to the same set of core principles, then there is an absolute meaning to life. If each entity comes to meaning through a extremely different outlook, then meaning is relative. This determination should not be made without regard to overall welfare, though. If the meaning that a miserable person ascribes to life is different to the happy and successful person, this does not imply that meaning is relative. Instead, we need to sort out if the unhappiness is under the person's control, and consider meaning and purpose only if they lead to a good outcome. If even in this case we still find that each person's meaning is unique or even highly differentiated by that person's makeup, then can we claim that meaning is relative. But if the core principles of an essential life change little from person to person, then meaning and purpose are absolute.