

The Power Utility

George F. Sowers Jr.
6346 Willow Springs Dr.
Morrison, CO 80465
USA
Tel: 303-697-8914
Email: gfsowers@comcast.net;
george.f.sowers@ulalaunch.com

The Power Utility

A normative theory of ethics is developed based on a utility of power, rather than the more typical utilities of happiness or preference. Conditions for a satisfactory normative theory of ethics are offered. A precise definition of power is introduced corresponding roughly to a measure of an agent's capability to affect a difference in the world. Societal or species power is defined as the aggregated power of individuals. The power utility is then defined as the total power of the human species in the long run. The utilitarian goal of maximizing power then leads to two main 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 so long as it does not come at the expense of the individual. These general laws can then be used to derive specific rules of behavior and address practical ethical issues. The power utility is justified via naturalistic and pragmatic arguments.

1. Introduction

One condition for truth is effectiveness: true beliefs lead to effective actions. If certain actions are effective, that effectiveness counts as evidence for the truth of the beliefs that lead 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 is a great intellectual value: that we should place the highest value in ideas that work, ideas that lead to outcomes we desire.

When applied to the field of normative 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 quantity that represents the goals of an agent or group of agents. Hence, these theories are also referred to as utilitarian theories. Utilitarian theories of ethics were first proposed by Jeremy Bentham (1748–1832) and refined by John S. Mill (1806–1873). Bentham and Mill proposed that aggregate happiness be used as the utility of the theory. In other words, utility is defined as the sum of happiness across all individuals within society. Then actions are judged on

their effect on this quantity. Actions that improve aggregate happiness are good; actions that degrade overall happiness are bad.

Since the days of Bentham and Mill, utilitarianism has enjoyed a fair amount of success as a normative theory of ethics. It is now a full blown school. But utilitarianism is bothered by a number of conceptual issues. Many of these issues revolve around the difficulty in defining happiness in a manner that allows its use as a quantifiable utility. These problem have led to the development of alternative utility measures. In vogue now are utilities that focus on the interests or preferences of individuals. Other issues arise in the practical difficulty in being able to calculate an action's effect on aggregate happiness. Much of the requisite information required to perform such a calculation would not be available to an individual contemplating a decision. This has led to rule utilitarianism, whereby the utilitarian calculus is used to devise a set of rules of behavior to be considered by individuals as either rigid, like laws, or simply guidelines or rules of thumb.

In this paper, I introduce a different utility measure, aggregate power at the species level. I will 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. Furthermore, a power based utility is justifiable in a broader scientific and evolutionary context, bridging the infamous is-ought gap.

I begin in Section 2 by introducing the logical foundations of decision theory. Utilitarianism is an application of this formalism. In Section 3, I propose some principles for a normative theory of ethics. What do we look for in such a theory? These principles then form a set of premises for subsequent logical development of the theory. Section 4 discusses the utilities typically used in utilitarian theories of ethics, happiness and preference, and some of the difficulties with these theories. Section 5 provides a formal definition of power as a utility. One of the advantages of this definition is its objectivity, avoiding the subjectivity of measures involving the happiness or preferences of individuals. 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. I explore briefly some of the consequences of these laws and the power theory in general. In Section 7, I provide a series of arguments to justify the use of

power in a utilitarian theory of ethics. These arguments are primarily naturalistic and might cause some to raise the standard objection of the naturalistic fallacy. My counter argument shows that the naturalistic fallacy is more of a semantic misunderstanding than a logical fallacy. Finally Section 8 contains my conclusions.

2. *Decision Theory*

Humans are agents—with free will. Agents are purveyors of actions; they are doers. As humans look to the future, they are faced with an infinite panorama of possible actions. How are they to choose? One *point* of philosophy is as a guide to action, a roadmap through the choices, big and small, faced on a continuous basis. Ethics is the branch of philosophy that deals with rules of action. Ethics, for individuals or societies, prescribes how to act, what to do. Ethics, then, are rules for decision making.

The branch of logic that deals with decision making is called decision theory. A large literature and powerful mathematical machinery have been assembled over the last half-century to deal with decision problems (Savage, 1954; Luce and Raiffa, 1957; Jeffrey, 1983; Bernardo and Smith, 1994). Utilitarianism is the branch of ethics that is aligned with the logic of decision theory and as such can take advantage of its power and coherence.

Decision theory is a theory of action by a rational agent. The agent has available a set of possible actions $\{a_i\}$, where i runs over some index set. These actions lead to a set of consequences, or future world states resulting from the actions. The central element of decision theory is the utility function, denoted by U . If s is a state of the world, embodying all information about the world relevant to the agent, then $U(s)$ is a real number. By convention, the value of the function U is contained in the real interval $[0 \leq U(s) \leq \infty]$. Conceptually the utility function $U(s)$ represents the preferability (according to the agent) of the state s . Thus, if $U(s_1) > U(s_2)$ then one would say that state s_1 is preferred to state s_2 . Since $U(s)$ is a real number, preferences are ordered per the natural ordering of the real numbers. This ordering satisfies the requirements of rationality. Hence, given any two states s_1 and s_2 , either s_1 is preferred to s_2 , s_2 is

preferred to s_1 or there is no preference. In terms of utility either $U(s_1) > U(s_2)$, $U(s_2) > U(s_1)$ or $U(s_2) = U(s_1)$, a condition clearly satisfied by the real numbers $U(s)$. In addition, one has the transitive property that if s_1 is preferred to s_2 and s_2 is preferred to s_3 then s_1 is preferred to s_3 . Again, this requirement is satisfied by the utility representation.

In a world without uncertainty, the theory would essentially stop here. Given a set of actions or theory choices $\{a_i\}$, the agent would simply choose the action which leads to the greatest utility—his most preferred state of affairs. In other words, there would be a specific state s_i associated with each choice a_i . The agent would evaluate the utility of each state, $U(s_i)$, then select (and execute) the choice associated with the greatest value. However in the real world, agents are not able to precisely predict the consequences of their actions, nor precisely evaluate the utility function of a given state. The mathematical details are not important here, but ultimately one obtains a probability distribution function of the utility for each action. This distribution codifies the uncertainty of the agent that the action a_i will actually lead to any particular value of the utility $U(s_i)$.

In these discussions I am assuming a subjective, epistemic interpretation of the probability measure: probability represents the rational degree of belief of the agent. More precisely, the integral of the probability distribution for $U(s_i)$ between two limits represents the degree to which the agent believes the true value of (future) utility to lie between those limits. In standard decision theory it is assumed that the agent bases his or her decision on the expected value of the utility distribution.

3. Principles for a Theory of Ethics

What kind of attributes should one look for in a normative theory of ethics? First, it should be a consequentialist approach based on an objective utility function. The utility should be an objective function of states of the world. This allows for a consensus, in principle, as to which actions are ethical. Consensus provides a crucial symmetry. For any given choice of actions by any given human agent, the ethical theory should yield a consistent result regardless of who does the moral reasoning. All should agree, in principle, on the ethical course of action for any agent.

Second, the utility should be universal across the set of all humans. In other words, the contribution to utility of all humans should be uniformly considered and equally weighted. This characteristic explicitly recognizes the evolutionary purpose of ethics—to foster cooperation and social harmony. Any bias in our theory toward one individual or another detracts from optimal conditions for cooperation. This puts the notion of fairness into the theory at the outset.

Universality applies only to members of the human species. Again, this feature is in recognition of the evolutionary purpose of ethics, social cooperation. It is just not possible to cooperate with other species, at least none that currently exist on earth. In other words, it makes no sense for humans to recognize the interests of non-human organisms because they have no way of recognizing human interests. This is not to say, however, that no consideration of other species is possible, just that any such consideration is derived from the utility of humans.

Third, 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 humans do not have just one nature. They have many natures that are often in conflict. Most fundamental is the natural 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. Beyond merely

aligning with human nature, a good normative theory would use human nature as a motive force.

Fourth, 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. There is clearly much that is right about that system.

Fifth, 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 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 issues. This is rule utilitarianism.

Sixth, and finally, the new theory should yield an ethical system that is optimal. I define optimal as that which provides the best probability of success. But how does one define success? An operational—pragmatic—definition hinges on competition. Which system provides a society with the best chance of winning in a competition with other societies? The nature of the competition is unspecified; it can be military, economic, cultural or in response to some natural phenomenon like an earthquake or hurricane. Pick one or all. This condition drives robustness into the system. The optimality condition is absolutely crucial. It is the crux of how the naturalistic fallacy is avoided, as I will show below. I have defined a goal, and that goal is winning.

These six features can be viewed as premises of the ethical theory. As I will argue in Section 6, they are justified by the facts of evolution and the higher value of pragmatism. They are also not independent. The sixth condition, the optimality condition, is the most fundamental and the others follow from it along with other facts of the natural world.

4. Utilitarianism

The main challenge in constructing a utilitarian systems of ethics is to choose the utility function. Bentham and Mill proposed a utility based on happiness or pleasure. In the words of Mill:

The creed which accepts as the foundation of morals, Utility, or the Greatest Happiness Principle, holds that actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness. By happiness is intended pleasure, and the absence of pain; by unhappiness, pain, and the privation of pleasure (Mill, 1861).

The pleasure-happiness utility has attracted a large following over the years. In many respects, it carries a compelling appeal. How can one argue against valuing what makes people happy? It gathers under its broad umbrella most of the qualities of life that are universally valued: love, affection, money, good food, beauty, etc. Anything that makes anyone happy qualifies. But this is also its downfall.

Preference utilitarians define utility around a person's preferences, or, in a more objective formulation, around a person's interests. 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). This reduces to the happiness utility with the reasonable assumption that satisfying one's interests tends to make one happy. A formal definition of a preference utility is given by Hugh Breakey based on a hierarchical preference structure defined for each individual (Breakey, 2009). Both happiness and preference utilitarians sum individual utility over all humans to derive a global utility value. (In fact, many present day utilitarians extend the concept of utility to all living beings.) This meets the universality condition established above.

Preference or happiness utilities are inadequate in two ways. First they are vague and subjective, violating my first condition above. Happiness and preference cannot easily be defined as an objective function of states of the world. To quantify such a utility, we could imagine a poll of everyone as to their personal degree of happiness. But it would be impossible to calibrate the responses, one to the other, in order to arrive at a global measure. In practical terms, the inability to define an objective, quantifiable

measure may not matter much. As practical utilitarians have shown, one can make progress with only vague concepts. For example, Singer goes far with the simple principle of minimizing suffering. Breakey simply assumes there is an objective measure of the degree to which an individual's preferences are fulfilled without specifying how such a measure is obtained.

Second, the happiness utility is not optimal in the sense of condition six above. 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. 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 again, I would argue that it is greatly sub-optimal.

This is where the slipperiness of the happiness criteria comes to bear. The happiness utilitarian may argue that the pastoral society should pay attention to safety and defense to insure its happiness in the long term. But this choice depends on an evaluation of risk and a view to the future. There may be other definitions of a happiness or preference utility that could avoid some of these issues, but I choose to take a different tack.

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.

There is some debate among utilitarians as to the merits of act utilitarianism versus rule utilitarianism. Act utilitarians believe that each act by a moral agent should

be judged relative to the objective standard of maximizing utility. Rule utilitarians, on the other hand, would devise a set of rules for behavior based on the utilitarian calculus and judge behavior relative to the rules.

To me, the debate is not one versus the other but when one should use rules versus when one should evaluate an individual act itself. Society needs rules—to include laws—that are specified and easily interpreted by individuals. The utilitarian calculus should be used in devising effective laws. In other words, a law (rule) should be such that it maximizes utility when adopted by the society at large.

But everything can't be specified by laws, and there will always be situations where the general legal framework provides no guidance. In those situations, individual acts must be judged on their own merits. Yet even here rules of thumb may be applied to reduce the calculation burden. The law in America does not mandate honesty except in certain circumstances like contracts or courtroom testimony. Nevertheless, most individuals adopt a rule like 'honesty is the best policy.' Rules of thumb are not absolute; the honesty rule is routinely violated, by telling white lies, for example. In this case the honesty rule may be superseded by a rule for avoiding embarrassment or conflict.

Rules are necessary and convenient. They are necessary to provide a consistent interpretation of morality to all members of society, which, in turn, is required to keep society harmonious and cooperative. That a satisfactory ethical theory should provide rules is condition five above.

5. The Power Utility

In this section I provide a formal definition of power. Conceptually, the measure of power relative to a particular rational agent, is the degree to which the range of actions available to the agent results in a range of different future physical states. Within that broad concept, various kinds of power can be identified, connecting my formal definition back to more familiar usages of the term. As I will show, the formal definition spans most of these usages, but without the kinds of pejorative connotations sometimes induced by the word "power." Personal power is the power accrued to a particular individual. 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. Technological power, the power of engineering, is an example of this.

Fundamentally, power represents our 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 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.

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). Typically, 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_{end} . Possible states are defined through a set of possible actions or sequences of actions available to an agent up through time t_{end} . A sequence of actions: $a_1, a_2, a_3... a_n$ is called a course of action and each course of action leads to a particular state at t_{end} , $S_{t_{\text{end}}}^i \in \{S_{t_{\text{end}}}\}$. The possible futures thus exhibit a state tree like structure, nodes

representing choices of the agent. The agent's power is then defined as a function of t_0 and t_{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.

I claim without proof that $POW(t_0, t_{\text{end}})$ is a decreasing function of t_0 (for fixed t_{end}) and an increasing function of t_{end} (for fixed t_0). That POW is a decreasing function of t_0 follows from the tree structure of the set of possibilities. As time passes and a specific course of action is followed, possibilities are eliminated. That POW is an increasing function of t_{end} is somewhat tougher to see. If additional actions are available to the agent after t_{end} , the conclusion is obvious. If not, then potentially chaotic physical divergence of states should cause POW to increase.

The function $POW(t)$ as derived above pertains to an individual agent (suppressing the dependence on t_{end}). As such it can be identified with the more familiar

concept of personal power. That is, the personal power of an individual at a particular time t is formally quantifiable by the mathematical function. The power of a group can be defined as the power derived from all possible actions of any of the members of the group. This is not simply the sum of the personal power of the group members, in contrast to happiness or preference utility. I call it the synergistic sum. (Note that the synergistic sum avoids the repugnant conclusion.) 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 backdrop 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 problem or task that cannot be accomplished alone, but is 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. All of the benefits of society accrue to groups. And by extension, the benefits accrue to the individuals who make up those groups.

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. The zenith in the hierarchy of group power is species power defined as the collective power of all the members of the human species. And it is long term species power I propose as the utility in my normative theory of ethics. By long term I mean one should set t_{end} sufficiently far into the future that its value makes no practical difference in our immediate choices.

It is useful to distinguish some additional flavors of power. 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.

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.

Physical power is the power of the engineer, the capability to change the physical state of the world. Here I am using the term 'physical' in a somewhat restricted sense to distinguish it from the previous qualifiers 'political' and 'economic.' Certainly political power is quantified by physical differences between states, but these differences are by and large relegated to differences in brain states, the minds of other agents. Pure physical power is the ability to manipulate non-mental aspects of the state of the world. 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 the 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 ingredient of species power.

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. The Two Moral Laws

Long term species power as the utility in a consequentialist theory meets all six criteria set forth in Section 3. First, long term species power is explicitly objective and, in principle, quantifiable. Second, since it is the aggregate of power across the species, it is universal introducing no preferences for any individual or group. Third, the power utility is directly aligned with human psychology as I will show in Section 7. Fourth, I will argue below that when practical ethics are derived from the power theory they are in general agreement with most of our ethical intuitions. I will also show that effective rules can be derived from the power principle, satisfying condition five. And finally, the power theory is explicitly optimal. By its very definition, power wins.

What kind of values can one derive from species power as a utility? What are the implications of the power theory? Species power can be viewed as a function of the individual powers of each member of the species (as well as factors like culture, science, and technology). It is more than a straight sum—it's a synergistic sum—but it has some of the properties of a sum. For example, the partial derivative of the function is positive with respect to each individual power. This means that, all other things being equal, species power increases when any individual's 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 synergistic sum of power 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 two ethical principles of the power theory lead to many of our familiar ethical rules. The law of the individual leads to concepts of individual rights and liberties. The law of cooperation leads to concepts of equality, diversity, reciprocity and charity.

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.

What can the power theory say about some of the pressing ethical issues of the present day? An entire book could be written, but I will address a few to illustrate the ideas and the process for connecting the power utility to practical ethics.

There are a number of controversial ethical issues revolving around life: when and how it begins and when and how it ends. In general, human life is precious. Life is the source of power and power is the ultimate value. But life in and of itself is not the goal. It is a means to power and achievement. The law of the individual implies that each individual has rights, among the most fundamental is right to life, but only insofar as to not infringe upon the rights of others. When one individual's right to life infringes on the rights of others, a fundamental conflict arises. This kind of conflict of rights is the crux of the issue of abortion.

Abortion is one of the most divisive issues facing our society today. On one side are those who believe any form of abortion is wrong—tantamount to murder. On the other side are those who focus on the rights of the mother and believe abortion is permissible as a choice of the mother. Much of the debate swirls around the moral status of the fetus. Many abortion opponents believe the right to life must be granted to the fetus based on a theory of sanctity of human life, often religiously inspired. This theory maintains that the intentional ending of a human life (or potential human life) is wrong no matter the circumstances. The power theory holds that the *quality* of life is as important as the *fact* of life. The intrinsic value in a human life is in the power and the prospects for power it delivers to the whole.

So how does one judge the abortion question? Following decision theory, one can cast the two sides as options or scenarios and evaluate the utility of each. Scenario one defines personhood as beginning at conception. This grants the right to life to the fetus. Scenario two defines personhood as beginning at birth. This denies the right to life to the fetus. What are the consequences of each scenario? First, scenario one has set up an irrevocable and continuing conflict between the rights of the fetus and the rights of the mother and thus her power. It is a fact that that the fetus inhabits the body of the mother until the time of birth. If the fetus is regarded as a moral and legal individual, the rights of the mother will be severely curtailed. The consequences of this inherent conflict are well known from the times when abortion was illegal. Illegal abortions are performed

with inherent health risks. Difficult pregnancies are carried to term with inherent health risks. Severely disabled babies are born. With abortion illegal, more babies would be born into the world to parents who are ill prepared or unwilling for the tremendous commitment required. Hence, the end effect of scenario one is a decrease in power for mothers, children and society in general.

In scenario two, most of the adverse consequences of scenario one are avoided. Since fetuses are not defined as persons, there is no inherent conflict with the rights of the mother and thus her rights prevail. With the option to terminate medically dangerous or unwanted pregnancies, the mother's health can be better addressed and babies would be born to mothers that truly desire to be parents and would presumably prepare for the challenge.

In my view, the power theory favors scenario two. A society where babies are universally wanted and mother's rights are unfettered will be more powerful than the alternative. It should be well understood that I am not recommending abortion. It would be true for both scenarios that effective birth control is the preferred method to regulate reproduction. The optimum situation is one where neither abortions nor unwanted pregnancies occur. Yet in the event there is an unwanted pregnancy, it is far better that the mistake be corrected rather than becoming amplified in a waterfall of ill consequences.

The rapidly expanding field of biotechnology also provides a rich harvest of ethical issues. I'll address two. First, is it ethical to use technology to enhance the capabilities of the human body? What about the mind? Second, are there ethical issues involved with using technology to reduce or eliminate the effects of aging?

In general, the power theory would fully sanction the use of technology to enhance human capabilities in any and all ways. Biotechnological enhancement of powers is just a specific example. Based on the law of the individual, any increase in the power of an individual gained through biotechnological means is good. What one must pay attention to is the harm an improvement may do to others as well as any harm that might come to our institutions of social cooperation.

Consider a specific example. A cochlear implant is a device containing an electrode array placed into the cochlea (inner ear) that allows severely hearing impaired

individuals to hear. It works in conjunction with an external microphone and signal processor, bypassing the ear's mechanical transmission of sound and directly stimulating the acoustic nerve endings in the cochlea. Tens of thousands of previously deaf individuals are now leading substantially normal lives. This is a therapeutic use of technology, correcting a defect. I know of no ethical objections to such usage.

But now suppose that the system is improved such that hearing acuity exceeds ordinary human capability. Normal individuals may be willing to pay to obtain the treatment and enhance their hearing. Should such treatment be allowed? Based on the power theory, the answer is a clear yes. I can imagine situations where enhanced aural capability would be quite useful. But the question can be extended. Enhancements in sight or strength or reflexes are right around the corner. And they may be produced genetically instead of electronically.

There are several ethical issues to address—more along the lines of cautions. First, does the technology work as intended or are there side effects or unforeseen consequences? This worry is especially acute with genetic enhancements. Genes are multi-functional and work in delicate harmony with other genes. What happens when one is changed? And how are the almost certain side effects handled? One way to deal with this kind of issue is through the current regulatory process. Any enhancement would be proven safe, with all side effects known, before allowed into the market. This will put a tremendous burden on genetic enhancements due to the limitations on clinical trials. The power theory would prefer a more balanced approach based on a risk to benefit assessment.

The second ethical issue is that of fairness. Will new class divisions emerge, an upper class of enhanced individuals versus a lower class of the unenhanced? It is a real possibility. Given the likely high cost of enhancements, they will be affordable by only a few, at least at first. Similar situations exist today. Cosmetic surgery is widespread and affordable to only a subset of society. There are at least two mitigating options. First, the social welfare system can make therapeutic enhancements widely available to those who can't afford them. And second, the economic system, if sufficiently robust, will create wealth so that the lower class continues to shrink. And as is the case with all technology,

the cost of enhancement technologies will decrease with time making them more and more widely affordable.

The question of life extension is answered similarly although additional ethical concerns arise. Over the last century the average life expectancy of an American has increased from 48 years to 78 years, over a 60% gain. An additional 60% gain would place it over 125 years. Again, based on the law of the individual, any life extension has to be viewed as good. And certainly no one would want to roll back the gains we have achieved. But there are several concerns.

First, an increase in average lifespan would necessarily be accompanied by an increase in population density unless birth rates were severely curtailed or additional living space obtained. A severely lowered birthrate is a danger. One source of power in human society is the constant refreshing of the population. New young people mean new energy, new ideas and fresh perspectives. In many branches of science, the best work is done by new faces, scientists in their twenties and thirties. Meanwhile, the role of the older generation is to teach and represent the mainstream. Increasing lifespan will reduce the proportion of the population in those young vital productive years.

Second, the overall social demographics will change. Due to the increase in life expectancy, today we have a much greater proportion of older people than ever before. Many of these people are retired and expect to be able to live at the expense of society, becoming net consumers, not net producers. The way to mitigate this effect is to increase the age of retirement as life spans increase, thus increasing the productive years proportionately.

This brings me to the final point. One must be certain that the quality of the extra life is on a par with what exists now. What benefit is 10 or 20 or even 100 extra years if one's faculties are severely diminished? One hundred years of feeble existence may not be worth having. On the other hand, an extra 100 years of vital life would be a blessing beyond estimation. Imagine knowing your great-great-grandchildren and witnessing the growth of multiple new generations. Imagine having multiple full careers. Who knows, perhaps the productivity of the 20 year old returns when entering a new field properly educated. And perhaps productivity in the new field would be enhanced by one's experience in another. It's a joy to imagine the possibilities.

Extending useful life has the potential to increase species power more than any other single technological advance. As such, it is worth almost any level of investment.

I have approached ethics as rules of behavior for individuals and guidelines for constructing social systems. But 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 the power 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. Destruction of the earth's environment has the potential to severely impact the power of the species. Although I believe it unlikely 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. Justifying Power

I have shown that the power theory meets all the features of a satisfactory normative theory of ethics. But there are deeper justifications. In this section, I offer four inter-related arguments for the power theory. First is the argument from evolution. Second is the argument from psychology. Third is the argument from ethics. Fourth and last is the argument from success. In these arguments, I make a number of far reaching claims without detailed evidence or explanation. That I leave for other books or papers. I end this section by considering potential objections to the power theory including the famous naturalistic fallacy.

The argument from evolution is straightforward and compelling. 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 50,000 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.

I claim that long term species power as I have defined it is tantamount to biological fitness for the human species. For our species, fitness is power—though this is not true for most other species. We have seen it throughout human history. Modern humans overcame the Neanderthals based on their greater power. Among humans, societies (and individuals) with greater power generally out-compete those with less power. Societies of greater power persist and grow; those of lesser power are overcome and subsumed. The conclusion is that aligning ethics with long term species power maintains the evolutionary vector of the species and best insures long term survival and well being.

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. The desire to control is one of the most fundamental features of human psychology. With no apologies to Nietzsche, I call this principle *the Will to Power*. The conclusion is that an ethical theory based on power is aligned with one of the most fundamental principles of human psychology, meeting condition three above. This is the argument from psychology.

The argument from ethics is as follows. Humans are social animals. Much of their fitness as a species comes from their ability to cooperate in groups. As shown above, the power of a group can greatly exceed the sum of the separate powers of its individuals. This increase in power arises through the mechanism of cooperation. The historical function—the evolutionary function—of ethics is to foster harmony and cooperation within human groups. Effective systems of ethics enable the group to persist and thrive. In short, effective ethics enhances the power of the group. This is why the power theory preserves, in fact justifies, many of our most cherished values: freedom, individual rights, equality, diversity and reciprocity.

The argument from success is the pragmatic argument; the argument that subsumes all the others. Power as a goal maximizes our evolutionary fitness and our probability of long term survival. It gives us the best chance of winning a no-rules conflict with another (alien) race. It gives us the best chance of avoiding or surviving a catastrophic event. Power allows us to become whatever we want to become, as individuals and as a society. In short, power is the optimal strategy for competition and survival.

Consider the following thought experiment. Imagine an ensemble of human societies. These societies are given properties and circumstances based on a statistical distribution of properties that represent the range of populations, technologies, social systems, cultures and external environments. This distribution covers every parameter that could possibly make a difference in the prospects of a society. Next, endow each society within the ensemble with the same goal. Then run the societies forward in time (think of them as simulations on a vast computer). Finally, gather statistics on survival

probability and fitness. The claim of optimality is simply this: The best results obtain when *power* is the goal.

One might object that power is not an end in itself but a means. The British philosopher Mary Midgley claims that

The desire for power is necessarily secondary to other desires, because power is *power to do* certain things, and valuing those things has to come first. Those who really pursue power just for its own sake are neurotics, entangled in confusion by habit and destroying their own lives. (Midgley, 1978)

What Midgley misses is that a desire for power is a desire for *all* things. Power does not discriminate. But Midgley's fundamental point is in error. It is imminently rational to pursue means as primary goals. Scientists often pursue knowledge for its own sake, without regard to its potential uses. Are they neurotic? Many people around the world pursue freedom as a primary goal even though freedom is merely *freedom to do*. Are they also neurotic?

The hierarchy of goals can be arranged in many ways. One might argue (as I have above) that winning is the goal and power is a means to that end. Or one might argue for power as the primary goal with winning as one justification. But power enables so much more than just winning. Physical power through technology can enable life, shelter, health, enhanced sensory perception and myriad other goods. The power of global cooperation can solve global environmental and security problems to guard against the extinction of our (and other) species. The list is endless. To me what is important is not which desires are primary versus secondary, but the end result.

There are several related philosophic arguments that make a great deal of the observation that statements of fact are qualitatively different than statements of value. These arguments are often used to attempt to invalidate naturalistic systems of ethics such as I have proposed here. The two main arguments are due to Hume and Moore. Hume made note of the is-ought gap and described the fallacy of arguments consisting of 'is' statements that lead to an 'ought' conclusion:

In every system of morality, which I have hitherto met with, I have always remark'd, that the author proceeds for some time in the ordinary ways of reasoning, and establishes the being of a God, or makes observations concerning human affairs; when of a sudden I am surpriz'd to find, that instead of the usual copulations of propositions, *is*, and *is not*, I meet with

no proposition that is not connected with an *ought*, or an *ought not*. This change is imperceptible; but is however, of the last consequence. For as this *ought*, or *ought not*, expresses some new relation or affirmation, 'tis necessary that it shou'd be observ'd and explain'd; and at the same time that a reason should be given; for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it. (Hume, 1739)

Hume's observation reveals the conservative property of formal deductive logic. Nothing can appear in the conclusion of a deductive argument that was not somehow present in the premises. Stated in this way, the 'is-ought' gap is not so remarkable, yet it has been used as an automatic refutation of naturalistic theories of ethics.

Let me examine the logic in a little more detail. Statements involving 'is' are statements of fact. For example, "The human species is a product of evolution." On the other hand, statements involving 'ought' are value laden—they imply a goal. Consider the statement: *Agent A ought to perform action x*. This statement is best interpreted with the additional context: *in order to achieve goal G*. In other words, once a goal for the agent is established, a course of action can be specified that will lead to that goal. In that case, we are justified in stating that the agent ought to execute that particular course of action. In other words, *if agent A desires goal G and action x leads to G, then A ought to perform x*. In short 'ought' implies a goal—that is the logical meaning of 'ought.' So the key to solving Hume's problem is to provide a goal. This goal should be viewed logically as a premise to any argument leading to statements involving 'ought'.

In this paper I have made a case for "maximizing long term species power" as the ultimate goal for the human species. That goal cannot be deductively arrived at from a series of statements of fact about the natural world. Logically, the definition of power as the utility in a decision theory will become a premise of the theory, not a deductive conclusion. As such, it should be justified via other arguments as I have done.

8. Conclusions

Adopting power as the utility in a consequentialist program leads to a satisfying normative theory of ethics. This theory meets all six conditions for a normative theory and it avoids many of the problems that have troubled standard happiness or preference

theories. 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 need a self consistent foundation. The power theory provides this. Most importantly, the power theory is 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

- 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.
- Hume, D. 1739: A Treatise of Human Nature.
- Jeffery, R. 1983: The Logic of Decision, Second Edition. University of Chicago Press. Chicago.
- Luce, R. D. and Raiffa, H. 1957: Games and Decisions. Dover Publications. New York.
- Midgley, M. 1978. Beast and Man. Routledge. New York.
- Mill, J. S. 1863: Utilitarianism.
- Parfit, D. 1984: Reasons and Persons. Clarendon Press. Oxford.
- Russell, B. 1938: Power. Routledge. New York.
- Savage, L. J. 1954: The Foundations of Statistics, Second Edition. Dover Publications. New York.
- Singer, P. 1993: Practical Ethics, Second Edition. Cambridge University Press. Cambridge.