The Power Theory of Scientific Progress

A normative theory of scientific progress is offered. Grounded in standard Bayesian decision theory, the theory employs a utility function which is a combination of explanatory and predictive power, called scientific power. Criteria are established for what we should expect from a normative theory of scientific progress. The utility function is defined on the overall scientific worldview and embodies both the explanatory and predictive aspects of scientific theories. Several justifications for the proposed theory are given. First, the power theory follows from the meta-scientific value of pragmatism, which in turn derives from biological and cultural evolutionary roots. In other words, the power theory is shown to be effective in that a scientific community and associated society which adopts it will out-compete one that does not. Second, the power theory provides a framework for reconciliation of other theories of scientific progress including realism, Laudan's problem solving effectiveness, and Giere's cognitive approach. Third, the power theory is consistent with scientific history in that science can be seen to have been progressive in the past in just those cases when it has conformed to the norms established herein.

1. Introduction and Background

Since the 1960's and the debut of Thomas Kuhn's famous work *The Structure of Scientific Revolutions* [Kuhn, 1970], the problem of scientific progress and the related problem of rational theory choice have been among the most widely discussed problems in the philosophy of science. Although Kuhn himself spent much of his remaining career backpedaling, the arguments he presented in his seminal work were generally viewed as having cast severe doubt that theory choice in science was rational or that progress could be said to occur in any meaningful sense. In the past decades many words and much emotion have been expended in trying to reestablish the fact and rationality of scientific progress. In this paper, I take those efforts one step further and propose a normative theory of scientific progress, which I call the power theory. My claim will be that theory choice in science is rational so long as it conforms to the norm. Furthermore, adherence to the norm gives rise to scientific progress. There is an analytic linkage between these two claims. Theory choice is rational if and only if it leads to progress. Conversely, progress is achieved through rational theory choice. This is the instrumentalist view of rationality.

The problem can be stated very simply: How ought science be conducted and what should be its overarching aims and goals? Once this question has been answered, several related problems are also solved. The problem of progress in science is answered by the simple statement: science will be progressive if it continuously moves toward its goals. Correspondingly, the rationality of science will be assured if its methods reliably produce results that are goal directed. Moreover, these methods should lead to optimal goal attainment in the sense that other competing methods can be shown less effective when measured against those goals.

We can identify several main schools of thought pertaining to the goals of science. Although much of the work of these schools has been descriptive, i.e., ‘what are the goals of science?’ or historical, i.e., ‘what have been the goals of science?’ each alternative can also be viewed as normative. The realist school has been most prominent in recent years. The realist view is that the goal of science is truth and that progress in science represents an increasingly closer approach to truth, perhaps asymptotically. Empiricists strive for theories that are empirically adequate: that explain or predict direct sensory experience. Pragmatists seek theories that can be predictably used to control the environment in support of other practical personal or societal goals. (My power theory is most closely aligned with this school.)

Each school has its proponents and critics. None has gained widespread acceptance. But each also contains a kernel of great value that I will incorporate into my proposed normative theory. With the realist, I acknowledge the notion of truth as an important value for science. Yet I also agree with Laudan (1984) in that absolute truth is not something we can be sure of, nor use as a metric for progress. With the empiricist, I acknowledge that the only direct access we have to the world is through our sensory experience. And finally, with the pragmatist, I assert that what matters, in the end, is what we can accomplish with our scientific knowledge, that the point of science is the decision making effectiveness it grants and the technological power it enables. As humans, as rational agents, as evolved living beings, what we seek most is control of the environment, control of the world—what we seek most is power.

What should we expect from a normative theory of scientific progress? A normative theory should specify the goals of science as well as a prescription for how to achieve those goals. A normative theory of scientific progress is tantamount to a theory of action for scientists to effectively pursue the aims of science. In turn, a theory of action should provide a means to choose among various possible actions available to scientists as they go about the business of science. To accomplish this, we should expect a normative theory to contain criteria to permit a scientist to judge between alternative scientific theories and the degree of progressiveness they entail.

The logical model I choose for the power theory is standard Bayesian decision theory. In this model, the goals of science are represented by a utility function. A rational agent—the scientist in this case—then acts in such a way as to maximize the expected value of the utility function at some future moment. Probability is interpreted in the Bayesian, subjective sense as representing the rational degree of belief of the agent. The utility of science should be the power (composed of equal parts explanatory power and predictive power) of the overall scientific belief system. This measure might be viewed as the norm of a two-component vector, one component being explanatory power, the other being predictive power.

Kuhn thought that an algorithmic theory of scientific progress was unachievable. His two main objections are best stated in his own words:

> The search for algorithmic decision procedures has continued for some time and produced both powerful and illuminating results. But those results all presuppose that individual criteria of choice can be unambiguously stated and also that, if more than one proves relevant, an appropriate weight function is at hand for their joint application. Unfortunately, where the choice at issue is between scientific theories, little progress has been made toward the first of these desiderata and none toward the second. Most philosophers of science would, therefore, I think, now regard the sort of algorithm which has traditionally been sought as a not quite attainable goal. [Kuhn, 1977, p326]

The power theory I propose addresses both of Kuhn's concerns. Scientific power as defined herein represents an unambiguous measure that contains inherently the relative weights of the more traditional scientific values. Where conflict between these values occurs, an appeal to the overall utility provides the means of resolution.

The organization of the rest of this paper is as follows. In Section 2, I briefly review the fundamentals of Bayesian decision theory focusing more on the concepts than the mathematical details. Section 3 provides an overview of my proposed utility function, scientific power, and relates it to other common conceptions of the goals of science. Section 4 explores how the utility might be quantified. Section 5 provides the pragmatic, meta-scientific justification for the normative theory and looks at how the theory reconciles several other views of scientific progress including realism, Laudan's problem solving effectiveness, Ronald Giere's cognitive model and Methodological Pluralism. Section 6 examines a few episodes in the history of science to show that science has been progressive in just those cases where the norm has prevailed, i.e., where power has been maximized. Finally, Section 7 presents my conclusions.

2. Decision Theory

A decision problem is any situation where a choice must be made by an agent from among a set of possible actions and the consequences of a particular action are uncertain. The conduct of science and, in particular, the problem of choosing among competing theories, falls under this broad definition. 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; Bernardo and Smith, 1994]. Here, I apply this machinery to the problem of scientific progress and theory choice. In so doing, I employ a version of standard Bayesian decision theory.

Decision theory is a theory of action by a rational agent—in our case, the scientist. 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 $[0 \le U(s) \le \infty]$. In our case, I restrict the scope of the world state *s* to represent only the entire accepted scientific worldview, the scientist's Zeitgeist.

Hence a choice between rival theories becomes a choice between different scientific worldviews denoted, for example, by s_1 and s_2.

Conceptually the utility function $U(s)$ represents the preferability (according to the agent or scientist) of the worldview *s*. Thus, if $U(s_1) > U(s_2)$ then we would say that worldview s_1 is preferred to worldview 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 worldviews 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 our 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, we have 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 scientist would simply choose the action which leads to the greatest utility—his most preferred worldview. In other words, there would be a specific worldview s_i associated with each choice a_i. The scientist would evaluate the utility of each worldview $U(s_i)$, 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 worldview. The mathematical details are not important here, but ultimately we obtain 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. Although it is possible to base a consistent decision theory on other statistics of the distribution, for our purposes, the expected value is adequate.

3. The Scientific Utility Introduced

Having dispensed with these preliminaries, creating a normative theory of scientific progress comes down to defining the utility function of science. Before diving in with a definition, let me first establish the characteristics we demand for an effective scientific utility. First, as was discussed above, the utility is a function of the complete scientific worldview or the Zeitgeist. It has been shown many times that individual theories cannot exist in isolation, that the context of a theory and its relations to other areas of scientific thought are crucial to its viability. This is the holistic picture advocated by Duhem (1906) and Quine (1953). Of course we can speak of the differences between rival theories with the assumption that the other components of the Zeitgeist remain unchanged. However, it is also important to consider how each rival may affect the rest of accepted science. My conception of the Zeitgeist is similar to Kitcher's (1993) concept of consensus practice extended to represent the whole of science.

Second, the utility should admit some concrete means of estimation. It is important that a normative theory allow, at least in principle, for rational debate to converge to an eventual consensus. That is the pragmatic point of a normative theory. On the other hand, it is inevitable, given the subjective nature of probability and the inherent difficulty of quantifying a worldview's utility, that differences will arise between scientists. Yet the utility should be quantifiable in terms sufficiently common and objective that convergence of rational opinion can be reasonably assured. This is the pragmatic resolution to the problem of incommensurability as defined by Kuhn (1970).

Third, the utility should embody those notions of scientific progress that are generally accepted throughout the scientific community. Values like Kuhn's simplicity, coherence, fruitfulness and scope [Kuhn, 1970], Laudan's problem solving effectiveness [Laudan, 1977], or Kitcher's unifying power [Kitcher, 1988] should be recognizable from within the broader utility framework. Furthermore, common sense and history demand that our theory finds that scientific progress has indeed occurred. The modern scientific worldview, based on, for example, the Standard Model of physics, should yield greater utility than the Aristotelian worldview of the Greeks; General Relativity should prevail over Newtonian Gravity; the Copernican solar system should have greater utility than the Ptolemaic solar system.

With these broad principles in mind, I propose that *scientific power* should be the utility of science. Scientific power I define to consist of two components: *explanatory power* and *predictive power*. In a general sense, power is capability, a pragmatic measure focused primarily on ends and effects. It is uniquely a property relating to agents or groups of agents, i.e., the power of an agent is the capability of that agent to achieve some aim through a course of action. Explanatory power is the capability of the scientific worldview (of an agent—or as we will be applying it, a group of agents) to reconcile and assimilate the facts of the cosmos. Predictive power is the capability of the worldview to generate accurate forecasts of future events. The overall utility can be regarded as the norm of a two component vector. We could inquire as to the relative weights of the two components. Yet so long as the utility is quantified consistently between the two, assigning any priority to one versus the other would be misguided. These two aspects of scientific power are compatible, indeed complementary, and operate at different phases within the methodology of science.

Explanatory power is the capability of a theory or model to explain the facts of the day. It is the power to reduce complicated phenomena to a simple, understandable basis: the capability to answer why-questions. It is the power to put facts, events, and phenomena into suitable relations with each other and with the other tenets of the belief system. In Kitcher's (1988) words, it is the unifying power of the explanatory store. The process of explanation is essentially deductive in that it begins with theories or models or other accepted beliefs and then deduces consequences, the facts to be explained. In part, explanatory power is a measure of reductionism—how completely we can reduce diverse experience to the effects of a small set of simple laws. Explanation is showing how one proposition is a logical consequence of all others in our belief system, that this proposition is necessary given all else we believe thereby integrating it into the whole. To use the marvelous word resurrected by E. O. Wilson (1998), explanatory power is a measure of the consilience of the worldview. It is the power to consiliate and re-conciliate our diverse experience.

Predictive power is the capability of a theory or model to make specific, accurate forecasts of future events. It is the power to look into the future and see the results of our actions. Predictive power fuels technology: the engineer takes the scientific theory at hand and creates artifacts to achieve his or her aims. Like explanation, the process of prediction is deductive in that it begins with theories or models and deduces consequences. Unlike explanation, the consequences are not facts to be explained but forecasts of future events. The predictive results are often in the form of if-then pairs or rules, like "if the earth-sun system is in so-and-so state at a given time, then it will be in such-and-so state a specific time hence." Or "if I undertake this specific action, then that will be the consequence." It is this last form that proves so valuable in the decision process. Predictive power is a measure of decision making effectiveness: to choose effective actions, an agent must be able to predict consequences.

Explanatory power and predictive power are closely related. They are both based on the scientific worldview, the theories and models that form the basis of deductive schemata. The conclusions of any explanatory deduction are beliefs already resident within the worldview. The conclusions of a predictive deduction are new beliefs, generated as a guide to action. In theory building, explanation and prediction are employed iteratively along with experiment to gain power. We look for a theory or model to explain known facts, especially puzzling ones, facts left out of a competing theory. Once we are successful there, we think of new facts that were unpredictable before and design and execute an experiment to test that prediction. Should our prediction come to pass, we search for new predictions; otherwise we go back to explanation—looking for new or modified theories which encompass the new facts.

4. The Scientific Utility Quantified

We can quantify the two components of scientific power in essentially the same way. Explanatory (or predictive) power is a measure of the accuracy of our explanations (or predictions), their scope and their ease of use. We want the scientific worldview to accurately and concisely explain all of the facts of the world. Similarly, we want the worldview to yield accurate predictions for all possible actions. The importance of accuracy in both explanation and prediction should be self evident. Theories that contain significant discrepancies cannot be relied upon as part of a larger Zeitgeist or, more critically, to make predictions. This is especially important in engineering or medical science where a misprediction can cost lives or treasure. The importance of scope is also evident. A Zeitgeist that explains some phenomena but leaves others unaddressed is less powerful than one that includes both sets. The issue of ease of use is an important pragmatic consideration. We should not lose sight of the fact that science is a human endeavor and that theories are constructed for human use. If a theory is too complex to yield useful explanations or predictions for human agents, its power is significantly decreased.

Accuracy is relatively easy to quantify. Several measures of the accuracy of a theory or model have been defined in the literature of statistics. The most familiar is the least square approach for fitting data. Given a set of data and a model with several adjustable parameters, the values of the adjustable parameters can be chosen so as to minimize the sum of the squares of the errors in the resultant fit. The sum of the squares of the errors is the accuracy measure. The least square approach is a special case of the more general approach of maximum likelihood whereby we would select the model which maximizes the probability of the data. That is, if we are to choose between two models, we assume in turn that each model is true and calculate the probability of obtaining the data we have in hand. Everything else being equal, we would prefer the model that produces the highest probability.

There are complications, however. We seek not only to minimize errors in fitting (explaining) the data we have today but also in predicting the data we may obtain tomorrow. This is where a crucial tension comes to bear. The sure-thing-theory, the theory that requires each existing data point to be as exactly as observed, maximizes likelihood—indeed, it results in a likelihood of one. But the sure-thing-theory depends absolutely on hindsight. It is generally useless in making accurate predictions. Thus we are pulled between opposite poles: accuracy in fitting existing data drives toward greater model complexity; accuracy in making predictions drives toward simpler—and presumably truer—models.

Is there a prescription for finding the right balance? Several have been proposed. Prominent among them are the Akaike Information Criterion, AIC [Akaike, 1973], and the Bayesian Information Criterion, BIC [Schwarz, 1978]. The Akaike criterion, for example, gives credit for likelihood but then subtracts for complexity, represented by the number of adjustable parameters in the model. Akaike's theorem proves that under certain conditions, the AIC represents an unbiased estimator of predictive accuracy. The BIC punishes complexity even more harshly than the AIC.

Simplicity, under the rubric "Ockham's razor," has long been touted as a virtue for scientific theories. The AIC and its kin provide a statistical justification based on predictive accuracy, a key component of scientific power. However, the statistical bias toward simplicity vanishes in the limit of large data samples. But there is the other—pragmatic—bias toward simple theories mentioned above: that theories and their extensions must be used by human agents to posit explanations and make predictions, both key elements in the day-to-day decision process. Thus, given the finite, limited nature of human mental processes and the limited timeframe in which many decisions take place, complicated theories would have reduced value. What we need to be effective decision-makers are theories that can be applied expeditiously. We need accuracy, but only insofar as it affects the outcome of the decision.

These considerations arise daily in the practice of engineering where cost is often an overriding concern. In designing a rocket, we build our models only to the level of detail necessary to achieve a particular aim, such as verifying the structural integrity of a strut or tank. Additional detail beyond what has been shown to lead to correct decisions only slows the design process and wastes money. Of course, there is judgment and experience required to determine just what level of detail is sufficient. Often, what is sought is a simple rule or set of rules which can be applied with confidence in certain circumstances. These rules of thumb can be dangerous, however, if used outside their domain of applicability. It is then when a more complex theory is required to locate the simple rules in the context of a more complete system and determine how they should be modified or replaced. The ascendance of the computer as a primary tool for engineering analysis has changed these considerations somewhat, moving the point of tradeoff by reducing the cost of complexity, but the consideration remains. As long as theories and models are built for a purpose, that purpose will dictate the level of complexity required. Hence, I state the engineer's version of Ockham's razor: entities should not be multiplied beyond what is necessary to achieve your purpose.

The other aspect of the explanatory-predictive power of a theory is scope. Scope is the breadth of applicability of a theory, its generality. Theories of broad generality are fruitful because they yield explanations or predictions in a wide range of situations. Consider Newton's theory of gravitation. Its scope far exceeded its predecessor for celestial phenomena, Ptolemy’s theory, in that it was as applicable here on earth as it was in the heavens. Newton's theory explains the trajectory of a baseball as well as it explains the orbit of Mars. By its universality—its scope—Newtonian gravity bestows great power. Scope can be measured by the quantity of different factual statements integrated into the Zeitgeist (for explanation) or the quantity of different predictive statements one could make.

For explanation or prediction, we can combine the three concepts, accuracy, simplicity and scope, into one utility by summing an AIC type measure over the set of all facts to be explained or predictions to be made. For explanation, the set is all empirical statements accepted by the scientific community. This is the means by which the power theory remains firmly grounded in the real world. It ensures the empirical adequacy of our worldview. We must be certain our measure weighs simplicity sufficiently to account for the ease of use consideration. Certainly, this set is difficult to define precisely, yet we have a good enough intuitive understanding of it to use the measure in a qualitative sense for comparative purposes.

For prediction, we encounter a minor difficulty. Conceptually, the set consists of all similar statements applicable to a future world state or currently unaccepted statements of present or past world states. Here is where we capture the value of fruitfulness. I.e., if a theory can generate significant statements about the world that can later be tested empirically, we say that it is fruitful. Unfortunately, we cannot determine the accuracy of a prediction unless we actually test it empirically. So, in practice, we must restrict our set to statements that have been predicted and tested, a subset of the statements included for explanation. Once the utility function is defined, all Kuhnian debates about conflicting values cease. The relative worth of competing scientific worldviews is settled by a comparison of their utilities.

Several higher level ways to compare the explanatory and predictive power of competing theories follow from the bottoms up definition given above. For example, problems or anomalies in the sense of Laudan (1977) subtract from explanatory power. The obvious, sometimes glaring, lack of explanatory power represented by a significant anomaly is often the impetus for the development of a new theory—more powerful since it lacks that particular anomaly. This consideration follows from the accuracy criterion in a trivial way. An anomaly is a serious defect in accuracy.

Another rule of thumb that follows from the accuracy and scope criteria is that we should prefer groups of theories that are consistent or coherent. Clearly, if two theories, which are elements of a complete worldview, fail to cohere, the accuracy of the worldview will suffer. Presumably, the lack of coherence will be manifest in some overlap region in either explanation or prediction. In that intersection, due to the inherent inconsistency, an explanation or prediction may not even be possible thus severely detracting from the power of the worldview. Historical examples of this abound. The anomaly of black body radiation which led to Plank's quantum postulate is one. The incoherence between Maxwellian electromagnetism and Newtonian mechanics that led to Einstein's special relativity is another. A current example of incoherence is the gulf between General Relativity (GR) and Quantum Mechanics (QM). The amount of intellectual resource presently devoted to solving that issue is truly staggering.

There are also higher level measures for predictive power. For example, technological power is strongly correlated with predictive power. Advanced technology is a result of a scientific worldview of great predictive power. In other words, being able to predict the results of actions accurately across a wide range of circumstances is requisite to executing an effective engineering design. For example, an airplane designer must be able to predict the lift forces that will act on the wing of an airplane across a variety of conditions described by parameters such as angle of attack and air speed. A valid theory of aerodynamics, i.e., the fluid dynamics of air, allows such a characterization. Generalizing, we can say that one theory has more predictive power than another if it allows a more capable technology.

Usually, theories with great predictive power contain equally great explanatory power. One notable exception is Quantum Mechanics (QM). Arguably, QM is the most predictively powerful theory in the history of science. It holds the record for the most accurate prediction ever made, that of the magnetic moment of the electron, predicted to better than one part in a quadrillion. The scope of QM predictions extends to the entire known universe, as far as we can tell. Yet the explanatory power of QM, in so far as we can reconcile its laws with the other tenets of the belief system, has up to now been found wanting. This is the famous foundational problem of QM. So far, any coherent interpretation of QM does major violence to other dearly held scientific beliefs like the principle of locality. In the end, something has to give. Believers in quantum philosophy are quite willing to toss out some of these other beliefs. Others continue the search for alternative interpretations. Overall, the tremendous predictive power of QM and its other explanatory successes dictates its central place in the scientific worldview of today—yet an uneasy tension remains.

5. Justification

The previous sections outline a proposal for how science ought to be pursued. In short, scientists ought to select theories that maximize the scientific power of the Zeitgeist. In this section, I present several justifications for my claim. The primary justification is a naturalistic argument: the power theory follows from a meta-scientific value of pragmatism which in turn has its roots in the evolutionary origins of humankind. In addition, I show that the power theory encompasses and reconciles other notions of scientific progress including realism. In the next section I will explore how some historical episodes can be viewed from this new perspective.

It is generally accepted that humans (and all other organisms on earth) are the result of 3.5 billion years of biological evolution. As such, the biological nature of humankind is largely adaptive—or at the very least, not maladaptive. More specifically, some six million years ago a twig on the bush of life split, one branch representing our lineage, the other branch representing the lineage of chimpanzees. Although the two species may seem quite different, humans still share 99% of their genes with chimps. Gradually, along the human branch, large brained forms emerged. Around 100,000 years ago, anatomically modern humans first walked the savannas of Africa. Sometime after 50,000 years ago, behaviorally modern humans, representatives of the species, *Homo sapiens*, emerged in equatorial East Africa. This new species had a very large brain (although Neanderthals had the largest brains of all hominids), and was certainly conscious in the modern sense of the word.

The original adaptive significance of our large brains is currently an open question. Many hypotheses have been proposed, none gaining widespread acceptance. One promising theory is that social intelligence was the driver, a sort of Machiavellian arms race [Mithen, 1996]. Another theory is that language was the driver [Deacon, 1997]. By no later than 50,000 years ago, consciousness emerged. Perhaps the key ingredient of consciousness is the construction of a mental model of the world, ourselves and our associates included. This mental model is the beginning of what has now become the worldview or Zeitgeist, the modern scientific version of which is the object of our scientific utility. The existence of these mental models in each member of a society of humans enabled the explosion of culture that began in Africa and rapidly spread throughout the world, reaching Europe by about 40,000 years ago. Essentially, a new mode of evolution emerged based on a new variety of replicator, now known as the meme [Dawkins, 1976; Blackmore, 1999]. Memes are pieces of cultural information copied by imitation. As replicators, memes can evolve according to the logic of universal Darwinism. Memetic evolution can be identified with cultural evolution—a much more rapid and fluid form of change than biological evolution.

As a feature of the *Homo sapiens* organism and a particular culture, the Zeitgeist is subject to selection pressure. In other words, a Zeitgeist that aids the survival and reproduction of the organism in which it resides will enhance its own replication by being passed on to new generations or by being copied by other organisms. Such a Zeitgeist will persist and multiply and thus possess a high fitness, to use a term from evolutionary biology. At this point, an important distinction must be made. The fitness of an individual meme, a (small) chunk of information, can in general be uncorrelated from the overall fitness of the society or culture in which it resides. In fact, even large pieces of a complete worldview might persist on their own through memetic evolution, yet detract from the overall fitness of the culture. It is the totality of the worldview—the complete memeplex—that can be related to the relative success of a society and thus contribute directly to the fitness of a society. We reach the conclusion that the adaptive purpose of the Zeitgeist is to maximally contribute to the success of its host society—the Darwinian fitness of that society.

The key point is this: I identify the fitness of the Zeitgeist in the sense described above, with scientific power. Clearly, both are properties of the Zeitgeist. But the best way to see the connection is to look to the pragmatic end. A society that possesses a science of great power—a science of great scope and accuracy—will be capable of developing a technology of great power, to improve its overall wealth, to improve the health of its citizens, to improve its power over nature and to improve its military capability. Such a society will out-compete its rivals, who possess a less capable science, and come to dominate. Its scientific worldview will persist and spread.

The argument can be summarized as follows. The fact and logic of evolution, both at a biological and at a cultural (memetic) level, leads to an overarching value of pragmatism: that we should value most that which contributes most to our success. When this value is applied to the problem of scientific progress, we conclude that we should prefer a Zeitgeist with the greatest potential to benefit society in terms of technology and general knowledge. In turn, this implies we should prefer a scientific worldview that maximizes scientific power as I have defined it, a combination of explanatory power and predictive power. In general, we should seek to maximize the *fitness* of our science.

There are many historical examples of this principle. In a conflict between cultures, it is almost always the one with the greater technology (hence greater scientific power) that prevails. The rare counter examples can be explained by other factors contributing to societal fitness, like resources, population or political organization. One such counter example is the defeat of Rome by the barbarian hoards. Other examples exist where a technological advantage overcame a huge disadvantage in numbers. The archetypical episode is Cortez's defeat of the Aztecs—a small band of Spanish adventurers conquering a vast empire.

How is the power theory of scientific progress related to other views? Let me first consider realism, the dictum that the goal of science is truth. My main claim here is that scientific power and scientific truth tend to be positively correlated. In other words, truer theories are more powerful, all else being equal. Again we must keep in mind that our measure of power is at the level of the complete Zeitgeist, not individual theories or practices. Certainly, in engineering, we regularly use approximations known to be strictly false (or less true) in order to expedite a solution. But it is the location of that approximation within a larger body of theory that justifies its use.

To understand the positive correlation between truth and power, consider a particular theory, Q. Suppose there is a discrepancy between Q and a hypothetical true theory, T. This discrepancy could be manifest in a series of predictions of Q that would differ from predictions of T by some amount over some domain. Now suppose we have another theory, R, which covers the same domain as Q. If the discrepancy between R and T is less than between Q and T, then presumably R would have greater predictive accuracy than Q and hence greater power. Theory R's greater predictive accuracy could be exploited to achieve a more powerful technology.

As an example, consider the Newtonian theory of gravitation (NG) versus General Relativity (GR). Across the entire domain of applicability, GR is more accurate than NG. This has been shown through countless experiments. Thus we would say that GR is closer to the truth than NG. However, in the domain of most ordinary engineering, including launching spacecraft into orbit, the difference between results computed using the two theories is negligible. Hence, due to their ease of use compared to GR, the equations of NG are generally used by engineers for those applications. Yet the engineers know that the more generally correct theory is GR and will resort to that when required. For example, the extreme precision required of the Global Positioning System (GPS) dictates that relativistic effects be accommodated.

Although truth can be related to power, by itself it is an insufficient goal for scientific progress. I will argue this point from several directions. First, truth is notoriously difficult to quantify. It is what Laudan (1984) calls a transcendental concept, in that it transcends our abilities for confirmation. For any given belief, the question of truth always boils down to questions of empirical support and logical consistency. There is just no other way to get at it. Hence, from a pragmatic point of view, we may as well dispense with claims to truth, and deal with those quantifiable attributes. In short, there is no quantifiable measure of truth, apart from the standard empirical criteria, with which we can compare alternative theories.

My second point is that we demand more of theories than mere truth. Truth resides at many levels and assumes many guises. For example, truth in the form of a large compendium of empirical facts will not necessarily impart any power at all. We can imagine a completely true theory of the world that consists of a vast library of facts, low level facts about the particulars of various features of the world. Such a theory would be devoid of both explanatory and predictive power despite being completely true and having great scope. Alternatively, we could imagine a true theory of such obscurity or complexity that it is impossible to use the theory to either reconcile experience or make predictions. This theory would also be completely devoid of power.

Modern examples of this second kind include the early days of Quantum Chromodynamics (QCD) and the current situation in string theory. For years after QCD was first proposed as a theory of the strong nuclear interaction, it was strikingly unproductive. Due to the non-linear nature of the quark interaction, it was initially impossible to extract any real predictions from it. It was the symmetries of the theory, matching some symmetries in the particle zoo, that attracted its early proponents. Some years later, Kenneth Wilson introduced the technique of lattice gauge theory [Wilson, 1974] which allowed real predictions to be derived from the theory. Currently, string theory has captured the imagination of many of the world's top theoretical physicists as a way to unify GR and QM. Yet the theory is presently unable to produce any useful predictions—the mathematics is just too difficult, almost intractable. It may be that a breakthrough like Wilson's is just around the corner, but at present string theory contributes little to scientific power. (Its contribution, if it has any, lies on the explanatory side as a potential unifying principle; at the moment, it is predictively sterile.)

We have seen how truth is a likely byproduct of the power theory of progress—but not its focus. Similarly, other proposed measures of progress can be understood from the power perspective. Laudan (1977) proposes that "the rationality and progressiveness of a theory are most closely linked—not with its confirmation or its falsification—but rather with its *problem solving effectiveness*." (Emphasis in the original.) Laudan's problem solving effectiveness can be seen as an aspect of explanatory power. In essence, a solved problem is an explanation: How does the specified phenomenon (or problem) fit within the broad scientific framework? How has it been assimilated? The capability of the Zeitgeist to assimilate and contextualize diverse problems is precisely what I mean by explanatory power.

Ronald Giere (1988) also makes use of decision theory in a general theory of science. He uses a cognitive approach to describe how real scientists actually do science. Within that overarching cognitive framework, Giere proposes a decision theoretic model for scientific judgment. Instead of the strict utility formulation proposed here, Giere uses the less precise satisficing approach first devised by Herbert Simon (1945). Giere also avoids the issue of defining a utility for science, falling back on the truth like measure of similarity between model and world for his examples. There is much to admire in Giere's cognitive theory as a description of science, and that is his aim. He makes no claim for his theory as a normative approach and expresses disdain for normative approaches in general:

> For too long philosophers have debated how scientists ought to judge hypotheses in glaring ignorance of how scientists in fact judge hypotheses. Before presuming to give advice an how something ought to be done, one should first find out how it is done. Maybe it is now being done better than one thinks. Indeed, attempting to follow the proffered advice might be detrimental to scientific progress. [Giere 1988, p. 149]

To answer Giere's criticism in light of my current endeavor, I would say this. There is a wide variation within the science community and without as to both the goals and rationality of science. The establishment of a clear consensus in those regards should facilitate the acceleration of scientific progress—and that is my purpose here.

Philip Kitcher (1993) defines several measures of progress that I believe are consistent with the power theory. Kitcher's concepts of conceptual progress, explanatory progress and erotetic progress are embedded in the explanatory component of scientific power, although a detailed reconciliation is beyond my scope here. Kitcher believes the enduring goal of science is to discover as much significant truth as possible for human beings. Kitcher is unabashedly realist and mounts a vigorous defense of his version of realism. In this regard, our approaches are similar, but with a difference in emphasis. Kitcher prefers to aim at truth and reap the side benefit of explanatory and predictive success, where I would measure progress as degree of success with all confidence that truth is a significant byproduct. In Kitcher's words, "…theories succeed…because they fasten on aspects of reality. If they did not, it would be "a miracle" if they were so successful." [Kitcher, 1993, p. 156] As I discussed above, there is a strong correlation between truth and power (success).

As a final example, consider methodological pluralism, a theory of scientific progress proposed by Howard Starkey (2000). Starkey has noticed—quite correctly, I think—that the methods of science are not fixed, but are variable in response to the situation at hand. Again, this view is accommodated within the power theory. Since the decision theoretic framework admits any action that maximizes the scientific utility, anything goes so long as it is goal directed. Methods are situational—it is the end, scientific power, not the means, that rules the day. And, in a stroke, we have co-opted the anarchist view of Paul Feyerabend (1975): theories, whole Zeitgeists, battling it out in the arena of scientific debate and societal competition—survival granted to the fittest, the most powerful.

6. Some Historical Episodes Examined

Beyond the justifications presented above, a good normative theory of scientific progress should be consistent with the history of science. We should find that historical episodes that are generally viewed as progressive are consistent with the norms established by the theory. In the case of the power theory, I assert that most, if not all, such progressive episodes can be shown to result in a significant increase in scientific power. I will illustrate my claim by briefly examining three famous examples all generally recognized as major progressive steps for science: the Copernican revolution, the triumph of general relativity (GR) over Newtonian gravity (NG), and the Quantum revolution.

In the year of his death, Nicholas Copernicus (1473–1543) published his life's masterwork, *De Revolutionibus Orbium Caelestium*. In it he detailed how the apparent motion of the sun, planets and celestial sphere could all be neatly and easily explained using a sun centered model of the solar system. Within a century, Copernicus' work had revolutionized astronomy, but its influence was also felt in physics, philosophy and religion. Precisely because of its very broad implications, requiring a revision of large tracts of the prevailing worldview of the age, the Copernican view took hold slowly, at times against a vigorous and powerful opposition.

The Copernican revolution ultimately succeeded because it afforded a gain in explanatory power as well as an eventual gain in predictive power. Yet those gains were far from immediate. The Ptolemaic cosmology had been accepted and refined since the Roman times. To explain the complex apparent motion of the planets, Ptolemy postulated a series of circles within circles, called epicycles. The more precise the planetary data were measured, the more complex the epicycle system became. Copernicus eliminated epicycles by a simple change of reference frame from earth centered to sun centered. In the new frame, the planets and the earth became unified in their (nearly) circular motion around the sun. This conceptual economy, i.e., explanatory power, as compared to the Ptolemaic system provided the initial appeal of the Copernican cosmology. Later, when augmented by Kepler's laws of planetary motion, it gained additional predictive power.

The great resistance to Copernicus came not so much from mathematical astronomy, but from religion and philosophy. Copernicus' system was simple and elegant, but it had profound ramifications for the Zeitgeist as a whole—a detraction from its explanatory power. By removing the earth from its privileged position at the center of cosmology, Copernicus also threatened the privileged position of mankind as the center of God's attention. Suddenly, our abode became "just" another planet circling the sun and the sun "just" another star in the cosmos. Psychologically, nothing is more severe than a challenge to the ego and the Church reacted accordingly. For example, one of

Galileo's sins was his acceptance of the Copernican doctrine—for this he was placed under house arrest. Nevertheless, over the course of a hundred years or so, the philosophical and religious positions softened and the Copernican view prevailed. The increase in explanatory power afforded by this view led directly to the next great advance in cosmology precipitated by Isaac Newton and his universal theory of gravitation, first published in 1687.

Newton's theory of universal gravitation afforded perhaps the greatest leap forward in scientific power seen up to that point in the history of science. Its scope was vast, spanning from the most trivial everyday phenomena of objects on earth to the motion of planets around the sun to the motion of the solar system within the galaxy. Its accuracy extended that scope to the degree to which experimental confirmation was available in those days. In fact, its accuracy is sufficient for most engineering applications even today. Its explanatory power essentially unified the entire physical world, from the very small to the very large.

Because of its great power, Newtonian Gravity (NG) reigned supreme for hundreds of years. It was not perfect, however. Conceptually, it relied on action-at-a-distance, a somewhat troubling mechanism whereby widely separated objects could seemingly communicate instantly. As time went on and measurement techniques improved, there began to accumulate a few anomalies, experimental results not accurately explained by the theory. Although these difficulties did not directly lead to the superceding of NG, they did provide a crack for the wedge of General Relativity (GR) to be inserted.

The motivation for GR came from a different direction. In 1901, in an attempt to unify the electro-magnetic theory of Maxwell with the field of general mechanics, Einstein proposed the Special Theory of Relativity. This theory unified the notions of space and time into one conceptual framework, a four dimensional space-time and provided the explanation for why the speed of light could be a universal constant in all inertial reference frames. Extending this new mechanics to the field of gravitation led to GR.

Does GR provide an increase in scientific power relative to NG? The answer is a clear "yes" although not so dramatic as the gain provided by NG some 250 years previously. First, GR eliminated the embarrassment of action-at-a-distance, vindicating all who argued against NG on that basis. Second, GR solved the few anomalies of NG, like the precession of Mercury's perihelion. Third, GR yielded many unique predictions that have since been verified. Examples include the slowing of time within a gravitational field, the existence of black holes, and most spectacularly, the Big Bang origin of the universe. Fourth, GR, despite its mathematical intractability, bestows great conceptual economy to our worldview. The identification of gravity with geometry lends a sublime wholeness to the Zeitgeist. And for most practical applications, the fact that NG can be recovered from the mathematical formalism of GR allows the simple Newtonian machinery to be employed with impunity. In summary, GR represents an unqualified increase in scientific power relative to NG.

The great scientific advance of the 20th century was quantum mechanics. But how does it stack up in terms of scientific power? By every measure QM adds scientific power, perhaps the most significant increase in the history of science. On the explanatory side of the ledger, QM can count the explanation of the periodic table of the elements underlying the field of chemistry and the unification of the fundamental forces (excepting gravity) among other significant successes. As discussed above, however, QM does contain some explanatory defects fueling the continued debate over quantum foundations. (I view this as very analogous to the action-at-a-distance defect of NG). But it is in the predictive arena that QM really stands out above any previous theory. QM has led to incredible advances in technology including modern electronics and information technology via the transistor and the laser. No aspect of modern life is untouched by the technology allowed by QM: computers, cell phones and television all owe their existence to the understanding of nature provided by QM. QM has also led to the fear of nuclear holocaust.

In general, I think it is clear that our everyday notion of scientific progress coincides quite closely with the concept of scientific power introduced here. Although I have here only briefly examined a few of the more prominent episodes in the history of science, I am quite confident that my conclusion extends to other episodes as well. Where the history of science can be seen to have been progressive, there has been a corresponding increase in scientific power.

7. Summary and Conclusions

In the preceding sections I have shown how a normative theory of scientific progress can be cast in the logic of Bayesian decision theory with scientific power as the utility function. Scientific power is defined relative to the complete scientific worldview and consists of equal measures of explanatory power and predictive power, which in turn are defined in terms of accuracy, simplicity and scope.

Proposing a normative theory is audacious, I know. No one likes being told how they should act or how they should make decisions, least of all scientists. But I firmly believe that philosophy should be more than an exegetical exercise. Philosophers should seek to transcend the supposed naturalistic fallacy; we should strive to obtain ‘ought’ from ‘is’.

The benefits of this theory to the scientific enterprise as a whole are immense. It can provide a philosophical foundation for all scientific inquiry. It can provide a solid defense against scientific relativism. It firmly establishes the rationality of science. It provides a new language for scientific discourse and debate: the language of power. It links the goals of science to the goals of the rest of society and ties science to the human condition as a whole. It casts historical episodes in a new and unifying light. Finally, it unifies many disparate views of the rationality and progressiveness of science showing them to be facets of a single principle.

Scientists should ask themselves these questions: Does my work contribute to an increase in scientific power? Does it unify previously disparate facts? Does it yield novel predictions? Does it contribute a unique point of view shedding light on previously obscure aspects of the worldview. Does it lead to new technologies? If so, scientists can be assured that they are following the normative, fully rational prescription, and that they are not only contributing to the power of science but to the power of society at large.

Lastly, it should be noted that what I am espousing here is not really new. Most scientists would admit that they seek to gain knowledge. It was Francis Bacon who recognized long ago that *knowledge is power*.

REFERENCES

Akiake, H. 1973: “Information Theory and an Extension of the Maximum Likelihood Principle,” in B. N. Petrov and F. Csaki (eds.). *2nd International Symposium on Information Theory*, Budapest, Akademiai Kaido, 267-81.

Bernardo, J. M. and Smith, A. F. M. 1994: Bayesian Theory. John Wiley & Sons. New York.

Blackmore, S. 1999: The Meme Machine, Oxford University Press, Oxford.

Dawkins, R. 1976: The Selfish Gene, Oxford University Press, Oxford.

Deacon, T. 1997: The Symbolic Species, W. W. Norton & Company, New York.

Duhem, P. 1954: The Aim and Structure of Scientific Theory. Princeton University Press. Princeton.

Feyerabend, P. 1975: Against Method. New Left Books, London.

Giere, R. 1988: Explaining Science. The University of Chicago Press, Chicago.

Kitcher, P. 1989: “Explanatory Unification and the Causal Structure of the World,” in P. Kitcher and W. C. Salmon (eds.), *Scientific Explanation, Minnesota Studies in the Philosophy of Science, Volume XIII*. University of Minnesota Press, Minneapolis.

——1993: The Advancement of Science. Oxford University Press. Oxford.

Kuhn, T. 1970: The Structure of Scientific Revolutions, 2nd Edition. University of Chicago Press. Chicago.

——1977: The Essential Tension. The University of Chicago Press. Chicago.

Laudan, L. 1977: Progress and Its Problems. University of California Press. Berekley.

——1984: Science and Values. University of California Press. Berekley.

Luce, R. D. and Raiffa, H. 1957: Games and Decisions. Dover Publications. New York.

Mithen, S. 1996: The Prehistory of the Mind. Thames and Hudson. London.

Quine, W. V. O. 1953: “Two dogmas of Empiricism,” in W. Quine, *From A Logical Point of View*. Harvard University Press. Cambridge, Mass.

Savage, L. J. 1954: The Foundations of Statistics, 2nd Edition. Dover Publications. New York.

Schwarz, G. 1978: “Estimating the Dimension of a Model,” *Annals of Statistics* 6:461-65.

Simon, H. 1945: Administrative Behavior. Free Press, New York.

Starkey, H. 2000: “Methodological Pluralism, Normative Naturalism and the Realist Aim of Science,” in R. Nola and H. Stankey (eds.), *After Popper, Kuhn and Feyerabend: Recent Issues in the Theories of Scientific Method, Australasian Studies in History and Philosophy of Science*, Volume 15. Kluwer Academic Publishers, Dordrecht.

Wilson, E. O. 1998: Consilience. Alfred A. Knoff. New York.

Wilson, K. 1974: “Confinement of Quarks”, *Physical Review D*, Volume 10 Number 8, 2445.