# Handout: Science, Prediction, and Risk: How Much Should We Worry About Impending Shortages?

Humanities 161 Global Crisis: Energy and Water Resources in the 21st Century

by Jonathan Gilligan
Reading for Mon. Jan. 17, 2005

## Contents

# 1 Introduction

This semester we will be looking at predictions of whether or not we will face serious shortages of energy and fresh water resources in the future. It's worth taking a little time at the beginning of the term to consider predictions in general. How do we make predictions about such things? What sorts of things might we do to decide whether or not to believe such predictions? How are predictions of coming dangers handled in our political system?

There is a long history both of false alarms that are taken far too seriously and serious warnings that went unheard. Years before September 11th Vice President Al Gore chaired a commission on airline security that recommended strengthening airline security measures, particularly screening passengers for weapons, because of the danger that hijackers would commandeer airplanes and crash them. Opponents of government regulation together with lobbyists for the airline industry convinced Congress that this would represent a costly and unnecessary measure.

On the other hand, in 1990 when a crackpot named Iben Browning proposed that there would be a large earthquake on the New Madrid fault, which lies along the Mississippi river near Memphis, many state and local government agencies, schools, universities, homeowners, and businesses took this warning seriously, setting up evacuation and rescue plans and buying earthquake insurance. Browning had no training in geology and no scientifically reasonable basis for his prediction, yet his warning was taken serously and millions of dollars were spent responding to it.

We would like to find a balance between worrying too much about any crazy idea that doomsday is just around the corner while also keeping our minds sufficiently open that we will not dismiss serious dangers out of hand simply because we haven't though about them before.

# 2 What is Science?

Much of this course is concerned with making predictions about oil and water, deciding how much to believe others' predictions, and deciding what to do about these predictions. Some of the predictions we will discuss are rooted deeply in the natural sciences. Others are more concerned with human behavior, and thus fall into the realms of politics, economics, or psychology. We will begin by considering scientific predictions.

In this section, we will consider three principal questions: What is scientific knowledge? How do we acquire it? How do we decide whether to trust it?

Perhaps the most important thing that distinguishes natural science from other disciplines is its ability to make detailed and accurate predictions of the future. Astronomers can tell us almost exactly where Mars will be one million years from now. Engineers can tell us almost exactly how much electricity we can get from burning one thousand cubic feet of natural gas in a gas-turbine generator. Not all scientific predictions are so accurate, though. Forecasts of next week's weather are useful but not terribly accurate. Still, we judge science largely on the success of its predictions.

Since we're concerned about running short of oil and water, we can begin by looking at what natural science can do to predict the amount of various resources that may be available to us in the future. Before we begin, let's step back and thing about what science is.

The philosopher, poet, and scientist Jacob Bronowski described science thus:

> *When [the poet] Coleridge tried to define beauty, he returned always to one deep thought: beauty, he said, is 'unity in variety.' Science is nothing else than the search for unity in the wild variety of nature—or more exactly in the variety of our experience.*[*]

---

[*] J. Bronowski, *Science and Human Values* (Harper Perennial, 1965) p. 16.

## 2.1 Vocabulary of science

To make predictions, we use scientific laws together with observations about the current state of the world to deduce what will happen at later times. To a scientist, observations are frequently called **data**. To philosophers of science, they're called **facts**. A fact is something you physically observe or measure, either with your own senses or with a measuring device, such as a camera or a volt-meter. The observation that a certain rock is brown would be a fact.

Facts alone don't make science. Science develops **theories** to explain the facts we observe. Theories consist of a number of rules, or laws, that facts must obey. To become part of a theory, a rule (or hypothesis) must be thoroughly tested by experiments or observations so that the scientific community will trust it.

A simple picture of how science proceeds is as follows: A scientist notices a pattern among events (facts) that he or she observes. The scientist proposes that this pattern is universal and puts it in the form of a **hypothesis**—a universal rule. Scientists test the hypothesis by looking for cases where it might not apply. If the hypothesis survives this testing, the community begins to trust it and incorporates it into theories. Actual science rarely follows this clean progression exactly, but most scientists do take this as a model of how science ought to work.

## 2.2 Inductive and Deductive Reasoning

Scientific reasoning as I have just described it has two components: inductive and deductive. **Inductive reasoning** occurs when we propose new hypotheses. We observe nature and generalize. We observe patterns in our observations and propose a hypothesis. From a number of specific observations, we jump to the idea that there is a universal principle involved. Galileo observed that many solid objects fell at the same speed when he dropped them, regardless of their weight, and came up with the general rule that gravity causes **all objects** to accelerate at the same rate.

There is no method for performing inductive reasoning correctly. Inductive reasoning is a creative activity—an intuitive leap that is not grounded in logic. Coming up with good hypotheses is a mysterious talent that has to do with insight and luck. Neither is there any way to prove that a rule arrived at by induction is correct. If we observe a phenomenon happen with perfect regularity 1000 times, it may not follow our rule on the 1001st time.

## 2.3 Falsifiability

The philosopher Karl Popper proposed that one thing that makes science different from other ways of acquiring knowledge is not that science can prove its hypotheses true, but that it can disprove incorrect hypotheses. What makes a hypothesis worthy of scientific study is its **falsifiability**.

When we use inductive reasoning to come up with a hypothesis, this hypothesis may be true or false. We have to test it. One way to do this is to use **deductive reasoning** to make predictions from the hypothesis and then see whether the predictions occur. If they don't, we can conclude that the hypothesis was incorrect. To Popper, the most important thing that distinguishes science from other disciplines is not that scientists are always right—in fact, you can never prove that scientific hypotheses are correct—but that when they're wrong there is a straightforward way to prove them wrong.

Galileo's hypothesis that gravity makes all bodies fall at the same rate in a vacuum is such as statement. If it were wrong, all we would need to to do disprove it would be to show one instance of two bodies falling through a vacuum at different rates.

A good scientist should always look for evidence that would falsify his hypotheses and he should always strive to make his hypotheses as precise as possible so that if they are wrong it will be easy to falsify them. Popper's criterion suggests that the most easily falsified hypotheses are the most trustworthy if experiments fail to disprove them.

In the sense we're discussing here, the term **scientific** does not mean correct. It simply refers to statements that can be tested in a certain way. A scientific statement may be true or false. It is important to realize that even if it is correct, this type of scientific statement can never be proved correct, but if it is false, it can potentially be proved false. Thus, testing scientific hypotheses consists of repeatedly trying to find flaws in them—places where their predictions do not come to pass. As a hypothesis is tested more and more times and is not falsified, we come to trust it more and more, but we can never become completely sure that it's true.

In this context, you may ask since a scientific statement can be either true or false, what would a nonscientific statement be? A nonscientific statement would be one that is stated in a way that makes it impossible to disprove it even if it's false. The statement "everything occurs according to God's will" is such a statement. It may be true or it may be false, but if it were false, you could never disprove it. How could you prove that something occurred not according to God's will. In general, scientific hypotheses allow us to make clear and precise predictions that can be compared to what actually comes to pass. Nonscientific hypotheses do not allow us to make these kinds of predictions.

## 2.4 Trans-science

There is another class of hypotheses, which we can call **trans-scientific**, which allow us to make clear and precise predictions, but where practical matters do not allow us to observe the predictions. A hypothesis that the sun will devour the earth four billion years from now is scientific, but since we won't be around to see whether this happens, we can't test it. A hypothesis that requires too much money or too much time to test in a reasonable amount of time is trans-scientific. Sometimes this definition is not absolute, but depends on the context. Many questions around the depletion of oil and changes to the earth's climate over the next several decades are trans-scientific from the perspective of making policy today but scientific in the sense that many of you will live to see whether or not they come to pass.

In discussing whether we face serious shortages of energy and water in the future, we will want to think about what the evidence is for and against different positions and whether this evidence is based on good science. We will want to distinguish science, non-science, and trans-science. We should recognize that non-scientific statements can be true and scientific ones can be false, but that the ways we go about deciding what to believe are different for the two types of knowledge.

## 2.5 Credibility

In an ideal world, each of us could test the scientific basis for these predictions ourselves, but the time and expense of doing testing everything are too great for this to be a practical method. Instead, at some level we must trust the work of others to be truthful and correct. For the most part science operates on a system of trust. Scientists check each other's work for honest errors, but mostly trust that no one is lying.

Science can go wrong in many ways. It is perfectly normal for hypotheses to be incorrect. Because there is no way to prove a hypothesis true, there will always be some number of incorrect hypotheses that slip through any battery of experimental tests. For this reason, there is no shame in proposing a hypothesis that turns out to be wrong.

There are two other sources of error in science: incompetence and dishonesty. An incompetent scientist may not carry out his tests properly and the results may be nonsense. In 1988, the French scientist Jacques Benveniste reported that a chemical dissolved in water produced an immune response in cells even after he had diluted the solution so much that there would be only one molecule of the chemical in more than 7000 gallons of water.* Later investigation of Benveniste's laboratory by John Maddox, the editor of the journal Nature together with James Randi, a professional magician and skeptic, revealed that Benveniste's experiments were "statistically ill-controlled, from which no

---

*Nature **333**, 816 (1988).

effort has been made to exclude systematic error, including observer bias, and whose interpretation has been clouded by the exclusion of measurements in conflict with the claim .... The phenomenon described is not reproducible."*

Jacob Bronowski wrote that science can only work in a community where scientists can trust one another. The very defining characteristic of the scientific community is its "habit of truth"

> *Individually, scientists no doubt have human weaknesses. ... But in a world in which state and dogma seem always either to threaten or to cajole, the body of scientists is trained to avoid and organized to resist every form of persuasion but the fact. A scientist who breaks this rule... is ignored.*†

As shameful as incompetence may be, the history of science is also rife with people baldly lying. In 1974, William Summerlin, the chief of transplantation immunology at the prestigious Sloane-Kettering Institute for Cancer Research in New York City, claimed to have made a breakthrough in transplanting skin grafts from black rats to white rats without the rejection that normally accompanies transplantation of tissue from one animal to another. It was later discovered that he had transplanted grafts from white rats and colored them with a black marker.

The scientific community depends on trust and reacts severely to cases of cheating. When a laboratory technician noticed that Summerlin's "black grafts" could be washed off with alcohol, this was reported quickly to the head of Sloane Kettering and by noon that day Summerlin had been fired.

In 2002 Jan Hendrik Sch́on, renowned young scientist at Bell laboratories was fired in a similar incident when it was discovered that he had not conducted the experiments on microelectronics that he had published, but had fabricated the data and lied to his co-workers.

Michael Crichton, the author of *Jurassic Park*, has just published a novel in which global warming is not really happening and the environmental disasters people blame on it are actually sabotage performed by wacko environmentalist terrorists who try to dynamite glaciers in Antarctica, create earthquakes and tsunamis in the South Pacific, cause flash floods, and steer hurricanes into Florida.

Because of the political and economic stakes in the debates over natural resources and environmental hazards, scientists on both sides of any issue are frequently subject to *ad hominem* attacks that question their motives and the integrity of their work. Over the course of this term, we will have to consider such charges and think about how we can judge among the different scientific positions we see.

Some of the things science does do to try to ensure integrity involve practices of peer review and openness. Scientific papers must be peer-reviewed before they are published. This means that the editor of a journal sends the manuscript out to one or more respected scientists who read the paper and provide anonymous reviews. These reviews look at the questions the paper asks, the methods it uses to answer the questions, and the logic of its arguments. If the methods are not reliable or appropriate to the questions asked, or if the data does not support the conclusions, the referees will recommend either rejecting the paper or fixing the problems. This sort of review can catch errors in approach, but not outright fraud.

Once a paper is published, if it is interesting or important other scientists will attempt to build on its results. If there are problems with the paper, these further investigations will reveal the problems and the community will work to understand and resolve them. In 1989 two scientists in Utah, Stanley Pons and Martin Fleischman, reported that they had discovered a cheap, clean, unlimited source of energy by creating a nuclear fusion reaction in a beaker full of what was essentially fancy salt water. As other physicists and chemists tried to study this phenomenon, they quickly discovered all manner of inconsistencies in the original report and were unable to duplicate Pons and Fleischman's results. Within a year the scientific community came to a solid consensus that the original experiment was riddled with errors and that Pons and Fleischman had lied to avoid admitting these errors.

---

*Nature **344**, 287 (1988).

†J. Bronowski, *Science and Human Values* (Harper Perennial, 1965) p. 59

Bjorn Lomborg's book, *The Skeptical Environmentalist* accuses many environmental scientists of similar mistakes and fraud, and the environmental science community has fired back with counter-charges that Lomborg committed unethical practices by misquoting and misrepresenting others' work.

It is worth noting that although it is much rarer than fraud and incompetence, there is another phenomenon we can see in science, where one person's new hypothesis, although correct, is so at odds with existing ideas that the community refuses to accept it. Copernicus's proposition that the earth orbits around the sun rather than vice-versa is a famous example of this.

Less well-known is the case if Ignaz Semmelweis, an Austrian physician in the mid-19th century. In Semmelweis's hospital, about one woman in eight who gave birth ended up dying of an infection known as puerperal (childbed) fever. Semmelweis noticed that when midwives instead of physicians attended the delivery, only one woman in 50 died of puerperal fever. Semmelweis also noticed that physicians frequently went directly from conducting autopsies to deliver babies without washing their hands in between. In 1947 Semmelweis proposed that somehow (he didn't know how because no one knew about bacteria and viruses yet) doctors were transferring infections from the corpses they were dissecting to the mothers. When he conducted an experiment by having physicians wash their hands in a disinfectant solution before attending births, the rate of death from puerperal fever dropped by a factor of seven.

However, the medical establishment ridiculed Semmelweis, pointing out that there was no theory that could explain why washing hands would have anything to do with disease. Semmelweis lost his job and his work lay largely unnoticed for almost two decades until Joseph Lister in England discovers the principles of antisepsis.

The lesson for us here is that in considering bizarre new ideas from science we must strike a balance between excessive gullability and excessive skepticism. The philospher Thomas Kuhn, who had studied many cases of good and bad science, described this balancing act as "*the essential tension*" that defines science. If we are too credulous, we will end up believing a lot of nonsense, but if we are too skeptical, we may end up like Semmelweis's superiors, rejecting true and useful ideas.

As we consider scientific treatments of energy and water resources, we will need to consider several questions: Are so-called scientific statements based on scientific (testable) hypotheses? If so, to what extent have the hypotheses been tested and confirmed? What are our standards of proof and how do we tell when they have been satisfied?

## 3 What Is Probability?

If we want to understand risks, we would do well to look at the insurance industry, which makes its money helping people deal with risk. If an insurance company does not understand the mathematics and psychology of risk, it will soon go out of business. In this section, we will look at the mathematics, specifically the question of how we learn about probabilities. In the next section, we will consider the psychology and economics of risks.

Insurance companies do their best work managing hazards that are random and recurring. The randomness assures that no one can predict who will be struck by misfortune and the recurrence enables the insurance company to determine the probability that the hazard will strike a given person.

### 3.1 The frequentist definition

How might we determine this probability. Let's consider the risk of dying in an automobile accident. There are about 45,000 deaths per year in the United States from car accidents and the total population is roughly 275 million. We can estimate the probability of a given person dying in a car accident in the next year as $P = 45,000/275,000,000$ or about one in six-thousand. We could refine this estimate by recognizing that some people are at greater risk than others, but this is a good starting point.

The method we just used to measure the risk of dying in a car accident is called the **frequentist** definition of probability. A frequentist believes that probability is defined as the fraction of times a given outcome occurs when you repeat a situation many times under identical conditions. To a frequentist, if you can't reproduce a situation over and over again and measure the number of times a given outcome occurs, then it is meaningless to talk about probablity. Thus, for a frequentist, it is meaningless to talk about the probability of the world running out of oil in the next fifty years because we can't endlessly reproduce the next fifty years on identical copies of the earth and count the number of times the world runs out of oil. Since we only have one earth, the probability that we will run out of oil in the next fifty years is either 100% or 0%. Any attempt to say otherwise is a misunderstanding of probability.

Insurance companies are not strictly frequentist, but they are most comfortable working as close as possible to frequentist definitions of probability—risks for which their actuaries can precisely measure probabilities. When they face situations they've never encountered before, such as the possibility of intense coastal storms due to global warming over the coming century, they become very nervous and often make bad business decisions.

## 3.2 The subjectivist definition

The principal alternative to the frequentist definition of probability is the **subjectivist** definition. Subjectivist accounts of probability are more mathematical than frequentist ones. Rather than determining probability from measurements of frequency, subjectivists attempt to deduce probability from what they know about the world. To a subjectivist, probability is not just a measure of the frequency with which certain outcomes occur when we repeat a process over and over. Probability is also a useful way to describe our state of knowledge or ignorance about the world. When a frequentist says that the probability that we will run out of oil in the next fifty years is either zero or one hundred percent, the subjectivist would say that we can't determine which of these answers is true until the fifty years are up. Before then, we can use probability to express how much we know about which answer is correct.

Consider a simple six-sided die with the sides numbered one through six. If you roll it, what is the probability that it will come up five? To a frequentist, you cannot answer this question except by rolling the die many times and counting how many times it comes up five. To a subjectivist, we might start with the observation that there is no particular reason for one side to come up more often than another, so we will assume that all six sides are equally likely to come up. Thus, a subjectivist would predict that the die would come up "five" one-sixth of the time, even before making a single experiment. The problem with this approach is that the die might not be fair. A cheat might have manipulated it or there might have been a defect in manufacturing it, so that it will actually come up "five" one quarter of the time or one seventh of the time. Frequentists criticize subjectivists for their need to make assumptions about the dice in order to tell us the probabilities.

Subjectivists retort that frequentists are useless for thinking about risks of things that have never happened before, or which have happened only rarely. A frequentist would assert that it is meaningless to discuss the probability of a terrorist group making a nuclear attack on the United States or the probability of a severe oil shortage in the next fifty years. A subjectivist would respond that if we are going to make sensible plans for dealing with either hazard, we must be able to say something about the probability that it will occur, even if our estimate of the probability is not perfect.

The principle difference between the frequentist and the subjectivist is that the frequentist believes that probability is strictly determined by measurements of frequencies in many repetitions of the same process, while a subjectivist attempts to determine probability mathematically from what he knows about the circumstances of the event. Subjectivists are willing to calculate probabilities of unique events, such as meltdowns of nuclear power plants, while frequentists believe that probability only has meaning for things that can be repeated many times under identical conditions.

For the frequentist, the question, "What is the probability that there will be a severe global oil shortage fifty years from now" is meaningless.

For a subjectivist, the question can be addressed by considering the factors that might contribute to such a shortage, estimating the probability of each factor, and then mathematically combining them to determine the overall probability. Such estimates are fraught with difficulties. Let's consider what might be involved. **Before you read further, take a few minutes and think of all the things you can that might affect the production of oil fifty years from now.**

## 3.3 Factors affecting oil production

A partial list of factors affecting oil production in the future might include:

- Demand for oil. Production of oil depends on many factors that affect the demand for oil (how much oil people want to buy at a given price). Here is a list of several factors that might affect the demand for oil:
  - The population of the earth. The more people there are, the more oil they will probably want to consume.
  - Energy intensity. Oil is one of our principal sources of energy. We use oil for transportation, heating and cooling our houses and offices, and for generating electricity. As a country's economy grows, its use of energy grows. The ratio of the amount of energy consumed to the size of the economy is called the **energy intensity** of a country's economy. If changing economic conditions (for instance a shift from manufacturing to processing data) were to reduce the energy intensity, we would be able to use less oil without sacrificing our standard of living.
  - Economic growth. If the world's economy grows faster than the energy intensity shrinks, demand for energy, and hence for oil, will grow.
  - Demand for petroleum as feedstock. We also use oil as a *feedstock*: a raw material for producing other goods, such as fertilizers, plastics, and medicines.
  - Alternate sources of energy. If we develop inexpensive alternate sources of energy, demand for oil may decline as more people choose the alternatives.
  - Environmental consequences. Burning fossil fuels has many bad environmental consequences, which range from the local (urban smog) to the regional (acid rain) to the truly global (climate change). If we decide that the environmental consequences are bad enough, we may choose to use less oil.
- Supplies of oil. Supplies of oil depend on two principal things
  - Size of known reservoirs.
  - Cost of extracting oil from existing reservoirs. Some oil fields are more expensive to exploit than others. The greater the price of oil, the more incentive there is to extract oil from the more difficult sources.
  - Discovery of new reservoirs. While there is a finite amount of oil in the ground, the higher the price, the more incentive there is to explore for new oil reservoirs.
  - Total amount of petroleum in the ground. This is fixed by nature and while we don't know the exact value, we will never be able to extract more oil than this. As we will learn, we are consuming oil millions of times faster than nature produces it, so eventually we will run out. The question is when.

This list is only a partial list of the various factors that can affect the supply and price of oil. You will notice that many of the items are not purely questions for natural scientists, but depend on political and economic conditions, which are very difficult to predict. Geologists may be able to learn about the amount of oil in the ground, but predicting how future politicians will think about pollution or what the price of alternate energy sources will be fifty years from now is far beyond the capabilities of even the best political scientists and economists. This makes it difficult to come up with any sensible kind of estimate.

To try to rescue some sort of scientific rigor from this difficult question, we can try to eliminate the parts of the question that depend on human choice and sociopolitical or economic factors and focus on the things that depend only on nature. We could refine the question by asking what the probability is that there is enough oil in the ground in reasonably accessible reservoirs to supply twice as much oil fifty years from now as we use today. This question eliminates the human factor and concentrates on the natural distribution of oil in the ground. The answer might not be particularly useful if the demand fifty years from now turns out to be different from twice the current demand, or if drilling technology changes radically in the next half-century, but at least the question can be answered with some level of scientific rigor.

This sort of approach was very useful in the early 1980s when the nations of the world were debating what if anything to do about the danger than a class of chemicals called chlorofluorocarbons, or CFCs (also known by the brand-name Freon) might destroy the so-called ozone layer in the stratosphere, which protects us from the most dangerous types of ultraviolet radiation in sunlight. Scientists just could not agree on the number of people who would die from skin cancer in the year 2100 if we did not eliminate CFCs. It turned out that most of the disagreements were over social and economic questions, such as the assumptions of what the industrial demand for CFCs would be and how quickly the world's economy would grow. When a question was asked: "What amount of CFCs in the atmosphere would reduce the ozone layer by half," the scientific community was quickly able to agree on an answer and this answer then proved quite useful to diplomats working on treaties to protect the atmosphere, even though the more interesting question about skin cancer in the year 2100 remained unanswerable.

Similarly, current concerns over global warming are difficult to answer because many of the questions we would like to answer, such as what would happen if we adopt the Kyoto treaty or if we do not, depend more than anything on the rate of growth of the world economy and on the details of how the economies of the industrial and developing nations evolve. The questions scientists can answer with reasonable confidence, such as how the temperature of the earth is likely to change if we quadruple the amount of carbon dioxide in the atmosphere, don't address the questions we would really like answered, such as how this would affect hurricanes in Florida or drought in Kansas.

Thus, while subjectivists can attempt to estimate probabilities when frequentists throw up their hands in frustration, we must always ask what assumptions the subjectivists are making and whether they're reasonable.

## 4 Perception of Probability

As hard as it is to manage the mathematics of probability and risk, we must also work to understand just what the numbers mean. In the 1970s and 80s, Daniel Kahneman, Paul Slovic, and Amos Tversky conducted a number of experiments on how people actually think about probabilities, statistics, and risks. To explain what they observed, these three psychologists developed something they called **prospect theory**. In 2002, Kahneman was awarded the Nobel Prize in Economics for this work. Before we go into a description of prospect theory, I would like to ask you to consider the following hypothetical questions:

1. Would you choose to make the following gamble?
    - 50% chance to win $150.
    - 50% chance to lose $100.
2. Would your answer to question 1 be different if you had just lost $100?

3. Which of the following alternatives would you choose?
    (a) Lose $100
    (b) Make the following gamble:
        - 50% chance to win $50
        - 50% chance to lose $200
4. Would your answer to question 3 be different if you had just received a gift of $100

We will return to these questions later.

Kahneman, Tversky, and Slovic found that the vast majority of people's intuitions are not well equipped to deal accurately with probabilities and statistics. Intuitive judgments about cause and effect, about probability, and about risk are often quite inaccurate. To make intuitive judgments, we rely on rules of thumb—what are called **heuristics** in psychology jargon. Heuristics are often drawn from experience, and good heuristics allow us to make decisions quickly by allowing us to ignore irrelevant information, avoid needless worrying about details, and go to the heart of the matter. Even if a heuristic misleads us some of the time, it may still be useful if it helps us make good choices most of the time and saves us a lot of time and effort.

The problems Kahneman, Tversky, and Slovic found with the heuristics people frequently use to manage probabilities is that they mislead people very often and lead people to seriously misunderstand the choices before them. Some of the effects of this are failure to perform sensible acts, such as wearing seat belts; bad choices in purchasing health insurance; invalid interpretations of medical tests (e.g., for cancer); and gross misunderstanding of risks from such things as floods and radioactive nuclear waste.

## 4.1 Affective heuristic

The principal assertion of prospect theory is that people's perception of probability and risk depends on the emotional context, even though the actual numbers do not. For instance, people worry much more about hazards that terrify them (so-called **dread risks**). Even though you are far more likely to die driving to or from the beach than to be killed by a shark, the thought of the shark probably makes your skin crawl in a completely different way from the thought of a car wreck. This **affective heuristic** leads most people to overestimate the danger of a shark attack.

Things that contribute to an increased affective emphasis on a particular type of risk include:

- the dread factor just discussed;
- whether we can control our exposure to the risk: Is it voluntary, like skydiving or smoking, or involuntary, like pollution or crime?
- whether it affects a few people at a time or many people all at once: Does it have potential to be a catastrophe?
- whether we can see it coming: Is it visible or invisible?
- whether it is well-understood or largely unknown

## 4.2 Availability and anchoring

We are also more likely to hear about shark attacks in the news than about car accidents. This tends to put shark attacks closer to our conscious imagination than dying in car accidents. This is an example of the **heuristic of availability**. Another aspect of this heuristic is that people tend to use available examples to **anchor** their perceptions of risk. People in a flood-plain tend to assume that the next large flood will be no larger than the last large flood.

Availability and anchoring might play into our thinking about oil and water shortages by connecting the prospect of future shortages to our experience of past ones (the Arab oil embargo of 1973–4 or the dust bowl drought of the 1930s). Such connections could lead us to either overestimate the probability of such an event or to underestimate its severity.

### 4.3 Prospect theory

There are many other phenomena that contribute to prospect theory and we do not have time to go into all of them. What you should be aware of is that people tend to misunderstand the numerical aspects of statistics and probability; that our intuitions (even those of experts) tend to do badly at estimating and reasoning with probabilities; and that there is an important emotional (affective) component to the way we perceive risks. This means that just presenting information may mislead more than inform. We need to consider not just the information we provide, but the context in which we place it.

In the gambles discussed above, most people asked would turn down the gamble in question 1 whether or not they had just lost $100. Most people asked would take the gamble in question 3 rather than giving up $100. However, taking the gamble in question 2 (when you just lost $100) would be equivalent to choosing the gamble in question 3. Taking the gamble in question 4, when you have just been given $100 would be equivalent to taking the gamble in question 1. If you are consistent, your choice to accept or decline the gamble in questions 1 and 4 should match and your choices in questions 2 and 3 should match. Most people's don't because the affective context is different for the two choices.

It is important to note that the affective heuristic is only a defect if we believe that people should make their choices on a rational numerical basis. If we believe that it's reasonable to prefer a larger risk of dying in a car wreck to a smaller risk of being eaten by a shark, then it's rational to give more weight to avoiding shark attacks. However we feel about this, if we are going to understand how people think about uncertain risks they face in the future we do need to understand something about how most people think about these things.

## 5 Social Institutions and Risk

Studies of professional risk managers show that they are susceptible to the same sorts of misperceptions that affect laypeople. Managing officials in charge of managing risks of floods and other natural hazards do not have time to examine every possible hazard in detail. If they spent all their time and money considering truly oddball possibilities they would never have resources to prepare for the hazards they know to be common and serious. Rather than examining all possible hazards and all possible responses, they tend to follow heuristics, or rules of thumb, that have worked in the past. Economists call this sort of behavior **bounded rationality** and business types call it **standard operating procedure**. Following procedures that have worked in the past allows us to concentrate our energies and money on solving problems rather than just thinking about them.

The problem with bounded rationality is that it does not work well when something completely new happens. There had not been a tsunami in the Indian Ocean for about a century, so no one made plans for what to do if one occurred and so millions of people were caught off guard two weeks ago and around 150,000 died. Our plans for dealing with airplane hijackers before 2001 did not consider the hijackers ramming airplanes into buildings, so we were unprepared for the events of September 11th.

In both cases, people managing risks tended to assume that the near future would be a lot like the recent past. In general, people's perception of the hazards that have a very low probability of occurring, but have catastrophic impact when they do occur are very distorted. Because of the low probability there may be very few, or even no historical instances, which makes these hazards difficult to study and makes their probabilities difficult to measure accurately.

People's response to such hazards tends towards either overreaction or underreaction. People may focus on the fears associated with the consequences and take very expensive precautions to avoid minuscule risks or else they may focus on the low probability and assume that the danger will never become real. Finding balance is difficult and science is often unable to provide clear answers to help us find a rational approach.

Moreover, when the issues are *trans-scientific*, as the questions we will be asking this term are, scientists' own biases and political beliefs may interfere with their ability to provide impartial advice. Consider the following three questions:

- How quickly will demand for electricity in the United States grow in the next twenty years?
- How will the cost of manufacturing photovoltaic solar cells (which generate electricity from sunlight) change in the next twenty years?
- Will uranium mines start to run out of uranium ore in the next twenty years?

There is no particular reason for any of these three questions to have anything to do with any of the others. Each is pretty much independent of the others and each is trans-scientific: it has a correct answer, but we will not be able to determine that answer for twenty years. Surprisingly, when a number of nuclear engineers were asked these three questions, their answers tended to form distinct patterns: those who thought demand for electricity would rise a lot also tended to think that photovoltaic cells would remain expensive and that we would run out of uranium. Those who thought that demand for electricity would not rise a lot tended to think that the cost of photovoltaics would drop and that we would not run out of uranium.

It turned out that those in the first group, who thought demand for electricity would rise and that we would also run out of options (photovoltaics, uranium) tended to be those who supported a technology called **breeder reactors**, a type of nuclear reactor that can produce nuclear fuel (plutonium) from a type of uranium that is plentiful, but ordinarily useless in reactors. Those in the second group tended to oppose building breeder reactors. The guesses about these three trans-scientific questions were largely predetermined by each engineer's personal opinion on the nonscientific question of whether it's a good idea to build breeder reactors.

In asking about scientific estimates of energy and water resources it's worth considering whether similar phenomena may be affecting scientists' best attempts to provide objective and impartial estimates of the future demand for these resources and the supplies we will have available.