October 17, 2005

Lecture notes Related to the Precautionary Principle

1 The Precautionary Principle

The Precautionary Principle says that the absence of scientific certainty should not be used as an excuse for inaction. The fact that we are not certain that a particular pollutant is dangerous does not mean that we should ignore the possible danger. This Principle is not a precise guide for action; it merely cautions us against a particular excuse for inaction: uncertain science.

The Precautionary Principle is relevant in circumstances when uncertainty is important. These notes discuss three aspects of economic modeling of uncertainty. The first of these involves our tolerance for risk. The second explains how anticipated learning should affect our choices. The third illustrates how we can model learning. It is possible to consider all of these aspects of the problem together, but for simplicity I deal with them separately.

2 Risk aversion

Economists model risk and uncertainty in a policy setting by supposing that a decision-maker maximizes the *expectation* of something. In this setting, a bad outcome that occurs with even very small probability can have a major influence on the optimal decision, if the decision-maker is risk-averse. Risk aversion makes us more cautious, and extreme risk aversion makes us very cautious. The degree of risk aversion is related to our "preferences", in particular, to our preferences regarding risk. Economists typically regard preferences as exogenous, that is, we take them as given. There need be no relation between an individual's risk preference and her rationality, just as there need be no relation between her preferences for ice cream and her rationality. There is no reason to suppose that an individual who behaves in an extremely risk averse manner is irrational.

For example, suppose that the benefit of GHG emissions, (x) is $x - \frac{x^2}{2}$ and the environmental damage associated with emissions is αx^2, where α is an uncertain parameter. The benefits minus costs are

$$y = x - \frac{x^2}{2} - \alpha x^2.$$

Since α is uncertain, we treat y as a random variable.

The magnitude of α determines the relation between the level of emissions and the magnitude of damages caused by GHG. We summarize our beliefs about the magnitude of α using a "subjective probability distribution" over the various possible values of this parameter. For example, suppose that we think that α is either 0 or 1, with equal probability; this is a particular subjective probability distribution. This subjective probability distribution is simply a way of formalizing the beliefs that we have about the different possibilities.

An example may help to distinguish between objective and subjective probabilities. Consider a coin which when tossed in a certain manner has probability p of coming up heads. If the coin is not "fair", then $p \neq 0.5$. p is an objective probability, in the sense that it is determined by physical characteristics of the coin and the tossing process. If the parameter p enters an optimization problem, but we do not know its true value, we might instead use an estimate (or guess) of its value. In this case we refer to the estimate as a "subjective" probability. Depending on the circumstances, e.g. whether we can perform experiments to learn the unknown value of p, we may be able to obtain a precise estimate, or we practically have to guess. Whichever is the case, the use of a subjective probability distribution is valuable because it forces us to be precise about the assumptions that we are using in making our decision.

If society is risk neutral, it chooses x to maximize the expectation of y. In the example above, where our subjective probability distribution states that α is either equal to 0 or to 1, with equal probability, the risk neutral planner chooses emissions, x to maximize.

$$E(y) = x - \frac{x^2}{2} - \frac{1}{2}\left(1x^2 + 0\right) = x - x^2.$$

The optimal value of x is $x = \frac{1}{2}$.

If society is risk averse, the planner can choose x to maximize the expectation of the *utility* of y (rather than choosing x to maximize the expectation of y). The more risk averse the planner it, the lower the optimal value of x. At the extreme case (an infinitely risk averse planner), the planner behaves in the most pessimistic manner, and acts as if the worst outcome ($\alpha = 1$) is certain. In that case he chooses x to maximize

$$x - \frac{x^2}{2} - 1x^2 = x - \frac{3}{2}x^2. \tag{1}$$

Here, the optimal value of x is $x = \frac{1}{3}$. As society's risk aversion goes from 0 (risk neutrality) to infinity, the optimal first period decision ranges between the levels $\frac{1}{2}$ and $\frac{1}{3}$. (At the other extreme, if society has "infinite risk preference", it acts as if $\alpha = 0$ with certainty. In that case, the optimal value of x is $x = 1$.)

The Precautionary Principle does not instruct us to assume the worst. It merely says that we need to take into account the possibility of bad outcomes. The Principle does not tell us exactly how to do this.

3 Decision-making with Anticipated Learning

A related question is how we should factor the prospect of *better information in the future* into our current decision. If we never expected to improve our understanding of this relation, we would maximize expected net benefits (or the expected utility of net benefits in the case of risk aversion), as above, using our subjective distribution. If we expect to improve our scientific understanding in the future, surely that should influence our current behavior. What is the direction of influence? That is, should the anticipation of better information (in the future) make us more cautious today or less cautious today?

The general answer is ambiguous – anticipated learning could make us either more or less cautious. I will give an example of both possibilities. The first example introduces the idea of an "option value". This example is important because it provides the intuition that many economists use to conclude that the prospect of better information in the future should make us behave more cautiously in the current period. The second example is a variation of the GHG model above; this example shows that the prospect of better information in the future might make us less cautious in the current period.

For both of these examples, I will assume that the decision-maker is risk neutral. He maximizes the expectation of net benefits, rather then the expectation of the utility of benefits. The effect, on a current decision, of risk aversion, and the effect of anticipated learning, are two distinct issues. So it is better to treat them separately, in an introduction.

Example 1 In the first example, the anticipation of learning makes us more cautious. Suppose that we can make an investment that costs \$60 and that returns either \$100 or \$50, with equal probability, in the next period. We have a one period discount factor of $\beta < 1$. We can make the investment either this period or the next period or not at all. First suppose that we do not expect to learn anything about the distribution of the payoff (the probabilities). If we invest today our expected payoff is

$$-60 + 0.5\beta\,(100 + 50) = 75\beta - 60.$$

Provided that $75\beta - 60 > 0$, i.e. that $\beta > \frac{60}{75} = 0.8$, the expected payoff from investing today is positive. The expected payoff of investing in the next period is also positive in this case, but because of discounting – and because of the assumption that we don't learn anything about the investment in the second period, it is lower than the expected payoff of investing in the first period. Therefore, for this example (given that $\beta > 0.8$) the optimal decision is to invest today.

Now suppose that in the first period we think that there is a 0.5 chance of the high or the low outcome (as before) but now we know that at the beginning of the second period we will learn whether the investment has the high or low payoff. If we invest in the first period, our expected payoff is $75\beta - 60$ as before. (Again, we assume that this amount is positive, i.e. that $\beta > 0.8$.) If we do not invest in the first period, we have the option of investing in the second period, when we will know the value of the investment. If we discover that the investment pays \$50, of course we choose not to invest, and if we discover that it pays \$100, of course we choose to invest, earning $-60 + \beta 100$.

The present value of the expectation of the payoff if we decide not to invest in the first period is β times the sum of payoffs in the two events multiplied by the probability of the two events. (Remember, we earn 0 if we discover that the investment pays \$50, because in that case we do not invest.) The present value of the expected payoff is therefore

$$0.5\beta\,(-60 + \beta 100)\,.$$

The difference between the expected value of not investing in the first period and investing in the first period is

$$\begin{aligned} & 0.5\beta\left(-60+\beta 100\right)-\left(-60+0.5\beta\left(100+50\right)\right) \\ = & -105.0\beta+50.0\beta^2+60.0>0. \end{aligned}$$

This amount is called the "option value", because it is equal to the additional expected value that we obtain from the project, due to the option of waiting to make our decision until we have better information. For example, if $\beta = 1$, the option value is $-105.0 + 50.0 + 60.0 = 5.0$.

Example 2 This example is a two-period version of the global warming model that I presented above. Here we ignore discounting, for simplicity. Suppose that in each of two periods we decide how much GHGs to emit, x_1 and x_2. In each period, our payoff from emitting GHGs (i.e., the decrease in the amount that we would have to pay in order to abate GHGs) is

$$x_i - \frac{x_i^2}{2}.$$

After the second period, the increased stock of GHG is $(x_1 + x_2)$. (I assume that there is no decay in the stock of GHGs.) Suppose that the damages associated with GHGs are

$$\alpha\left(x_1+x_2\right)^2.$$

In period 1 we do not know the value of α but we think that it is either 0 or 1; each of these possible values has equal probability.

As before, first consider how we would behave in the first period, if we do not expect to improve our information during the planning horizon. In that case, we can decide, in the first period, on emissions in both periods. Because of the simplifying assumption that there is no discounting, the optimal value of emissions in each period is the same. We want to maximize the value of emissions in the two periods, minus the expected value associated with damages. The problem is

$$\max_x 2\left(x-\frac{x^2}{2}\right)-0.5(1)\left((x+x)^2\right)=2.0x-3.0x^2. \tag{2}$$

The optimal first period emissions is $x = \frac{1}{3} = 0.333$. (For this problem, the expected value of the program is also equal to 0.333.)

Now consider how we would behave if we expect to learn the value of α before we need to make our second period decision. We solve this problem by "working backwards", i.e. first figuring out what would be optimal in the second period, depending on the action that we took in the first period and the information that we acquired about the value of α. Suppose that in the first period the level of emissions was y. If at the beginning of the second period, we discover that $\alpha = 0$, the the value of the optimal decision solves the following problem:

$$V_0(y) = \max\left(x - \frac{x^2}{2}\right).$$

It is easy to show that the optimal level of second period emission is 1 and that $V_0(y) = .5$. If at the beginning of the second period we discover that $\alpha = 1$, the problem is

$$V_1(y) = \max\left(x - \frac{x^2}{2}\right) - 1\left(x + y\right)^2.$$

After a bit of algebra, you can show that the optimal level of the second period emissions is a decreasing function of y,

$$x_2^* = \frac{1}{3} - \frac{2}{3}y \tag{3}$$

and that the optimal value of the second period payoff is

$$V_1(y) = \frac{1}{6} - \frac{2}{3}y - \frac{1}{3}y^2.$$

In the first period, the optimal choice of first period emissions (given that we expect to learn the value of α before we make our second period decision) is the solution to

$$\max_y\left(y - \frac{y^2}{2} + (0.5)(0.5) + (0.5)\left(\frac{1}{6} - \frac{2}{3}y - \frac{1}{3}y^2\right)\right)$$

The optimal expected value of this program is 0.75; in the case without learning, we saw that the optimal value was 0.333. Thus anticipated learning more than doubles the expected value of the program. The optimal first period level of emissions is $y = 0.5$. Recall that in the case without anticipated learning, the optimal first (and second) period level of emissions is 0.333. Thus, anticipated learning increases the first period emissions by a factor of $\frac{.5}{.333} = 1.5$, i.e. by 50%.

In the second period, if the decision-maker learns that $\alpha = 0$, it is optimal to emit $x_2 = 1$ in the second period. If he learns that $\alpha = 1$, it is optimal to use the decision rule in equation (3), together with the first period level of emission ($y = x_1 = .5$) to obtain the second period level of emissions

$$\frac{1}{3} - \frac{2}{3}y = \frac{1}{3} - \frac{2}{3}(.5) = 0.$$

In this example, the anticipation of future learning makes us *less cautious* in the first period (that is, we increase emissions relative to the no-learning case). If we know that we will learn, in the second period, about the damage associated with GHG emissions, then we will adjust our second period decision in light of this information. This ability to make future adjustments (when we have better information) decreases the expected marginal damage of current emissions, thereby increasing the optimal first period level of emissions.

4 Statistical Learning

This section of the notes explains how we can use information to update our beliefs about an event. Suppose that we are uncertain whether a "project" (or a commodity, such as hormone-fed beef) is harmful or benign. Our problem is to pick a rule that determines when it is permissible to ban the project (or commodity). For example, we might decide that it is legitimate to ban the project when we believe that the probability that it is harmful exceeds a certain threshold. We will assume that there is agreement about the level of the threshold, so our problem is to decide how to use information to determine whether that threshold has been crossed. In order to do this, we need a formal model, which I now describe.

We begin with the *subjective* belief that the project is harmful (B for bad) with probability α; it is not harmful (G for good) with probability $1 - \alpha$. We undertake experiments ("trials") that provide information about the true nature of the project. The nature of the experiment/trial depends on the specific problem. For example, an experiment might involve giving a laboratory animal a chemical, or it might involve observing the outcome of consumption in a human population. We can think of consumption in the US of hormone-fed beef as a kind of experiment, that might be useful to the EC in deciding whether hormone-fed beef is harmful.

A "success" is defined as a trial that provides a piece of evidence that the activity is harmful. A test produces a success for a harmful project with probability p and a test produces a success for a not-harmful project with probability q. Thus, the probability of getting a false negative (a result that says "not harmful" when the project really is harmful) is $1-p$ and the probability of getting a false positive (a result that says "harmful" when the project is really not harmful) is q. This model is extremely simple, because the outcome of each trial is either a success or a failure, i.e., the outcome is binary. In a more general model, many different outcomes might be possible.

The probability of M successes with N tests when the project is actually harmful is given by the binomial formula:

$$s\left(M,N;B\right)=\frac{N!}{M!(N-M)!}p^{M}\left(1-p\right)^{N-M}.$$

(The notation $x!$, read as "x factorial" means $x\left(x-1\right)\left(x-2\right)...1$. For example, $5!=5\left(4\right)\left(3\right)\left(2\right)\left(1\right)=120$.) The first argument in the function $s\left(M,N;B\right)$ indicates the number of success, the second indicates the number of trials, and the third (the B) indicates that the project is actually harmful (bad). Similarly, the probability of M successes with N tests when the project is actually not harmful is given by the binomial formula:

$$s\left(M,N;G\right)=\frac{N!}{M!(N-M)!}q^{M}\left(1-q\right)^{N-M}.$$

Suppose that we begin with the subjective probability α that the project is harmful, and we conduct N tests, of which M result in a "success" (indicating that the project is harmful). The posterior probability, λ, that a project is harmful given M successes with N tests is

$$\lambda=\frac{\alpha p^{M}\left(1-p\right)^{N-M}}{\alpha p^{M}\left(1-p\right)^{N-M}+\left(1-\alpha\right)q^{M}\left(1-q\right)^{N-M}}. \tag{4}$$

If our decision rule states that we ban the project only if the subjective probability that it is harmful exceeds a threshold level, we can use equation (1) to determine the combination of M and N that result in a ban.

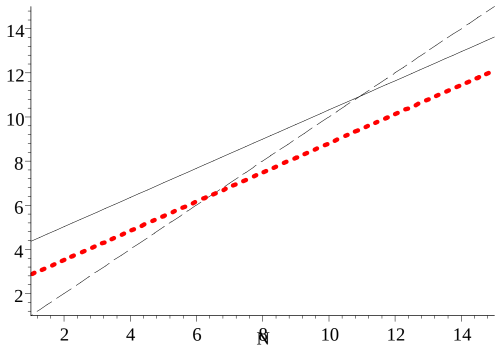

Figure 1: Solid curve $\lambda = .95$; dotted curve $\lambda = .7$ $(\alpha = .1, p = .8, q = .5)$

Figure 1 shows the combination of N and M needed to provide either a .95 or a 7 probability of concluding that the activity is harmful (the solid and the dotted lines, respectively). This figure assumes that.$\alpha = .1$, $p = .8$, $q = .5$. Since it is not possible that $M > N$, only the parts of the curves below the 45 degree line (the dashed curve) are relevant.

Figure 2 shows the combination of N and M needed to provide a .95 probability of harm, for $p = .8$, $q = .1$, with different priors, $\alpha = .5$ (solid) and $\alpha = .1$ (dotted). Once again, since $M \leq N$, the set of feasible outcomes lie below the (dashed) 45 degree line. The minimum number of experiments, N^*, needed to establish the hypothesis that the project is harmful, with a probability at least λ, is the smallest integer N greater than the intersection between the graph of equation (**??**) and the 45 degree line. This intersection, denoted N^*, is given by

$$N^* = \frac{\ln \lambda + \ln(1-\alpha) - \ln(1-\lambda) - \ln \alpha}{-\ln q + \ln p}.$$

If I ignore the "integer constraint", I can think of N^* as the minimum number of trials needed to establish the "standard of proof", λ. If each trial is expensive, I would like to have as few trials as possible in order to reach a certain standard of proof. Consequently, to reach the standard of proof, I would require N^* consecutive successes.

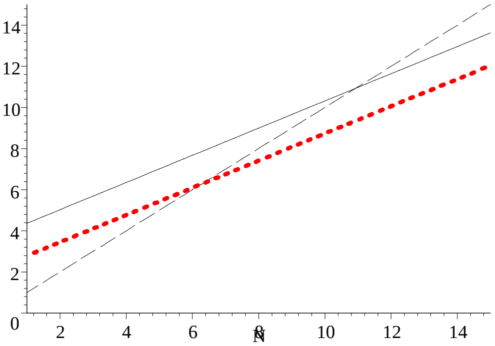

Figure 2: Solid curve $\alpha = .5$, dotted curve $\alpha = .1$ ($p = .8, q = .5, \lambda = .95$)

This model provides a simple way to show how the number of experiments affects the degree of certainty. It also show the role of the prior probability in determining beliefs after a particular set of outcomes. Reasonable people could examine the same evidence and reach different conclusions, because of different priors. Figure 3 shows the number of consecutive successes that are needed in order to establish the hypothesis that the project is harmful, with probability $\lambda = .95$, when $p = .8$ and $q = .5$. When the initial assessment is α close to 0 (so that we think it is very unlikely that the project is harmful, we need over 15 consecutive successes in order to establish harm (with probability $\lambda = .95$). When our initial degree of certainty that the project is harmful is quite close to the threshold $\lambda = .95$, we need only one successful trial to reach the threshold. The point of this example is that if you have two people who have exactly the same model (except for α) and agree on the same critical threshold (the value of λ), they may require vastly different amounts of evidence in order to agree on a course of action (either banning or accepting the project).

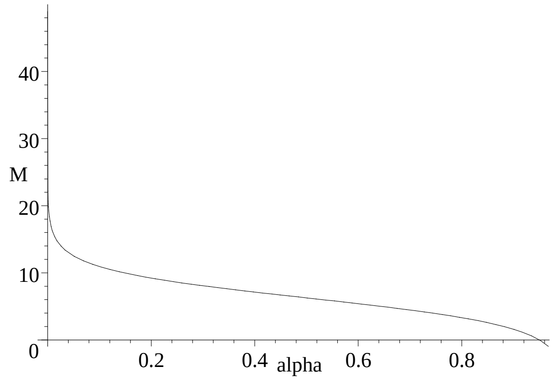

Figure 3: Number of consecutive successes needed, as a function of α, when $\lambda = .95$, $q = .5$ and $p = .8$.