

# *Power Based Consequentialism*

*George Sowers*

A normative theory of ethics is developed based on a utility of species wide power, rather than the more typical utilities of happiness or welfare. Conditions for a satisfactory normative theory of ethics are developed by considering the overarching goal of long term viability of human society. A precise definition of power is introduced corresponding roughly to a measure of an agent's capability to change the world. Societal or species power is defined as the aggregated power of individuals. A utility is then defined as the total power of the human species in the long run. This definition of power is compared to others found in the social science literature. Various kinds of power are discussed. The utilitarian goal of maximizing power then leads to two primary moral principles which I call the law of the individual and the law of cooperation. The law of the individual states that any increase in the power of an individual is good. The law of cooperation states that any increase in the means of cooperation is good. These general laws can then be used to derive specific rules of behavior and address practical ethical issues. Power based consequentialism avoids many of the issues surrounding traditional utilitarianism and can be extended beyond ethics to address policy issues in other fields like environmentalism, science, technology and international affairs.

## *1. Introduction*

One condition for truth is effectiveness: true beliefs lead to effective actions. If a certain action is effective, that effectiveness counts as evidence for the truth of the beliefs that led to the action. This is the pragmatic view of truth. Truth is what works. In the words of William James, truth has cash value. But beyond a theory of truth, pragmatism can be viewed as an overarching intellectual value: We should place the highest value in ideas that work, ideas that lead to outcomes we desire.

When applied to ethics, the pragmatic approach leads to consequentialist theories. These theories are goal oriented, focused on ends and results. Often, the desired outcome

is expressed in terms of a utility, a numerical quantity representing the goals of an agent or group of agents. Hence, these theories are also known as utilitarian theories. Utilitarian theories of ethics were first proposed by Jeremy Bentham (1748–1832) and refined by John S. Mill (1806–1873). The principle of utility, as proposed by Bentham and Mill, states that actions or behaviors are right in so far as they promote *happiness* or *pleasure*, wrong as they tend to produce *unhappiness* or *pain*. Utility is defined as the sum of happiness across all individuals within society. Then actions that improve overall utility are good; actions that degrade utility are bad.

Since the days of Bentham and Mill, utilitarianism has enjoyed a fair amount of success as a normative theory of ethics. But utilitarianism is bothered by a number of conceptual issues. Early issues included the difficulty in defining happiness in a manner that allows its use as a quantifiable utility. This problem led to the development of alternative utility measures. In vogue now are utilities that focus on the interests or preferences of individuals or other measures of individual well being. Other issues arise in the practical difficulty of calculating an action's effect on utility. Much of the requisite information required to perform such a calculation is not available to a decision making agent. This has led to rule utilitarianism, whereby the utilitarian calculus is used to devise a set of rules of behavior to be internalized by all members of society. These rules can be rigid, like laws, or simply guidelines or rules of thumb. Modern rule consequentialism (Hooker, 2000) relies on a version of the principle of utility whereby some measure of welfare, summed over individuals, represents the overarching goal of the theory, the quantity to be maximized by the rules of behavior. Yet there is little to justify this principle other than bare intuition.

In this paper, I take the pragmatic approach to another level. I use a pragmatically motivated thought experiment to derive a utility fundamentally different than welfare. From a functional perspective ethics are rules for human action and interaction. Their purpose is to enable the harmonious and effective organization and operation of human society. By thinking broadly about how different systems of ethics contribute to long term viability of human society, I derive several conditions for the optimal utility. This leads to the introduction of a different measure of utility, aggregate power at the species level. I argue that the power utility leads to a satisfying and coherent theory of ethics while avoiding many of the issues that hinder other utilitarian theories.

In Section 2, I use a thought experiment to develop principles for a normative theory of ethics. These principles provide guidance for the subsequent logical development of the theory. Section 3 examines several candidates for utility settling on power as most suitable. Section 4 provides a formal definition of power as a utility. I define power with enough mathematical rigor to enable some derivative properties to be developed. In Section 5, I compare this new definition of power to others found in the social sciences literature. In Section 6, I take the power utility and derive, in the spirit of rule utilitarianism, two fundamental moral laws which I call the law of the individual and the law of cooperation. Finally Section 7 contains my conclusions. I show how the power theory meets the considerations derived in Section 2 and how it addresses some of the issues of traditional utilitarianism.

## *2. Principles for a Normative Theory of Ethics*

What kind of attributes should we look for in a normative theory of ethics? Fundamentally, ethics are rules for human action and interaction. Their purpose is to enable the harmonious and effective organization and operation of human society. A normative theory, then, would provide rules that are optimal for that purpose. The means to judge optimality is to examine the effectiveness of different ethical systems if implemented into a human society encountering real world issues and problems.

The problem then becomes one of defining optimal in a way that allows us to make progress. Imagine an ensemble of human societies, each equipped with a system of ethics generally accepted and internalized by nearly all members of society. Each society is faced with the sort of challenges real societies have encountered or might encounter in the future. These range from internal problems like providing for the daily needs of its members to natural disasters to competition with other societies (including potentially a hostile alien race) and environmental degradation.

Next imagine a simulation of these societies from the present to some point sufficiently far into the future to capture meaningful performance data. Several hundred years is the sort of span I have in mind. We perform a Monte Carlo analysis varying all of the internal and external conditions, but keeping the system of ethics invariant. In practice, this would require a vast computer far beyond the capabilities of today's fastest supercomputer. That's why it's a thought experiment. Finally we collect statistics on the outcomes. Here we have a choice. What parameters will we measure? Candidates include population, living standard, economic output or one of the welfare measures of traditional consequentialism. But choosing any of these will lead to further debate. A simple measure that should garner nearly universal assent is basic viability: Did the

society survive? Did it avoid assimilation by a competitor or, even worse, extinction. The optimal ethics then is the ethical system that produces the highest viability rate in the long run.

Given this starting point, further attributes of our theory can be derived. We are optimizing our theory to produce a specific outcome or consequence, namely long term viability. Therefore it is natural that our theory should be consequentialist. To enable the mathematical machinery of modern decision theory to be brought to bear, I assume a utility function can be defined. And since the desired outcome is objective and determined by the physical state of affairs in the future, the utility should be an objective function of states of the world. An objective utility allows for consensus, in principle, as to which actions are ethical. For any given choice of actions by any given human agent, the theory will yield a consistent result regardless of who does the moral reasoning. All will agree, in principle, on the ethical course of action for any agent. Hence the theory is agent neutral.

Next, the optimal utility will be universal across the set of all humans in the society. Maximizing future viability will entail maximizing the contributions of all members of society toward that end. In other words, the contribution to utility of all humans should be uniformly considered and equally weighted. Any bias in our theory toward one individual or another detracts from optimal conditions for cooperation. Hence, the notion of fairness is built into the theory at the outset.

Further, since competition among human societies poses a significant risk to any particular society, cooperation among societies is encouraged. In fact unification across the species is truly optimal. In this case, the utility becomes universal across all humans.

It would also exclude non-intelligent species here on earth. They are not agents whose actions can contribute to the long term viability of society. This is not to say, however, that no consideration of non-intelligent species is possible, just that any such consideration is derived from the utility of humans, i.e., how other species contribute toward the long term viability of human society.

Let me summarize this important point. The question of boundaries (or scope) of the ethical theory is decided by our initial desideratum. Given the goal of long term viability of the species, ethical coverage is extended to those agents (humans until we encounter another intelligent species) with whom cooperation enhances long term viability.

A few other considerations are worth mentioning. First, the normative theory should not run counter to human psychology—basic human nature. Humans have a nature, a fundamental psychology, as a result of their evolutionary past. Any theory of ethics that requires humans to consistently act against their natures will result in conflict and unhappiness. This is a tricky condition to satisfy because human nature is complex. Different aspects of human nature are often in conflict. Most fundamental is the tension between cooperation and competition. The answer is balance; the theory should seek the equilibrium point between conflicting interests within a particular individual and among different individuals. But, beyond merely aligning with human nature, a good normative theory would use human nature as a motive force.

Second, the normative theory should agree—in general—with standard ethical intuitions. Current ethical systems have allowed human societies to be successful. For example, the Judeo-Christian system has permitted, even fostered, the unprecedented

success of the western world. In a pragmatic sense, there is clearly much that is right about that system. This criterion is often referred to in the literature as reflective equilibrium (Rawls 2001), that a moral theory should cohere with the moral convictions we have after careful reflection (Hooker 2000).

Third, the theory should be able generate rules of behavior. A utilitarian theory is unwieldy on a day to day basis. One cannot expect people to calculate or even estimate an objective utility function on states of the world as they go about their daily lives. The requisite information just does not exist and predicting the future is problematic even for meteorologists. The theory should provide a theoretical basis for simple, accessible rules for daily use. It should also provide a foundation for more global decisions involving difficult moral or policy issues. This is rule utilitarianism. In what follows, I take the structure and methodology of rule utilitarianism developed by Hooker to be substantially correct (Hooker, 2000).

These last three features are not independent of the primary desideratum. They represent my expectations of the result of the optimization procedure.

### *3. The Optimal Utility*

Now the question becomes one of defining a utility function that optimally leads toward the goal of long term viability of the species. We already have a number of conditions defined. The utility is to be an objective function of states of the world, agent neutral and universal across all humans. Furthermore, it should be aligned with human nature as far as possible, agree with most standard moral judgments and permit the definition of rules

for daily use. But most importantly, it should provide the optimum set of ethics for long term viability of human society. How will we judge this? Certainly we don't really have access to a vast computer nor the software required to run the requisite Monte Carlo simulations. What we do have are examples from history. History is replete with examples of societies that have failed, and for a myriad of reasons. Societies have been conquered from without (Carthage, the Inca, the Aztec), fallen victim to revolutions from within (Czarist Russia, Imperial China), been decimated by natural disaster (the Maya, the Anasazi) or disease (North American Indians). Others have been quietly assimilated into a technologically more advanced society (all modern primitive tribes).

Let's examine some of the standard utilities proposed in the literature. Might they be credible candidates for optimality? I start with the happiness utility of Mill (Mill, 1863). How will such a society fare? Will a society whose ultimate value is happiness win a competition with another society with a different value? History is replete with examples of (seemingly) idyllic societies being overcome by outside forces. A pastoral state may contain maximally happy citizens, but unless attention is paid to safety and defense, it may be vulnerable to natural disaster and conquest. In modern times, citizens of the Scandinavian states seem to score highest in happiness polls (Forbes, 2010), but they were swiftly overcome by Hitler's armies in World War II. One can imagine situations where happiness can be induced technologically. Larry Niven, in his *Ringworld* series of novels, described a technique where a wire is connected directly into the pleasure center of the brain. Addicts of this procedure, called wireheads, would die through neglecting their bodily needs, being unwilling to disconnect from the wire even to eat or drink. A society of wireheads, kept alive by mechanical means, would seem to



be the happiness utilitarian utopia. But it would be hard to imagine such a society having any chance at long term viability.

Most modern utilitarians define utility around a person's preferences, or, in a more objective formulation, around a person's interests or well being. To Peter Singer, a prominent contemporary utilitarian, "best consequences" is understood as meaning what, on balance, furthers the interests of those affected, rather than merely what increases pleasure and reduces pain (Singer, 1993). A formal definition of a preference utility is given by Hugh Breakey based on a hierarchical preference structure defined for each individual (Breakey, 2009). Individual utility is summed over all humans to derive a global utility value. This meets the universality condition established above and is agent neutral.

But I claim that a preference utility fails to be optimal for the simple reason that most humans seem to prefer short term gain over long term viability. In the micro sense this effect is well documented in the psychological literature on delayed gratification (e.g. McClure, 2004). In the macro sense, it is playing out today in Europe and America as many nations attempt to institute austerity measures to deal with massive debt. These debts were incurred to implement programs to satisfy the preferences (or interests) of the citizens. The funds to pay for these program were, in effect, borrowed from future generations, at the obvious and direct expense of long term viability. Yet even when faced with the impending default of their government, Greek citizens took to the streets to protest the elimination of programs that served their direct interests (Reuters, 2011). A very poignant example from history is provided by the Easter Islanders who completely

denuded their once forest covered island, even though the trees were their only means to build canoes for deep water fishing, a key to their prosperity (Diamond, 2005).

Any other utility measure based on individual welfare fails for similar reasons. Maximizing welfare in the present is not really conducive to long term viability. The United States, for example, has implemented welfare programs like Social Security and Medicaid that threaten to engender huge debts and bankrupt the country in the next several decades. But attempts to curtail it so as to provide longer term fiscal viability run into arguments that doing so comes at the expense of the immediate welfare of millions of individuals.

Welfare utilities being inadequate, I need to take another tack. I need to aim more directly at the goal. As discussed above, long term viability requires the capability for a society to cope with many kinds of challenges, both internal and external. But the problem is that we can't predict in advance the exact nature of the difficulties that will be encountered in the future. Hence what is needed is capability to do anything required in service and defense of the society. What is needed is capability, or to use a more edgy term, power. The social sciences have produced a rich literature on power (of which more below), but what I am referring to here is *power-to* in the broadest sense of capability or capacity. Is it possible to define power in a manner that will allow its use as a utility? And furthermore will such a utility be optimal in promoting long term viability of human society? The answer to both questions is yes as I will show presently.

#### *4. Power Defined*

In this section I provide a formal definition of power. I start with defining the power of an individual agent and then use that definition to extend the concept to groups and ultimately the entire species. I have attempted to introduce enough mathematical rigor so that the concept is precise and can be used, in principle, as a utility function in a formal decision theory. This allows the impressive mathematical machinery of that well developed subject to be brought to bear (Jeffrey, 1983; Bernardo and Smith, 1994).

Fundamentally, power represents an agent's capacity to make a difference in the world. A completely powerless individual is one who can effect no change in the future state of reality regardless of his actions. A maximally powerful individual can determine the future absolutely, constrained only by physical law. With this in mind, one can see that power depends on possible courses of action, and the possible future states of reality as determined by those courses of action. Furthermore, the degree of change an agent can effect is dependent on how far in the future one looks. For example, no matter what I do, what action I choose, the state of reality an instant from now will look substantially the same. However, I can affect reality one year from now to a much greater extent. I could father a child or build a house or explode a bomb.

These considerations lead me to the following definition: *the measure of power of a particular agent is the degree to which the range of actions available to the agent results in a range of different future physical states.* I call this the personal power of the agent. Within that broad concept, various kinds of power can be identified, connecting this definition back to more familiar usages of the term.

Key to quantifying power is a means to identify differences between states of reality and to quantify the degree of difference. What this means in mathematical terms

is that states of reality are elements of a vectorspace (so one can add and subtract them). And the vectorspace is equipped with a norm (to determine the magnitude of states and the magnitude of differences of states). Happily, in modern physical theory, states of the world are indeed elements of a vectorspace. To simplify the discussion, I will ignore the uncertainty in predicting future states, whether the uncertainty is epistemic (from lack of knowledge) or physical (for example from quantum mechanics).

Now I can state the formal definition of power. Let  $S_0$  be the current state of reality (at time  $t_0$ ). Further, let  $\{S_{t_{\text{end}}}\}$  be the set of possible states of reality at a future time,  $t_{\text{end}}$ . Possible states are defined through a set of possible actions or sequences of actions available to an agent up through time  $t_{\text{end}}$ . A sequence of actions:  $a_1, a_2, a_3 \dots a_n$  is called a course of action and each course of action leads to a particular state at  $t_{\text{end}}$ ,  $S_{t_{\text{end}}}^i \in \{S_{t_{\text{end}}}\}$ . The possible future states thus exhibit an event tree like structure, nodes representing choices of the agent. The agent's power is then defined as a function of  $t_0$  and  $t_{\text{end}}$ :

$$POW(t_0, t_{\text{end}}) = \max_{i,j} |S_{t_{\text{end}}}^i - S_{t_{\text{end}}}^j|,$$

where the superscripts  $i$  and  $j$  range over courses of action. The vertical bars represent the norm or magnitude of the states of the world, or in this case, the magnitude of the difference between states. The power of an agent thus represents the maximum difference between any two states of the world achievable via action of that agent.

Let's examine some consequences of this definition. First, suppose an agent has only one possible course of action available. Then, as determined by the laws of physics there will only be one possible future state and the agent's power will be zero. This

situation describes the powerlessness imposed by a lack of options. Similarly, consider the situation where the agent has many possible actions at his disposal, yet the resulting future states are all the same. Again, the agent's power is zero. This would be the situation of a man in prison without access or contact to the outside world. No matter what he does up to and including suicide, the world continues as before, unaffected by and oblivious to his actions. (In fact, one of the main objectives of this punishment is to deprive the prisoner of his personal power.) At the other extreme is an agent who has all the resources of his society at his personal disposal—like a king. He has a span of actions available to him ranging from squandering the society's resources to using them for destruction or oppression or charity.

The function  $POW(t)$  pertains to an individual agent (suppressing the dependence on  $t_{\text{end}}$ ). It is the personal power of that agent because the courses of action are restricted to those of the agent alone. In an analogous fashion, the power of a group can be defined as the power derived from all possible actions of all of the members of the group. This is not simply the sum of the personal powers of the group members (in contrast to the welfare based utilities). I call it the synergistic sum. In most cases the group power will far exceed the sum of personal powers, a consequence of the possibility for cooperation. The state tree for a group includes branch points for all the possible actions of any member of the group, hence the complexity of the tree, the number of branch points and branches, increases exponentially with the number of group members. Contrast this with the much simpler state tree for the individual. There the rest of the universe, including all other agents, was part of the environment for the decisions of the

individual. Now all the possible actions of all the group members are considered simultaneously.

The mere fact that the number of possibilities increases speaks nothing about the real difference between any of them and thus nothing about the power of the group. All these myriad states might be very similar. My claim is that the power of a group can be far greater than the sum of the powers of each individual. The gain comes from cooperation. One is often confronted by problems or tasks that cannot be accomplished alone, but are simple and easy with a helper. A large group of engineers and technicians can build a rocket to send spacecraft to other planets. This obviously cannot be done by any single individual, no matter how powerful. One could argue that the ruler of a sufficiently powerful society could order it done, but in that case, the ruler is leveraging the power of his whole society, achieved by cooperation between individuals.

The power of a society or the power of a nation can be defined in a straightforward manner based on the power of a group. Societal power is the formal power based on all the members of the society. Similarly, the power of a nation is the power based on all the individuals of that nationality. And species power is defined as the collective power of all the members of the human species. It is long term species power I propose as the optimum utility for our new consequentialist theory. By long term I mean one should set  $t_{\text{end}}$  sufficiently far into the future that its value makes no practical difference in our immediate choices.

The greatest good then is not to promote collective happiness or welfare, but collective power—not just in the here and now—but in the long term. The implications of this change in perspective are wide ranging and many of the issues facing standard

consequentialism are swept away. But first let me dig a bit deeper into the conceptual fabric of power.

### *5. Kinds of Power*

Power as a concept has enjoyed a prominent place in the history of ideas. Hobbes believed the pursuit of power to be one of the prime motivators of mankind. Spinoza made the distinction between "power-to" in the sense of capacity or capability and "power-over" as a relationship between agents. This distinction is crucial in understanding the modern social science literature in power, which is mostly concerned with power-over. Nietzsche explored the psychology of the pursuit of power and codified it with his famous phrase "the Will to Power."

More recently, power has been a topic of much discussion and debate in the social science literature. The main issues are how one group attains and maintains power over another group and how compliance of the disaffected group is secured. For example, Feminists are concerned about the role of power in the male domination of society. Marxists are concerned how ruling elites acquire and maintain power. However, there has been some recent work with regards to understanding power in the sense of power-to, notably Morriss (Morriss, 2002) and Lukes (Lukes 2005). Chapter 2 of Lukes contains an excellent up-to-date survey of the different power concepts extant in the literature. Lukes believes power-over is a special case of power-to. But even he views power in the restricted sense of a dynamical force between groups and not something to be aspired to at the level of human society as a whole. Hannah Arendt, perhaps, comes closest.

According to Arendt, "power corresponds to the human ability not just to act but to act in concert" (Arendt, 1970).

My conception of power presented in the previous section is thoroughly power-to. It is a dispositional concept; that is, it represents a disposition or potential for an agent or group to affect the future state of the world. I have defined it in terms of possible courses of actions that lead to different possible future states. Different kinds of power can be obtained by restricting our consideration variously to whose actions, the kinds of actions, or different aspects of the possible future states.

Personal power is the power accrued to a particular individual agent. It's obtained by considering all possible courses of action of the agent where other agents are viewed as part of the environment. Societal power is the aggregate power across all the members of the society determined by the range of options available to all individuals that constitute the society; it is here one gains the tremendous synergistic benefits of cooperation between members of the society. Species power extends societal power to the entire human species.

Physical power is the ability to affect the physical state of reality—that is, outside the social milieu. Physical power is the capability to construct cities: roads and bridges and skyscrapers and water systems and electrical power grids. It is the capability to move around the world in cars or airplanes or send a man to the moon or a spacecraft to Mars. It is the capability to rearrange matter to construct micro-chips and computers. And it is the capability to manipulate the genome of living things to create plants and animals of certain desired characteristics. Physical power is manifested in technology.



A specialized form of physical power is military power, the ability to prevail in armed conflicts between societies or nations. Military power is represented by the size and capabilities of the armed services, but more importantly by the quality and technology of their armaments.

Another type of power is scientific power, made up of equal parts explanatory power and predictive power. Explanatory power is the enabler, providing the context, the consilience leading to a worldview underlying all rational decision making. Predictive power fuels the engine of technology leading to physical power. The engineer uses the predictive capability of science to lead him to the physical designs of his artifacts. He relies on scientific power. Scientific power is a critical component of overall societal power. The society with the best science generally wins. Consequently, scientific power is a key to long term viability.

Political power is the ability to make a difference in the political state of a society or nation. Political power is the capability to persuade and convince other agents to adopt a course of action. It is the capability to foster cooperation, concerted and unified pursuit of a goal by a group. Political power is the power of the charismatic leader to bend the will of the people to his will, to unify his people to a common purpose. Political power perhaps comes closest to the concept of power-over that concerns social scientists.

Economic power is the ability to change the economic state of a society, the capability to make a difference in the society's system of goods and services, production and distribution. Economic power can often be identified with wealth, wealth being commutable into economic influence and difference making. In the United States, great economic power is held by the government with its ability to set interest rates, tax rates

and overall economic policy. Large chunks of economic power are also wielded by corporations, legal entities formed to achieve some economic goal, or the few very wealthy individuals.

Bertrand Russell drew the analogy between power as a fundamental concept in social science and energy as a fundamental concept in physics (Russell, 1938). Let me expand a bit on this analogy. Just as one form of energy can be converted into another, different kinds of power are also convertible. For example, scientific power can be converted into physical power or military power through the development of technology. Economic power can lead to scientific power through funding of scientific research programs. Conversely, scientific power can lead to economic power again through technology and the development of new goods and services. This closed loop process is responsible for the incredible advances experienced in the West over the last several hundred years. The scientific revolution begun by Bacon and Descartes led to the industrial revolution which led to even greater science which led to more powerful technology and so forth.

## *6. Two Moral Laws*

Having settled on species power as the utility, the next step is to use that utility to derive moral laws which will form the basis for a moral code in the rule utilitarian sense. According to Hooker, the ideal moral code is the one whose internalization (acceptance) by the vast majority of the population leads to the best consequences (Hooker 2000). We

have determined that the best consequences is the maximization of power. Hence, we seek a moral code that maximizes the power of the human species.

To determine laws, it is helpful to think back to the mathematical definition of the human society and borrow a technique from mathematical physics. Species power is obtained by considering the collective courses of action of all humans. Hence it can be viewed as a mathematical function of the personal powers of each human (as well as factors like culture, science, and technology). As such, species power can be expanded in a *power series* of personal powers. (The italicized usage of power represents the mathematical term meaning exponentiation.) The first order term is then the sum of the personal powers of the members of society. Higher order terms contain products of personal powers which capture the effect of cooperation. Hence we have:

$$\text{Species Power} = (\text{sum of personal powers}) + (\text{power due to cooperation}).$$

Considering the first term it is clear that, all other things being equal, species power increases when any individual's personal power increases. The caveat means that the power of an individual can increase, but if that increase comes at the expense of a decrease in the power of another individual, species power might decrease. This observation leads me to the first of two ethical principles of the power theory.

*The increase in power of any individual or group of individuals is good so long as the increase is not exceeded by a concomitant decrease in the power of others.*

I will call this the law of individual achievement, or the law of the individual, for short. The converse of this law also holds: Any decrease in the power of an individual or a

group is bad (again so long as the decrease is not exceeded by a concomitant increase in the power of others).

The second term is dependent on the effectiveness of the means of cooperation among individuals. The greater the effectiveness of cooperation, the greater will be the aggregate power. Thus I come to the second ethical principle of the power theory:

*The increase in the effectiveness of the means of cooperation between individuals and groups of individuals is good.*

I will call this the law of cooperation. As before, the converse of the law of cooperation also holds. A decrease in the effectiveness of the means of cooperation is bad.

These two laws go together to maximize overall power, but they can seem to be at odds. This is the fundamental tension between competition, by which individuals seek to increase their own power, and cooperation, by which individuals work together to increase group power. The phrase “so long as the increase is not exceeded by a concomitant decrease in the power of others” is my attempt to relax the tension by proscribing destructive competition. But the real key to the dilemma is balance: to value the achievements of individuals, but not at the expense of others and to value the group and its institutions that foster cooperation, but not at the expense of individuals.

The give and take between individual achievement and group achievement shows up at many levels. Within each individual there is a tension between the internal drive for personal gain and the externally enforced constraints for social cooperation. For example, sports teams seek the most talented individual players but also have to be mindful of team chemistry and effective cooperation. Many teams rich in talent do not succeed because the talented players put personal glory ahead of the team’s goals.

Playing like a team is a cliché describing groups that achieve effective cooperation. The same drama plays every day on the national political stage. The political left emphasizes the group—the socialist ideal. The political right focuses on the individual. Neither side is interested in balance. Nevertheless, some modicum of balance is achieved, at least in America, because neither side can gain absolute control. The principles and institutions of government are constructed to maintain balance and stability.

The law of the individual states that we should value, as a highest value, the achievement of power by individuals. It follows, then, that we should also value those principles that allow and encourage individual achievement. Among the most treasured of these are individual rights and freedoms. Freedoms include freedom of thought and speech, freedom of association and assembly and freedom of religion. Rights include the right to life, the right to security, and property rights.

The law of cooperation states that we should value, as a highest value, those things that enhance cooperation among individuals. From this value other values can be derived including many that have become central to modern life. The principles of equality and diversity state we should value the contributions of all individuals regardless of race or gender or other characteristics. This implies that any system of governance have no preference—explicit or implied—for some groups over others. Principles of reciprocity and charity enable a social system in which all individuals can feel secure, which, in turn enhances their value in cooperative ventures.

As indicated, these two laws can be used to fully develop a code of ethics which can then be used to address issues in practical ethics. For now, I leave that task for the future. Let me instead shift the focus to the national level. The current world is

composed of nations: self contained societies, each with its own social system and set of ethical principles. Is it meaningful to speak of ethical conduct between nations? Historically, nations have tended to compete more than cooperate and often that cooperation took the form of temporary alliances to defeat other nations. But I deliberately defined utility at the species level. To maximize species power, nations must find a way to cooperate.

The first step is to develop principles for cooperation at the national level. One can start with international analogues of the two fundamental principles. Thus one has that the increase in power of a nation is good so long as the increase does not come at the expense of other nations. And the increase in the effectiveness of the means of cooperation between nations is good. As principles, these will do nicely, but what is lacking on the international level is a means of enforcement. In the meantime, much progress has been made in globalization through free market mechanisms. International cooperation has occurred at the level of businesses and individuals resulting in great economic gains.

The benefits to species power from global cooperation will be immense. There are critical issues facing the species that can only be addressed globally. Foremost is security. Many individuals in today's world live in constant fear for the physical well being of themselves and their families. This drastically curtails their personal power. Environmental degradation is another critical issue. Harm to the earth's environment has the potential to severely impact the power of the species. Although I believe it uncertain that the current phase of climate change is human caused, it is imperative that climate

dynamics be much better understood so effective action can be taken if needed. And there are many other real environmental issues to tackle.

## *7. Conclusions*

I have proposed species power as the optimal utility for a consequentialist theory of ethics. I have defined optimal as that which leads to the highest long run viability rate for society in particular and the human species in general. But is power optimal? My claim is that it is optimum or very nearly optimum. I will provide an intuitive argument and a historical argument.

Intuitively, power is optimal because it directly aims at the goal of long term viability. It is impossible to predict what exactly society needs to survive in the long run. Power provides general capabilities to counter a wide range of dangers. Military power provides the capability to resist aggression by hostile external societies or even a hostile alien race. Scientific power through medical technology provides the capability to resist decimation by pandemics. Green technologies provide the means to cope with or avoid environmental disasters. Through engineering advances we can design buildings to resist earthquakes and deploy satellites to predict and track hurricanes. And space technology may allow the future dispersal of the species throughout the solar system and beyond, providing the capability to survive many global scale catastrophes.

Historically, we see powerful societies prevail time and again. The most successful, long lived society in antiquity was the Roman Empire lasting over one thousand years. It was also the most powerful. Western societies fueled by the

tremendous increase in power provided by the Enlightenment and the advance of science, prevailed in encounters with weaker societies from the Aztecs to the Inca to the Maori to the Japanese and Chinese. In modern times, it took the scientific, technological and industrial power of America and her allies to defeat the threat of Nazi Germany and Imperial Japan in World War II. Similarly, superior economic and military power allowed America to prevail in the Cold War against the former Soviet Union.

What about the other considerations I developed in Section II? How does the power theory with its two moral laws fare here? The first consideration was that the moral theory should not run counter to human psychology. My claim is that an ethical theory based on power is aligned with one of the most fundamental principles of human psychology, which (borrowing the phrase from Nietzsche) I call *the Will to Power*. Justifying this claim deserves a book length argument, but let me briefly summarize it here. Humans are products of biological evolution. The processes of evolution have created the human organism to be optimized for survival and reproduction. These attributes are captured in an overall measure called fitness. The fitness level achieved by the human species is the highest yet seen in the 3.5 billion years of existence of the biological world. In the mere 100,000 years since the first modern humans appeared in Africa, they have spread over the entire globe and reached a population level in excess of six billion individuals.

The key attribute of the human organism that has led to this incredible level of fitness is, I claim, conscious intelligence. The adaptive function of intelligence is prediction and control; i.e. power. The mental mechanism that enables prediction is representation. The capacity for the brain to build a mental representation of the external



world, both social and physical, was the adaptation that set modern humans apart from their less capable forbearers. The capacity to represent the external world enables the human organism to predict the results of its actions and thus gain a measure of control over its surroundings. In other words, representation begets prediction; prediction begets control; and control is power. Hence, the desire to control (the Will to Power) is one of the most fundamental features of human psychology, and an ethical theory based on power is aligned with the psychology of the species.

A similar conclusion can be reached by considering the two moral laws as compared with more detailed psychological factors. The law of the individual captures our needs to advance our own individual interests while the law of cooperation captures our needs to be a part of a larger social group. Exploring these connections is again beyond the scope of this paper, but I think the reader can readily see the plausibility of the claim.

The second consideration was that the theory should be in general agreement with our moral intuitions. My claim is that the two moral laws derived from the power theory do indeed lead to rules of conduct that are in general agreement with our intuitions. Again this could be the subject of a book length exposition. But in general, our moral proscriptions against harming others are built into the law of the individual. Our moral prescriptions for helping others in need are built into the law of cooperation.

The third consideration is that the theory should be useable to generate rules of behavior in the rule consequentialist sense. This consideration was met explicitly by constructing the two moral laws. Other cherished rules follow from these two. Rules

protecting individual rights and freedoms follow from the law of the individual. Rules for charity and equality follow from the law of cooperation.

I also believe the power theory avoids some of the problems that have plagued other forms of consequentialism. To me, the most glaring issue with traditional utilitarianism is the difficulty in defining the utility. Whether happiness or preferences or welfare, we are left with the question of what counts as happiness, or which preferences or aspects of welfare to include. To me, there is no means—even in principle—to resolve these questions, so authors are left to appeal to their own intuitive biases. The power theory, in contrast is based on an objective measure defined with mathematical precision that is based on the single goal of preserving the viability of the human species in the long run.

Another issue that troubles traditional utilitarianism theories is that they seem to be overly demanding with respect to charitable contributions by individuals. If the welfare of each individual counts equally toward utility, and based on the principle of marginal utility, then traditional utilitarianism would imply that the relatively well off should give nearly everything to the poor and starving. Some authors embrace this conclusion (Singer, 1993), others find it troubling and at odds with our moral intuitions (Hooker, 2000).

In the power theory, the issue is not so clear cut. The power theory does encourage charitable giving because such giving enables destitute individuals to enhance their own contributions to the power of society. This is one of the effects of the law of cooperation. On the other hand, a wealthy person might contribute more to the power of society by investing in new technology or creating industries that employ thousands or

contributing to scientific research. The power theory seeks to find the optimum balance in such choices.

One other issue with traditional utilitarianism deserves mention. Due to the linear combination of individual utilities into overall utility, it follows that we should prefer a society of very large population where each individual has a low but positive utility to a much smaller society where each individual has a very high utility. Called the repugnant conclusion, this objection was first posed by Derek Parfit (Parfit, 1984), and has resisted all attempts at resolution by traditional utilitarian theories. The power theory avoids this pitfall for the simple reason that the power of society is not a simple linear sum of the powers of individuals. Cooperation plays a critical role. The population of a society and the powers of its individuals is not enough to settle the issue. Equally important is how the society is organized for cooperation and its access to resources, science and technology.

In summary, adopting power as the utility in a consequentialist program leads to a satisfying normative theory of ethics. This theory meets all the requirements for a normative theory laid out in section 2 and it avoids many of the problems that have troubled traditional utilitarianism. It leads to moral rules that preserve and justify many of our cherished values. Having this theory in hand yields many benefits. It provides an overarching theoretical framework from which to address ethical issues. Since power is an objective measure, it provides, in principle, a means to reach consensus on these issues. Society needs laws and other rules to foster cooperation and maintain harmony, but the laws and rules deserve a self consistent foundation. The power theory provides

this. Most importantly, the power theory is explicitly geared toward the long term success of our species.

Beyond ethics, the power theory can be expanded to a general theory of action. My proposal for the ultimate utility is long term species power. That is, each individual, each group, each state and organization of states should act in a way so as to maximize the long term power of our species as a whole. The ultimate goal of humanity is thus ultimate power.

The fundamental principle is simple: Any increase in the power of humanity is good. Any decrease in power is bad. This broad principle can guide policy decisions at all levels. Specific policies that would be advanced by the power theory include promoting science and technology, political globalization and human exploration and expansion into space. Viewed in this way, the power theory offers a positive, hopeful and optimistic vision for the future.

## REFERENCES

- Arendt, H. 1970: *On Violence*. Harcourt Books. Orlando.
- Bernardo, J. M. and Smith, A. F. M. 1994: *Bayesian Theory*. John Wiley & Sons. New York.
- Breakey, H. 2009: The Epistemic and Informational Requirements of Utilitarianism. *Utilitas* Vol. 21, No. 1. Cambridge University Press. Cambridge.
- Diamond, J. 2005: *Collapse*. Penguin Books. New York.
- Hooker, B. 2000: *Ideal Code, Real World*. Oxford University Press. Oxford.
- Jeffery, R. 1983: *The Logic of Decision, Second Edition*. University of Chicago Press. Chicago.
- Lukes, S. 2005: *Power, a Radical View, Second Edition*. Palgrave Macmillan. New York.
- McClure, S. et. al. 2004: Separate Neural Systems Value Immediate and Delayed Monetary Rewards. *Science* Vol. 306.
- Mill, J. S. 1863: *Utilitarianism*.
- Morriss, P. 2002: *Power, a Philosophical Analysis, Second Edition*. Manchester University Press, Manchester.
- Parfit, D. 1984: *Reasons and Persons*. Clarendon Press. Oxford.
- Rawls, J. 2001: *Justice as Fairness, a Restatement*. Harvard University Press. Cambridge, MA.
- Russell, B. 1938: *Power*. Routledge. New York.
- Singer, P. 1993: *Practical Ethics, Second Edition*. Cambridge University Press. Cambridge.
- Forbes, 2010: <http://www.forbes.com/2010/07/14/world-happiest-countries-lifestyle-realestate-gallup-table.html>
- Reuters, 2011: <http://www.reuters.com/article/2011/10/20/us-greece-idUSTRE79H1FI20111020>