Oxford Economic Papers 39 (1987), 190–222

ISSUES IN THE APPLICATION OF COST-BENEFIT ANALYSIS TO ENERGY PROJECTS IN DEVELOPING COUNTRIES*

By SUDHIR ANAND *and* BARRY NALEBUFF

1. Introduction

THE THEORY behind cost-benefit analysis is relatively well developed [see Little and Mirrlees (1968), (1974); Dasgupta, Marglin and Sen (1972); Mason and Merton (1984)], but its application to real world problems is not always easy. This paper considers the application of cost-benefit analysis to energy problems in oil-importing developing countries (OIDCs).

The results are divided into four parts: Section 2 presents a methodology for handling uncertainty; Section 3 studies the implications of exhaustibility; Section 4 examines the externalities associated with exploration; Section 5 applies these tools in an analysis of the costs associated with importing oil.

In Section 2 we begin with a discussion of the theoretical issues involved in cost-benefit analysis in the presence of uncertainty. Problems arise when there are conflicting estimates of a project's chance of success. The analysis must consider how the range of possible outcomes affects both the total risk borne in the economy and the income distribution. Uncertain future returns have to be appropriately discounted to the present. Projects which deplete exhaustible resources are irreversible and therefore the cost-benefit analysis must take into account the value of current reserves.

In Section 3 Hotelling's rule is presented as the starting point for forming expectations about the price path of oil and other exhaustible energy resources. All project appraisals for programs ranging from new energy development to conservation or stockpiling are strongly influenced by the expected future price of energy, and oil in particular. Energy prices today and in the future depend on estimates of total world hydrocarbon reserves, their extraction costs, and the predicted availability of backstop technologies. Uncertainty about future reserves leads to variability in supplies and price, and causes an inefficiency in the intertemporal allocation of oil; depletion takes place more slowly in order to maintain flexibility in the event of worse-case scenarios. The gains from developing new energy supplies (and conservation) include lower prices, reduced vulnerability, and

* We are particularly grateful to Harinder Kohli for motivating this work and for providing continuous support and encouragement. We appreciate discussions with, or comments on earlier drafts from, Panos Cavoulacos, Angus Deaton, Kemal Dervis, William Hogan, Harinder Kohli, Deepak Lal, Alan Manning, Ricardo Martin, James Mirrlees, John Page, D. C. Rao, Paul Seabright, Edilberto Segura, Robert Stobaugh, Yannis Stournaras, Michael Webb, and two anonymous referees. Audrey Hiscock, Jean Ponchamni and Tony Tenorio typed the different versions of this paper, and we are extremely grateful to them. Financial support for the research was provided by the Industry Department of the World Bank. An expanded version of this paper appeared as World Bank Staff Working Paper Number 738 [Anand and Nalebuff (1985)].

greater competition in the world market. How countries respond to the expected price path of oil will depend on their discount rate. Oil-importing developing countries frequently have above average discount rates due to below average creditworthiness and constraints arising from imperfect capital markets or restrictions on foreign investment; from their perspective, the high price of oil today relative to the future should lead them to deplete their own resources at a faster rate.

Section 4 widens the scope of cost-benefit analysis to capture the externalities involved when research strategies for exploration are coordinated. The uncertainty concerning the total size of oil, gas and other energy reserves hurts all oil-importing countries. Because information is a public good, the gains to information gathering extend beyond a country's border. Projects that reduce the uncertainty associated with reserve levels should be evaluated from an international perspective. There may also be economies of scale in gathering information; a little knowledge may be useless but it is still costly to acquire. Related is the question of economies of scope: to develop renewable energy resources, how many competing directions of research should be financed? Since only the best solutions will be implemented, this reduces the advantages to diversification. With regard to diversification, there is an important distinction between risk-sharing and risk-pooling: sharing involves spreading risk among a larger number of individuals, while pooling involves accepting a larger number of less than perfectly correlated projects. Pooling together independent projects increases both the total expected return and the total variance. When the risks are already spread as thinly as possible, there is no safety in the law of large numbers.

Section 5 examines the options for an oil-importing developing country to reduce both its vulnerability and dependence on imported oil. Stockpiles can be an effective and inexpensive tool to lower a country's vulnerability to short-run supply disruptions. Other strategies which reduce dependence on foreign oil in the long-run include (i) increasing conservation, (ii) promoting renewable energy projects and, in the medium-run, (iii) depleting their stocks of exhaustible energy resources. Because of capital constraints and the difficulties associated with financing private contracts, there may be too little energy development taking place in oil-importing developing countries. The evidence presented by Blitzer *et al.* (1984) is quite striking. While 70% of worldwide exploratory wells in 1980 were drilled in the United States compared with about 3% for the entire group of oil-importing developing countries, the success rates were almost equal at 30%. During the 1970s the cost-effectiveness of developing countries was four times greater: 1.6 barrels of hydrocarbon reserves were generated per dollar of investment in developing countries compared with 0.4 barrels of reserves in the United States.

Section 6 brings together the theory and applications in a brief conclusion that summarizes the results.

2. Cost-benefit analysis under uncertainty

The first and most obvious problem in evaluating a project is the fact that nothing is certain. As a starting point the Expected Present Monetary Value (EPV) is used as a proxy for the expected welfare value of a project. If in period t the probability of an output worth Z is $p_t(Z)$, and r_t is the appropriate average discount rate between periods 1 and t, then

$$\mathrm{EPV} = \sum_t \left[\int Z p_t(Z) \, \mathrm{d}Z \right] \Big/ (1 + r_t)^t. \tag{1}$$

However, the expected present value index is only the first step in project appraisal. The following subsections discuss how this criterion must be adjusted to take account of: conflicting probability estimates $p(Z)$; the costs of uncertainty; income distribution effects; the appropriate discount rate for the project; the value of the status quo, and project timing.

2.1. *Evaluating conflicting probability estimates*

Probabilities are only estimates, and in practice they may be no better than rough guesses. Thus, there may be several conflicting estimates of the likelihood of any outcome. This uncertainty is contrasted in the next subsection with the different type of uncertainty caused by the possibility of multiple outcomes; for example, in a coin toss, the outcome, either heads or tails, is uncertain but the probability of heads is known to be 0.5 with certainty. Here we focus on how to evaluate EPV when the possible outcomes are known (e.g., heads or tails) but the probability of each event is uncertain.

Consider the following two situations. In Case A, there are two estimates: both show that the probability of a 1 million rupee return is 0.25. In situation B, there is uncertainty about the estimates. Estimate 1 shows the probability of a 1 million rupee return to be 0.10; estimate 2 shows the probability to be 0.40. In the second case there are conflicting estimates of the chance of success, one putting it at only 10% while the other more optimistically expecting a 40% chance of success. The probabilities are shown in the table below.

Case	A	B
Estimate 1	0.25	0.1
Estimate 2	0.25	0.4
Average	0.25	0.25

Since both estimates agree in Case A, it is straightforward that the expected present value should be calculated using a 0.25 probability of a 1 million rupee return. For Case B, assume that the two conflicting estimates

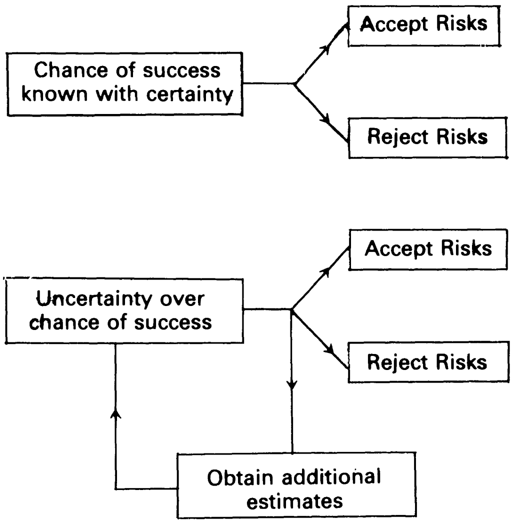

FIG. 2.1. Decision chart with uncertain probabilities of success.

are equally precise.[1] Both estimates are then given equal weight and the expected chance of a 1 million rupee return in Case B is

$$p(1{,}000{,}000) = (0.5 \times 0.1) + (0.5 \times 0.4) = 0.25,$$

which is the same as in Case A. The expected return, the variance, and all other moments are equal in the two situations.

The project in situation B should *not* be penalized because there are conflicting estimates about its actual chance of success. There is no risk aversion to uncertainty about the estimated probability of success. Indeed there can even be a benefit arising from uncertainty in the estimated probability of success.

The benefit from uncertainty becomes clearer in the context of sequential decision-making [Rothschild (1971)]. When the probability of success is known with certainty, there is no more information that can be found out. For example, the probability of heads when flipping an unbiased coin is known to be 1/2. With this information, the only remaining decision is whether or not to take the gamble. Alternatively, when the probability of success is unknown, there is the additional possibility of gaining further information (see decision chart in Fig. 2.1). This new information will lead to new probabilities and thus a potentially different decision. Of course, gathering this information may be expensive and thus it is not necessarily worthwhile.

[1] Technically, the probability estimate is the mean of some distribution of probabilities. The precision is inversely related to the variance of this probability estimate.

The value of the additional information depends on the extent to which the estimates diverge. When the probability is known (as in the case of the unbiased coin) there is no value to obtaining additional estimates. To the extent that the conflicting estimates are farther apart, there is more uncertainty to be resolved.[2] In our example, there is greater confidence that 0.25 is the actual chance of success in Case A since both estimates agree; in Case B at least one, and perhaps both, of the estimated probabilities must be wrong.

This point may be summarized by noting that once the decision is made to accept the risks, all that matters is the expected probability of each possible outcome. There is no reason to be risk averse due to divergent probability estimates. But, divergent estimates suggest that there may be a value to obtaining further information.

2.2. *The costs of uncertainty*

A cost from uncertainty can arise when a project has several possible outcomes. Individuals are averse to fluctuations in their income. The cost from bad outcomes is greater than the benefits from correspondingly good outcomes. Decision-makers prefer a sure thing to any pure gamble with the same expected outcome. And, except in the hypothetical case of complete contingent claims markets, it is not possible to purchase complete insurance and eliminate the effects of uncertainty [Arrow and Lind (1970)].

Risk aversion becomes important when the range of the possible income levels for *individuals* is more than trivially affected by the outcome of the project. For Bangladesh, with its population of about 100 million, even a large project having returns between $50 million and $150 million results in an income variation per person of at most a dollar. This may turn out to be a small risk relative to even a low per capita national income of $100.

There are two factors that complicate this straightforward example of a risk-sharing argument. Project outcomes are correlated with the state of the economy; because there are real costs to the risks associated with fluctuations in the economy, adding *incremental* risk will be important. Secondly, project risks may not be spread evenly across the population; appraisals must include consideration of the differential effects of a project on different groups of people. These points are taken up in Sections 2.3 and 2.5.

2.3. *Valuation of project output under uncertainty*

This subsection demonstrates the importance of taking into account the correlation between fluctuations in per capita income and project output when calculating the costs of uncertainty.

[2] While the value of additional information may be high, the existence of conflicting probability estimates does not in itself imply anything about the costs of additional information. As always, benefits must be compared with costs.

Let Y = national income per person; Z = net project output per person; $V(Y)$ = social valuation of representative person's income, where $V'(Y) > 0$ and $V''(Y) < 0$. A project Z should be done if it raises expected social welfare, i.e., $E[V(Y+Z)] > E[V(Y)]$, where $E[\ ]$ denotes the expectation operator.

Define the certainty equivalent z of a project as that income which the government would consider equally satisfactory as a project Z whose net output is random. In other words, z is defined through the equation:

$$E[V(Y+Z)] = E[V(Y+z)]. \tag{2}$$

Letting $\bar{Y} = E[Y]$ and $\bar{Z} = E[Z]$, we can expand the left-hand and right-hand sides of equation (2) around $(\bar{Y} + \bar{Z})$. The Taylor series expansions to second order of the left-hand and right-hand sides, respectively, are:

$$\begin{aligned} E[V(Y+Z)] &= V(\bar{Y}+\bar{Z}) + V'(\bar{Y}+\bar{Z})E[Y-\bar{Y}+Z-\bar{Z}] \\ &\quad + (1/2)V''(\bar{Y}+\bar{Z})E[(Y-\bar{Y})+(Z-\bar{Z})]^2 \end{aligned} \tag{3}$$

$$\begin{aligned} E[V(Y+z)] &= V(\bar{Y}+\bar{Z}) + V'(\bar{Y}+\bar{Z})E[Y-\bar{Y}+z-\bar{Z}] \\ &\quad + (1/2)V''(\bar{Y}+\bar{Z})E[(Y-\bar{Y})+(z-\bar{Z})]^2 \end{aligned} \tag{4}$$

Equating the right-hand sides of equations (3) and (4) yields

$$\begin{aligned} &(1/2)V''(\bar{Y}+\bar{Z})[\text{Var}(Y) + \text{Var}(Z) + 2\,\text{Cov}(Y, Z)] \\ &\quad = (z-\bar{Z})V'(\bar{Y}+\bar{Z}) + (1/2)V''(\bar{Y}+\bar{Z})[\text{Var}(Y) + (z-\bar{Z})^2]. \end{aligned} \tag{5}$$

Hence,

$$\begin{aligned} z &= \bar{Z} + (1/2)\frac{V''(\bar{Y}+\bar{Z})}{V'(\bar{Y}+\bar{Z})}[\text{Var}(Z) + 2\,\text{Cov}(Y, Z) - (z-\bar{Z})^2] \\ &= \bar{Z} - (1/2)\frac{R(\bar{Y}+\bar{Z})}{(\bar{Y}+\bar{Z})}[\text{Var}(Z) + 2\,\text{Cov}(Y, Z)] \end{aligned} \tag{6}$$

where $R(y) = -yV''(y)/V'(y)$ is the relative risk aversion (or elasticity of marginal valuation of income) at income level y, and $(z-\bar{Z})^2/(\bar{Y}+\bar{Z})$ is assumed negligible compared with $(z-\bar{Z})$.

The costs of uncertainty can be broken up into two terms: the variance effect and the covariance effect. We look at these two terms sequentially in order to gauge their relative importance. As will be seen below, the cost due to the variance term will typically be significantly smaller than the benefit associated with the (negative) covariance term for energy projects in OIDCs.

Both terms are equally affected by the choice of the relative risk aversion parameter, denoted by R. The literature on alternative estimates of the elasticity of social marginal valuation of income has been reviewed by Stern (1977) [see also Anand (1973)], where values between 1.5 and 2.5 are suggested as being reasonable and broadly acceptable. We assume a constant value of $R = 2$ in the relevant range around $(\bar{Y} + \bar{Z})$, and hence the valuation function specified as $V(y) = a - by^{-1}$.

First we focus on the variance term. If national income per person is itself certain (i.e., Y is not a random variable), then $\mathrm{Cov}\,(Y, Z) = 0$. In this case, equation (6) reduces to:

$$z = \bar{Z} - (1/2)R\frac{\mathrm{Var}\,(Z)}{(\bar{Y} + \bar{Z})}. \qquad (7)$$

Consider a large government project in which per capita income can vary by 1% over the range of possible outcomes. For the example of Bangladesh, this would perhaps be a project with an outcome that is uncertain over a $100 million range. Even for countries with smaller populations such as Mali, Burkina Faso, or Botswana where energy projects of $150 million would not be unrealistic, the *uncertainty* associated with the project outcome will almost always be less than 5% of per capita income.

Starting with a per capita income of $100, an example of a project with a 1% range is:

$$Z = \begin{cases} \$0.5 \text{ with probability } 1/2 \\ \$1.5 \text{ with probability } 1/2 \end{cases} \qquad (8)$$

The project has an expected per capita return $\bar{Z} = 1$ and a variance $\mathrm{Var}\,(Z) = 0.25$. The risk premium or proportion by which the expected monetary value of the project has to be deflated because of the uncertain return is:

$$\frac{(\bar{Z} - z)}{\bar{Z}} = (1/2)(R)\frac{\mathrm{Var}\,(Z)}{\bar{Z}(\bar{Y} + \bar{Z})}$$
$$= (1/2)(2)\frac{0.25}{101} = 0.0025. \qquad (9)$$

Thus, the risk premium is just under one-quarter of one percent—a miniscule amount which, for practical purposes, may be ignored.

When a per capita income of $100 is subject to a 5% risk, the risk premium is twenty-five times larger, and a downward adjustment of about 6% is required. Although this is now relevant, we see below that it may still be only one-eighth as large as a realistic correction for the corresponding *covariance*.

Economists are concerned with the value of the energy project conditional on the state of the economy and of the world. Knowing the correlation between the economy and the project is important both for determining the *value* to the economy of money from the project and for determining the *amount* of money the project is likely to generate. A small variation of the earlier examples shows how the *correlation* between national income and the project's return can have a much more significant effect on the certainty equivalent monetary value of the project.

The first example demonstrated that the risk adjustments associated with even large projects are relatively small (with the qualification that income distribution effects must also be considered). But to determine the monetary value of a project, it is essential to realize that the price of the output, especially in the case of energy, will in general be related to the state of the economy. The welfare value of the project will also depend on the state of the economy. We are not worried about causality, whether the project affects the economy or the state of the economy affects the project. Here, we are emphasizing that the swings in the economy may be *correlated* with the output of the project and this should be taken into account. For energy projects in oil-importing developing countries it is very likely that the marginal welfare value of the project's output, $ZV'(Y)$, and national income per person in the economy as a whole, Y, are *negatively* correlated since they are oppositely affected by variations in the world price of oil. An oil embargo will make the development of internal energy sources (including energy conservation) particularly attractive both because the economy will be depressed due to the embargo and because of the high price of the remaining available oil; the marginal welfare value of income will be high when the project output is high. That is, for an economy dependent on oil imports, oil will usually be expensive when the economy is depressed and vice versa. These effects would be further amplified by changes in the shadow price of foreign exchange, which are also likely to be negatively correlated with the economy.[3]

We can now make illustrative estimates of the extent to which the expected monetary value of output may have to be adjusted to take account of fluctuations in the marginal value of income. Consider the earlier example of a 1% variation again, but now suppose that when the output of the project is \$0.5, national income per head is \$110, and when the output of the project is \$1.5, national income per head is \$90. Again suppose that the two events occur with probability 1/2. We continue to assume an elasticity of marginal valuation of income of 2 in this range, i.e., $V'(y) = by^{-2}$. In this case, there is a *negative covariance* between Y and Z, given by

$$\begin{aligned}\text{Cov}(Y, Z) &= E[(Y - \bar{Y})(Z - \bar{Z})] \\ &= -5. \end{aligned} \tag{10}$$

The certainty equivalent monetary value for this project, z, is *greater* than its expected monetary value, $\bar{Z}$, and the proportional upward adjustment

[3] A project appraisal has to consider all the general equilibrium effects of changes in the economy. It is necessary to estimate changes both in the price of energy and in the prices of all other commodities (i.e. foreign exchange, interest rates, transportation, etc.). To calculate the EPV, the expectation of the shadow value of output has to be taken across all states of nature corresponding to oil price (or supply) shocks. Strictly speaking, this requires working out all the shadow prices in each state of nature contingent on assumptions about government policy. This is not an easy task!

can be obtained from equation (6) straightforwardly as

$$\frac{(z-\bar{Z})}{\bar{Z}} = -(1/2)(R)\frac{[2\,\mathrm{Cov}\,(Y, Z)]}{(\bar{Y}+\bar{Z})\bar{Z}}$$

$$= (1/2)\frac{(2)10}{101} = 0.099. \qquad (11)$$

This ignores the (downward) adjustment for Var (Z) which was shown to be negligible (0.0025). The negative correlation between Y and Z therefore calls for an upward adjustment in $\bar{Z}$ of 9.9%. For $R = 4$ the increase is twice as large, almost 20% of the expected monetary value of the project. Such upward adjustments imply corresponding upward changes to the expected present value of the entire stream of returns from the project.

When the project outcome ranges over 5% of per capita income (e.g., $Z = -\$1.5$ when income is \$110 and $Z = \$3.5$ when income is \$90) then the covariance correction is five times larger at almost 50%. It is now important to include the 6% reduction due to the variance term, which still leaves a net upward correction of approximately 44%.

Our conclusion from this analysis is that it may be necessary to make significant adjustments to account for uncertainty. When calculating an example of the size of the effects associated with the uncertainty in the price and availability of oil imports, oil dependency can lead to a negative correlation between the output of an energy project and national income per head which is forty times more important than the variance in income from the project for the average individual. This ratio of 40 to 1 does not depend on the magnitude of the relative risk aversion, as changes in R affect both terms proportionally. However, the ratio does depend on the size of the project relative to national income; doubling the size of the project (or considering a country with half the per capita income) also doubles the *relative* importance of the variance effect. Even in the largest projects (5% case), though, the covariance term is still eight times larger and of the opposite sign.

A negative correlation between the return on asset i and the portfolio as a whole is beneficial because it reduces the variance of the total return.[4] Two important examples of energy projects negatively correlated with the economy—conservation for the long-term and stockpiling for the short-term—are considered in Sections 3.6 and 5.1, respectively.

The beneficial effects of this negative correlation are analogous to those associated with asset pricing in the literature on modern finance. Let r_i and r_m denote the random rates of return on stock i and on a portfolio of all

[4] If the negative correlation is sufficiently large, people would be willing to hold an asset which *loses* money on average. House insurance is such an asset; the expected return to the owner is negative (otherwise insurance companies would be unprofitable) but the payoff occurs precisely when the marginal utility of income is very high.

stocks in the market, respectively. Under the assumptions of the capital asset pricing model (CAPM), the *expected risk-adjusted* rate of return on stock i, $\tilde{r}_i$, in an efficient market satisfies the equation

$$\tilde{r}_i = r + \beta_{im}[\tilde{r}_m - r] \tag{12}$$

where

$$\beta_{im} = \frac{\text{Cov}(r_i, r_m)}{\text{Var}(r_m)} \tag{13}$$

is a measure of risk, r is the risk-free interest rate, and $\tilde{r}_m$ is the expected rate of return on the market portfolio [Sharpe (1964)]. In countries with well-developed capital markets, $\tilde{r}_i$ is the appropriate risk-adjusted discount factor to use in evaluating a project's returns. This approach has the advantage that it is not necessary to estimate risk-aversion coefficients needed in equation (6); the cost of risk is determined by the required return on the market portfolio. In developing countries, it may be easier to estimate risk-aversion parameters and thus calculate certainty equivalents directly through equation (6) than to estimate the return on the market portfolio (even if it exists).

2.4. *Correcting for uncertainty through the discount rate*

Two caveats should be noted in applying CAPM to make risk adjustments in the required rate-of-return or discount rate in calculating EPV. First, the adjustments are made period-by-period. Second, discounting for risk continues only until the uncertainty is resolved.

A difficulty arises in attempting to aggregate the effects of multi-period risk into a single adjustment factor. Consider a typical investment project which becomes profitable after a set-up period in which there are losses. Initially, when the expected return in a period is negative, the marginal valuation of income increases, so the direction of the appropriate adjustment in the discount rate applied to the expected return is to *lower* the rate. Later, when the expected returns are positive, diminishing marginal valuation of income requires a *higher* discount rate. It follows that no single adjustment in the discount rate can be made to reflect risk correctly *at every point in time* if the expected return is positive in some periods and negative in others [Lind (1982)]. Even if there is some risk-adjusted rate $\tilde{r}$ which equates the present value of expected returns ($\bar{Z}_t$) discounted at $\tilde{r}$ to the present value of the certainty equivalent of these returns (z_t) discounted at the risk-free rate r, there will be no guarantee that $\tilde{r} > r$.

Another, quite separate, point concerning the discount rate arises if it is known that the project risk will be resolved in the course of the project's life [Wilson (1982)]. Assume that the values of the uncertain return in each period are, in fact, correctly captured by discounting the expected returns by a risk-adjusted rate, $\tilde{r}$. It is only correct to use the rate $\tilde{r}$ to discount the expected returns *until* the uncertainty has been resolved. A return twenty years in the future that is uncertain today but will become known in five years should only be discounted for five years of risk. This can be seen by adapting our earlier example.

Suppose the project Z at the end of each period yields a return of \$0.5 with probability 1/2 and \$1.5 with probability 1/2. After the first period, the uncertainty becomes resolved one way or the other forever. Thus at the end of the first and all subsequent periods, the returns are either (\$0.5, \$0.5, \$0.5,...) with probability 1/2 or (\$1.5, \$1.5, \$1.5,...) with probability 1/2. For example, the project might be an oil-drilling project whose returns depend on the size of the reserves which become known only at the *end* of period 1.

For this project at the *beginning* of period 1, the expected present value of all future returns discounted at the rate $\tilde{r}$, $\mathrm{EPV}(\tilde{r})$, is given by:

$$\mathrm{EPV}(\tilde{r}) = \sum_{t=1}^{\infty} \frac{1}{(1+\tilde{r})^t} = \frac{1}{\tilde{r}}. \tag{14}$$

On the other hand, the present value of the project at the *end* of period 1 will be either $(0.5)(1+r)/r$ with probability 1/2 or $(1.5)(1+r)/r$ with probability 1/2. Because this return is risky, it has to be discounted back to the beginning of the first period by the rate $\tilde{r}$. Hence, the correct measure of expected present value is:

$$\mathrm{EPV}(\tilde{r}, r) = \frac{1+r}{r(1+\tilde{r})} = \frac{1+r}{r+r\tilde{r}}. \tag{15}$$

When $\tilde{r} > r$, this yields a higher present value because

$$\mathrm{EPV}(\tilde{r}, r) = \frac{1+r}{r+r\tilde{r}} > \frac{1+r}{\tilde{r}+r\tilde{r}} = \frac{1}{\tilde{r}} = \mathrm{EPV}(\tilde{r}).$$

Using a higher discount rate to value uncertain future returns underestimates their value if the uncertainty is resolved or diminished at some intermediate point in the life of the project.

2.5. *Income distribution considerations*

Project risks are not borne equally across the population. Although the earlier analysis in Section 2.3 with the representative individual assumes that the government *can* and *will* spread risks evenly over the entire population, in practice it may turn out that certain groups bear a disproportionate burden of the variations associated with the price of energy. Government policy could in principle correct the distributional consequences of price changes, but this cannot be taken for granted. On balance, and without strong evidence to the contrary, it is safer to assume that government policy will remain unchanged. Project choice must take into account sub-optimal policies of government, whatever the reason for their existence. However, identifying government policies is not always straightforward. For example, when energy import costs rise the government's financial position is affected. If its policy is to maintain a balanced budget then the government may be forced to raise taxes; if its policy is to maintain tax rates then the government may be forced towards a budget deficit. In advance, it is difficult to know whether government policy is to maintain tax rates or a balanced budget.

If government policy fails to spread risk evenly then the use of a representative individual can underestimate the costs associated with energy price fluctuations, and the appraisal of domestic energy projects should take such distributional implications into account.[5]

2.6. *Project appraisal in relation to the status quo*

Cost-benefit analyses often neglect the costs associated with depleting a finite reserve of an exhaustible resource. The price of oil in the ground is called its royalty value. This royalty value should be *subtracted* from the market (i.e. well-head) price of oil in a project appraisal. The reason is that oil extracted now cannot be extracted later: the total supply is fixed. This opportunity cost is a real cost. In particular, instead of extracting the oil it may be possible to sell it in its current form by auctioning the drilling and extraction rights. That is why it is important to emphasize that costs and benefits should be measured relative to the *status quo*. Currently the oil is in the ground. When developing exhaustible resources, the relevant measure of benefit is the *value added*, the difference between the price of oil and the royalty value of deposits.

The royalty value of an oil reserve is based on the market price of oil in the ground. A problem arises in countries whose reserves are not traded. Domestic conditions may prevent many LDCs from trading their coal, gas and oil reserves internationally. In the absence of complete markets, countries can at least begin to place bounds on the value of their reserves by comparing the relative profitability of their reserves with that of reserves traded in the market, such as off Alaska or in the North Sea.[6]

Project appraisals which include in the benefits the value of extracted oil but neglect estimated royalty costs can easily have *internal rates of return* of several hundred percent. Using this method of appraisal, however, can lead alternative projects to have still higher rates of return. For example, a country that sells its oil reserves in the ground to a foreign oil company receives a large positive payment and apparently incurs no costs. This leads to a practically infinite internal rate of return. This example helps illustrate

[5] The analyses in this paper consider the sum of a project's effects across a population and are thus cast in the framework of utilitarianism. Other ethical considerations may affect the results of cost-benefit analysis [Schulze and Kneese (1981)]. A Rawlsian criterion would focus on the utility of the worst-off individual. Interpersonal comparisons of utility can only be avoided if the project improves everyone's expected welfare (Pareto criterion).

[6] The value of the reserves depends on extraction costs, transportation costs to points of sale, and degree of uncertainty about reserve levels.

both the weakness of using the internal rate of return criterion and the mistakes that may be made if the royalty cost of depleting the resource is left out of the calculation.

It is also important to consider the timing of energy projects. Since oil reserves are exhaustible, complete extraction today precludes extraction tomorrow; that is, alternative programs for oil extraction are *mutually exclusive*. Even with renewable energy projects, timing decisions require comparisons of mutually exclusive options; the decision to proceed with a synthetic fuels program today precludes postponing the project for one year. Thus, even when the expected discounted value of one project option is positive, it may not be the optimal alternative. Project appraisal should choose the program with the *highest* (positive) *expected present value* among all of the mutually exclusive options.

3. Hotelling's rule and the price path of oil

Oil is an exhaustible resource; a barrel extracted today results in one barrel less for extraction tomorrow. When the timing of oil extraction and sales is chosen to maximize discounted profits, this leads to an expected price path due to Hotelling (1931). The technique for calculating an optimal extraction rule is first demonstrated in a general model, and then specialized to cases examining perfect competition, extraction costs, and imperfect competition.

Consider an entrepreneur who owns S barrels of oil reserves. If he sells at a rate of q_t barrels at time t, he expects to receive revenue at the rate of $R(q_t)$; in the case of perfect competition, $R(q_t)$ will be equal to expected price p_t times quantity q_t. His costs of extraction, $C(q_t, \dot{q}_t, t)$, depend on the time t at which the extraction occurs, and on both the extraction rate (q_t) and the change in the extraction rate ($\dot{q}_t$); extracting a larger quantity is costlier as is the attempt to *speed up* the extraction rate.[7] The entrepreneur is constrained to sell no more than his total reserves, S.

Formally, the problem can be expressed as[8]:

$$\underset{q_t}{\text{Max}}\left[\int_0^{\infty}[R(q_t) - C(q_t, \dot{q}_t, t)]\mathrm{e}^{-r_t t}\,\mathrm{d}t\right], \tag{16}$$

subject to

$$\int_0^{\infty} q_t\,\mathrm{d}t \leqslant S, \tag{17}$$

[7] The problem can be further complicated by permitting extraction costs to depend on the size of remaining reserves and by permitting storage (instead of sales) after extraction.

[8] We use r_t to represent the continuous time analogue of the average discount rate between 0 and t, i.e.,

$$r_t = \left[\int_0^t r(s)\,\mathrm{d}s\right]\Big/ t$$

where $r(s)$ is the instantaneous discount rate at time s.

and

$$q_t \geq 0. \tag{18}$$

Letting subscripts denote partial derivatives with respect to the relevant argument, the Euler condition for an optimal path can be written as

$$[R_1(q_t) - C_1(q_t, \dot{q}_t, t)]\mathrm{e}^{-r_t t} - \lambda \leq -\mathrm{d}[C_2(q_t, \dot{q}_t, t)\mathrm{e}^{-r_t t}]/\mathrm{d}t \tag{19}$$

with strict equality whenever the extraction rate is positive, and where λ is chosen to ensure that the solution satisfies the exhaustion constraint (17).

3.1. *Perfect competition, costless extraction, and a constant discount rate*

The simplest and most idealized special case involves entrepreneurs acting under perfect competition with costless extraction and a constant discount rate, r. The optimality condition (19) reduces to

$$q_t > 0 \text{ and } p_t\mathrm{e}^{-rt} = \lambda; \quad or \quad q_t = 0 \text{ and } p_t\mathrm{e}^{-rt} \leq \lambda \tag{20}$$

where p_t is the market price. Thus oil is extracted and sold only in the period(s) in which its present discounted value is highest. The Lagrange multiplier is the maximum discounted price, $\lambda = \max_t [p_t\mathrm{e}^{-rt}]$.

Unless the present discounted price for every period is the same, the market will not be in equilibrium. Thus if the price of oil is expected to rise more quickly than the nominal interest rate, then all suppliers will hold on to their oil anticipating higher profits. This raises the price of oil today relative to its price in the future and restores equilibrium. If the price of oil is expected to rise slower than the interest rate, suppliers will try to sell more oil today and invest their profits to receive the nominal interest rate. This flooding of the market leads to a lowering of the price today relative to the future and restores the equilibrium path.

Oil prices cannot continue to rise at the nominal interest rate if there is a backstop technology which can provide an unlimited supply of energy at a constant price. In this case, the oil price path must also be at the right level to ensure that the stock of depletable resources will be exhausted just at the time when the backstop technology is *expected* to become competitive. If entrepreneurs thought that an alternative supply would arrive before their stocks are exhausted, they would increase their oil supply now in order to prevent being stuck with excess oil in the future. This brings the market back into equilibrium as it results in a lower price and higher demand in each period.

Because of the boundary condition that leads to exhaustion, the price of oil today is sensitive to changes in long-run interest rates. A higher interest rate results in a price path of oil where prices rise faster over time. If today's price remains constant then all future consumption will be reduced. To ensure exhaustion, the equilibrium price today must fall to compensate for its swifter rise. Higher interest rates shift consumption from the future to the present. Similarly, because of the boundary condition, if new reserves are found or an increase in conservation takes place, the whole equilibrium price path must be lowered to stimulate the additional demand needed to ensure exhaustion; thus, today's price must fall.

Note that if different suppliers have differing discount rates then those with above average discount rates (typically OIDCs) will be the first to extract their reserves. The price will rise at the discount rate of the marginal supplier.

The recent dramatic rises and falls in the price of oil have caused many to doubt the validity of Hotelling's implied steady rise in the price of oil. But this conclusion depends on there being perfect competition, costless extraction and a constant discount rate. Even under these very stylized assumptions, Hotelling's rule only requires that the price of oil is *expected* to rise at the interest rate. The price path of oil will be changed by variations in the interest rate, by new discoveries that affect expected world oil reserves, by new technologies for the use of oil substitutes (coal, synthetic fuels, etc.) or for interfuel substitution, and by changes in consumption patterns (due for example to faster than anticipated conservation efforts). The volatility of nominal interest rates and fluctuations in the world's proven reserves are reflected in a high variance of oil prices around their expected (or average) growth path.

Several other factors can also interfere with the stylized Hotelling solution:

(1) Extraction is not costless. As illustrated in expression (16), costs depend both on the rate of extraction and on *changes* in the rate of extraction;

(2) Oil for the future is held by a monopolist who in the interest of profit maximization chooses not to sell today;

(3) Oil for the future is held by a cartel with conflicting objectives different from *joint* profit maximization;[9]

(4) Oil for the future is held by both a cartel and a competitive fringe, and either (i) the competitive fringe chooses to save rather than sell its oil for strategic reasons [Nichols and Zeckhauser (1977), Crawford and Sobel (1982)]; or (ii) the competitive fringe tries to sell oil as fast as it can but is constrained by a physical limit to the rate of extraction; or (iii) the relative market shares of the cartel and competitive fringe change.

The first two issues above are discussed in the subsections that follow. In the short run, the market may be out of equilibrium because most oil is located far away from where it is consumed; transportation is slow and expensive; there are lags in production; demand is volatile; and expectations

[9] For example, current revenue requirements or *individual* profit maximization may lead some members of the cartel to overproduce, thus disrupting production quota agreements and lowering price.

may turn out to be wrong. When prices are thought to be too high, competitive producers who try rapidly to increase their extraction rate often must pay much higher marginal costs and may even sacrifice some of their recoverable reserves (due to technological inefficiencies involved in excessively rapid extraction).

3.2. *Extraction costs*

Even in a perfectly competitive market with a constant discount rate, extraction costs can fundamentally alter the expected price path of oil. With constant unit extraction costs, $C(q_t, \dot{q}_t, t) = cq_t$, the Euler conditions lead to an equilibrium in which

$$(p_t - c)\mathrm{e}^{-rt} = \lambda. \tag{21}$$

Price net of unit extraction costs rises at the discount rate; this implies that price rises *slower* than the discount rate. When unit extraction costs are constant with respect to output but fall over time with technical progress, the corresponding equilibrium is

$$(p_t - c_t)\mathrm{e}^{-rt} = \lambda.$$

Here, falling marginal costs c_t further slows the rate of price appreciation. Only if marginal costs actually rise faster than the discount rate will price appreciation be faster than the discount rate. Significantly more complicated is the solution when extraction costs depend on changes in the rate of extraction.[10]

Extraction costs may also vary across different suppliers. For any price path, those with smaller unit extraction cost will choose to extract first. Marginal revenue net of marginal cost appreciates faster for suppliers with high marginal costs; high-cost suppliers have a relatively greater benefit from postponing their extraction costs into the future. The shift over time from low to high cost extractors (or from low to high cost techniques) will be partially offset through improvements from technological innovations.

3.3. *Imperfect competition*

Compared with the competitive solution, the price path of exhaustible resources may be different when the oil reserves are held by a cartel or monopolist. A monopolist will choose an intertemporal distribution of

[10] If speeding up extraction is expensive, it becomes necessary to evaluate the term

$$d[C_2(q_t, \dot{q}_t, t)]/\mathrm{d}t \quad \text{from equation (19).}$$

The solution remains tractable in the special case where the cost function takes the linear form:

$$C(q_t, \dot{q}_t, t) = \gamma_0 + \gamma_1 q_t + \gamma_2 \dot{q}_t.$$

Then, along an equilibrium path,

$$p_t = \gamma_1 + r\gamma_2 + \lambda \mathrm{e}^{rt}.$$

supply to maximize total discounted profit; this requires the present discounted value of marginal revenue minus marginal cost to be the same for any period. Unlike the competitive solution, the monopolist may choose *not* to exhaust the entire stock of resources by the time the backstop technology is expected to become competitive. Greater profits may result if the monopolist initially follows a higher price path and later when the competitive backstop technology constrains his ability to raise prices further he exhausts all remaining stock at the constant backstop price [Stiglitz and Dasgupta (1980)].

There are two potential differences to the solution of Hotelling's rule under a monopoly: (i) the price path may be different to the extent that marginal revenue is different from price; (ii) the whole price path may be higher if the monopolist finds it optimal not to exhaust supply by the stage at which the backstop technology becomes competitive. Under monopoly the rate of oil price rises can be either faster or slower than under competition. If extraction is costless and the price elasticity of demand is *constant,* then the price of oil rises at the monopolist's discount rate [Stiglitz (1976)]. Marginal revenue is given by $p(1-1/\varepsilon)$, where p is price and ε is the elasticity of demand. With costless extraction, marginal revenue and therefore price (which is a fixed multiple of marginal revenue) rises at the monopolist's discount rate. In this case, competition and monopoly imply the same rate of price increase if the discount rates in the two situations are the same.

The expected price path of oil may differ from Hotelling's rule to the extent that oil reserves are controlled by a monopoly and the price elasticity of demand changes over time. For example, an increase in the price elasticity brought about by gradually increased substitution possibilities will imply that prices rise slower than the monopolist's discount rate.

3.4. *The discount rate in developing countries*

The appropriate social discount rate for developing countries may be higher than the interest rates which prevail in the rest of the world. The suboptimality of domestic savings [Little and Mirrlees (1968), (1974)] and restrictions on foreign investment suggest that rates of return to domestic projects in developing countries may be higher than those obtainable elsewhere. Thus, these countries attempt to borrow funds from international development agencies and capital markets until they run up against lending constraints and large risk premia. In contrast, the discount rate for oil-exporting, capital-rich countries is likely to be lower than in the rest of the world.

There are several implications of an OIDC's discount rate being higher than that of oil-exporting countries. First, given any path for the world price of oil, it will pay an OIDC to deplete its own reserves at a faster rate than the oil-exporting countries if it has the same (or lower) extraction costs. To the extent that its marginal extraction costs rise with both the amount and the speed of extraction, this depletion policy will be moderated. Such a policy does have the side effect of increasing the OIDC's future dependence on oil imports.

A higher discount rate also raises the costs of stockpiling. Whatever the capital gain (or loss) from changes in the world price of oil over time, the annual opportunity cost of the capital tied up in the oil stockpile is higher. Other opportunity costs of storing oil will also be higher because the capital tied up (in extraction and storage facilities—see Section 5.1) has to be valued at the higher interest rate.

3.5. *The cost arising from uncertain reserves*

Unless future supplies of oil are known with certainty, it is impossible to allocate oil consumption efficiently over time [Gilbert (1976)]. If an unexpected bonus of oil is found, then too little consumption will have occurred; if a well runs dry unexpectedly soon, then too much consumption will have occurred. A cost to uncertainty arises because we assume the loss from a shortage is greater than the benefit from an equivalent surplus.[11] If it is increasingly difficult to find replacements for oil as shortages become more severe then prices must rise more steeply in the event of a shortage than they fall when there is a surplus. Resolving the uncertainty over oil reserves sooner will on average increase today's consumption and lower the price.[12]

Consider a two-period model in which the total quantity of oil is either 0.5 or 1.5 (each with chance 1/2). Extraction is costless. At the time when first-period consumption is decided, total supply is uncertain. Before the second period, the remaining supply becomes known. In a competitive market, second-period prices will lead to the resource just being exhausted. To capture the increasing difficulty of finding substitutes for oil, the demand function must be assumed convex in price. For example, let demand as a function of price be

$$q(p) = 1/p. \tag{22}$$

If consumption in the first period is q_1, then the first-period price is

$$p_1 = 1/q_1. \tag{23}$$

Second-period consumption q_2 will either be $(0.5 - q_1)$ or $(1.5 - q_1)$. Second-period prices $p_2(q_2)$ may be written as a function of the total stock Z

[11] Devarajan and Fisher (1982) highlight a benefit from uncertain reserves. When there are economies of scale, expected extraction costs of a random reserve are lower than for a certain reserve equal to the expected size of the random reserve.

[12] Resolving the uncertainty will not always result in oil stocks being equal to the expected reserves; oil stocks may be higher or lower than the average of all possibilities. The benefit of resolving the uncertainty is that we know the actual stocks sooner and can thus plan better.

and first-period consumption q_1 as follows:

$$p_2(q_2) = 1/q_2 = 1/(Z - q_1), \qquad \text{since } Z = q_1 + q_2. \tag{24}$$

Hotelling's rule can be used to determine the equilibrium. The *expected* discounted second-period price must equal the first-period price. Thus, if r is the risk-free interest rate, this implies

$$p_1 = [(1/2)p_2(1.5 - q_1) + (1/2)p_2(0.5 - q_1)]/(1 + r). \tag{25}$$

Hence first-period consumption q_1 is given as the solution to the equation:

$$2/q_1 = [1/(1.5 - q_1) + 1/(0.5 - q_1)]/(1 + r). \tag{26}$$

In the two-period model, it is reasonable to think of the periods as each being ten years. Then if $r = 100\%$, this gives

$$q_1 = 0.39, \qquad p_1 = 2.5;$$
$$p_2(0.5 - q_1) = 9.1, \qquad p_2(1.5 - q_1) = 0.9; \qquad \text{hence } E[p_2] = 5.0.$$

If the stock Z were known with certainty at the *beginning* of the first period, the outcome would be either

$$\text{for } Z = 0.5\text{:} \quad p_1 = 3, \quad p_2 = 6; \quad q_1 = 0.33, \quad q_2 = 0.17$$

or

$$\text{for } Z = 1.5\text{:} \quad p_1 = 1, \quad p_2 = 2; \quad q_1 = 1.0, \quad q_2 = 0.5.$$

The average price and consumption levels are:

$$E[p_1] = 2, \; E[p_2] = 4;$$

and

$$E[q_1] = 0.67, \; E[q_2] = 0.33.$$

In this example, the cost of uncertainty in the first period results in a 25% higher price compared with the average of the paths when the reserves are known.

The inefficiency in intertemporal allocation arises because of the need to maintain flexibility in facing the uncertain second-period supply; this inefficiency is relative to the optimal depletion path in a world with no uncertainty. Caution in the first period reduces the costs of adjusting to a worse than expected outcome in the second period. When the elasticity of demand is lower, the cost of adjusting to a shortage is greater and more flexibility is needed; first-period consumption must be reduced even further which leads to greater inefficiency and higher prices.

This subsection illustrates an important externality associated with oil exploration projects. The evaluation of such projects should include the benefit of reducing the misallocation arising from uncertainty about the size of global reserves.

3.6. *The effects of new energy supplies*

In evaluating the benefits from the technologies of producing energy, it is important to make a distinction between the production of energy from exhaustible and from renewable resources. The development of known exhaustible resources in a world of certainty should not affect the world price because everyone knows that they will be extracted sooner or later.[13] Prices are lowered with the development of renewable energy resources (e.g., solar, hydro and nuclear) since this increases the world's total supply of energy.[14] This may also be accomplished through energy conservation. By using the available oil or coal more efficiently, the same outputs can be produced with less oil or coal [for estimates of this potential, see World Bank (1982)]. When conservation results in a 25% greater efficiency this is equivalent to an effective increase of 25% in the entire stock of reserves. This benefit must be weighed against the costs of conservation.

Reducing energy import needs either through developing renewable resources or conservation decreases dependency and hence vulnerability to oil-price shocks. Less dependency creates greater competition in world oil markets. This should benefit all oil-importing countries.[15]

A related issue is: when should importing countries extract their limited stocks of exhaustible resources [see Crawford and Sobel (1982) and Gilbert *et al.* (1978)]? The cartelization of the current market is moderated by the fact that many countries at present have sizeable reserves, i.e. there is a large *potential* competitive fringe. Without these reserves, there would be a much greater potential for cartelization in the future. But, this issue arises only if developing countries increase their production of *depletable* resources. Conserving energy and developing renewable energy resources have no such negative externality: the world's energy supply is increased.

4. Coordinating strategies

The advantage to coordinating strategies for information gathering and developing alternative energy sources has resulted in oil-importing countries forming the International Energy Agency (IEA).[16] This is only a first step in the right direction. Countries (or multinational oil companies) acting alone

[13] Other oil-exporting nations can cut back current production knowing that their supply will be more in demand later when other countries' exhaustible resources run out. This assumes that they can still meet their current revenue requirements by borrowing against their now larger future incomes; some of the effects of borrowing constraints are discussed below.

[14] Similarly, a new technology that improves the cost-effectiveness of alternative exhaustible resources (e.g., gas or coal) creates additional supplies of energy.

[15] The development of both renewable and non-renewable energy resources shifts revenues away from oil-exporters, but in the case of non-renewable resources these revenues are simply postponed until the future.

[16] Participants in the International Energy Plan hope to minimize the costs of a supply interruption by formalizing sharing rules in advance. However, Hogan (1981) argues that the sharing rules are poorly designed and may be expected to break down in the event of a major disruption.

collect too little information and wastefully duplicate or excessively diversify their research strategies. This section discusses externalities associated with exploration and research diversification.

4.1. *Information externalities in exploration*

The first step in developing energy resources is the gathering of information. There are important reasons why this should be done by a central authority rather than the free market (if it exists). Information is perhaps one of the only true public goods. It can be shared by an unlimited number of people and across countries. Thus everyone has an incentive to wait for others to gather the information rather than to pay the costs and duplicate efforts. There is an inefficiency when more than one group pays the costs to find out the same information. While there can be private advantages when information is withheld, the sum of the benefits is generally not reduced when the information is shared. Duplicating efforts to gather information frequently occurs in zero-sum games such as futures markets, where each trader tries to take advantage of another. Competitive information gathering affects the distribution of income but not the size of the pie.

At present, oil-importing developing countries all pay a risk premium because the world's total supply of oil (and its substitutes) is not known precisely. As demonstrated in Section 3.5, uncertainty leads to an inefficient intertemporal allocation of resources. No one country can eliminate this uncertainty by investing in exploratory drilling. Yet each country that investigates its level of reserves contributes to a reduction of the total uncertainty. Although developing countries might not be willing to do exploration on their own, the IEA could promote a quid-pro-quo policy where oil-importers agree to find out more about their own reserves provided others find out about theirs.

For oil-exporting countries, information gathering provides both advantages and disadvantages. More information about the level of a country's reserves decreases the risk of unanticipated future price variability for its own oil [Gilbert *et al.* (1978)]. However, each individual oil-exporting country will prefer greater uncertainty about others' reserves since its oil will then command a higher risk premium.

4.2. *Diversification and multi-project analysis*

The discussion has focused on oil but the issue is clearly one of energy generally. A country considering research and development strategies should examine the various approaches to producing both oil and its substitutes. For example, hydroelectric power and shale oil may prove equal to oil in their importance for many South American countries. How much should research be diversified, how should the eggs be spread among the baskets?

When there are several projects, all of which have goals that are substitutes for one another, it is important that they be evaluated simultaneously. Independent project evaluation can lead to excessive diversification; too many projects could be accepted.

Uncertainty usually motivates diversification. Even though research and exploration are inherently risky propositions, diversification may not be advantageous. There are two important special features of research projects that motivate specialization even in the presence of uncertainty. First, there may be economies of scale in producing information. Second, there may be diseconomies of scope; when there are several alternative solutions, only the best one is implemented.

A little information costs more than it is worth [Radner and Stiglitz (1975), Dasgupta (1982)]. Gathering information has a direct economic cost. The benefits can only be calculated indirectly. Actions are based on the information available. The information's value is then determined by the expected return from the ensuing action. When actions are chosen optimally (conditional on the information), a slight perturbation will not change the expected return. The value of incremental information arises from the distribution of actions becoming more appropriate given the outcomes. At zero information, only a single action is taken and hence there are no gains from redistributing actions more appropriately. Starting from ignorance, there is initially a cost to obtaining information but no benefit. Thus, research projects initially have increasing returns to scale.[17]

An example may help illustrate this point. Imagine that the information concerns the probability of finding oil in a given geological structure. In the extreme case of perfect information the probability of finding oil is either 0 or 1; drilling when undertaken will always be successful. At the other extreme of zero information, there is still a prior probability of finding oil. If this probability is above some critical value, then drilling for oil is expected to be profitable. Gathering only a very small amount of information can only make a small change from the prior probability of success. A small amount of information cannot be sufficiently favorable to induce drilling when, based on prior beliefs, drilling was inadvisable. Since the distribution of actions is the same, there is no expected gain from only a marginal amount of information. There is an analogy between fixed costs in production and the ineffectiveness of a small amount of information: the fixed cost associated with gathering information is the cost of obtaining the minimal amount of information needed to have a potential effect on the decision.

Diseconomies of scope may also provide an argument for specialization. If society is only interested in the single best technique for producing energy, and its research projects do not suffer from decreasing returns to scale, then as shown below it is indeed appropriate to put all the money into

[17] This argument relates to projects starting from scratch: small additions to ongoing projects may well be worthwhile.

one project [Nalebuff and Varian (1983)]. While this result is obvious if there is no uncertainty about the returns from the project,[18] its interest lies in the fact that it is also true when the fruits of research are uncertain. One goes with the expected winner. This result does not depend on the correlation between projects (positive, negative, or zero).

To formalize this proposition, consider a government allocating funds to various projects. Each project produces output $f_i(X_i, e_i)$ where X_i are the funds allocated to project i and e_i is a random variable determining the project's success. The government has a budget constraint that $\sum X_i = B$. The random variables may have any general joint probability density function $g(e_1,\dots, e_n)$. Society only cares about the winner. The value of outputs $(f_1,\dots,f_n)$ is $\text{Max}\,(f_1,\dots,f_n)$. Conditional on any realization of $\boldsymbol{e} = (e_1,\dots, e_n)$, the value of each research project, $f_i(X_i, e_i)$, is assumed to be convex; as argued above, there are increasing returns to scale in gathering information. The expected return from following research strategy $\boldsymbol{X} = (X_1,\dots, X_n)$ is

$$E[\text{Max}\,[f_i(X_i, e_i)]] = \int\dots\int \text{Max}\,[f_1(X_1, e_1),\dots, f_n(X_n, e_n)]g(e_1,\dots, e_n)\, de_1 \dots de_n.$$

For each realization of $\boldsymbol{e}$, the valuation function $\text{Max}\,[f_i(X_i, e_i)]$ is the composite of two increasing and convex functions and is thus convex. The expectation of $\text{Max}\,[f_i(X_i, e_i)]$ is just a weighted average of the valuation function conditional on e_i; a weighted average of convex functions is still a (weakly) convex function. Because the maximum value of a convex function occurs at a boundary, it is optimal to fund only one project: for some i, $X_i^* = B$ and $X_j = 0$ for $j \neq i$.

A government does not care whether there is a *second* firm to discover a new technology; only the first discoverer counts. Specialization may hasten the time and improve the quality of the discovery. Diversification is only justified if the leading project encounters diminishing returns to scale or if there is diminishing marginal valuation with respect to increases in the value of the best outcome.

Initially it may seem counter-intuitive that specialization is optimal even in an uncertain world. On reflection it becomes clear: diversification is implied by the convexity of the constraints or the concavity of the objective function. If these conditions are not met, diversification is not necessarily optimal. The standard assumptions about concavity are likely to be violated when considering the returns from several simultaneous research projects.

4.3. *Risk pooling and multi-project analysis*

Outside of research and development, the total value of several projects is generally equal to the *sum* of the outputs. Then, one justification for accepting several simultaneous projects is the advantage from risk-pooling.

[18] With no uncertainty, the best project is known and only *it* should be funded.

As discussed earlier in the context of the capital asset pricing model (CAPM), when there are several projects which are negatively correlated there is a reduction in risk because the sum of the outputs becomes less risky. Often this intuition is mistakenly extended to include projects which have a less than perfect positive correlation, and the claim is made that accepting several *independent* projects also results in lower risk.

As the number of independent and identically distributed assets increases, the *average* return converges in probability to a constant. Thus, adding an additional asset *and* an additional shareholder (to an existing portfolio of n assets held by n people) results in lower risk to all; the additional asset increases the variance by a factor of $(n+1)/n$ while the additional shareholder lowers the portfolio's variance by a factor of $[n/(n+1)]^2$. But a government may not be able to increase its number of shareholders beyond its existing population. Investments will already be spread as thinly as possible. Given that the maximum risk-spreading has already taken place, a government cares about the *sum* of the returns rather than the *mean* return. Accepting additional independently distributed projects linearly increases both the total expected return *and* the total variance. This may or may not be beneficial.

The fallacy of large numbers [Samuelson (1963)] is designed to contrast with the law of large numbers; it shows that investing in a large number of independent and identically distributed assets may be undesirable. Consider an asset with a positive expected return but with sufficient variance so that an investor would choose not to hold this asset at *any* income level. The "fallacy" shows that holding any number of independent and identically distributed *replicas* of this asset is also undesirable. Thus, if the government of Monaco is sufficiently risk averse that it never wants to be the house for a single bet at Monte Carlo, then it will also refuse to play the house over a period of a day, year, or century. The proof follows by induction. Since by assumption accepting a single bet is never desirable, expected welfare is improved by eliminating the final replication. Continuing to remove each remaining final replication implies that expected welfare is maximized when none of the bets is left.

5. The costs of importing oil

There are two interrelated costs facing an economy which is heavily dependent on imported oil: the danger from disruptions in supply and a higher variance in the economy's performance due to fluctuations in the price of oil. Our discussion starts at supply disruptions.[19]

[19] Developing countries have at least two important reasons why they should be concerned about their supply of oil being cut off: embargos and wars. The history of 1973 proves the real possibility of an embargo. Wars have disrupted the flow of oil either by cutting off the producers (as in Iran–Iraq) or by blockades around the consumer (as in Malawi during the Tanzania–Uganda war and in Argentina during the Falklands crisis). The costs of a disruption include threats to national security (defence), massive inefficiencies from disturbances in production, and hardships suffered by individuals unable to obtain or to afford fuel for cooking and heating [Deese and Nye (1981), Plummer (1982)].

The responses to this danger include: (1) reducing domestic consumption through tariffs and quotas and through increased conservation; (2) buying on the spot market; (3) investing in backstop technologies; (4) strategic stockpiling of crude oil; and (5) carrying excess domestic production capacity. We consider each of these options.

The market price fails to reflect all the externalities associated with importing an additional barrel of oil. To reduce demand to the appropriate level, it may be optimal for a government to combine the use of tariffs and quotas. If demand is restrained solely by the use of quotas, then the oil-exporting country would raise its price until market demand is at the quota level. The importing country could capture this price increase by imposing a tariff on oil imports [Hogan (1982)]. Domestic consumers would be better off when part of the price of oil is paid to their own government rather than to sellers of oil abroad.

In the presence of uncertainty, there is an advantage to supplementing tariffs with quotas. A quota helps to limit a country's foreign exchange commitment in the importation of energy. In countries where oil imports form a very large proportion of export earnings, small fluctuations in the volume of energy imports can cause large changes in the balance-of-payments. With their capacities to borrow already stretched (and their foreign reserves exhausted), many OIDCs face the risk of sharp rises in the exchange rate. In such situations the costs to the non-traded sectors of the economy can be excessive.

Ideally, countries can reduce their risk by investing in assets that are negatively correlated with the supply, or positively correlated with the price, of oil (e.g., purchasing shares of oil companies). Subsidizing conservation has a high payoff in the event of a supply disruption. Similarly, the development of a contingent rationing scheme creates another negatively correlated asset; it is valuable only in the event of a supply disruption. Ration coupons which are tradable help protect the poor against oil-price shocks while maintaining the incentive for efficient allocation.

To a greater or lesser extent, all countries respond to threatened disruptions by buying on the spot market. There are two disadvantages to this approach: (i) the spot market has accounted for a relatively small proportion of the total oil traded, so that individual countries may face high prices if they all move to the spot market together to meet disruptions; (ii) oil purchased in spot markets may be at great distance from the home country and transport costs and time lags could entail considerable losses to the economy.

Reliance on renewable resources (such as alcohol from biomass) is not economic at present world oil prices but may become so in the future with a long-term increase in real prices. Because oil is an exhaustible resource whereas most synthetic fuels are not, we would expect the real price of oil to rise in the long-run but that of most synfuels to remain more or less constant. Thus, in due course, synfuels should become competitive in relation to oil supplies that are exhaustible. Some development of renewable resources at higher cost today may in any case be warranted as a way of reducing the uncertainty in supplies faced by oil-importing developing countries. The extent of such development would depend on the premium attached to security of supplies; this varies according to a country's circumstances and the technological possibilities open to it.

A further protection against interruptions of oil supply is for countries to keep a strategic stockpile of oil reserves.[20] The argument for public stockpiling must depend on an implicit belief that the private market does not have sufficient incentives to maintain an adequate stockpile [Wright and Williams (1982)]. There are several potentially large externalities that are not captured by the private incentives to stockpile. In the event of a shortage, private sellers expect (with good reason) that the government will impose a price ceiling. This cuts off part of the favorable tail of benefits. The expected value to society of the stockpile exceeds its market price when the government limits prices. A second factor mitigating competitive stockpiling arises when the size of the optimal stockpile is large enough to affect prices. Purchasing a large stockpile may force oil prices to rise if it restricts the supply available to consumers. The higher oil prices resulting from stockpiling creates a comparative advantage for the oil producers to hold the stockpiles. Having a stockpile may also provide strategic advantages not captured by the market. Hogan (1982) stresses the advantage of a stockpile in reducing the probability of a disruption; as the impact of an embargo is diminished, it is less likely to be deployed.

Countries that extract and refine some of their own oil have the option of keeping a safety reserve in the form of untapped wells that can be brought into immediate production (i.e. carrying excess capacity). The economic theory of exhaustible resources shows that this is essentially another form of stockpiling.

5.1. *The costs of stockpiling*

The true costs of stockpiling an *exhaustible* resource are likely to be lower than those associated with carrying inventories of ordinary (renewable) goods. The reason is that exhaustible resources such as oil are expected to reap capital gains from price rises along a Hotelling equilibrium path. Although there is a notional interest loss on the capital tied up in stored oil, this is offset by the extent to which the real price of oil rises over time. With perfect competition and zero extraction costs, the price of oil would be expected to rise at exactly the rate of interest. In this case there is no opportunity cost to the capital invested in the stockpile—it is simply another form of holding savings. Authors such as Samouilidis and Berahas [1982, p.

[20] There are several ways in which these reserves can in principle be held. They can be kept above ground in storage tanks, or underground—in leached salt caverns or abandoned salt, hematite, limestone, granite, chalk, or coal mines, or indeed in existing or disused oil-wells. Some new suggested systems for crude-oil storage include rubber bags and artificially created "lagoons", sometimes lined with impermeable synthetic plastics of polyethylene products. Idle tankers are also obvious potential storage facilities.

569] underestimate the returns to stockpiling because they attribute an interest cost to the stockpile but do *not* take into account the expected gains in its value.

The opportunity cost of capital tied up in stockpiling must be evaluated at the country's *own* internal discount rate. Yet the price path for oil is largely determined by the *lower* discount rates prevailing in capital-rich, oil-exporting countries (Section 3.4). Together with the effect of positive extraction costs in moderating the rate of price increase, this suggests that capital gains in the value of the stockpile are unlikely to offset the full opportunity cost of capital tied up in the stockpile.

There are additional costs of stockpiling. Building storage facilities for oil can be expensive[21] and there is generally some loss from evaporation (of the order of 1% to 2% per year). When extraction costs are incurred for oil extracted from own reserves, these could have been delayed until the oil was actually needed. The money spent on early extraction leads to an opportunity cost that should be evaluated at the country's discount rate. Extraction costs may also be falling over time with improvements in technology. The lost savings in extraction costs should be attributed to the costs of maintaining the stockpile. A developing country may not have adequate refining capabilities. Hence, it may need to stockpile refined instead of crude oil. The additional expense of purchasing refined oil could also have been postponed until the oil was just ready for use. Although refining costs do grow to keep up with inflation, real interest is lost on the capital spent on refining.

5.2. *Inventories and flexibility in production*

Even a country which produces some of its own energy has reasons to carry stockpiles. These inventories can be used to counter the disruptive effects of the second major problem associated with importing oil, viz. its price variability.

Inventories can be used to accommodate fluctuations in supply while reducing price variability. If inventories are not used (or run out), the price mechanism will act as a brake; during peak periods of demand, price will rise to choke off some of the excess demand. The optimal stabilization results in neither a constant price nor a band-width rule; it is the solution demonstrated by Gustafson (1958) to a dynamic optimization problem that turns out to be similar to Hotelling's rule [Newbery and Stiglitz (1982)]. The costs of oil supply shocks can be reduced but not eliminated through an optimal management of strategic oil reserves.

Note that price variation is in itself not always harmful. Random variation in input costs is *beneficial* because cost functions are concave in input prices.

[21] Costs presently range from $3 per barrel in existing salt mines to $16 per barrel in above-ground tanks.

Firms can take extra advantage of oil when it is a bargain and substitute away from it when it is expensive. The use of oil inventories accentuates this advantage if stockpiles can be built up when the purchase price is low.

Inventories illustrate one of the advantages of flexibility. This argument has been extended to claim that the choice of energy production techniques should be relatively more capital-intensive so as to have greater flexibility in meeting variations in demand. By paying a higher fixed cost, firms hope to exploit the opportunity of relatively efficient production over a range of output levels. *Yet, it may be preferable to use inventories rather than excess capacity to meet these types of unexpected fluctuations.* Making relatively large fixed cost investments takes away an even more important aspect of flexibility, the option of shutting down [see Mason and Merton (1984)]. If the price of imported oil becomes very low (for example, because of a technological break-through in offshore development) then one of the production options is to shut down and import oil.

The fact often overlooked is that for projects with a high marginal (i.e. variable) cost, there is a greater probability of the economy being able to take advantage of cheaper alternatives. The choice of a technology usually involves making trade-offs between high fixed costs and high marginal costs. It is optimal to choose projects with a relatively small fixed-cost component (and correspondingly higher marginal costs) when taking into account the often desirable option of temporarily shutting down and relying on imports if they become cheap.[22]

5.3. *Financing energy development in* OIDCs

A primary concern of many oil-importing developing countries is the financing of their imports. Countries that are constrained in their ability to borrow may also find it difficult to undertake the large investments needed to develop their own domestic energy sources.[23]

A loan to pay for current consumption of energy is an especially risky proposition. Moreover, the risk premium with loans may be expected to rise as the total outstanding debt grows. Countries such as Brazil are already beginning to find themselves in situations where borrowing money to buy oil has become prohibitively expensive. But at the same time that it is negotiating crippling loans, Brazil also has a large savings account in the

[22] For example, in the production of ethanol from corn, the marginal input costs (corn @ $2.50/bushel) form over 85% of the total production costs of a project with an output of 50 million gallons per year [Manassah (1981), Part A, Table I, p. 335]. Similarly, in the production of ethanol from sugarcane in South-east Brazil, biomass inputs made up almost 65% of the production costs while the levelized investment cost was less than 20% of the total expenditures [Gray (1981), p. 298]. These types of projects have the flexibility to be expanded or contracted depending on whether the world oil price is high or low.

[23] The cost of importing energy has formed a significant proportion of many developing countries' exports (in 1980 it was more than 50% for India, Brazil, etc.). The true cost of these imports is particularly high when account is taken of their overvalued exchange rates. Large expenditures on imported oil add pressure to the exchange rate and increase the risk premia on loans to such countries (e.g. Brazil, Turkey, and Jamaica).

form of shale oil deposits. Borrowing is one form of dissaving. Brazil has started "borrowing" from its future self by dissaving (i.e., depleting) its stocks of non-renewable resources. Although future generations are left with less savings in the ground, they will also have less external debt to carry. When servicing the debt costs more than the return on reserves, future generations are better off in net terms. A second advantage of this type of "borrowing" is that a risk premium is avoided; a corresponding disadvantage is that the country cannot default (on itself). Investors and banks realize that the risks associated with developing a country's supply of exhaustible resources are far smaller than financing a loan to pay for current energy consumption. Thus, developing countries often contract with multinational oil companies both for their access to capital and for their technical expertise. There are advantages and disadvantages to this approach [see Blitzer *et al.* (1982)].

Selling a contract for oil exploration and development can help a small country spread the risk associated with the uncertain size of its reserves and the uncertain cost of developing them. But, in addition to sharing risks, the contract should provide the outside developer with incentives to minimize costs and to explore efficiently. For example, efficiency is promoted if a developer could pay a *fixed sum* in return for the rights to all the costs and benefits; in this contract, the developer is also accepting all of the risks.

A difficulty arises if the contract can be broken. Most multinational oil companies have reason to fear that they will be nationalized or heavily taxed if oil reserves or oil prices prove much higher than expected. This uncertainty reduces their incentives to finance the exploration and development costs. This effect is especially severe if oil companies depend on their few large successes to cover their more frequent small losses.

A valuable area of ongoing research examines optimal contract design in a constrained environment: how should energy development contracts make tradeoffs between risk-sharing and incentives if the contract may be broken? Blitzer *et al.* (1984) compare four commonly used contract formulas: service contracts, toll or fee per barrel contracts, production-sharing contracts, and royalty contracts. They emphasize that the contracts should attempt to spread the different risks in proportion to each party's comparative ability to accept these risks. Consider the production-sharing contract; the developer is rewarded with a fixed fraction of the benefits for assuming a fixed fraction of the costs. If the fraction of benefits equals the fraction of costs then the developer has the correct incentives for efficient exploration and production. However, developers often do not expect to receive benefits above some maximum level (due to taxes, nationalization, currency restrictions). In this case, efficient incentives are provided only if the developer's fraction of benefits is larger than his fraction of costs to compensate for his loss of the upper tail of benefits. This uneven sharing of costs and benefits results in the host country bearing more risk—a result of its inability to commit itself credibly to contractual agreements.

6. Conclusions

The expected present value (EPV) of domestic energy production or conservation programs does not reflect their true social welfare benefits. The EPV must be adjusted upwards to account for the reduced variability in national income, and improvements in income distribution. A reduced dependency on foreign imports may also create externalities which help reduce the world price of oil.

The appraisal of domestic non-renewable energy projects must take account of the same insurance benefits, but there is also a large adjustment in the opposite direction. Calculations must include the opportunity cost of depleting the exhaustible resource. A royalty value per barrel should be subtracted from the market price of oil. With extraction, we are interested in the *value added* by the project as opposed to the *value* of the project. By the same token, when calculating the true costs of stockpiling, the opportunity cost of capital invested in the stockpile will be lower since it must be adjusted to include expected appreciation in the value of the stockpile.

No correction for depletion is necessary for renewable energy resources. Given a long-run increase in the expected price of exhaustible resources under Hotelling's rule, renewable energy projects will eventually become economic. When production technology has low fixed costs and high marginal costs, this allows greater flexibility in responding to the variability of world prices; there is the option of shutting down temporarily and relying on imports when these are cheap. A low fixed-cost technology becomes more attractive with a higher discount rate; owing to capital shortages, the discount rate is likely to be higher in OIDCs. Other things equal, this will also imply a faster depletion rate of its exhaustible energy resources.

The development of domestic exhaustible energy resources is a way of borrowing from future generations without paying a risk premium. In situations where many OIDCs are overexposed to foreign borrowing, a reduction in the risk premium on loans can represent a considerable gain.

Even after all the corrections have been made that translate expected present value into a measure of social welfare, this does not guarantee that the appropriate social welfare criteria will be used in the decision-making process. Cost-benefit analysis cannot be applied in isolation of the incentives of those who have the decision-making power. To implement decisions according to cost-benefit criteria, the decision-makers must either be benevolent or they must be rewarded according to the gain in social welfare. Designing incentive schemes that motivate implementation of cost-benefit rules is an important area for future research.

Cost-benefit analysis is an essential part of development economics. Energy is an essential ingredient. In order to implement the optimum development of OIDCs' exhaustible and renewable energy resources, it is first necessary to specify and then to use appropriate project evaluation criteria/guidelines. The special features associated with energy projects in oil-importing developing countries require significant adjustments from expected present discounted value.

St. Catherine's College, Oxford, and Princeton University, USA.

REFERENCES

ANAND, SUDHIR (1973): "Distributional Weights in Project Analysis", mimeographed, Development Research Center, World Bank, Washington, D.C., July.

ANAND, SUDHIR and NALEBUFF, BARRY J. (1985): "Issues in the Appraisal of Energy Projects for Oil-Importing Developing Countries", World Bank Staff Working Paper Number 738, Washington, D.C., July.

ARROW, KENNETH J. and LIND, ROBERT C. (1970): "Uncertainty and the Evaluation of Public Investment Decisions", *American Economic Review,* Vol. 60 (June), pp. 364–378; reprinted in Richard Layard (ed.), *Cost-Benefit Analysis,* Penguin Modern Economics Readings, Penguin Books, England, 1972.

BLITZER, CHARLES R., LESSARD, DONALD R. and PADDOCK, JAMES L. (1982): "Risk Bearing and the Choice of Contract Forms for Oil Exploration and Development", Working Paper No. MIT-EL 82-053WP, International Energy Studies Program, Massachusetts Institute of Technology, Cambridge, September.

BLITZER, CHARLES R., CAVOULACOS, PANOS, LESSARD, DONALD R. and PADDOCK, JAMES L. (1984): "An Analysis of Financial Impediments to Oil and Gas Exploration in Developing Countries", Working Paper No. MIT-EL 84-005, Center for Energy Policy Research, Massachusetts Institute of Technology, Cambridge, January.

CRAWFORD, VINCE and SOBEL, JOEL (1982): "Strategic Considerations in the Stockpiling of Oil", mimeographed, University of California at San Diego.

DASGUPTA, PARTHA S. (1982): "Environmental Management Under Uncertainty", in Smith and Krutilla (eds.) (1982).

DASGUPTA, PARTHA S., MARGLIN, STEPHEN A. and SEN, AMARTYA K. (1972): *UNIDO Guidelines for Project Evaluation,* United Nations, New York.

DEESE, DAVID A. and NYE, JOSEPH S. (eds.) (1981): *Energy and Security,* Ballinger Publishing House, Cambridge, Massachusetts.

DEVARAJAN, SHANTAYANAN and FISHER, ANTHONY C. (1982): "Measures of Natural Resource Scarcity Under Uncertainty", in Smith and Krutilla (eds.) (1982).

GILBERT, RICHARD J. (1976): "Optimal Depletion of an Uncertain Stock", SEER Technical Report No. 10, University of California, Berkeley.

GILBERT, RICHARD J. and STIGLITZ, JOSEPH E. (1978): *Effects of Risk on Prices and Quantities of Energy Supplies,* Volumes I to IV, Report EA-700, EPRI, Palo Alto, California.

GILBERT, RICHARD J., NEWBERY, DAVID M. G., and STIGLITZ, JOSEPH E. (1978): *An Overview of the Economic Theory of Uncertainty and its Implications for Energy Supply,* EA-586-SR, EPRI, Palo Alto, California.

GRAY, DALE F. (1981): "Methods for the Economic Evaluation of Petroleum Exploration and Synthetic Fuels Production: An Application to Brazil", Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.

GUSTAFSON, R. L. (1958): "Carryover Levels for Grains", U.S. Department of Agriculture, Technical Bulletin No. 1178, Washington, D.C.

HOGAN, WILLIAM W. (1981): "Energy and Security Policy", Energy and Environmental Policy Center, Discussion Paper E-81-09, John F. Kennedy School of Government, Harvard University, Cambridge, Massachusetts, September.

HOGAN, WILLIAM W. (1982): "Policies for Oil Importers", Energy and Environmental Policy

Center, Discussion Paper E-82-02, John F. Kennedy School of Government, Harvard University, Cambridge, Massachusetts, February.

HOTELLING, HAROLD (1931): "The Economics of Exhaustible Resources", *Journal of Political Economy,* Vol. 39, pp. 137–75.

LIND, ROBERT C. (1982): "A Primer on the Major Issues Relating to the Discount Rate for Evaluating National Energy Options", in Lind *et al.* (1982).

LIND, ROBERT C. *et al.* (1982): *Discounting for Time and Risk in Energy Policy,* published by Resources for the Future, Inc., Washington, D.C., The Johns Hopkins University Press, Baltimore, Maryland.

LITTLE, IAN M. D. and MIRRLEES, JAMES A. (1968): *Manual of Industrial Project Analysis in Developing Countries, Vol. 2, Social Cost-Benefit Analysis,* Development Center, Organization for Economic Cooperation and Development, Paris.

LITTLE, IAN M. D. and MIRRLEES, JAMES A. (1974): *Project Appraisal and Planning for Developing Countries,* Basic Books, New York.

MANASSAH, JAMAL T. (ed.) (1981): *Alternative Energy Sources, Parts A & B,* Academic Press, New York.

MASON, SCOTT P. and MERTON, ROBERT C. (1984): "The Role of Contingent Claims Analysis in Corporate Finance", mimeographed, Sloan School of Management, M.I.T., Cambridge, Massachusetts, January.

NALEBUFF, BARRY J. and VARIAN, HAL (1983): "Some Aspects of Risk Sharing in Non-Classical Environments", in (ed.) Dale W. Jorgenson, *Proceedings of the Arne Lynde Symposium on Social Insurance,* North Holland.

NEWBERY, DAVID M. G. and STIGLITZ, JOSEPH E. (1982): "Optimal Commodity Stock-Piling Rules", mimeographed, Princeton University, New Jersey, February.

NICHOLS, ALBERT L. and ZECKHAUSER, RICHARD J. (1977): "Stockpiling Strategies and Cartel Prices", *The Bell Journal of Economics,* Vol. 8, No. 1, Spring.

PLUMMER, JAMES L. (ed.) (1982): *Energy Vulnerability,* Ballinger Press, Cambridge, Massachusetts, U.S.A.

RADNER, ROY and STIGLITZ, JOSEPH E. (1975): "Fundamental Non-Convexities in the Value of Information", mimeographed, Stanford University, California.

ROTHSCHILD, MICHAEL (1971): "A Two-Armed Bandit Theory of Market Pricing", Working Paper No. 10, University of California, Berkeley.

SAMUELSON, PAUL A. (1963): "Risk and Uncertainty: A Fallacy of Large Numbers", *Scientia,* Vol. 6, April/May.

SAMOUILIDIS, J. E. and BERAHAS, S. A. (1982): "A Methodological Approach to Strategic Petroleum Reserves", *Omega,* Vol. 10, No. 5, pp. 565–574.

SCHULZE, WILLIAM D. and KNEESE, ALLEN V. (1981): "Risk in Benefit-Cost Analysis", *Risk Analysis,* Vol. 1, No. 1, March, pp. 81–88.

SHARPE, W. F. (1964): "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk", *Journal of Finance,* Vol. 19, September, pp. 425–442.

SMITH, V. KERRY and KRUTILLA, JOHN V. (eds.) (1982): *Explorations in Natural Resource Economics,* published for Resources for the Future, Inc., The Johns Hopkins University Press, Baltimore, Maryland.

STERN, NICHOLAS H. (1977): "Welfare Weights and the Elasticity of the Marginal Valuation of Income", in M. Artis and A. R. Nobay, (eds.) *Current Economic Problems,* Basil Blackwell, Oxford.

STIGLITZ, JOSEPH E. (1976): "Monopoly and the Rate of Extraction of Exhaustible Resources", *American Economic Review,* Vol. 66, pp. 655–61.

STIGLITZ, JOSEPH E. and DASGUPTA, PARTHA S. (1980): "Market Structure and Resource Depletion: A Contribution to the Theory of Intertemporal Monopolistic Competition", Econometric Research Program Research Memorandum No. 261, Princeton University, Princeton, New Jersey, March.

WILSON, ROBERT (1982): "Risk Measurement of Public Projects", in Lind *et al.* (1982).

World Bank (1982): *Potential for Industrial Energy Conservation in Developing Countries,* by Edilberto Segura and staff of the Industry Department, Washington, D.C., October.

Wright, Brian D. and Williams, Jeffrey C. (1982): "The Roles of Public and Private Storage in Managing Oil Import Disruptions", mimeographed, Yale University and Brandeis University, March.