

RESEARCH 101 FOR ENGINEERS

George A. Hazelrigg¹
National Science Foundation

1. Introduction

About 25 years ago, National Science Foundation (NSF) director Erich Bloch made the comment, “Isn’t engineering research an oxymoron?” This comment is particularly relevant in the context of NSF, which funds fundamental research and shuns the funding of developmental efforts. Indeed, the NSF Strategic Plan for 2011-2016 notes that the NSF mission is “*to promote the progress of science*” through its role in research and education. In keeping with this mission, NSF seeks high-quality research proposals across virtually all areas of science, engineering and mathematics.² The funding of research projects is also in keeping with NSF’s mission of education, as research projects form the basis of master’s and PhD theses and promote the training of a skilled STEM workforce. These concepts seem to be quite well infused into the NSF culture.

Over the course of my 30 years at NSF, I have had the opportunity to run and to sit in on many hundreds of panels and to oversee the review of thousands of proposals. My observations in these panels lead me to conclude that proposal reviewers and panelists have quite a clear understanding of the mission of NSF with respect to research, and they rate proposals accordingly. Proposals that have a focus on research have a clear advantage in the review process over those that focus on other activities, such as development, design, and so on. Given my experience with the NSF proposal review system, I have no evidence to indicate that it is a random process (although there is always some variability in the measurement of anything). For this reason, it is my belief that it pays to learn how to write a research proposal.

2. What is Research?

If you are going to write a research proposal, it seems reasonable that the first thing you might want to know is, what is research? Although I do not think that it is important that you agree with my definition of research, I do think it is important that you have thought about it long enough and hard enough to have your own definition. My definition is the following:

Research is the process of finding out something that we don’t already know.

First, it is important to recognize that research is a process. That is why you get grants to conduct research. If, instead of being a process, it is, say, an end state, then there would be no point in funding it. Research funding is provided to conduct the research process. Second, the purpose of research is to find out something that we don’t already know. This means that the focus of research is on knowledge, it is not on artifacts. A proposal that focuses on an artifact, for

¹The views expressed here are strictly those of the author, and they do not necessarily reflect those of the National Science Foundation or the Federal government.

²Notable exceptions include weapons research, which Department of Defense and Department of Energy fund, research in space, which is funded mainly by NASA, and clinical research, which is funded by the National Institutes of Health.

example, the development of a sensor, is most likely not a research proposal. Rather it likely is a development proposal. If your proposal discusses artifacts and not knowledge, it is likely to get lower ratings in the review process.

Not all discovery of new knowledge is through a process of scientific research. However, NSF strives to fund scientific research. Thus, we need to distinguish between scientific research and other types of research. I believe that scientific research has three distinguishing characteristics:

1. *Scientific research is methodical.* By this, I mean that, in advance of conducting the research, it can be planned. This is an incredibly important concept, for it is the research plan that enables evaluation of the proposed research. We will soon show that it is not difficult to state a research objective, say an hypothesis test, for which we cannot write a plan. A very simple example of such an hypothesis is, *God exists*.
2. *Scientific research is repeatable.* By this, I mean that, if reasonably competent people follow the research plan reasonably well, they get reasonably similar results. If this were not the case, the results would be random and, thus, no new knowledge would be gained.
3. *Scientific research is verifiable.* By this, I mean that the research conclusions are based on tangible evidence. For example, suppose my hypothesis is, *ghosts exist*. My research plan is to come into my office late at night with my camera. I turn off the lights and wait until it gets really spooky. Then, say at about 3 AM, I take my camera to the elevator, get in and press the basement button. When the doors open, the ghosts will be there, and I'll take their picture. As we all know, ghosts' images do not appear on film. Thus, when the ghosts do not appear in the pictures, I will have proven their existence. This conclusion does not derive from tangible evidence.

When you frame research for NSF, you need to have an understanding of what scientific research is, and you have to frame it with these characteristics in mind.³

What is Science?

To be sure, our goal here is to understand engineering research and to learn how to frame an engineering research project. But first, we have to understand what science is. Again, I present my perspective on this. I believe that science is our understanding of nature and, thus, science research (as distinguished from scientific research) is research aimed at improving our understanding of nature. This definition immediately begs the question, does science exist, namely, is it possible to understand nature? Most of us believe that, indeed science does exist. This belief rests on the assumptions that all the laws of nature apply everywhere all the time, and they are invariant over all space and time. As such, they are there for us to discover. Were this

³Of course, if you don't agree with my definition and characteristics, use your own. But be sure that what you present is consistent with your definition and characteristics.

different, they might be constantly changing and, hence, elusive and impossible to discover.⁴ Or, after their discovery, they could change, thus invalidating our prior knowledge. In such a case, science could not build upon itself. Now, the fact is that we have no proof that the laws of nature are fixed over all time and space. But we do have a fair bit of evidence, with data spanning some 10^{30} cubic light years of space and about 13 billion years in time and, with the possible exception of the big bang itself, we have yet to discover a violation of this belief.

If we accept these premises of what science is and what comprises scientific research, we can then proceed to build a body of scientific knowledge. The manner by which we do this is to present hypotheses that are testable and falsifiable, and to conduct tests of these hypotheses. For example, let's consider a test of the hypothesis $F=ma$. To test this hypothesis, we compare experimental data to the model. We could plot the model as shown in Figure 1. This is a representation of the hypothesis. Now we look to see if this representation corresponds to nature. The only way we have to do this is through experiments. This is where it becomes a bit tricky to do science research. We hope to make observations that relate to our hypothesis in such a way as to be uncorrupted by exogenous factors. But every test result is a measurement, and every measurement is itself dependent upon a model. So it is very easy indeed for this process to become corrupted. This is the main reason why science must proceed by expanding our frontiers of knowledge on the margin rather than by taking giant leaps into the unknown.

Back to our hypothesis that $F=ma$. Suppose we now conduct a very large number of experiments, say 10,000. And suppose all our results are in near perfect agreement with our hypothesis. Does this prove our hypothesis to be true? What level proof would make you comfortable in the assertion that the hypothesis is true? We would probably agree that, if we test every conceivable case covered by the hypothesis and get agreement across the board, we could assert the truth of the hypothesis. But what if we test only half the possible cases? Or 1 percent of the possible cases, or 10^{-10} percent? What we need to realize is that our hypothesis includes an infinity of possible cases. So even 10,000 test cases represents precisely 0 percent of the possible cases. Clearly, we have not proven the hypothesis to be true, even with a huge number of test cases. On the other hand, one valid test case that violates the hypothesis will falsify it. Ergo, science does not build on validated principles and laws. Rather it builds on principles and laws that we have, after extensive testing, failed to falsify. Of course, in the case of $F=ma$, we have literally billions of test cases.

Thus we see the need for a valid scientific hypothesis to be falsifiable under a finite plan of testing. Indeed, in the case of $F=ma$, it is conceivable that we could falsify our hypothesis with the very first test. But we cannot prove its truth with anything short of an infinity of tests, so the truth of a scientific hypothesis always remains a bit elusive. We accept the truth of such




Fig. 1 Test of an hypothesis

⁴Note that, if a law such as $F=ma$ is not fixed across time and space, but there exists a more fundamental law that governs its change, then we would have to find the underlying law to understand nature. At the most fundamental level, however, the laws must remain fixed in order that we “know” them.

hypotheses only after extensive testing fails to falsify them. This leads us to see science as the *pursuit of disbelief*. The scientist is always striving to prove his or her theories false, and accepts them as true only after extensive tries to falsify them fail. This is in contrast to religion, wherein practitioners accept a premise to be true until overwhelming evidence is provided to the contrary.

What is Mathematics?

Mathematics also plays a key role in engineering. So it is equally important that we understand what mathematics is. Whereas the sciences deal with nature, namely with things that can be observed and measured (seen, felt, heard, smelled, tasted, weighed, etc.), mathematics deals with the mind. The mind is known only to its owner; it produces thoughts, and thoughts cannot be seen, felt, tasted, smelled or otherwise observed and measured. Mathematics is about logic. Indeed, there is no agreement that mathematics is a science at all. But, using the word “science” loosely, mathematics is the science of consistency.⁵ Mathematics is not connected to the physical world in any particular way, rather it’s all about what is in our minds. As engineers, we make decisions, and we want our decisions to be consistent with our beliefs, our preferences and the available alternatives. Mathematics provides the framework for doing this.

At the most fundamental level, we ask the question, how do I know what I know? For example, can I say that I know that $2+2=4$? The answer is that I can say that I know something if that something is entirely consistent with everything else that I would say that I know. Suppose I have three rocks, A , B and C , and I know that A is heavier than B , and I know that B is heavier than C . Consistent with this knowledge, I can say that I know that A is heavier than C . However, I cannot say that I also know that C is heavier than A , as this presents an inconsistency. Adding this statement to my knowledge base regarding the weights of the rocks destroys my knowledge base altogether. For the scientist, then, mathematics plays a very important role in understanding nature.⁶ Each new item of knowledge must be entirely consistent with all other items of knowledge.

Let’s go back to our example hypothesis, $F=ma$. We begin with a suspicion that there is a relationship between F , m and a . We hypothesize the form of the relationship, then design a set of experiments, collect data and compare the data and the hypothesis. The data by themselves have no meaning. In the absence of the hypothesis, the data are just a random collection of numbers. They provide no new knowledge. The hypothesis transforms the data into knowledge in the form of a mathematical relationship, which is consistent with the data. It is the mathematical relationship (namely a logical construct), in our heads, that comprises new knowledge. If this relationship is inconsistent with any other relationship that we would say we know, then there is at least something in the set of things in our heads that we cannot know. In fact, until we resolve the inconsistency, we risk destroying not only the new item of knowledge, but all other things we hold as knowledge as well.

This brings us to another interesting concept. Suppose you and I have mutually contradictory items of knowledge that each of us would claim we know. I might say A is B , whereas you might say A is *not* B . Mathematically, I would say $A=B$, whereas you would say $A=\neg B$. The question is, can we both legitimately claim our statements as knowledge, that is, can

⁵An excellent book for understanding mathematics is Douglas R. Hofstadter, *Gödel, Escher, Bach*, Basic Books, New York, 1979.

⁶Indeed, employing mathematics to understand nature is the major contribution made by the great scientist Galileo.

I say that I know that $A=B$ and you say that you know $A=\neg B$, and we both are right? Contradictory as this may sound, it is possible that we are both right. Remember that we know something when it is entirely consistent with everything else that we know. And since this places a new item of knowledge into the context of all previous items, and since you and I will not share the same set of previous items, it is indeed possible. We create a logic system by stating a (minimal) set of beliefs—we call them axioms. We then build upon this set of axioms by creating theorems, which become the operative elements of our logic system. If we hold different beliefs, that is, if our logic system is built on different axioms, our theorems will be quite different, and what we see as consistent with these will differ also. There is no underlying set of axioms that are the right axioms in the sense that all others are wrong. Axioms pertain to each of us separately as individuals. On the other hand, most of us share parts of our logic system, such as numbers and arithmetic, as they are so fundamental to our everyday existence.

Recognizing that mathematics provides a framework for thinking logically, we can look at the mathematical subdisciplines:

Arithmetic is a framework for thinking logically about numbers.

Algebra is a framework for thinking logically about relationships between objects.

Geometry is a framework for thinking logically about lines and shapes.

Calculus is a framework for thinking logically about change.

Statistics is a framework for thinking logically about the past.

Probability theory is a framework for thinking logically about the future.

Decision theory is a framework for thinking logically about decisions.

Game theory is a framework for thinking logically about interactions among two or more people.

And so on. One thing that we see here, which has been problematic in many engineering curricula, is the difference between statistics and probability theory. Statistics deals with data. It asks the question: what conclusions can one draw that are logically consistent with a set of data. All data come from the past. Ergo, statistics is about the past. Probability theory deals with beliefs about events that may occur in the future. There are no data from the future. Ergo, probability theory deals with the question: what conclusions can one draw that are logically consistent with a set of beliefs? Probability theory has its major application in decision making as an aid to enabling a decision-maker to make decisions that are consistent with his/her beliefs, preferences and available alternatives, recognizing that the future can never be predicted with both precision and certainty.

What is Engineering?

Engineering is many things to many people. It involves analysis. To some it is all about problem solving. It certainly involves meetings. But, on the bottom line, what defines engineering as distinct from the sciences and other disciplines is that engineers do design. Design is a process that begins with choosing what to design, how to configure the device or system, and determination of its dimensions and tolerances, materials, fabrication and manufacturing techniques, packaging, rules for use of the product or system, and even its eventual disposal, recycling and reuse. In all these elements of design, the engineer is in the role

of decision maker. So a key element of design is decision making. Through his/her decision making, the engineer is manipulating nature to benefit at least some segment of society.⁷

Recognizing that a key function of engineering is to make decisions, it is appropriate to have some understanding of what defines a decision. Decisions all share a set of properties:

1. As distinct from problem solving, decision making always involves a commitment of resources. Resources can be in the form of money, equipment, materials, time, labor, energy, and so on. Indeed, one can view a decision as the irrevocable allocation of resources. If the decision involves resources that could be later revoked, the revokable resources are not really part of the decision.
2. All decisions have outcomes. Outcomes may be good or bad, they are not right or wrong as is the case with problem solutions. The degree of good or bad associated with a particular outcome is determined by the preferences of the decision maker.
3. Decisions are always optimizations. That is to say, the decision maker will always choose the decision that is best in his/her eyes at the moment of the decision.
4. Decisions cannot be made in the absence of preferences, namely how the decision maker would rank order outcomes. Should a decision maker assert that he/she has no preference over outcomes, it is equivalent to saying that the decision maker is indifferent across all outcomes. In this case, the decision maker would accept a random choice.
5. Decisions are always made in the present (the past is gone—it cannot be changed, and the future isn't here yet), and outcomes are always in the future. Ergo, decision making always demands that the decision maker predict the future.
6. Predictions are never both precise and certain. Thus, all real decisions involve some degree of uncertainty and risk.⁸
7. Decisions are made only by individuals. Groups have an emergent behavior that is the subject of game theory, they do not make decisions.

As a key element of engineering is decision making and since good decision making demands good prediction, engineers must be good at prediction. In designing a system, for example, the engineer must be able to predict the behavior of the system as a function of the design choices made regarding the system. In order to affect good prediction, engineers invoke the belief that all laws of nature apply everywhere all the time, and it is the laws of nature that

⁷The concept of providing benefit to at least some segment of society is key here. Without it, the engineer would have no incentive to design (nor would he/she be paid). On the other hand, design clearly need not benefit all of society. For example, Hitler's engineers designed gas chambers, which the majority of us would be quick to assert did not benefit a very wide circle of people. Yet, without benefit to at least someone, they would not have been designed. A proper theory must accommodate even these egregious cases.

⁸A detailed exposition of engineering decision making is provided by George A. Hazelrigg, *Fundamentals of Decision Making for Engineering Design and Systems Engineering*, 2012, www.engineeringdecisionmaking.com.

determine the system's behavior.⁹ But there are lots of laws of nature, and everywhere is a lot of places. Let's consider a very simple law, $F=0$. This law governs the behavior of a static structure.¹⁰ We'll apply this law to a structure, such as a bridge. To do this, what we teach is that the structure consists of a set of lines and nodes, and we balance the forces at the nodes. But no real structure consists of lines and nodes. A real structure has volume, which is to say that if we think of $F=0$ being applied to lines and nodes, for a real structure, the structure contains an infinity of nodes. We have no solution method for an infinity of nodes. So we must resort to finite approximations for the solution. Hence, we see that, in predicting the behavior of a system, the engineer faces two key questions: which laws of nature dominate the behavior of the system, and what computational algorithm will enable enforcement of these laws with adequate precision?

As we noted above, no prediction is both precise and certain. Thus, uncertainty is a key element of all prediction. Indeed, mathematically speaking, a prediction is defined as a probability distribution on the outcome of an event. Failure to treat a prediction in this mathematically correct manner will result, most often, in substantial error. A simple example will help to understand these mathematics. Suppose an engineer is engaged in the design of an airplane. It is important that the engineer know the weight of the airplane; if it is too heavy, it won't fly. Now, if the airplane were built, the simple thing to do would be to weigh it. But we cannot build the airplane until it is designed, and we cannot design the airplane until we know what it will weigh. So, what to do? The answer is that we use a mathematical model to predict the weight of the airplane. A seemingly appropriate model would be to assume that the weight of the airplane will equal the sum of the weights of the parts of the airplane,

$$W_A = \sum_{i=1}^n W_i \quad (1)$$

Of course, it would be reasonable to test the model against data. We could obtain data by experimenting with airplanes that are already built. To do this, we would gather a collection of built airplanes, weigh them carefully, then disassemble them, weigh their individual parts, sum the weights of the parts and compare the sum to the originally measured weight of the assembled airplane. No doubt, if done carefully, we would get a result that corresponds well with this model. But now to predict the weight of the airplane under design. Where do we get the W_i ? We could weigh the parts only if they exist. But they will exist only after they are built, which is after they are designed. So we are almost back to our original problem. Here we do something different, however. We estimate the weights of the parts (perhaps using simple models—volume times density, for example). This means that the W_i are random variables. But arithmetic (in this case, addition) does not apply to random variables. So we know that the above equation cannot be valid for predicting the weight of the airplane. We have to correct the equation so that we are adding cardinal numbers. Expected values are cardinal numbers, thus we can write,

$$E\{W_A\} = \sum_{i=1}^n E\{W_i\} \quad (2)$$

⁹Many systems involve human input. For these system, engineers predict the behavior of the non-human components by relying on the laws of nature.

¹⁰I never cease to be amused by the fact that we teach a three-credit course, *Statics*, on the solutions to this equation.

Eq. 2 is a predictive model. It is very important to recognize that it is not the same as Eq. 1, and it cannot be tested against data. It is a model that applies to a specific event, and events occur only once. Any attempt to test it using a statistical approach would rely on the untestable assumption that outcomes of repeated experiments are drawn from independent and identically distributed (IID) distributions. Moreover, even in the case that a clairvoyant could provide us in advance with the outcome of the event, this knowledge would not enable a test of the distribution on W_A and, hence, Eq. 2.

The key thing to recognize about the difference between Eqs. 1 and 2 is that they can produce very different results. Eq. 1 tempts us to use for the W_i our best guesses (most likely values) for these quantities. However, Eq. 2 notes that we should be using the expected values. As weights are never negative, the distributions on all W_i are skewed to the right. Ergo, use of best guesses underestimates every single term on the right hand side of Eq. 2, leading to substantial underestimates of W_A . Thus, we see that engineers must deal properly with uncertainty, and they need computational procedures to aggregate uncertainty estimates on components to an uncertainty estimate on the overall system.

What is Engineering Research?

Given that the distinguishing difference between science and engineering is that scientists seek to learn about nature, while engineers seek to manipulate nature via the act of design, and noting from the arguments above that design decision making demands an ability to predict the behavior of a system and to quantify the uncertainty associated with such predictions, we can now distinguish research objectives that are uniquely engineering.

First, to predict the behavior of a system as a function of design decisions, an engineer must understand which laws of nature dominate the behavior of the system. For example, for a bridge, in the short term, we might hypothesize that the behavior of the bridge is dominated by the balance of forces ($F=0$) and Hooke's law (strain is proportional to stress). Of course, in the longer term, other laws of nature might also be important, such as laws regarding corrosion and fatigue. Research regarding which laws of nature dominate the behavior of a system is framed as an hypothesis denoting a conjecture to be tested. We observe that such an hypothesis forms the basis of research that is uniquely engineering. It is not science because it is not seeking to understand fundamental laws of nature. Rather, it takes these laws as given and applies them to a specific case. Furthermore, the case to which they are applied relates to the design of an artifact, a system or a process, and that is uniquely engineering. Finally, it clearly fits within our definition of research—to find out something that we don't already know, namely which laws of nature dominate behavior in specific situations.

Second, recognizing that, in general, we cannot apply even the most simple laws of nature (such as $F=0$) with total precision in any real case, a second area of research that is uniquely engineering is finding computational algorithms that provide adequate representations of the real cases and their limitations. Finite element analysis is a good example of such a computational algorithm. Research on this topic can also be framed as an hypothesis, notably that a particular algorithm provides a representation of reality within a specified error bound. And, again, this area of research fits our criteria as uniquely engineering.

Note that, for these areas of engineering research, framed as hypothesis testing, it is not possible to provide a finite plan to prove the truth of the hypotheses. However, they are both testable and falsifiable. Thus, they qualify as valid scientific hypotheses.

Third, we noted above that no prediction is both precise and certain. Ergo, consideration of uncertainty is mandatory and, in fact, a prediction is defined as a probability law on the outcome of an event. Thus, all predictions are stochastic in nature. This means that we need methods to obtain uncertainty data and to aggregate these into probability distributions on the behavior of a total system. To do this correctly, one needs to understand that a probability is a belief held by an individual. Thus, it would be inconsistent to aggregate uncertainties (probabilities) obtained from more than one individual. Rather, the aggregation method needs to make use of Bayes Theorem to enable the decision maker to properly account for inputs from others. The necessary computations can be somewhat complex.

Fourth, also noted above, all decision making is optimization. This poses two research issues that are uniquely engineering. Optimization is possible only in the presence of an objective function. Further, all objective functions represent preferences, namely the preferences of the decision maker. Though it is not widely recognized, objective functions may be valid or not valid. They are valid only if they rank order all outcomes in precisely the same order as would the decision maker. Otherwise, they are not valid. Optimization using an objective function that is not valid is every bit as wrong as any other mathematical or theoretical mistake. Furthermore, as preferences are in the mind of the decision maker, we can assume that they are known with precision and certainty. Any other assumption invalidates all theory of optimization. For example, the CEO of a profit-making company might have the preference *to make money, more is better*. This preference rank orders all profit measure outcomes. However, it may be quite difficult to map engineering decisions, even down to the level of sizing a bolt or setting a tolerance, into this preference. Therefore, preference mapping between engineering decisions and realistic preferences is yet another area of research unique to engineering.

The other researchable problem that optimization presents that is relevant (although not necessarily unique) to engineering is that of dimensionality. A typical system may have upwards of tens of thousands of design variables subject to optimization. Some of these are continuous and some are discrete. Most relate to the underlying preference nonlinearly. And, since there is uncertainty in the system performance as measured by the objective function, we are faced with a stochastic optimization problem. Such problems can be enormously difficult to solve. Yet, advances toward optimal solutions often have considerable payoff, and they are worthy of pursuit.

The fifth area of research that is uniquely engineering is one that has received very little rigorous research to date. On the one hand, we recognize that optimization of designs is highly preferred for two reasons, most obviously because optimized designs provide more of what is intended than non-optimal designs (profit, for example). Secondly, however, optimal designs provide assurance of convergence to desirable outcomes. In the absence of a rigorous decision-making environment, such as provided by an optimization framework, it is possible and even likely that decision making will be quite poor. So, the problem is this: an optimization framework is feasible only when there is a single design decision maker. The moment that two or more people make design decisions for a given system, which is to say in virtually every case, it is no longer possible to provide a rigorous optimization framework for design decision making, and this opens the door for very bad decisions. As a result, there is a true need to do “damage control” in the case of all designs that involve more than one designer. This may involve the mathematics of game theory and social choice theory.¹¹

¹¹One should consult the extensive work of Donald G. Saari for insights into problems of choice.

Framing Research

It should now be evident that there are indeed areas of research that are uniquely engineering. So the question is, how should we frame research objectives for these areas? In my experience, I have seen only four proper ways to frame a research objective. This is not to say that, for certain, there is no other way. But, in the several thousand proposals I have seen, I have yet to come across one that could not have framed its proposed research, probably better, in one of these four ways. Briefly, they are:

1. The research objective of this project is to test the hypothesis H .
2. The research objective of this project is to measure parameter P with accuracy A .
3. The research objective of this project is to prove the conjecture C .
4. The research objective of this project is to apply method Q from field F to solve problem R in field G .

As noted in the previous section, much of engineering research can be framed in the form of an hypothesis test. However, it is most important that you frame your research in the form of a valid scientific hypothesis. Recall that a valid scientific hypothesis is both testable and falsifiable. A valid hypothesis generally takes the form A is B . For example, " $F=0$ and Hooke's law dominate the behavior of a static structure." One could paraphrase this as "these laws are the dominate factors in the behavior of the system," which is the form A is B . You must avoid the use of words such as "can." These words make an hypothesis unfalsifiable with a finite plan. For example, suppose we have the hypothesis, the addition of nitrous compounds can improve the wear resistance of machine tools. The problem is that there exist an infinity of nitrous compounds, and falsification would require testing them all (note that one case that does not improve wear resistance does not falsify the hypothesis as another case could prove it true), and this is not possible under a finite plan as it requires an infinite amount of testing. Of course, it is possible that the hypothesis would be found true if, by luck, a case is found that works. But requiring luck is not an acceptable form for framing research.¹²

It is also important to frame a measurement correctly. It is not adequate to say simply, "the research objective of this project is to measure the distance from earth to moon." This objective is entirely different if the objective is to make the measurement with an accuracy of 0.01 meters as opposed to 1,000 km. Both the method of the measurement and the value of the measurement will differ widely depending upon the desired accuracy. Ergo, if the objective of your research is a measurement, it is necessary that you include the accuracy of the measurement as a part of your objective statement.

The proof of a conjecture is generally a mathematical objective. For example, to prove the four-color conjecture (that a map can be colored with a maximum of four colors with no two abutting regions of the same color).

Finally, we observe that much of the advancement in science comes from the application of methods or principles taken from one field and applied in another. We sometimes consider this to be multidisciplinary research, as it requires knowledge of both fields. Underlying the concept of this as research is the fact that the notion of disciplines is human. Nature does not recognize disciplines. Yet nature is entirely self-consistent. Thus, if our theories regarding nature are inconsistent across disciplines, something must be wrong. By resolving

¹²The reader is referred to the work of Karl Popper.

inconsistencies across disciplines, we advance our knowledge of nature and, hence, we advance science. This is the basis of the fourth method of framing research.

Lessons Learned for Proposal Writing

A proposal is an offer to perform some task or act for receipt of some compensation. A research proposal is typically an offer to perform certain research for receipt of a monetary payment. To be a proposal, a document must make a statement about what is to be done and what will be the compensation. In the case of a research proposal, it must include a clear statement of a research objective and a budget. In the absence of these two things, a document is not a proposal. In fact, in the absence of a clear statement of an objective, it may well seem that the proposer does not know what he/she intends to do. Furthermore, proposals submitted to the NSF are subject to a merit review process. Over the years, NSF has varied its review criteria. However, it has always come down to five basic questions:

1. What does the Principal Investigator (PI) propose to do?

This question is answered by the statement of the research objective. If the proposal does not have a clearly stated objective, the reviewers are left to guess what the PI has in mind. This is not a good place for the PI to be. Keep in mind that a reviewer cannot evaluate a proposal until he/she knows what it is about. For this reason, the statement of the research objective should be the first sentence of the first paragraph of the first page. Whatever may be said prior to this statement may be torn off and discarded. It will not be used in the reviewers' evaluation, and it is a waste of space that could be used for meaningful statements that actually add to the proposal. Further, NSF funds fundamental research, so a proposal to NSF should not focus on development. A clear test is to note whether the proposal discusses mainly an artifact or knowledge. If the proposal is about an artifact (a device, a product, a sensor, a system, etc.), then it most likely is a developmental proposal. Developmental activities tend to be open ended (they are never complete), they are omni-disciplinary (which is to say they involve all conceivable disciplines), and they are generally rather expensive. Developmental proposals tend not to fare well in the NSF review process. If the focus of the proposal is on knowledge, such as the result of an hypothesis test, then the proposal is a research proposal. Research projects tend to be well defined, with a beginning and an end, they can have a clear disciplinary focus, and they tend to cost much less than developmental efforts.

2. If the research plan is carried out competently, will it accomplish the stated research objective?

A proposal to NSF must include a plan. Too many proposals go on at length about background and theory, and neglect to present a specific plan of research. The plan obviously should be a plan to accomplish the stated research objective (too many proposals miss this simple concept). The reviewers will assess the efficacy of the plan as the primary element of their review process. Further, without a plan, it would not be possible to perform a merit review of a proposal. One would simply have to take it on faith that the objective would be

accomplished.¹³ In the case of an hypothesis test, the plan needs to address the nature of the tests to be performed, their relevance to the hypothesis, the number of tests needed to comprise a reasonable test, the parameters that will be varied and the range over which they will be varied for the tests to be done. For a measurement, the plan needs to discuss the measurement apparatus, and it needs to provide evidence that the proposed apparatus can be obtained and will yield the proposed accuracy. For the proof of a conjecture, the approach to the proof needs to be outlined. For the application of method Q from field F to solve problem R in field G , the proposal needs to identify the known inconsistencies between the fields or the barriers to success and provide insights as to how they will be resolved or overcome.

3. Can the PI (or PIs) competently carry out the proposed plan?

The proposal should offer evidence that the PI (or PIs) can carry out their research plan. This means showing that they have the knowledge, skills and expertise to conduct the research, and that they also have access to all needed equipment and supplies. Most PIs are oblivious to the need to offer this evidence, and merely fill out their proposals in accord with the mandated sections. However, several parts of the proposal provide opportunity to offer such evidence explicitly: the PIs' bios with publications carefully chosen to illustrate expertise in the specific areas of need, results from prior research, preliminary results and facilities. All should be crafted around the need to support the PIs' ability to carry out the proposed research plan.

4. Is the research worth doing?

For proposals submitted to NSF, this question should be addressed under the headings of Intellectual Merit and Broader Impact. The Intellectual Merit is the contribution that the research makes to the scientific or engineering field. How does the research extend the frontiers of knowledge and what new research might it facilitate? The Broader Impact is the contribution that the research makes to society at large. The proposal should make a cogent argument that the proposed research will benefit society, and it should detail the nature of that benefit. It is not, and never was, intended that the requirement of a Broader Impact statement necessitate additional activities specifically designed to engender benefit to society. Rather it is intended to be a statement that argues that the proposed research itself benefits society at large.

5. Is the budget reasonable?

The budget should be appropriate to the work proposed, neither more nor less. If more, the reviewers will feel that the PI is greedy. If less, the reviewers will feel that the PI is naive and incapable of performing the research. As NSF's mission includes education, it is important that the budget provide educational support, usually in the form of graduate student support. If support for a post doc is requested, it should be evident that this is requested in support of the NSF mission of education, and not in support of "empire building" for the PI.

¹³I have actually received proposals with no plan. One, requesting \$5 million, had the simple objective, "to prove the existence of God." The technical section stated, in toto: "*So, how will we do it? Well, it won't be easy. But surely with modern computers and perseverance, we shall succeed.*"

Summary

To summarize, a well written proposal provides clear narrative to the reviewers for evaluation of the above five questions. It begins with a clearly written statement of a research objective, and follows with a plan to accomplish that objective. It offers sufficient background to assure the reviewers that the proposed objective moves the frontiers of knowledge forward, framing the proposed research in the context of and giving acknowledgment to prior work by others as well as the PI. It also offers clear evidence of the PI's capability to carry out the plan and to show that the proposed work is worth doing.

Proposals that do not have a clearly articulated research objective leave the reviewers feeling that the PI doesn't really know what he/she intends to do. Further, without a clearly stated objective, it is not possible to do a comprehensive review of the research plan. Obviously, one needs to know what the plan is intended to accomplish in order to assess its efficacy. Furthermore, without a clear objective, even the PI is at a loss to propose a coherent research plan, and the proposal tends to fall apart. It is equally important that the objective be a research and not a developmental objective. Too often, engineering research proposals actually turn out to be developmental proposals. Developmental objectives lead to open-ended projects that always leave loose ends, and these are picked up by the reviewers to criticize the work. It may not be easy to formulate a clear research objective. But if you don't take the time, the rest of the work you put into the proposal is most likely a waste. Think of someone coming to you asking for a quarter of a million dollars or more. You would be very unlikely to give it to them if they didn't know what they wanted it for. This is how you will be treated if you don't have a clear objective.

In the end, good proposal writing is just common sense. But you have know your field, you have to be an expert in the area of research you are proposing, you have to know where the frontiers of the field lie and where extensions of these frontiers make sense, you have to have a very clear idea of what will be your contribution, and you have to have a comprehensive plan to make your contribution.

Good luck!